



The Virtualization Cookbook: Day 2 - Linux

Michael Maclsaac - mikemac@us.ibm.com
Thursday, August 17th, 4:30 PM
Session 9217



Abstract

An updated redbook is available: 'z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9'. The goal of the book is to allow you to install and configure z/VM, install and configure Linux and be cloning Linux in two working days. On the second day, you install Linux twice onto the same virtual machine. The first install is called the 'master image' which will be cloned from. The second install is called the 'controller' which will do cloning and other tasks. The data is copied from the NFS server to the controller so the PC can be retired. Then cloning is addressed and the supplied clone script is described. You should be ready to configure clones into Web, LDAP, file and print, development or other servers. Also backup and restore issues are addressed.

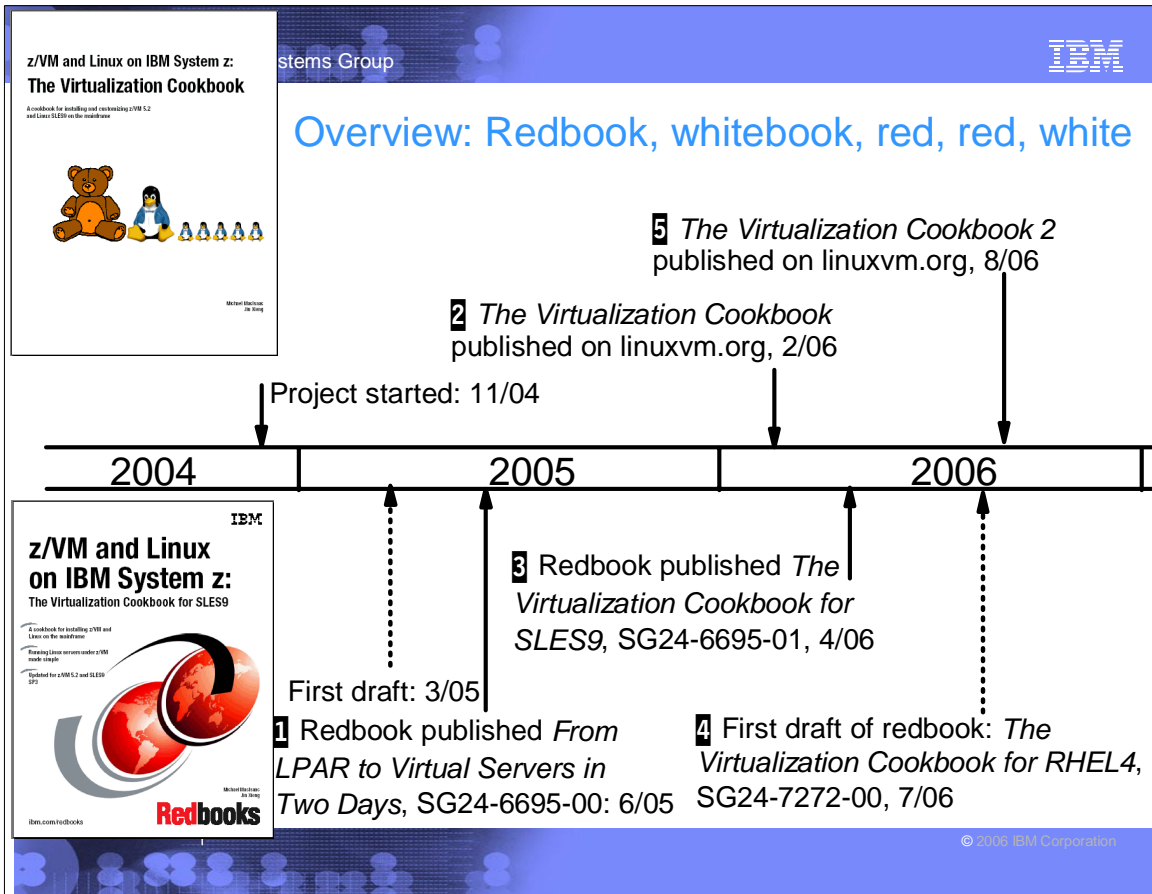
Who am I?, who are you?

