From LPAR to Virtual Servers in Two Days: Day 1 - z/VM

Michael MacIsaac - mikemac@us.ibm.com
Thursday August 25rd, 1:30 PM
Session 9216

Abstract

There is a new redbook with a similar title: *z/VM and Linux on zSeries: From LPAR to Virtual Servers in Two Days*. As the title suggests a goal of the redbook is to allow you to install and configure z/VM, install and configure Linux and be cloning Linux in 2 working days. Rexx EXECs and scripts are provided with the book to make this aggressive goal reality. This talk focus on the first day which includes planning and desktop tools, and then gets into a z/VM 5.1 install from DVD. Once installed, the following topics are addressed:

* The SYSTEM CONFIG file
* TCP/IP configuration
* Paging volumes are configured
* A user ID for common files is created
* A highly available VSWITCH is created
* System startup and shutdown is addressed
* Security issues and backup are touched upon

While z/VM is being installed, the task of setting up a PC Linux NFS server is addressed. This allows you to supply the EXECs and scripts associated with the book and a Linux install tree. One additional topic of monitoring z/VM is briefly addressed.

Who am I?, who are you?

- Mike MacIsaac, mikemac@us.ibm.com
  - 18 years at IBM in Kingston and Poughkeepsie, NY
  - VM+Fortran+REXX
  - RS/6000, C, ksh
  - SP/2, C++
  - OS/390, USS, redbooks
  - z/VM + Linux + bash
  - Wrote a lot of *z/VM and Linux on zSeries: From LPAR to Virtual Servers in Two Days*
    - http://www.redbooks.ibm.com/abstracts/sg246695.html
- Who are you?
  - Have you tried the steps in this redbook?
  - Are you thinking about using this redbook?
  - Have you never heard about this redbook?
  - ...
  - Any other categories of attendees?
  - Any VM gurus?

Overview

- This presentation is based on a redbook and associated files
  - ***z/VM and Linux on zSeries: From LPAR to Virtual Servers in Two Days***
    - http://www.redbooks.ibm.com/abstracts/sg246695.html
  - Authors: Michael MacIsaac, Jin Xiong and Curtis Gearhart
  - Written: Nov 1 2004 - March 31 2005
- Overall goal: approach the concept of *Linux appliances*
  - Previous titles (rejected by IBM reviewers :)
    - The *zSeries Linux Appliance Cookbook*
    - *From LPAR to Virtual **Blades** in Two Days*
- Tangible goal: Install and configure z/VM, then install, configure and clone Linux ***in 2 working days***
- 4 usability tests were conducted - 3 achieved the goal
  - The other guy was too busy to work 2 days straight :))

Overview (cont'd)

- Philosophy:
  - "*Everything should be made as simple as possible, but not simpler*" -Albert Einstein
- Choices made in keeping with this philosophy:
  - "Roll your own" cloning rather than other products
    - Aduva's Onstage, IBM TPM, Levanta are all valid and more sophisticated, complex solutions
  - USER DIRECT file over directory maintenance products
  - z/VM user ID must be predefined in order to clone
  - ECKD DASD - no SCSI/FCP disks
  - Read/write /usr/ file system over shared read/only
    - Wanted to document a DCSS with XIP2 file systems
      - Ran out of time
  - Cloning and manual install hinge on CMS parameter files
    - Cloning done from Linux, no VM service machine needed

Outline for the two presentations

- **Planning**
- **Configuring a Windows desktop**
- **Installing and configuring z/VM**
- **Configuring an NFS server**

- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- **Monitoring z/VM and Linux**
- Backup and restore
- **Appendix - Relabel z/VM system volumes**

Planning

- Bill of materials
  - Hardware
    - zSeries LPAR - 2 IFLs recommended
      - 3GB central:1GB expanded storage (1.5G:512M OK)
      - 24 3390-3 DASD or more (ask for 32 :))
      - Two OSA cards for HA VSWITCH (One is OK)
    - Temporary Linux PC for NFS server (or equivalent)
  - Software
    - z/VM 5.1 on DVD (tape is OK)
    - Linux distro ISO images (CDs can be converted)
    - Code associated with redbook
      - ftp://www.redbooks.ibm.com/redbooks/SG246695/
  - Networking resources
    - TCP/IP address for z/VM
    - One TCP/IP address for each Linux (ask for 16 or more :))
    - DNS names (helpful but not required)

Planning (cont'd)

