

# The Cloud Computing Cookbook: The Hypervisor Side

Michael Maclsaac  
IBM

March 12, 2012  
10174



## Abstract

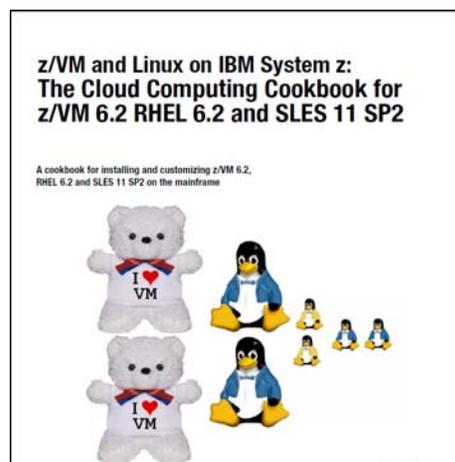
The "Virtualization Cookbook" for System z, usually in the form of a Redbook, has been a popular reference for many years. It has been updated for 2012 and renamed "The Cloud Computing Cookbook". This presentation will focus on the latest function provided in the z/VM hypervisor. New sections of the book, including enabling and using DirMaint and SMAPI, will also be addressed. Session, "The Cloud Computing Cookbook: The Linux Side", will focus on the Linux end of the cloud. A live demo will be included.

## Introductions

- Mike MacIsaac, mikemac@us.ibm.com
  - ▶ 25 years at IBM in NY
  - ▶ Programmer, z/OS USS, Redbook project lead
  - ▶ Marketing technical support of z/VM, Linux, IBM software, ...
  - ▶ z/VM development manager
  - ▶ Lab-based Services (New)
- Who are you?
  - ▶ Experience with z/VM and Linux:
    - z/VM?
    - Linux?
    - Other?
    - None of the abovehis book?
  - ▶ IT status:
    - Do you have Linux and z/VM in production?
    - In test?
    - Planning a proof of concept?
- Something you are hoping to get out of the next three days?

## Overview

- The **Virtualization Cookbooks** and now the **Cloud Computing Cookbook** have always had the same goal in mind: to be a single source for installing and customizing z/VM, installing and customizing Linux, and getting to the point of cloning and making appliances of Linux virtual servers. Over the years, commonly used **Miscellaneous Recipes** have also been documented.



See: <http://www.vm.ibm.com/devpages/mikemac/>

L P A R 1	LPARs BVM1 and BVM2: z/VM 6.2 on z114	L P A R n
	IDENTITY MAINT: z/VM system administration	
	IDENTITY TCPMAINT: TCP/IP administration	
	IDENTITY TCPIP: TCP/IP stack	
	IDENTITY AUTOLOG1: z/VM configuration at IPL	
	IDENTITY DTCVSW1/DTCVSW2: VSWITCH controllers	
	USER LNXMAINT: CMS files common to Linux systems	
	IDENTITY LNXADMIN: Linux system administration	
	USER RH62GOLD - RHEL 6.2 golden image	
	USER S112GOLD: SLES 11 SP2 golden image	
	USER LINUX153: Linux virtual server 1	
	USER LINUX157: Linux virtual server 2	

## Overview of entire system

### Resources:

CPU: 2 IFLs, shared  
 Memory: 6GB/2GB or more  
 Disk: 42 3390-3 DASD  
 Network: 16 OSA-E addresses  
 TCP/IP 6 TCP/IP addresses