- Mike Maclsaac, mikemac@us.ibm.com
 - ▶ 20 years at IBM in NY
 - ▶ z/VM and Linux evangelist
 - ▶ Wrote much of *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9*
- Who are you?
 - ▶ No Linux on zSeries?
 - ▶ Testing Linux/PoC?
 - ▶ Linux in production?
 - ▶ =====
 - ▶ Have you tried the steps in this book?
 - ▶ Are you thinking about using it?
 - ▶ Have you never heard of it?

Outline for the two presentations

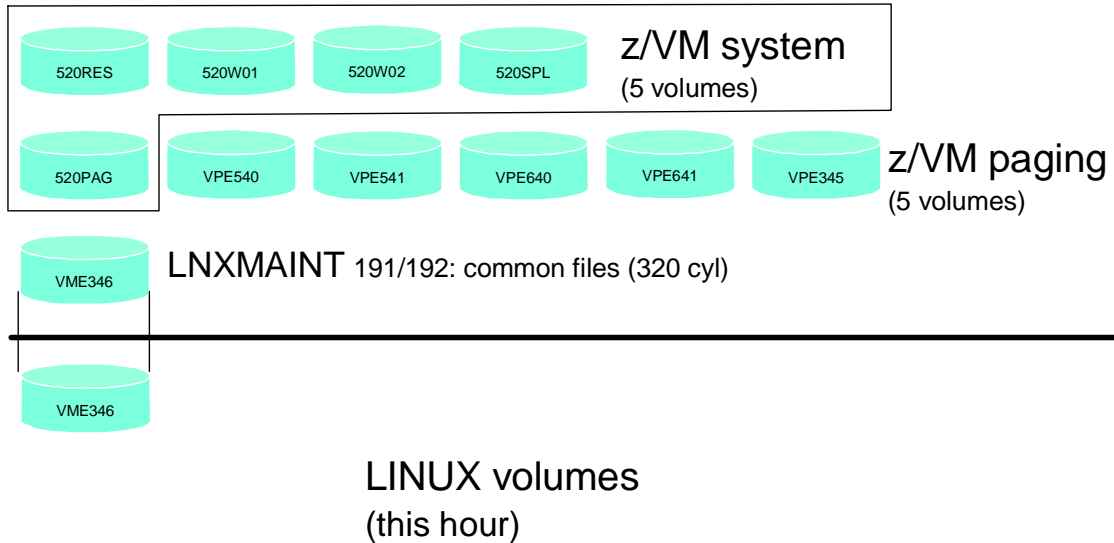
Focus is on Linux this hour

- **Overview**
- Planning
- Configuring a Windows desktop
- Installing and configuring z/VM
- Servicing z/VM
- Configuring an NFS server
- ---
- **Install and configure Linux**
- **Configure NFS on controller**
- **Configure Linux for cloning**
- **Cloning basic virtual servers**
- **Cloning IBM Middleware virtual servers (NEW!)**
- Monitoring z/VM and Linux
- **Backup and restore**