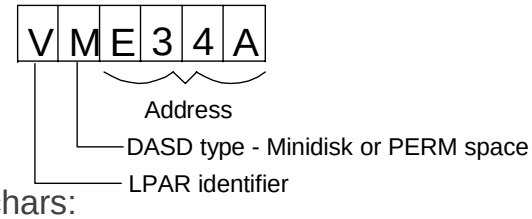

- Conventions
  - Volume labeling convention
    - Volume labels are only 6 chars
    - Using device address in last 4 chars:
      - Guarantees unique labels
      - Is convenient because of the z/VM label<=>address disconnect
      - First character is LPAR identifier
      - Second character is function (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 (1) z/VM resources, (2) Linux resources, (3) z/VM DASD, (4) Linux user IDs

Outline

- ~~Planning~~
- **Configuring a Windows desktop**
- Installing and configuring z/VM
- Configuring an NFS server
- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- Monitoring z/VM and Linux
- Backup and restore
- Appendix - Relabel z/VM system volumes

User interface

Configuring a Windows desktop

- SSH client
  - PuTTY is described
    - Set SSH protocol to "2 only"
    - Add rows and columns
    - Save sessions
  - For Linux, x3270 is most popular
- 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
- VNC client or X server
  - Not described, but RealVNC client seems to be most popular
  - For Linux desktop - you will have an X server

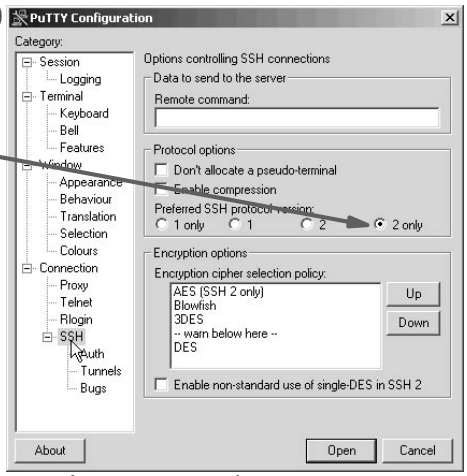

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- **Installing and configuring z/VM**
- Configuring an NFS server
- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- Monitoring z/VM and Linux
- Backup and restore
- Appendix - Relabel z/VM system volumes

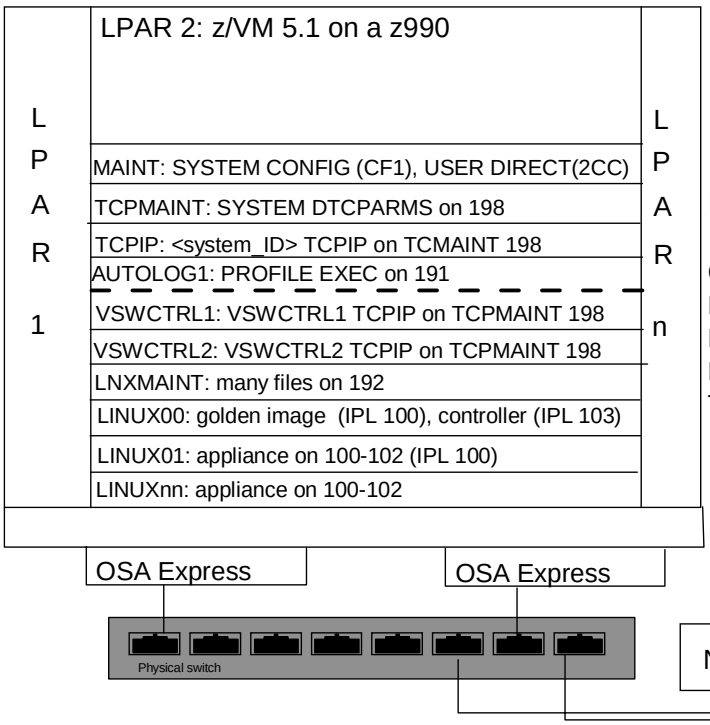

Block Diagram

Resources:

CPU: 2 IFLs, shared
Memory: 3GB/1GB
Disk: 32 3390-3 DASD
Network: 16 OSA-E addresses
TCP/IP 13 TCP/IP addresses

Install and configure z/VM

- Install z/VM from DVD
  - Install from DVD is documented in some detail
  - Use the *Integrated 3270 console* on HMC
  - Bug in INSTDVD EXEC is fixed, not worked around
- Configure z/VM - Overview
  - Customize SYSTEM CONFIG
  - Customize TCPIP with IPWIZARD
  - Add paging volumes
  - Create a common CMS user ID - LNXMAINT
  - Create a highly available VSWITCH
  - Customize startup and shutdown
  - Address security issues
- Backup system to tape
  - Modify system labels
  - Restore system

