



IBM Systems Group

The z/VM and Linux Virtualization Cookbooks for SLES 10 and RHEL 5

Michael Maclsaac - mikemac@us.ibm.com
Wednesday July 25th, 11:45 AM

© 2007 IBM Corporation



IBM Systems Group

Abstract

Updated books are available with titles: "z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES 10" and another "for RHEL 5". The goal of these books is to allow you to install and configure z/VM, install and configure Linux and be cloning Linux in two or three working days. Beyond the initial configuration, there are many other sections on customizing virtual servers and putting together solutions. This presentation will describe the two books, and will also give a live demonstration of a Linux cloning process and a new simple Web status application.

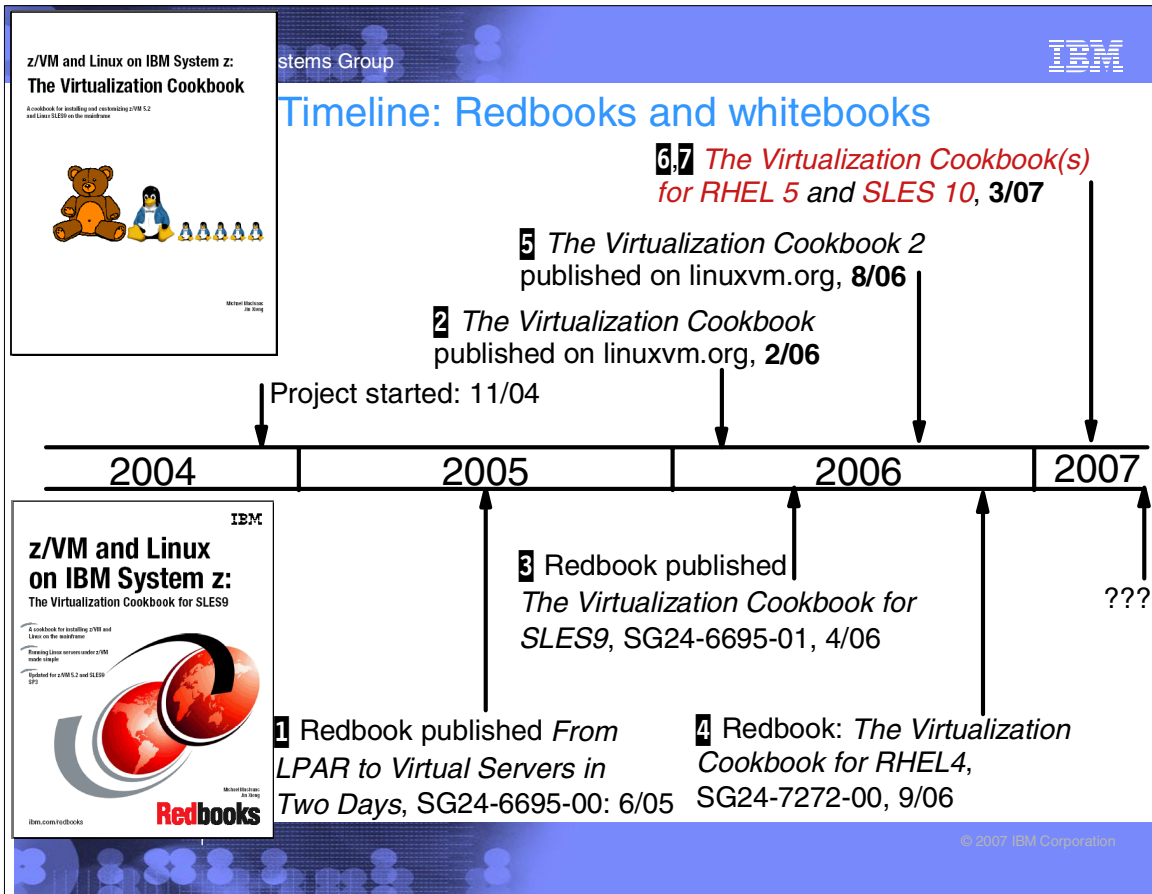
© 2007 IBM Corporation

Who am I?, who are you?

- Mike MacIsaac, mikemac@us.ibm.com
 - ▶ 20 years at IBM in NY
 - ▶ Programmer
 - ▶ OS/390, USS, Redbook project lead
 - ▶ Marketing technical support of z/VM, Linux, IBM software, ...
 - ▶ Wrote much of the *z/VM and Linux on zSeries: Virtualization Cookbooks*
- Who are you?
 - ▶ Experience with this book:
 - Have you tried the steps in this book?
 - Thinking about using this book?
 - Never heard of this 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 this talk?