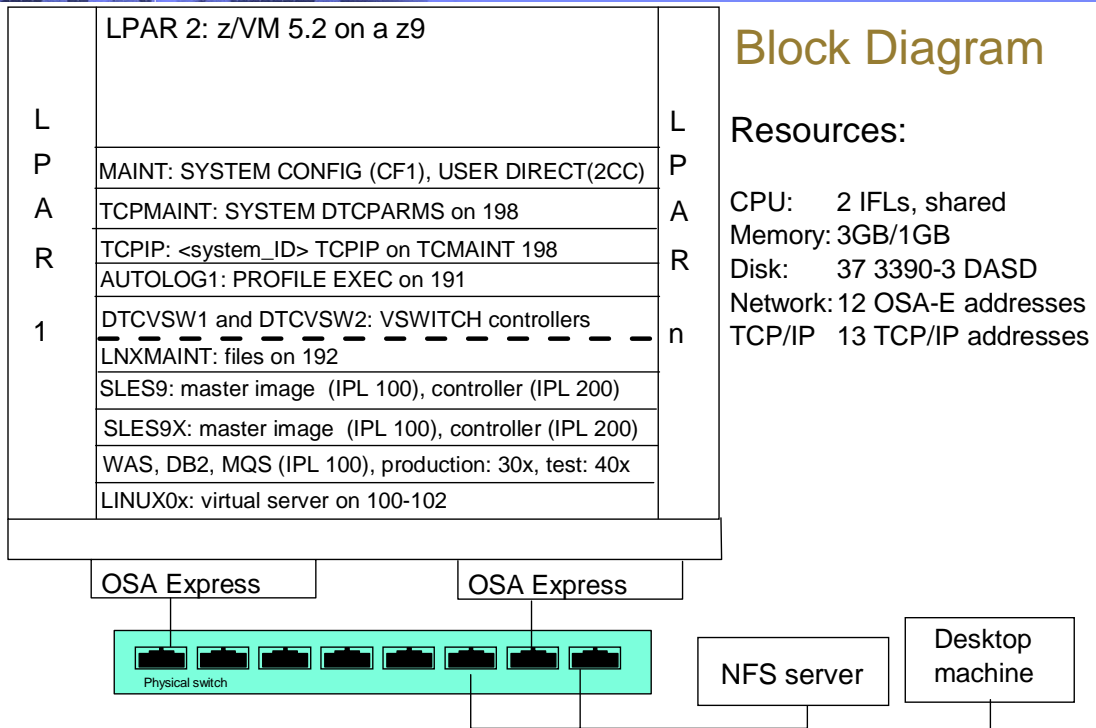


- IBM Systems Group
- ## Overview: how to get it (let me count the ways :))
- Redbook version (SLES9 SP3, VM 5.2)
 - ▶ *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9*
 - <http://www.redbooks.ibm.com/abstracts/sg246695.html>
 - ▶ Associated files:
 - <ftp://www.redbooks.ibm.com/redbooks/SG246695/>
 - "Whitebook" version (about same as redbook)
 - ▶ *z/VM and Linux on IBM System z: The Virtualization Cookbook*
 - <http://linuxvm.org/present/misc/virt-cookbook-1.pdf>
 - ▶ Associated files
 - <http://linuxvm.org/present/misc/virt-cookbook-1.tgz>
 - Redbook draft (RHEL4 U3, VM 5.2)
 - ▶ *z/VM and Linux on IBM System z: The Virtualization Cookbook for Red Hat Enterprise Linux 4*
 - <http://www.redbooks.ibm.com/redpieces/abstracts/sg247272.html>
 - ▶ Associated files
 - <ftp://www.redbooks.ibm.com/redbooks/SG247272>
 - Whitebook-2 draft
 - <http://linuxvm.org/present/misc/virt-cookbook-2.pdf>
 - ▶ Associated files
 - <http://linuxvm.org/present/misc/virt-cookbook-2.tgz>
- © 2006 IBM Corporation