OSA Express

OSA Express

PC Linux  
NFS serverDesktop  
machine

© 2011 IBM Corporation

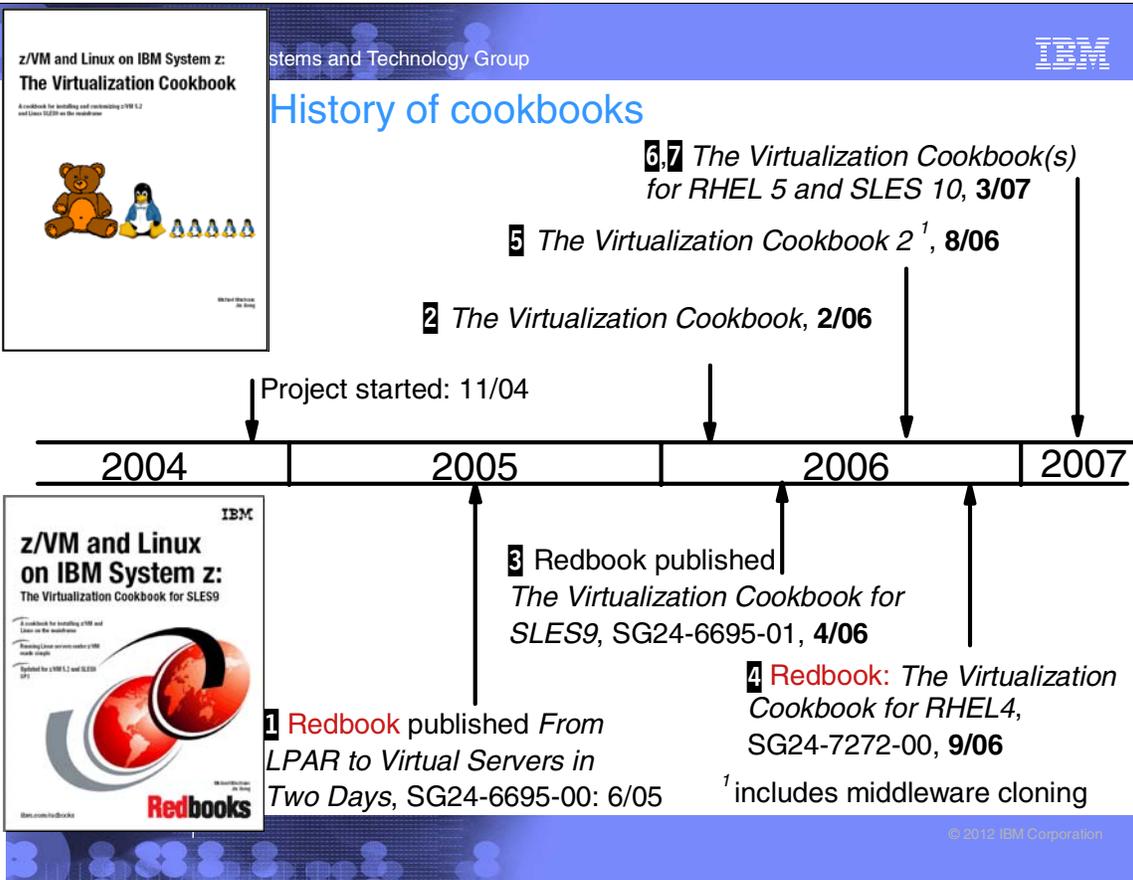
## Outline of current book

1. **Introduction and z/VM** - introduces z/VM 6.2, discusses planning, then installation and configuration into a two member SSI with z/VM 6.2.
2. RHEL 6.2 Linux - install, customizing and clone Red Hat Enterprise Linux (RHEL)
3. SLES 11 SP2 Linux - install, customizing and clone SuSE Linux Enterprise Server (SLES)
4. **Other topics** - includes chapters on:
  - a. **Live Guest Relocation (LGR) between SSI members**
  - b. **Configuring DirMaint and SMAPI**
  - c. Monitoring z/VM and Linux
  - d. Miscellaneous "recipes"
  - e. xCAT - the eXtreme Cloud Administration Toolkit
5. Appendices - includes references, cheat sheets and lists the source code

© 2012 IBM Corporation

## Introduction and z/VM

1. History of cookbooks (Preface)
2. Introduction to z/VM and Linux (Chapter 1)
3. Planning (Chapter 2)
4. Configure a desktop machine (Chapter 3)
5. Configure an NFS/FTP server (Chapter 4)
6. Install a z/VM SSI cluster (Chapter 5)
7. Service z/VM (Chapter 6)