Outline - by book chapters (*common*)

1. *Introduction to z/VM and Linux*
2. *Planning*
3. *Configuring a desktop machine*
4. *Installing and configuring z/VM*
5. *Servicing z/VM*
6. Configuring an NFS server
7. Installing and configuring Linux
8. Configuring Linux for cloning
9. Installing Linux with kickstart (**RH5 book only**)
10. Servicing Linux with Red Hat Network (**RH5 book only**)
11. Cloning open source virtual servers
12. Miscellaneous Recipes
13. *Monitoring z/VM and Linux*
14. *Backup and restore*



IBM Systems Group

Introduction:

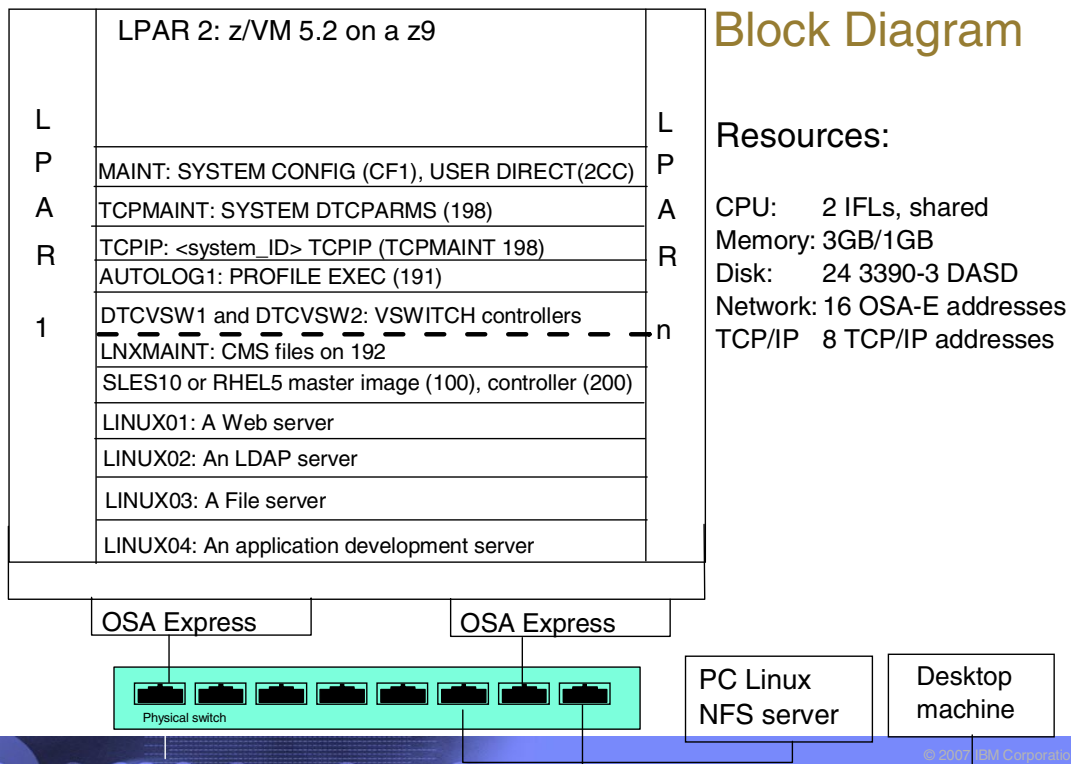
- History: project started 2004: impetus from a zBLC working group
 - ▶ Wanted Linux on System z to be appliance-like
- Philosophy
 - ▶ Cookbook to install/customize z/VM, install/customize Linux, and clone virtual servers
 - ▶ *Everything should be made as simple as possible, but not simpler.* -Albert Einstein
 - ▶ Reader (sysadmin) wants to understand all steps, takes ownership
 - ▶ Open "source" - free as in beer, free as in liberty
- What is new?
 - ▶ Two new cookbooks for RHEL 5, SLES 10
 - ▶ 2 VDISK swaps/server => larger root file system
 - ▶ Associated controller files are an RPM
 - ▶ New section *Centralizing home directories for LDAP users* - brings together:
 - Cloning, LDAP + PAM + NSS, NFS + automount of /home/
 - ▶ New section *Rescuing a Linux system*
 - ▶ Web application to monitor/log system: Data About z/VM and Linux (DAZL)
 - ▶ **clone.sh** script (now in /sbin/) is updated:
 - Added code to create one line description file
 - Removed hard-coded reference to the master image's mount point
 - Master image minidisk is now a variable (\$disk)
 - ▶ Installing z/VM onto 3390-9s is addressed

© 2007 IBM Corporation

Overview (cont'd):