Configure z/VM

- Customize SYSTEM CONFIG
  - Change system name
  - Allow users to create VDISKs
  - Define a highly available VSWITCH named VSW1
  - Run CPSYNTAX
- Customize TCPIP with IPWIZARD
  - Use this tool ONCE - after that, edit files
  - Configure TCP/IP to start in AUTOLOG1's PROFILE EXEC
  - Rename PROFILE TCPIP to <system_ID> TCPIP
  - Configure FTP server
  - Shutdown and ReIPL
- Add paging volumes - 5 new volumes are recommended
  - Format the volumes - CPFORMAT EXEC is written
    - Wrapper around CPFMTXA
    - Tries to alternate between FLASHCOPY and CPFMTXA
  - Update the SYSTEM CONFIG file and ReIPL

CPFORMAT and CPLABEL EXECs - details

```
==> cpformat
Synopsis:
  Format one or a range of DASD as page, perm, spool or temp disk space
  The label written to each DASD is V<t><xxxx> where:
    <t> is type - P (page), M (perm), S (spool) or T (Temp disk)
    <xxxx> is the 4 digit address
Syntax is:
                                              .-PAGE-.
   >>--CPFORMAT--.-rdev--------------.--AS---+-PERM-+--------><
                 | <---------------< |       '-SPOL-'
                 '-rdev1-rdev2-------'

==> cplabel
Synopsis:
  LABEL and ALLOCATE DASD as page, perm, spool or temp disk space
  The label written to each DASD is V<t><xxxx> where:
    <t> is type - P (page), M (perm), S (spool)
    <xxxx> is the 4 digit address
Syntax is:
                                              .-PAGE-.
   >>---CPLABEL--.-rdev--------------.--AS---+-PERM-+--------><
                 | <---------------< |       '-SPOL-'
```

CPFORMAT EXEC - details

```
firstChar = 'V'
parse upper arg dasds "AS " type
if ((dasds = '') | (dasds = '?')) then call help
labelPrefix = getLabelPrefix(firstChar type)
numDasd = parseDasd(dasds)
answer = areYouSure(type)
if (answer = 'Y') then  /* the user is sure */
do
  retVal = doFormat(labelPrefix numDasd type)
  call doReport
end
else
  retVal = 2
exit retVal
doFormat: procedure expose dasdList.
  arg labelPrefix numDasd type
  'CP TERM MORE 1 1'
  do i = 1 to numDasd
    retVal = 1
    label = getLabel(labelPrefix dasdList.i)
    if (i // 2 = 0) then
      retVal = tryFLASHCOPY(dasdList.1 dasdList.i label)
    if (retVal <> 0) then
      call formatOne(dasdList.i type label)
  end /* do i = */
  'CP TERM MORE 50 10'
return 0 /* from doFormat */
```


Configure z/VM (cont'd)

- Create a common CMS user ID - LNXMAINT
  - LNXMAINT 192 is a minidisk for common CMS files
  - Define user ID in USER DIRECT
  - Define dummy minidisks of cylinder 0 in USER $ALLOC$
  - Logon and copy files associated with the book
- Create a highly available VSWITCH
  - Define two VSWITCH controller user IDs
  - Configure files on TCPMAINT 198 disk
  - Shutdown and ReIPL
- To speed running the DISKMAP command DISKMAPP is written
  - Simply an EXEC/XEDIT pair that:
    - Runs DISKMAP
    - Edits USER DISKMAP with subcommand "ALL /gap/|/overlap/"
  - You should always see only one gap of 501 cylinders

Configure z/VM (cont'd)

- Customize startup and shutdown - you have a choice
  - Configure AUTOLOG1's PROFILE EXEC
    - **XAUTOLOG** TCPIP and two VSWITCH controllers
    - Disable minidisk cache in XSTOR - **set mdc xstore 0m 0m**
    - Overcommit memory - **set srm storbuf 300% 250% 200%**
    - Send signals to Linux to shutdown - **set signal shutdown 180**
    - Grant each Linux ID access to the VSWITCH, e.g.:
      - **set vswitch vsw1 grant linux00**
    - **XAUTOLOG** each Linux ID (if desired)
  - Sine Nomine's SYSVINIT package is also documented
- Security issues, relabel system labels, back up system
  - CHANGEPW XEDIT is written to change USER DIRECT passwords
  - Backup system to tape
  - Modify 510xxx labels? (more later)
  - Restore system from tape