Review - End of 1 day - DASD view



Block Diagram



Outline

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

Install and configure Linux - define user ID SLES9X

```

PROFILE LNXDFLT
  IPL CMS
  MACHINE ESA 4
  CPU 00 BASE
  CPU 01
  NICDEF 600 TYPE QDIO LAN SYSTEM VSW1
  SPOOL 000C 2540 READER *
  SPOOL 000D 2540 PUNCH A
  SPOOL 000E 1403 A
  CONSOLE 009 3215 T
  LINK MAINT 0190 0190 RR
  LINK MAINT 019D 019D RR
  LINK MAINT 019E 019E RR
  LINK LNXMAINT 192 191 RR
  LINK TCPMAINT 592 592 RR
  ...
USER SLES9X LNX4VM 256M 1G BEG
  INCLUDE LNXDFLT
  OPTION LNKNOPAS APPLMON
  MDISK 100 3390 0001 3038 MMA751 MR LNX4VM LNX4VM LNX4VM
  MDISK 102 3390 3039 0300 MMA751 MR LNX4VM LNX4VM LNX4VM
  MINIOPT NOMDC
  MDISK 200 3390 0001 3338 MMA752 MR LNX4VM LNX4VM LNX4VM
  MDISK 203 3390 0321 3018 MMA753 MR LNX4VM LNX4VM LNX4VM
  MDISK 204 3390 0001 3338 MMA754 MR LNX4VM LNX4VM LNX4VM
  MDISK 205 3390 0001 3338 MMA755 MR LNX4VM LNX4VM LNX4VM
  MDISK 206 3390 0001 3338 MMA756 MR LNX4VM LNX4VM LNX4VM
  MDISK 207 3390 0001 3338 MMA757 MR LNX4VM LNX4VM LNX4VM

```

← VSWITCH "adapter"

← master image

← controller

← /backup/

← /nfs/ LV

Install and configure Linux (cont'd)

- Prepare bootstrap files
 - ▶ FTP from SLES9 install tree on Linux PC to LNXMAINT 192:
 - Kernel SLES9X KERNEL
 - Initial RAMdisk SLES9X INITRD
 - ▶ Install the master image on 100-102
 - Trickier install panels are documented well
 - ▶ Configure the master image
 - Apply service if necessary - Y.O.U. is documented
 - Copy associate files from NFS server via "scp -r"
 - Remove unneeded/add additional RPMs
 - Turn off unneeded services
 - Configure rsyncd
 - Configure sitar
 - Configure VNC

Install and configure Linux (cont'd)

- Set the software clock from the controller (**/usr/sbin/ntpdate -q**)
- Turn off the hz_timer
- Harden the system with Bastille Linux
- ▶ Reboot system
- Install the controller - same as master except
 - ▶ Minidisk 200 is the root file system
 - ▶ Minidisk 100 is /sles9master (Don't format!!)
 - ▶ Mindisk 203 becomes a file system over /backup
 - ▶ Mindisks 204-207 become a logical volume mounted over /nfs
- Configure the controller
 - ▶ Apply service, if necessary
 - ▶ Copy files from NFS server
 - ▶ Remove some RPMs
 - ▶ Add some RPMs

Install and configure Linux (cont'd)

- ▶ Install cmsfs package by Rick Troth
- ▶ Turn off unneeded services
- ▶ Turn on the NFS server
- ▶ Turn on the NTP server
 - Get accurate time from two Stratum-2 servers on the Internet
- ▶ Turn off hz_timer
- ▶ Configure SSH keys for key-based authentication to clones
- ▶ Harden the controller with Bastille Linux
- ▶ Reboot
- You now have a *dual-boot* Linux ID
 - ▶ Controller normally runs
 - ▶ Golden image can be brought up (or can be **chroot**'ed into)
 - ▶ Controller can get to clone via keys (without SSH password)**

Outline

- Planning
- Configuring a Windows desktop
- Installing and configuring zVM
- Servicing zVM
- Configuring an NFS server
- Install and configure Linux
- **Configure NFS on the controller**
 - Configure Linux for cloning
 - Cloning basic virtual servers
 - Cloning IBM middleware virtual servers
 - Monitoring zVM and Linux
 - Backup and restore
 - Appendix—Relabel zVM system volumes