- What has been removed?
 - ▶ Cloning IBM middleware
 - Hope to add back to SLES 10 book soon, RHEL 5 book eventually (???)
 - ▶ DCSS/XIP2
- Choices made:
 - ▶ "Roll your own" cloning rather than other products
 - Other solutions are all valid, more sophisticated, more complex
 - ▶ 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
 - ▶ Cloning and manual install hinge on CMS parameter files
 - ▶ Cloning done from Linux, no VM service machine needed
- Many usability tests conducted
 - ▶ Completion now takes about 4 days
- Get the books free on the Web at:
 - ▶ <http://linuxvm.org/present/>

Block Diagram



Planning - bill of materials

- Hardware
 - ▶ zSeries LPAR
 - 2 IFLs recommended (One is OK)
 - 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 5.2 on DVD (tape is OK)
 - ▶ SLES-10 or RHEL 5 Linux distro ISO images
 - ▶ Code associated with book - Tar files:
 - <http://linuxvm.org/present/>
- 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
 - 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

M	M	E	3	4	A
---	---	---	---	---	---



Configuring a desktop machine

- SSH client
 - ▶ PuTTY is described
 - Set SSH protocol to "2 only"
 - Add rows and columns
 - Save sessions
- VNC client
 - ▶ 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

Installing and configuring z/VM

- Install z/VM from DVD
 - ▶ Install from DVD is documented in some detail
 - ▶ Use the Integrated 3270 console on HMC
- Customize TCPIP with IPWIZARD
 - ▶ Also configure FTP server
- Customize SYSTEM CONFIG
 - ▶ Define a VSWITCH
- Add 5 paging volumes
 - ▶ Use supplied CPFORMAT EXEC to format
- 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
- Address z/VM security issues
 - ▶ Passwords in USER DIRECT
- Backing/restore system to tape
 - ▶ No recipe
- Relabeling system volumes

CPFORMAT EXEC

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-----'
    
```


Example:

```

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

Why relabel z/VM system volumes?

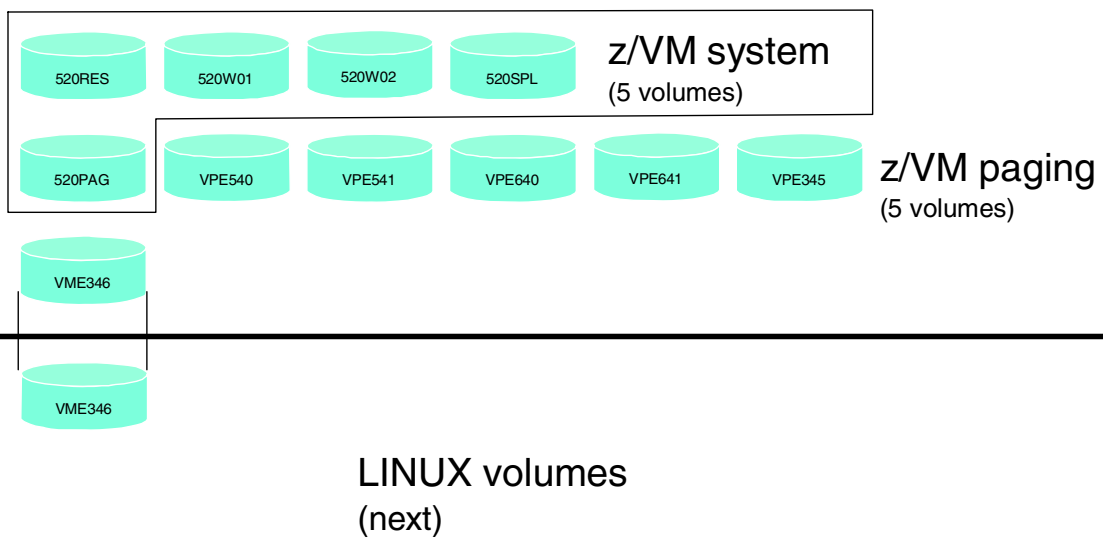
- Two reasons:
 - ▶ Avoid duplicate labels (esp. 520RES) if z/OS can see all DASD
 - ▶ Avoid z/VM systems from cross-referencing PAGE and SPOOL



Servicing z/VM

- Apply a Programming Temporary Fix (PTF)
- Apply a Recommended Service Upgrade (RSU)
 - ▶ Getting service via Internet FTP
 - ▶ SERVICE ALL
 - ▶ PUT2PROD
- Determining z/VM's service level
 - ▶ Adapted from ibm.com/vm pages