End of 1 day - DASD view

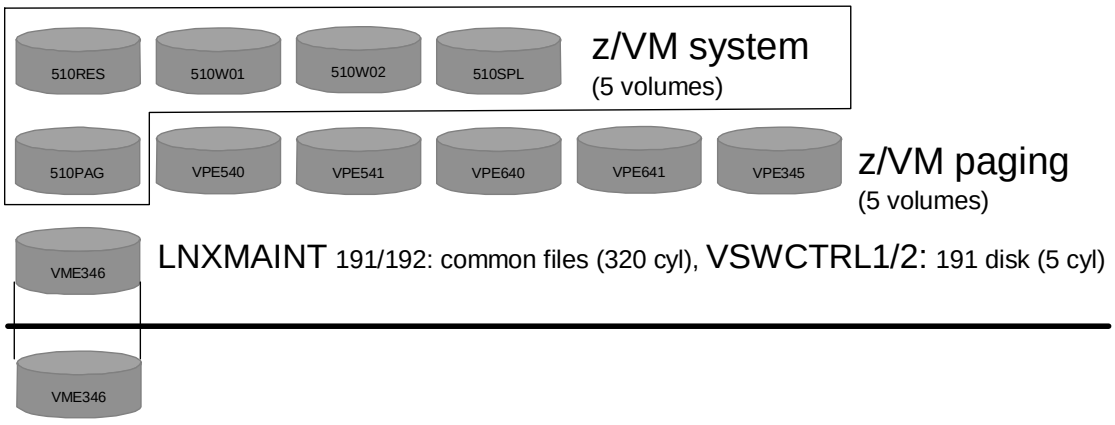

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- ~~Installing and configuring z/VM~~
- **Configuring an NFS server**
- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- Monitoring z/VM and Linux
- Backup and restore
- Appendix - Relabel z/VM system volumes

Configure an NFS server

- Install Linux on an Intel PC - same distribution recommended
- This server is temporary - replace contents to controller
- Copy files associated with this book to this new server
  - ftp://www.redbooks.ibm.com/redbooks/SG246695/
    - lpar2vs-small.tgz ~1MB - all associated files
    - lpar2vs.tgz ~50MB - includes z/VM and Linux manuals
- Create a SLES9 install tree
  - mksles9root script - looks for 6 SLES9 ISO and 2 SP1 ISO images

```
# ./mksles9root s390
Making a 31-bit SLES9 tree ...
Making the directory structure ...
Copying SLES9 ISO images ...
  Mounting and copying SLES-9-s390-RC5a-CD1.iso ...
  ...
  Mounting and copying SLES-9-s390-RC5-CD6.iso ...
Copying SLES9 SP1 ISO images ...
  Mounting and copying SLES-9-SP-1-s390-RC5-CD1.iso ...
  Mounting and copying SLES-9-SP-1-s390-RC5-CD2.iso ...
Removing temporary mount point ...
Making symbolic links ...
Creating yast/*order files ...
```

Configure an NFS server

- Creates a SLES9 install tree (cont'd)

```
# cd /nfs/sles9root
# ls -l
lrwxrwxrwx root root  boot -> SUSE-SLES-Version-9/CD1/boot/
lrwxrwxrwx root root  content -> SUSE-SLES-Version-9/CD1/content
lrwxrwxrwx root root  control.xml-> SUSE-SLES-Version-9/CD1/control.xml
lrwxrwxrwx root root  driverupdate-> SUSE-SP1-Version-9/CD1/driverupdate
lrwxrwxrwx root root  linux -> SUSE-SP1-Version-9/CD1/linux//
lrwxrwxrwx root root  media.1 -> SUSE-SLES-Version-9/CD1/media.1/
drwxr-xr-x root root  SUSE-CORE-Version-9/
drwxr-xr-x root root  SUSE-SLES-Version-9/
drwxr-xr-x root root  SUSE-SP1-Version-9/
drwxr-xr-x root root  yast/
# cat yast/instorder
/sles9/CD1      /sles9/CD1
/core9/CD1      /core9/CD1
```

- Configure NFS
  - Modify /etc/exports

```
# tail -2 /etc/exports
/nfs/sles9root              *(ro,sync)
/nfs/lpar2vs                *(ro,sync)
```

  - Configure NFS to start with **chkconfig** command

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- ~~Installing and configuring z/VM~~
- ~~Configuring an NFS server~~
- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- **Monitoring z/VM and Linux**
- Backup and restore
- Appendix - Relabel z/VM system volumes