IBM Systems and Technology Group

## History of books (cont'd)

See: <http://www.vm.ibm.com/devpages/mikemac/>

**12** *The Cloud Computing Cookbook for z/VM 6.2, RHEL 6.2 and SLES 11 SP2, 1/12*

**9** *The Virtualization Cookbook for SLES 11, 2/10*

**8** *Redbook: The Virtualization Cookbook for SLES 10 SP2<sup>2</sup>, 10/08*

**10** *Redbook: The Virtualization Cookbook for SLES 11 SP1, 1/11*

**11** *Redbook: The Virtualization Cookbook for RHEL 6, 2/11*

See: <http://www.redbooks.ibm.com/> <sup>2</sup> includes travelling /home

© 2012 IBM Corporation

IBM Systems and Technology Group

## Changes in the Jan 1, 2012 book

- *z/VM and Linux on IBM System z: The Cloud Computing Cookbook for z/VM 6.2 RHEL 6.2 and SLES 11 SP2* has many new sections:
  - ▶ z/VM sections are updated for 6.2 with a two member SSI setup
  - ▶ Linux sections are updated for both RHEL 6.2 and SLES 11 SP2, combined in one book
  - ▶ NFS-exported files are stored in /var/nfs/ rather than /nfs/ in keeping with Linux FHS
  - ▶ Use of both layer 2 and layer 3 virtual switches
  - ▶ VSWITCH authorization granted through COMMAND statements in user directory profile
  - ▶ Section on relabelling z/VM system volumes removed
  - ▶ New chapter (17) on Live Guest Relocation (LGR) between SSI members
  - ▶ New chapter (18) on how to install and configure z/VM's DirMaint and SMAPI
  - ▶ New chapter (21) on how to install and configure xCAT
  - ▶ New section (19.4) on how to install and configure sysstat on Linux
  - ▶ Title is buzzword compliant :))

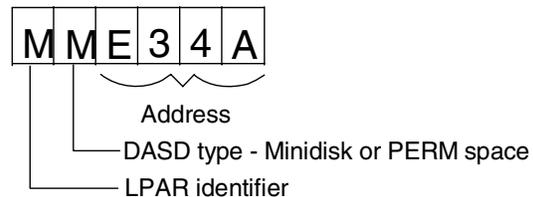
© 2012 IBM Corporation

## Introduction - Planning - bill of materials

- Hardware
  - ▶ System z LPARs (2 or 4 for SSI)
    - IFLs
    - Memory (aka *storage*)
    - DASD (aka *storage* :))
    - Two OSA cards for HA VSWITCH (One is OK)
  - ▶ Temporary Distributed server
- Software
  - ▶ z/VM 6.2
  - ▶ Linux
    - SLES-11 SP2
    - RHEL 6.2
  - ▶ Code associated with book: <http://www.vm.ibm.com/devpages/mikemac/CKB-VM62.tgz>
- Networking resources
  - ▶ TCP/IP address for z/VM
  - ▶ One TCP/IP address for each Linux
  - ▶ DNS names

## Introduction - Planning (cont'd)

- Conventions
  - ▶ Volume labeling convention
    - Volume labels are only 6 chars
    - Using device address in last 4 chars:
      - Guarantees unique labels
      - First character is LPAR identifier
      - Second character is function (P=page, S=spool, M=minidisk)
  - ▶ File naming convention
    - File that is shipped with VM/Linux - ORIG or .orig suffix
    - File that was last working - WRKS or .works
- Password convention - z/VM admin, Linux admin, Linux users
  - ▶ Worksheets - 2 sets of 4 worksheets
  - ▶ Populated set of worksheets for examples used in the book
  - ▶ Blank set of worksheets for
    - z/VM resources
    - Linux resources
    - z/VM DASD
    - Linux virtual machines



## Introduction - Configure a desktop machine

- SSH client
  - ▶ PuTTY is described
    - Set SSH protocol to "2 only"
    - Add rows, columns, scrollback buffer
    - Save sessions
- VNC client
  - ▶ Recommended for install of Linux, some software
  - ▶ RealVNC is described
- 3270 emulator
  - ▶ Set Enter and Clear key if possible
  - ▶ Set to use 43 lines
  - ▶ Set to Reconnect after logoff
  - ▶ For Linux, x3270 is most popular