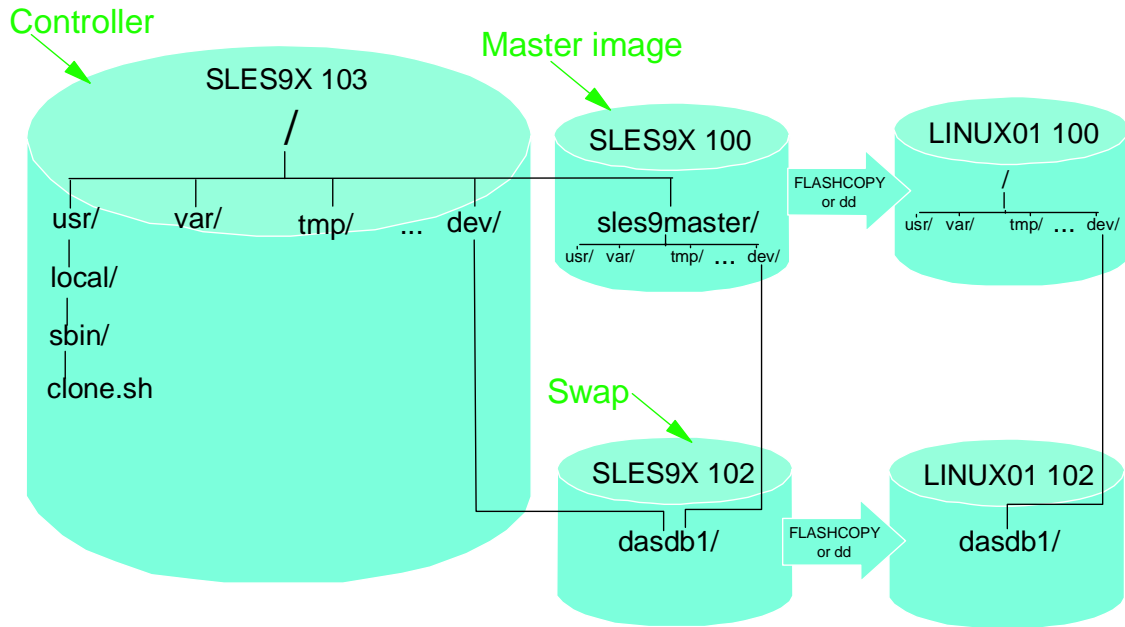
Configure NFS on controller

- **Copy files from PC NFS server**
 - # `scp -rp pc.nfs.server.ip:/nfs/sles9x/sles9sp3root /nfs`
 - # `scp -rp pc.nfs.server.ip:/nfs/virt-cookbook-2 /nfs`
- **Configure NFS server**
 - ▶ Same steps as on PC server
- **Change YaST install tree location**
 - ▶ Change source on the golden image
 - # `chroot /sles9master`
 - `yast => Software => Change Source of Installation`
 - # `exit`
 - Somewhat of a paradox, but useful on clones
 - ▶ Change source on controller
 - ▶ Change source in CMS parameter file
- **Retire PC NFS server (if desired)**