DASD view of z/VM

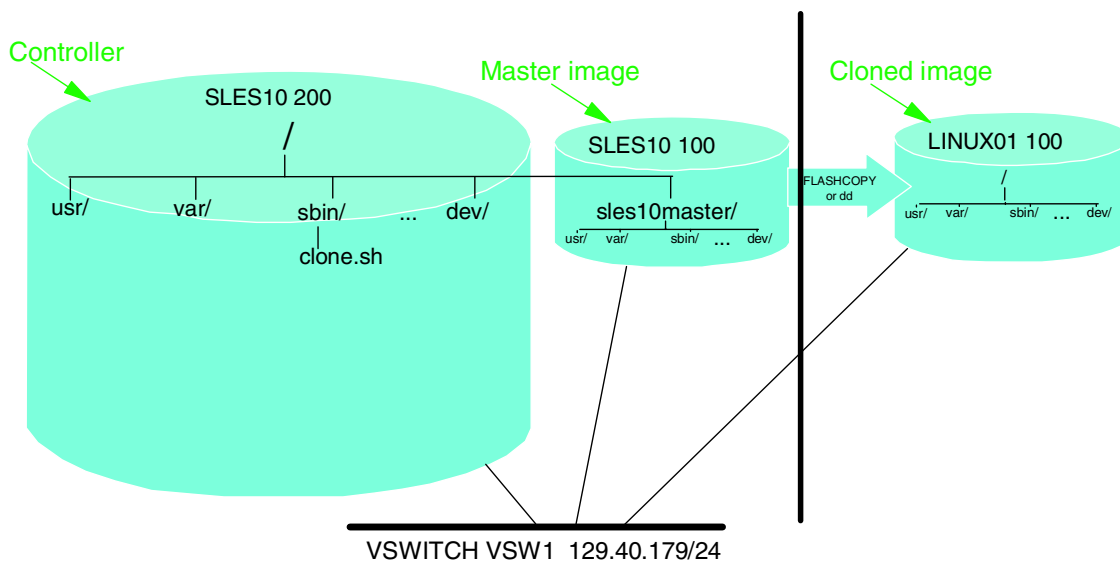


Configure a PC NFS server

- Installing Linux on zSeries is a chicken and egg problem
- Recommendation: install Linux on an Intel-architecture PC
- Server is a temporary NFS server (retire it after chapter 8)
- Steps:
 - ▶ Install Linux onto a PC
 - ▶ Copy files associated with this book to this NFS server
 - <http://linuxvm.org/present/misc/virt-cookbook-RH5.tgz> (for RHEL 5)
 - <http://linuxvm.org/present/misc/virt-cookbook-S10.tgz> (for SLES 10)
 - untar to /nfs/virt-cookbook-<distro>/
 - ▶ Set up an install directory under /nfs/<distro>/
 - ▶ Configure the NFS server to export these two directories

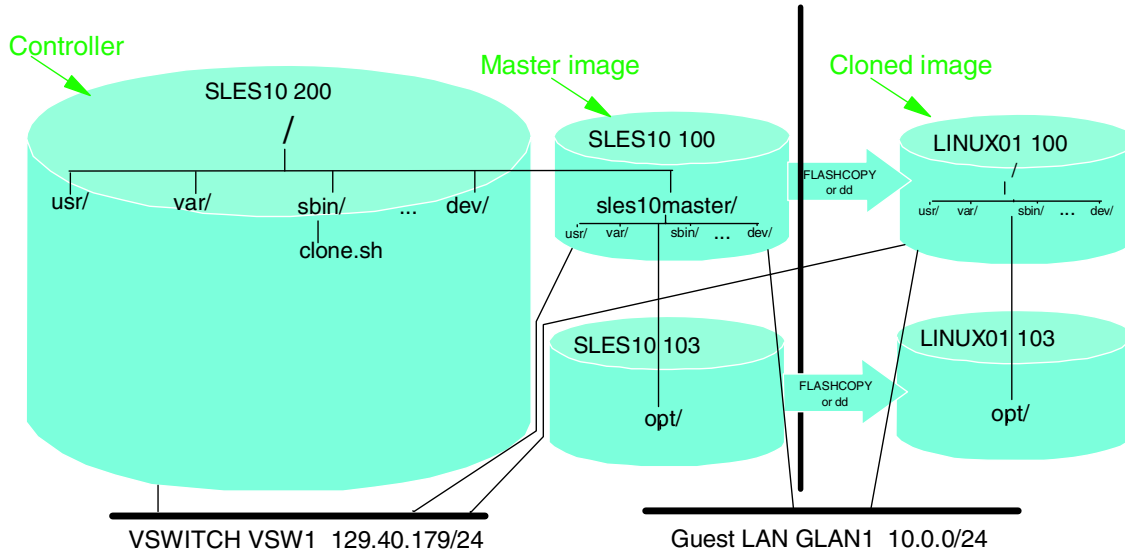
Installing and configuring Linux

- First: a conceptual diagram:
 - ▶ Controller/master user ID is dual boot
 - ▶ The `clone.sh` script copies the 100 minidisk to target user ID