## Introduction - Configure a PC server

- Installing Linux on zSeries is a chicken and egg problem
- Recommendation: install Linux on an Intel box as a temporary NFS server:
  - ▶ Install Linux onto a PC
  - ▶ Copy files associated with this book to this NFS server
  - ▶ Untar to `/var/nfs/CKB-VM62/`
  - ▶ Set up an install directory under `/var/nfs/<distro>/`
  - ▶ Configure the NFS server to export these two directories

## Installing and configuring z/VM

- Obtain z/VM through electronic download
- Configure an FTP server for z/VM installation
- Install z/VM from DVD or FTP server
- Customize TCPIP - z/VM stack, FTP server
- Customize SYSTEM CONFIG
  - ▶ Define VSWITCHes, other configuration
- Add volumes for paging and minidisks
  - ▶ CPFORMAT EXEC is included
- Create LNXMAINT for common CMS files- kernels, RAMdisks, PARMfiles, etc.
- Customize system startup and shutdown
  - ▶ SHUTDOWN z/VM signals Linux servers to shutdown
  - ▶ IPL of z/VM autologs (boots) important Linux servers
- z/VM security issues

## Obtain z/VM through Electronic Download

- Go to the z/VM service page:  
<http://www.vm.ibm.com/service/>
- Click on the link **IBM Shopz** in the section *IBM Support Portals*
  - ▶ Sign in by clicking on the link *Sign in for registered users* in the upper right
  - ▶ Click on the link **create new software orders**
  - ▶ On *Step 1*, click on the radio button **z/VM Products** and choose **VM SDO version 6** in the dropdown menu to the right. Click **Continue**.
  - ▶ On *Step 2*, select a hardware system on which you plan to run z/VM
  - ▶ On *Step 3*, first filter, select **VM - VM Base Product**, second filter, select **Show all products** then click **Show catalog**
  - ▶ Select **z/VM V6 3390 System DDR** and click **Continue**
  - ▶ On *Step 4*, verify the order and click **Continue**
  - ▶ On *Step 5*, verify the entitlements and click **Continue**
  - ▶ On *Step 6*, for the *Preferred media*, select **Internet** and click **Continue**
  - ▶ On *Step 7*, review and click **Submit**

## Configure an FTP server for z/VM installation

- Prepare the z/VM product install files
- Install the FTP server
- Configure the FTP server
  - ▶ Anonymous or not?
- Test the anonymous FTP server

## Install z/VM from DVD or FTP server

- Start the z/VM install
  - ▶ Important screens (below and next chart)
- Copy a vanilla z/VM system to DASD
- IPL the first SSI member
  - ▶ New IPL Params:
    - ==> **q iplparms**
- IPL remaining SSI members
- Verify the installation
- Configure TCP/IP

FN=SYSTEM      FT=CONFIG      PDNUM=1      PDVOL=D964

```

MrMutt: Integrated 3270 Console for MR28SE:BVM1
File Keys Font Help
*** z/VM INSTALLATION PLANNING PANEL 2 ***

Would you like to have your system automatically configured to be
managed by the Unified Resource Manager or some other SMAPI client
for system management? (Y/N)

Keep The Following in Mind:

If you say YES, you should not attempt to manage your system in
any other way.

If you'd like to manage your own system, or use a purchased
external security manager or a purchased directory manager say NO
  
```

### z/VM install screens (cont'd)

MrMutt: Integrated 3270 Console for M28:BVM1

File Keys Font Help

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

SSI Cluster Name: POKSSI

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	POKDEV62	BVM1
2	POKTST62	BVM2