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

Cloning - block diagram



© 2006 IBM Corporation

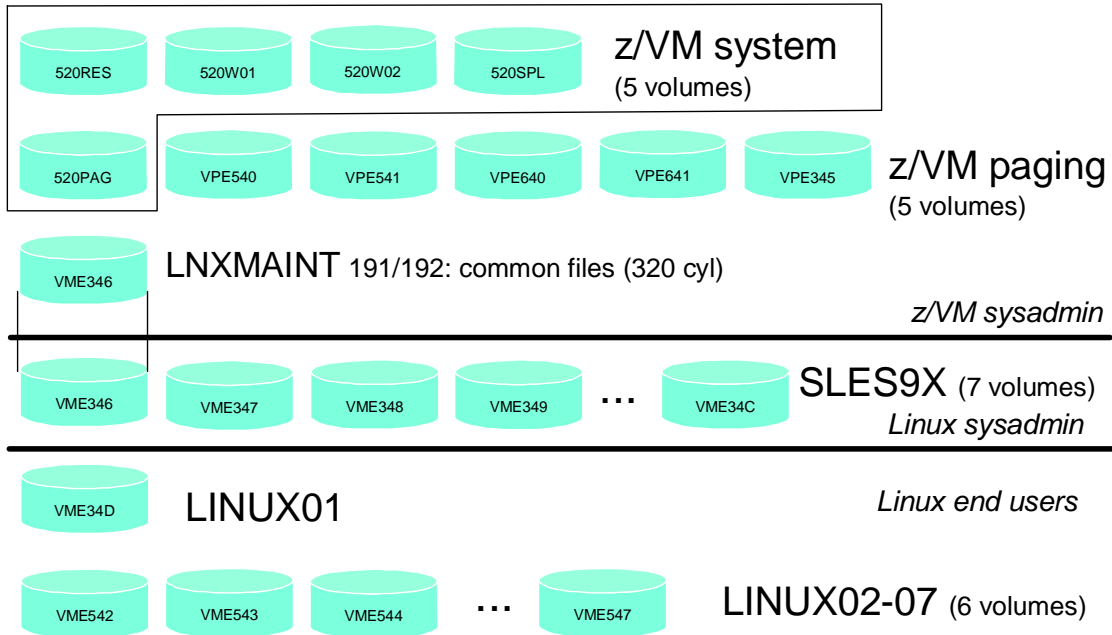
Cloning details

```
# main() # the first three lines are some important global variables
master_mnt_pt="/sles9master"           # set directory of master root file system
cloned_mnt_pt="/mnt/sles9cloned"       # set directory of temporary mount point
backup_dir="/backup/linux"            # set directory of Linux backups

process_arguments $@                   # process arguments passed by user
check_target_id                         # be sure user ID exists and is logged off
get_parmfile_info                       # get info from source and target parm files
ask_are_you_sure                        # confirm that disks will be overwritten
copy_system                             # copy 100 and 102 disks to target ID
modify_cloned_image $clone_id 100 1100 # modify newly copied system
cp_cmd XAUTOLOG $clone_id              # bring new clone to life
make_backup_dir                         # make a backup directory
echo "Successfully cloned $master_mnt_pt to $clone_id"
echo "You should be able to ping $target_IP within one minute"
exit 0
```

© 2006 IBM Corporation

Review - DASD and role view



© 2006 IBM Corporation

Outline

- Planning
- Configuring a Windows desktop
- Installing and configuring z/VM
- Servicing z/VM
- Configuring an NFS server
- Install and configure Linux
- Configure NFS on the controller
- Configure Linux for cloning
- **Cloning basic virtual servers**
- Cloning IBM middleware virtual servers
- Monitoring z/VM and Linux
- Backup and restore
- Appendix – Relabel z/VM system volumes

© 2006 IBM Corporation

Basic Linux virtual servers

- Basic virtual server types
 - ▶ A virtual Web server
 - ▶ A virtual LDAP server
 - ▶ A virtual file and print server
 - ▶ A virtual application development server
- Clone a virtual server, with exception of LDAP server
- Install scripts to install/setup virtual server
- Procedures on how to bring them to life
- Also: Sharing executables via DCSS/XIP2

A virtual Web server

- Clone a server and install Apache2 RPMs
- Testing Apache
 - ▶ Start Apache: **rcapache2 start**
 - ▶ Point Web browser to server address
- Configure HTTPS in Apache
 - ▶ Create a server certificate using openssl
 - self-signed
 - generate request for CA-signed
 - ▶ Configure Apache for HTTPS
 - runs on a different port, another Web server
 - create a virtual host
- Populate your Web site
 - ▶ Web pages located in /srv/www/htdocs/ by default

A virtual LDAP server

- OpenLDAP, pretty standard with most distros
- Centralized login authentication and user and group ID resolution
- You must have DNS for LDAP to work successfully
- Install Linux image manually
 - ▶ During the 2nd half of the YaST install, setup LDAP in the Service Configuration Panel
- Customize after install
- Add a new user using YaST and login using new credentials
- Setting another virtual server to use LDAP server for login authentication
 - ▶ Add 2 RPMs: pam_ldap and nss_ldap
 - ▶ Use YaST to modify LDAP client configuration
 - ▶ Modify the OpenLDAP client configuration file
 - ▶ Test that it's working by logging in using a user defined on the LDAP server

A virtual file and print server

- Samba
 - ▶ Allows Windows clients to map Linux file systems as shared drives
 - ▶ Acts as middle-man between Windows clients and a Linux print server
- Clone a virtual server and install necessary RPMs
- Tweak Bastille configuration to allow NFS to run
- Configure the Samba configuration file to indicate which directories to share and how
- Add a Samba user
- Start Samba: rcnmb start, rcsmb start
- Test the changes
 - ▶ Map network drive on Windows client to the Samba share
- Print server configuration: *Printing with Linux on zSeries Using CUPS and Samba*, REDP-3864.