Installing and configuring Linux

- Modified cloning configuration:
 - ▶ Add a 103 minidisk (3390-3) mounted over /opt
 - ▶ Add a interface to a private Guest LAN



Installing and configuring Linux (cont'd)

- Create new user ID - SLES10 or RHEL5 - with 7 3390-9s
- Add to z/VM startup - AUTOLOG1's PROFILE EXEC
- Prepare bootstrap files (kernel, RAMdisk, parmfiles) on LNXMAINT 192
- Install master image onto 100 with 101/102 VDISK swaps
- Configure master image
 - ▶ Create nightly.sh script
 - ▶ Adding additional RPMs
 - ▶ Configuring the VNC server
 - ▶ Preparing for Online Update
 - ▶ Removing unnecessary RPMs
 - ▶ Turning off unneeded services
 - ▶ Configuring rsyncd
 - ▶ Applying service - online update
 - ▶ Configuring sitar
 - ▶ Setting the software clock accurately
 - ▶ Setting system to halt on SIGNAL SHUTDOWN
 - ▶ Turning off the hz_timer
 - ▶ Modifying zipl.conf
- Install controller onto 200
 - ▶ 100 disk is /sles10master/, /backup/ file system, /nfs/ is 4 volume LVM (9GB)

Installing and configuring Linux (cont'd)

- Configure controller
 - ▶ Copying files to the controller
 - ▶ Adding additional RPMs
 - ▶ Configuring the VNC server
 - ▶ Removing unnecessary RPMs
 - ▶ Turning off unneeded services
 - ▶ Applying service if necessary - online update
 - ▶ Configuring sitar
 - ▶ Installing the cmsfs package
 - ▶ Turning on the NFS server
 - ▶ Turning on the NTP server
 - ▶ Enabling the vmcp module
 - ▶ Setting system to halt on SIGNAL SHUTDOWN
 - ▶ Turning off the hz_timer
 - ▶ Configuring SSH keys
 - ▶ Configuring Apache for DAZL
 - ▶ Setting ownership of Linux backup directories

Configuring NFS on the controller

- Copying files from NFS server to controller
 - ▶ Copying the SLES10 ISO images
 - ▶ Copying the files associated with this book
- Configuring the NFS server
- Changing the YaST installation location
- Retire the PC NFS server

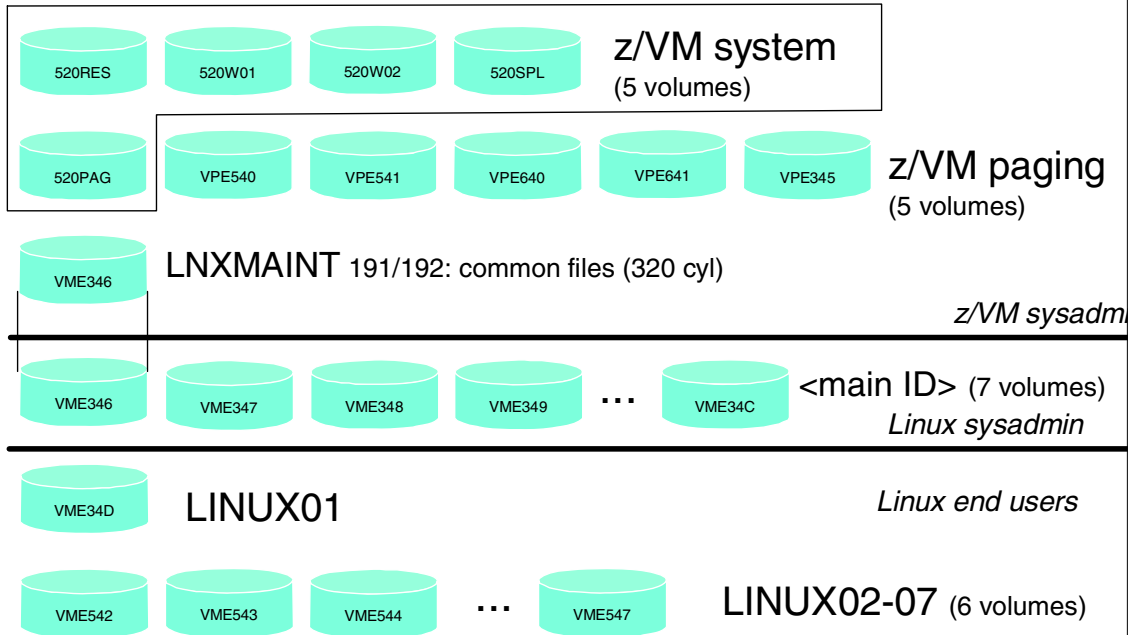
Configuring Linux for cloning

- How to clone manually
- How to use the clone.sh script
- Both processes do about the same tasks:
 - ▶ Link target ID as 1100
 - ▶ Copy from source (100) to target (1100) - use FLASHCOPY if you have it
 - ▶ Mount copied file system
 - ▶ Modify networking info - usually just IP@ and hostname
 - ▶ Detach target disk
 - ▶ IPL new clone
 - ▶ Modify SSH keys