MrMutt: Integrated 3270 Console for MR28SE:BVM1

File Keys Font Help

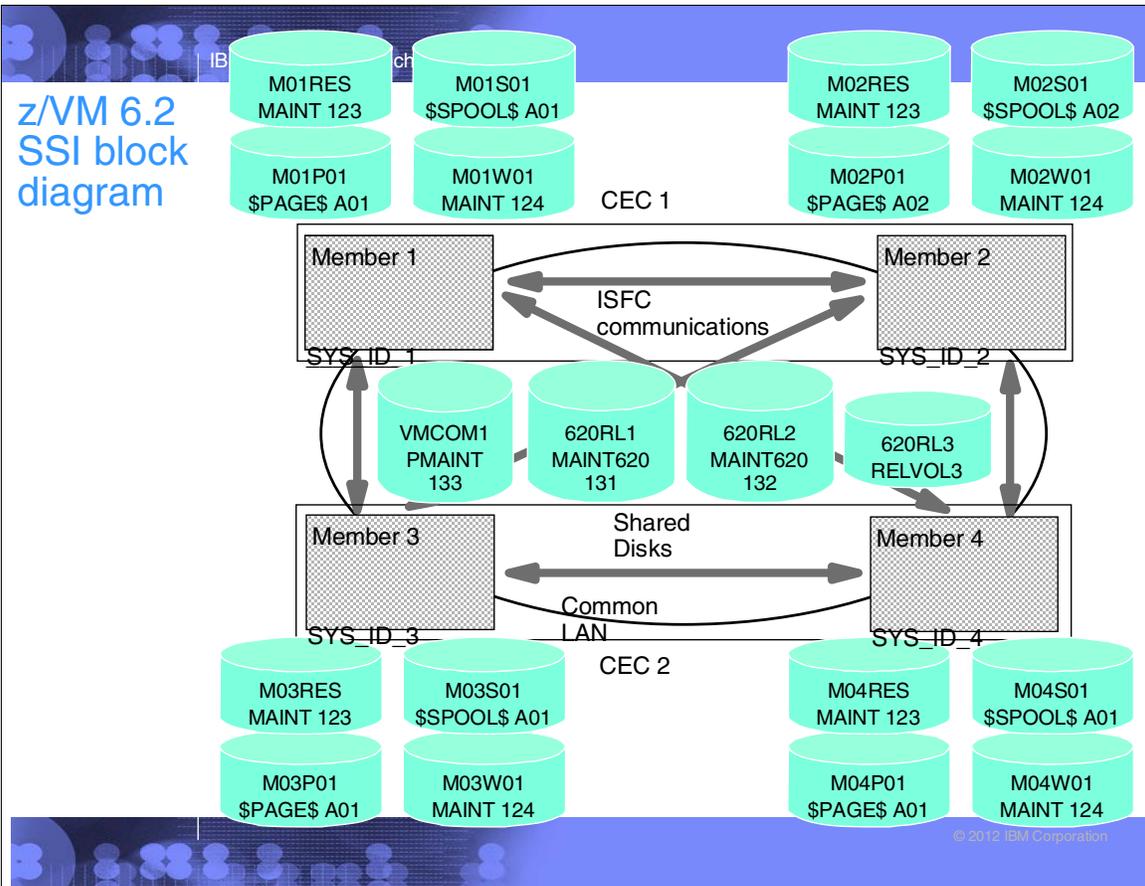
\*\*\* z/VM INSTALLATION VOLUME DEFINITION \*\*\*