A virtual application development server

- "Pure" development environment
 - ▶ Most distributions come with a basic set of development tools
 - C/C++ development environment
 - Java development environment
 - Scripting languages
 - ▶ Web dev platform LAMP; Linux, Apache, MySQL, and Python/Perl
 - ▶ Basic C/C++ and Java development techniques

Using DCSS/XIP2

- Reduce memory requirements of Linux servers
 - ▶ One copy of shared data loaded in memory per VM
 - ▶ DCSS is z/VM Discontiguous Saved Segment
 - ▶ XIP2 is Linux file system that allows a DCSS to be mounted
- Creating a DCSS/XIP2 shared file system:
 - ▶ Determine the size of the DCSS
 - ▶ Determine start and end addresses of the DCSS
 - ▶ Prepare Linux ID for DCSS creation (only for first guest)
 - ▶ Create a file system image for the DCSS in Linux
 - ▶ Create a DCSS from the file system image in VM
 - ▶ Change Linux to use DCSS
 - ▶ Test DCSS using XIP2
 - ▶ Activate XIP2 file system at boot time
 - ▶ Modify other Linux IDs to use DCSS - and save memory!
- Section based on work by Carsten Otte and Carlos Ordonez


Outline

- Planning
- Configuring a Windows desktop
- Installing and configuring z/VM
- Servicing z/VM
- Configuring an NFS server
- Install and configure Linux
- Configure NFS on the controller
- Configure Linux for cloning
- Cloning basic virtual servers
- **Cloning IBM middleware virtual servers (NEW!)**
- Monitoring z/VM and Linux
- Backup and restore
- Appendix – Relabel z/VM system volumes