Cloning open source virtual servers

- Clone and customize 4 open source virtual servers
 - ▶ LINUX01 - Web server
 - apache2 RPMs
 - Turning on a firewall (RH5 book only)
 - Creating a certificate (S10 book only)
 - ▶ LINUX02 - LDAP server
 - openldap RPMs
 - Migrate /etc/ users and groups via PADL migration tools
 - ▶ LINUX03
 - File server - Samba
 - Create one new Samba user, one new file share
 - ▶ LINUX04 - Application development server
 - Python, Perl, tcl, PHP
 - C/C++
 - Java
 - etc.

Review - DASD and role view



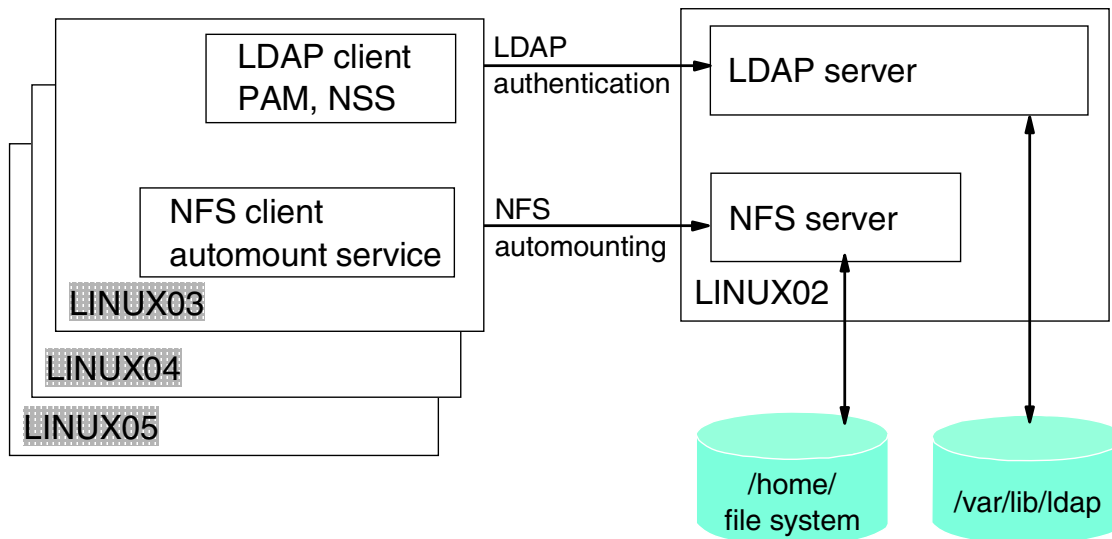
© 2007 IBM Corporation

Miscellaneous Recipes

- Other tasks you might want to do:
 - ▶ Adding a logical volume (LVM) via line commands
 - Aside: a rule of GUI administration tools:
 - *First learn line commands to perform a task and know what files are changed. Then use a GUI tool to do the same task if it is faster or more usable.*
 - ▶ Extending a logical volume via line commands
 - ▶ Centralizing home directories for LDAP users - brings together
 - LDAP in LINUX02
 - Logical volume just created
 - PAM, NSS for authentication
 - Automount NFS for "traveling" /home/
 - ▶ Rescuing a Linux system

© 2007 IBM Corporation

Miscellaneous recipes (cont'd)