Monitoring z/VM and Linux

- Using the INDICATE and other basic commands
  - Using INDICATE written by Bill Bitner
  - Using other QUERY and HELP commands
- z/VM Performance Toolkit
  - How to configure
  - How to use (brief)
- Monitoring Linux - two options
  - With the Linux RMF data gatherer (aka rmfpms)
    - Tar file downloaded from IBM FTP site
    - Modify a configuration file and run
    - Monitor data can be viewed via a browser
  - With APPLMON data gatherer built into SLES9 kernel
    - Must load three modules
    - Must turn on via kernel /proc variables
- Linux images can be registered with the Performance Toolkit

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- ~~Installing and configuring z/VM~~
- ~~Configuring an NFS server~~
- Install and configure Linux
- Configure NFS on controller
- Configure Linux for cloning
- Basic Linux virtual servers
- A virtual communications server
- A virtual communications controller server
- ~~Monitoring z/VM and Linux~~
- Backup and restore
- **Appendix - Relabel z/VM system volumes**

Relabel z/VM system volumes

- What happens when a second z/VM system can be seen by the first?

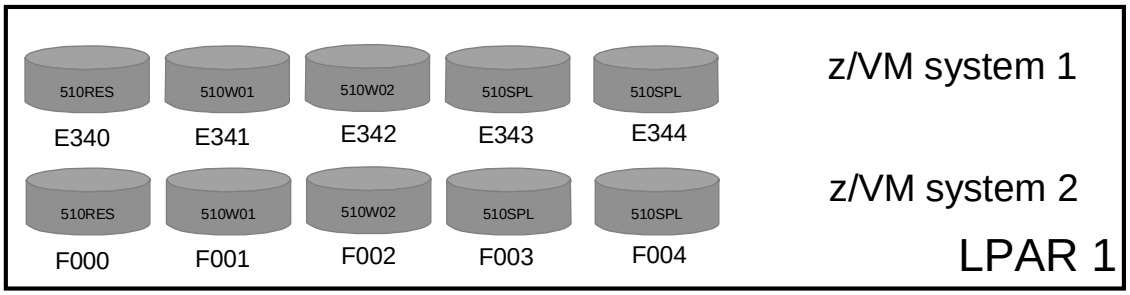

  - What happens with "IPL E340"?
  - What happens with "IPL F000"?
- Steps to relabel a z/VM system:
  - Modify labels in SYSTEM CONFIG file
  - Modify labels in USER DIRECT file
  - Change labels of the 5 volumes
  - Run DIRECTXA (will give a return code of 5)
  - SHUTDOWN REIPL (and cross fingers)
- LABEL510 EXEC and LABEL510 XEDIT written to help prevent mistakes

IBM Systems Group

Relabel z/VM system volumes

- What happens when a second z/VM system can be seen by the first?

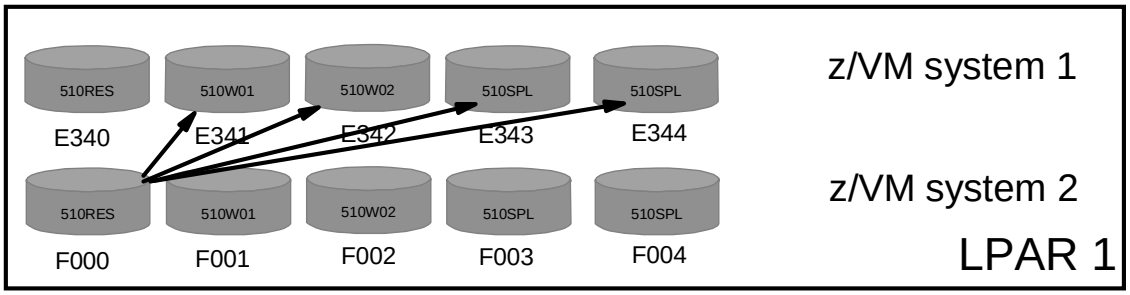

  - What happens with "IPL E340"?
  - What happens with "IPL F000"?
- Steps to relabel a z/VM system:
  - Modify labels in SYSTEM CONFIG file
  - Modify labels in USER DIRECT file
  - Change labels of the 5 volumes
  - Run DIRECTXA (will give a return code of 5)
  - SHUTDOWN REIPL (and cross fingers)
- LABEL510 EXEC and LABEL510 XEDIT written to help prevent mistakes

Resources

- Redbook *z/VM and Linux on zSeries: From LPAR to Virtual Servers in Two Days*
  - http://www.redbooks.ibm.com/abstracts/sg246695.html
- Files associated with the redbook
  - ftp://www.redbooks.ibm.com/redbooks/SG246695/
- *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
- SUSE LINUX Enterprise Server 9 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 - ???