TYPE	LABEL	ADDRESS	FORMAT (Y/N)
COMMON	JV6284	6284	Y
RELVOL	CV6285	6285	
RELVOL2	CV6286	6286	
RELVOL3	CV6287	6287	

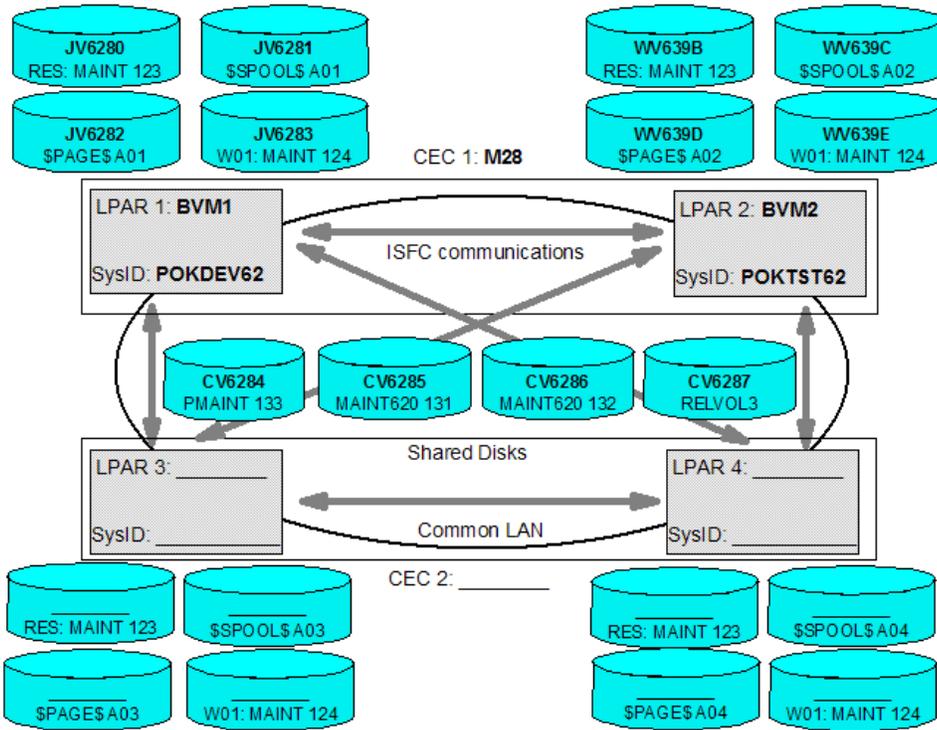
  

TYPE	LABEL	ADDRESS	TYPE	LABEL	ADDRESS
POKDEV62			POKTST62		
RES	JV6280	6280	RES	WV639B	639B
SPOOL	JS6281	6281	SPOOL	WS639C	639C
PAGE	JP6282	6282	PAGE	WP639D	639D
WORK	JV6283	6283	WORK	WV639E	639E

© 2012 IBM Corporation



## SSI block diagram - values used in book



## Customize z/VM TCP/IP stack and FTP server

- Recommend IPWIZARD for TCP/IP configuration
  - ▶ Run once for each SSI member
  - ▶ Configure XEDIT profile on TCPMAINT
- Recommend turning on z/VM FTP server
  - ▶ Run once for each SSI member

## Customize SYSTEM CONFIG file

### ■ Recommendations

- ▶ Increase retrieve key capacity - from 20 to 99
- ▶ Allow VDISKS to be created for swap spaces
  - Using SWAPGEN EXEC is common to create in-memory Linux swap spaces
- ▶ Turn off the Disconnect Timeout feature
  - So Linux virtual machines are not forced off by SYSTEM
- ▶ Define layer 2 and 3 virtual switches
  - Layer 2 is becoming more common - for DHCP
- ▶ Set up "Equivalency IDs" - new for z/VM 6.2

```
/* Add EQID statements for OSA addresses and unique MAC IDs */
POKDEV62: begin
  rdev 4200-420f eqid osaset1 type osa
  rdev 4300-430f eqid osaset1 type osa
  vmlan macprefix 02000b
POKDEV62: end
POKTST62: begin
  rdev 4200-420f eqid osaset1 type osa
  rdev 4300-430f eqid osaset1 type osa
  vmlan macprefix 02000c
POKTST62: end
```

## CPFORMAT EXEC

==> **cpformat**

Synopsis:

Format and label DASD as page, perm, spool or temp disk space  
 The label written to each DASD is W<t><xxxx> where:  
 <t> is type - P (page), M (perm), S (spool) or T (Temp disk)  
 <xxxx> is the 4 digit address

Syntax is:

```
<-----<
>--CPFORMAT--vdev-----AS--PERM----->
          '-vdev1-vdev2-'          '-PAGE-'
                                   '-SPOL-'
                                   '-TEMP-'
```

Example:

```
==> att <a775-a779> *
A775-A779 ATTACHED TO MAINT
==> cpformat <a775-a779> as page
...
```