© 2007 IBM Corporation

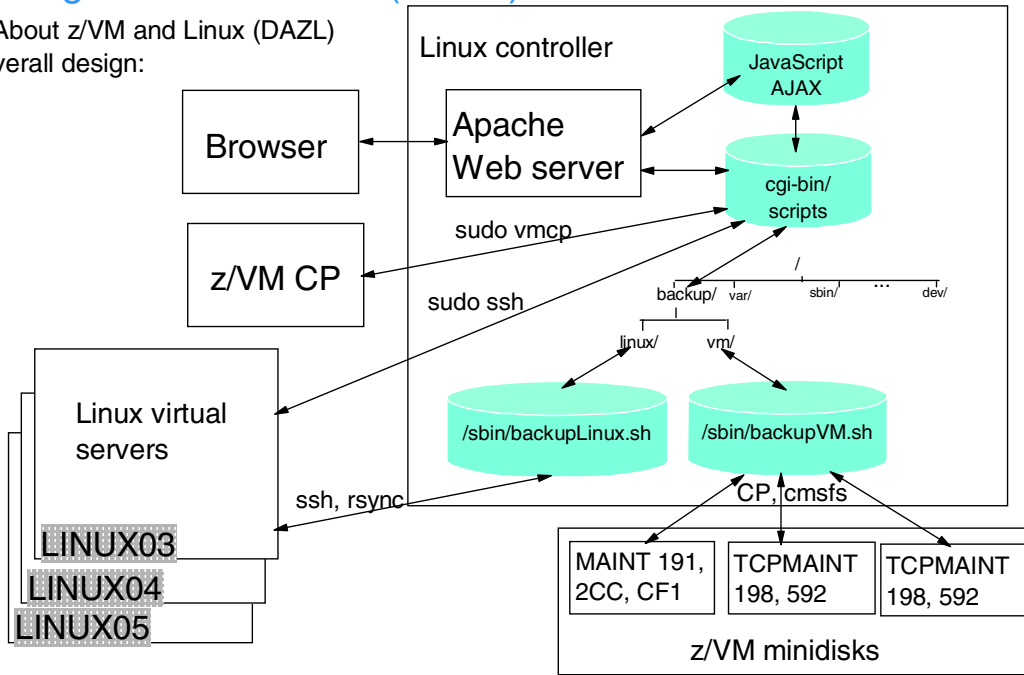
Monitoring z/VM and Linux

- 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 basic and as a Web server
 - ▶ How to use (brief)
- Monitoring Linux - two options
 - ▶ With the Linux RMF data gatherer (aka rmfpm)
 - ▶ With APPLMON data gatherer built into SLES9 kernel
- Linux images can be registered with the Performance Toolkit
- New Web app: Data About z/VM and Linux
 - ▶ Disclaimer:
 - Not formally supported
 - Not heavily tested
 - "Quick and dirty"
 - ▶ Does not replace IBM Director!
 - See "*IBM Director 5.2 on System Z with z/VM Center extensions*"
 - Session 9219, Friday at 9:30

© 2007 IBM Corporation

Monitoring z/VM and Linux (cont'd)

- Data About z/VM and Linux (DAZL)
 - Overall design:



Monitoring z/VM and Linux (cont'd)

- DAZL Screen shots:

+1 page

+3

Volume Label	Device Address	Minidisks (first three)	Status	Cylinders used/total
MMA711	A711	\$ALLOC\$ A04 SLES9X 100 SLES9X 102	CP SYSTEM	3339/3339
MMA712	A712	\$ALLOC\$ A05 SLES9X 200	CP SYSTEM	3339/3339
MMA713	A713	\$ALLOC\$ A06 LNVMMAINT 0191 LNVMMAINT 0192 More	CP SYSTEM	3339/3339
MMA714	A714	\$ALLOC\$ A07 SLES9X 204	CP SYSTEM	3339/3339
MMA715	A715	\$ALLOC\$ A08 SLES9X 205	CP SYSTEM	3339/3339
MMA716	A716	\$ALLOC\$ A09 SLES9X 206	CP SYSTEM	3339/3339
MMA717	A717	\$ALLOC\$ A0A SLES9X 207	CP SYSTEM	3339/3339
MMA718	A718	\$ALLOC\$ A0B LINUX01 200 LINUX01 202	CP SYSTEM	3339/3339

Monitoring z/VM and Linux (cont'd)

- DAZL Linux report:

z/VM User ID	Linux Host name	IP address	Sitar data	Description/Log
DB2	lat133.pbm.ihost.com	129.40.178.133	Cron New	descLog.txt file not found
LINUX01	lat121.pbm.ihost.com	129.40.178.121	Cron New	This is a Web Server
LINUX02	lat122.pbm.ihost.com	129.40.178.122	Cron New	This is an LDAP Server
LINUX03	lat123.pbm.ihost.com	129.40.178.123	Cron New	This is an application development server
LINUX04	lat124.pbm.ihost.com	129.40.178.124	Cron New	This is a Samba Server
LINUX05	lat125.pbm.ihost.com	129.40.178.125	Cron New	linux05 - WAS clone
LINUX06	lat126.pbm.ihost.com	129.40.178.126	Cron New	linux06 - DB2 clone
LINUX07	lat127.pbm.ihost.com	129.40.178.127	Cron New	MQ clone

Monitoring z/VM and Linux (cont'd)