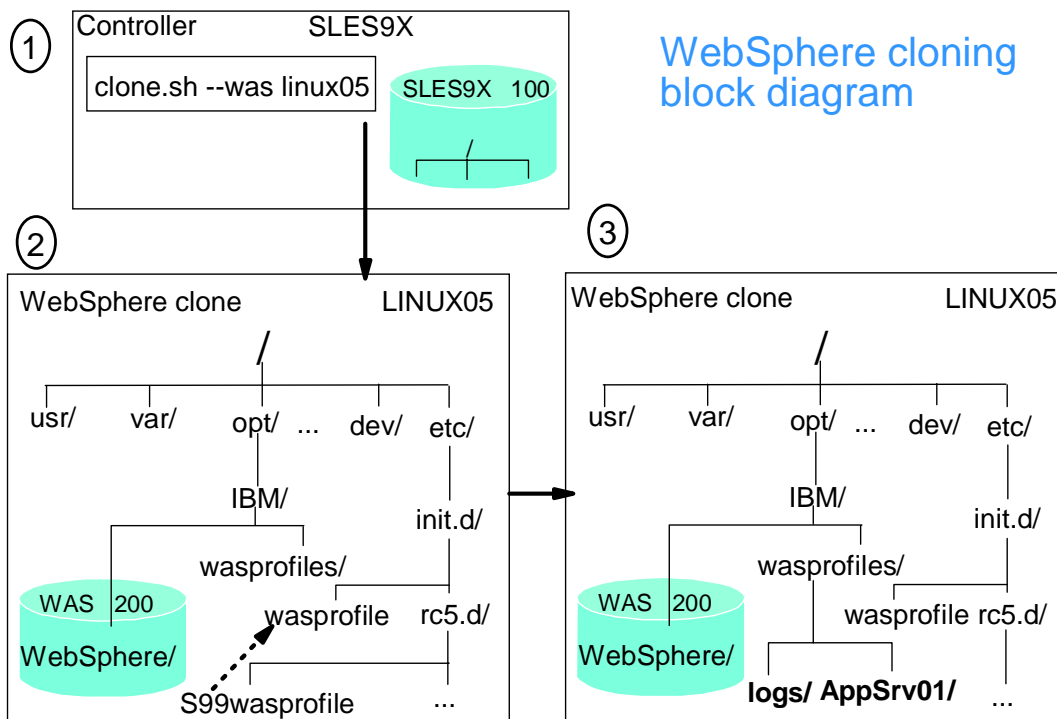
Cloning IBM middleware virtual servers

- Discusses some IBM middleware products
 - ▶ WebSphere Application Server V6.0.2 (V6.1?)
 - ▶ DB2 V8.2
 - ▶ MQ Series V6.0
- Install binaries on a master server
- Link binaries R/O and create application data R/W on clones
 - ▶ WebSphere uses profiles
 - ▶ DB2 uses instances
 - ▶ MQ Series uses queue managers
- Use 2 disks on each of the masters
 - ▶ Production - stable version
 - ▶ Test - version.next
- Clones can alternate between production and test binaries

Overview: Cloned system block diagram



WebSphere cloning block diagram



Outline

- Planning
- Configuring a Windows desktop
- Installing and configuring z/VM
- Servicing z/VM
- Configuring an NFS server
- Install and configure Linux
- Configure NFS on the controller
- Configure Linux for cloning
- Cloning basic virtual servers
- Cloning IBM middleware virtual servers
- Monitoring z/VM and Linux
- **Backup and restore**
- Appendix – Relabel z/VM system volumes

Backup and restore

- Overview
 - ▶ Incremental backup of z/VM
 - ▶ Incremental backup of Linux
 - ▶ Disaster recovery
- Incremental backup of z/VM
 - ▶ There are 6 important configuration files:
 - SYSTEM CONFIG (on MAINT CF1)
 - USER DIRECT (on MAINT 2CC)
 - <system_ID> TCPIP, SYSTEM DTCPARMS (TCPMAINT 198)
 - TCPIP DATA (on TCPMAINT 592)
 - PROFILE EXEC (on AUTOLOG1 191)
 - ▶ If you have copies of these files, you can restore a z/VM system
 - ▶ Recall that the /backup file system was created on the controller
 - ▶ These files can be backed up nightly - here is one method ...

Backup and restore (cont'd)

- ▶ On controller: to backup Linuxes a script, **backup_linux**, is written:

```
# tail -7 /etc/cron.daily/backup_linux
backup_dir="/backup/linux"
cd $backup_dir
for i in LINUX*-on-* # iterate through directories starting with LINUX*-on-*
do
  IP_addr=${i#LINUX*-on-} # this chops the head off and grabs the IP address
  cd $backup_dir/$i      # change directory
  rsync -r --timeout=30 $IP_addr:/etc . # use rsync to back up the /etc dir
done
```

- ▶ On clones, sitar is run nightly

```
#!/bin/bash
sitar --format=html --outfile=/etc/sitar.html
```

- By writing output of sitar to /etc/, it is also backed up

- Disaster recovery

- ▶ Follow good s/390 procedures (but for a different backup ...)
- ▶ Consider that three volumes contain
 - LNXMAINT 192 (Common CMS files)
 - LINUX00 100 (Golden), 200 (Controller), 203 (/backup)
 - Along with the 5 vanilla z/VM volumes, a system could be rebuilt

© 2006 IBM Corporation

Live Demo!

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




© 2006 IBM Corporation

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 - ???



- Mike's email: mikemac@us.ibm.com
- Whitebook: <http://linuxvm.org/present/misc/virt-cookbook-2.pdf>
- Associated files: <http://linuxvm.org/present/misc/virt-cookbook-2.tgz>
- This session: 9217