## Add volumes for paging and minidisks

- Copy the CPFORMAT EXEC
- Format volumes for page space
  - ▶ Use the CPFORMAT EXEC with "for page"
- Format DASD for minidisks
  - ▶ Use the CPFORMAT EXEC with "for perm"
- Update the SYSTEM CONFIG file. e.g.:

```
POKDEV62: BEGIN
  CP_Owned Slot 251 JP628A
  CP_Owned Slot 252 JP6288
  CP_Owned Slot 253 JP6233
  CP_Owned Slot 254 JP6232
  CP_Owned Slot 255 JV6282
POKDEV62: END

POKTST62: BEGIN
  CP_Owned Slot 251 WP633E
  CP_Owned Slot 252 WP633C
  CP_Owned Slot 253 WP633B
  CP_Owned Slot 254 WP628B
  CP_Owned Slot 255 WV639D
POKTST62: END
...
User_Volume_List CV6285 CV6286 CV6287
User_Volume_Include JM6*
```

## Create LNXMAINT for common CMS files

- Define virtual machine
- Customize virtual machine
- Copy files
  - ▶ 191 disk: PROFILE EXEC, PROFILE XEDIT
  - ▶ 192 disk: Common Linux files
    - PROFILE EXEC
    - PROFILE XEDIT
    - SAMPLE CONF-RH6
    - SAMPLE PARM-S11
    - SWAPGEN EXEC
    - RHEL62 EXEC
    - SAMPLE PARM-RH6
    - SLES11S2 EXEC
    - <Linux> RAMDISK
    - <Linux> KERNEL

## Customizing z/VM startup and shutdown

- Add a minidisk link to AUTOLOG1 user directory entry
- Call a startup EXEC common to all SSI members
 

```
/* Common code to be run at SSI IPL time */
"CP XAUTOLOG TCPIP" /* Autolog TCPIP */
"CP SET MDC STOR 0M 128M" /* Limit minidisk cache in CSTOR */
"CP SET MDC XSTORE 0M 0M" /* Disable minidisk cache in XSTORE */
"CP SET SIGNAL SHUTDOWN 600" /* Allow guests 10 min to shut down */
```
- Start Linux virtual machines on appropriate SSI members
 

```
/* Start Linux systems on SSI member 1 */
"CP XAUTOLOG LINUX01"
"CP XAUTOLOG LINUX02"
```
- Test a SHUTDOWN REIPL

## SSISHUTD and SSICMD EXECs

==> **ssishutd help**

Synopsis:

SHUTDOWN or SHUTDOWN REIPL an SSI cluster

Syntax is:

```
>>--SSISHUTD-----,-----><
          '--REIPL--'
```

==> **ssicmd**

Synopsis:

SSICMD cmd

cmd is a command to be issued on each of the members  
in the SSI cluster using the AT command.

Example:

==> **ssicmd q proc**

POKDEV62:

```
PROCESSOR 00 MASTER CP
PROCESSOR 01 ALTERNATE CP
```

POKTST62:

```
PROCESSOR 00 MASTER CP
PROCESSOR 01 ALTERNATE CP
```

## z/VM security issues

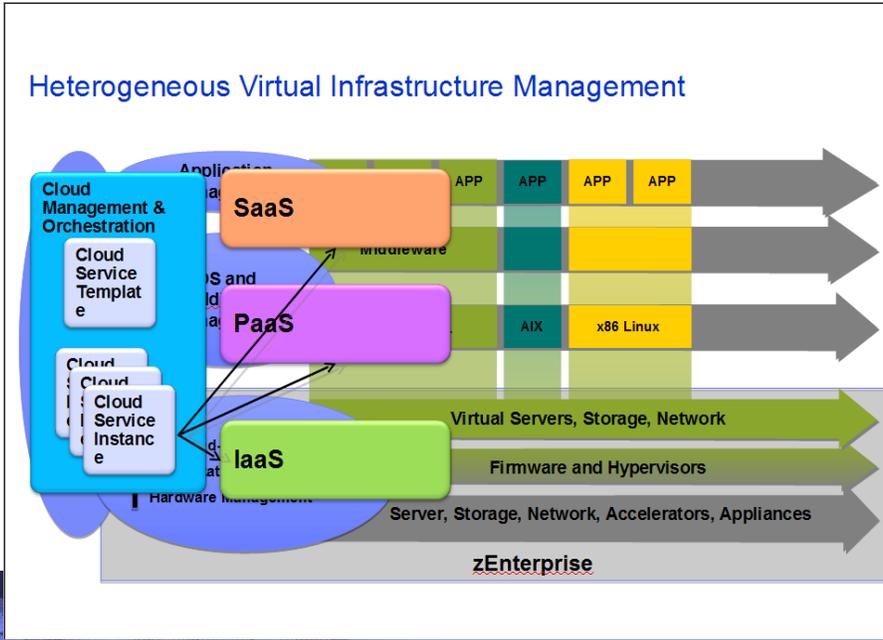
- Change passwords in USER DIRECT
- Use a z/VM Security product?
  - ▶ IBM RACF
  - ▶ CA VM:Secure
- The paper *z/VM Security and Integrity*
  - ▶ <http://www.vm.ibm.com/library/zvmsecint.pdf>