- DAZL editing in place:

z/VM User ID	Linux Host name	IP address	Sitar data	Description/Log
DB2	lat133.pbm.ihost.com	129.40.178.133	Cron New	descLog.txt file not found
LINUX01	lat121.pbm.ihost.com	129.40.178.121	Cron New	This is a Web Server It was updated with the test Web Site fubar on Nov 21, 2006.
LINUX02	lat122.pbm.ihost.com	129.40.178.122	Cron New	This is an LDAP Server
LINUX03	lat123.pbm.ihost.com	129.40.178.123	Cron New	This is an application development server

Monitoring z/VM and Linux (cont'd)

- DAZL z/VM report:

Indicate Load: IND

```

AVGPROC-000% 04
XSTORE-00001/SEC MIGRATE-0000/SEC
MDC READS-000000/SEC WRITES-000000/SEC HIT RATIO-000%
PAGING-1/SEC STEAL-000%
Q0-00000(00000)          DORMANT-00020
Q1-00003(00000)          E1-00000(00000)
Q2-00001(00000)  EXPAN-001 E2-00000(00000)
Q3-00001(00000)  EXPAN-001 E3-00000(00000)

PROC 0000-000%          PROC 0001-001%
PROC 0002-000%          PROC 0003-000%

LIMITED-00000

IND QUEUES EXP
    
```

Display system allocation: QUERY ALLOC MAP

VOLID	RDEV	EXTENT START	EXTENT END	TOTAL IN USE	HIGH USED	%
MVA700	A700	1	20	20	1	5%
MVA701	A701	1	3338 600840	90673	94486	15%
MVA702	A702	1	3338 600840	86981	181893	14%
MVA705	A705	0	3338 601020	89740	181789	14%
MVA706	A706	0	3338 601020	89940	212788	14%
MVA707	A707	0	3338 601020	90479	182157	15%
MVA708	A708	0	3338 601020	90978	190033	15%
MVA709	A709	0	3338 601020	88533	197445	14%

Display who is logged on/disconnected: QUERY NAMES

```

LINUX07 - DSC , MQS      - DSC , LINUX06 - DSC , LINUX05 - DSC
LINUX04 - DSC , LINUX03 - DSC , LINUX02 - DSC , LINUX01 - DSC
FTSPSERVE - DSC , DTCVSW2 - DSC , DTCVSW1 - DSC , TCP/IP - DSC
OPERSYMP - DSC , DISKACNT - DSC , EREP      - DSC , OPERATOR - DSC
SLES10   - DSC
VSM      - TCP/IP
    
```

Display system service level: QUERY CPLEVEL

```

z/VM Version 5 Release 2.0, service level 0601 (64-bit)
Generated at 02/13/06 11:45:25 EDT
IPL at 08/14/06 08:51:11 EDT
    
```

Display central/expanded storage: Q STOR/XSTOR

```

STORAGE = 1536M
XSTORE = 512M online= 512M
XSTORE = 512M userid= SYSTEM usage= 99% retained= 0M pending
XSTORE MDC min=0M, max=0M, usage=0%
XSTORE = 512M userid= (none) max. attach= 512M
    
```

Display virtual switch info: QUERY VSWITCH

```

VSWITCH SYSTEM VSW1      Type: VSWITCH Connected: 9      Maxco
PERSISTENT RESTRICTED  NONROUTER                      Accou
VLAN Unaware
State: Ready
IPTimeout: 5          QueueStorage: 8
Portname: UNASSIGNED RDEV: 3004 Controller: DTCVSW2 VDEV
Portname: UNASSIGNED RDEV: 3008 Controller: DTCVSW1 VDEV
    
```

z/VM Version 5 Release 2.0, service level 0601 (64-bit)
Generated at 02/13/06 11:45:25 EDT
IPL at 08/14/06 08:51:11 EDT

Live Demo

Remember:
If it's not working,
just pretend it is



Resources

- Book *z/VM and Linux on System z: The Virtualization Cookbook for SLES 10*
 - ▶ <http://linuxvm.org/present/virt-cookbook-S10.pdf>
- Files associated with the SLES 10 book
 - ▶ <ftp://linuxvm.org/present/virt-cookbook-S10.tgz>
- Book *z/VM and Linux on System z: The Virtualization Cookbook for RHEL 5* (publish Feb 28?)
 - ▶ <http://linuxvm.org/present/virt-cookbook-RH5.pdf>
- Files associated with the RHEL 5 book (Feb 28?)
 - ▶ <ftp://linuxvm.org/present/virt-cookbook-RH5.tgz>
- *The Linux for zSeries and S/390 portal*
 - ▶ <http://linuxvm.org/>
- The linux-390 list server
 - ▶ <http://www2.marist.edu/htbin/wlindex?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 - ???

Q: What is the answer
to The Ultimate Question
Of Life, the Universe
and Everything?
A: 42

But what is the
ultimate
question?