## Servicing z/VM

- Apply a Programming Temporary Fix (PTF)
  - ▶ Get service from Internet
  - ▶ Receive, apply and build
  - ▶ Put into production
- Apply a Recommended Service Upgrade (RSU)
- Determining z/VM's service level

## Tangent - <X>aaS

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS) - aka "container"



## Tangent - virtualization terminology

- User ID
- Virtual machine
- Guest
- Container

## z/VM Live Guest Relocation (LGR)

- LGR considerations
  - ▶ USERS are relocatable, not IDENTITYs
  - ▶ Memory size (central, expanded)
  - ▶ Link and resource contention
  - ▶ Add `OPTION CHPIDV ONE` to the Linux PROFILE in user directory
  - ▶ Linux must not have CMS disks at relocate time
    - Disks can be detached at Linux boot time
    - Added to `/etc/rc.d/rc.local`:
 

```
...
chshut halt vmcmd logoff
chshut poff vmcmd logoff
modprobe vmcp
vmcp det 190
vmcp det 191
vmcp det 19d
vmcp det 19e
rmmod vmcp
```
- Relocate a Linux system
  - ==> `vmrelocate test <user ID> <target system ID>`
  - ==> `vmrelocate move <user ID> <target system ID>`
  - ▶ Demo?

## Configure DirMaint and SMAPI

- Configure DirMaint
  - ▶ Enable DirMaint
  - ▶ Tailor DirMaint
  - ▶ Customize the EXTENT CONTROL file
  - ▶ Start DirMaint
  - ▶ Test DirMaint
  - ▶ Test DirMaint at IPL time
- Configure SMAPI
  - ▶ Set up basic SMAPI configuration
  - ▶ Turn off ensembles
  - ▶ Start SMAPI at IPL time
  - ▶ Test SMAPI
- Some common DirMaint tasks
  - ▶ Update a user directory entry
  - ▶ Edit the EXTENT CONTROL file
  - ▶ Get a copy of the user directory
  - ▶ Add an IDENTITY

# DASD view of the system



# Live Demo



## Resources

- All *Virtualization Cookbooks* and other papers:
  - ▶ <http://www.vm.ibm.com/devpages/mikemac/>
- *The Linux for zSeries and S/390 portal*
  - ▶ <http://linuxvm.org/>
- The linux-390 list server
  - ▶ <http://www2.marist.edu/htbin/wlvindex?linux-390>
- Linux for zSeries and S/390 developerWorks®
  - ▶ <http://awlinux1.alphaworks.ibm.com/developerworks/linux390/index.shtml>
- Red Hat Enterprise Linux evaluation
  - ▶ <http://www.redhat.com/rhel/server/mainframe/>
- SUSE LINUX Enterprise Server evaluation
  - ▶ <http://www.novell.com/products/linuxenterpriseserver/eval.html>
- z/VM publications
  - ▶ <http://www.vm.ibm.com/pubs/>
- z/VM performance tips
  - ▶ <http://www.vm.ibm.com/perf/tips/>

## Questions

- Are there any questions?