The Virtualization Cookbook: Day 1 - z/VM

Michael MacIsaac - mikemac@us.ibm.com
Thursday August 17th, 3:00 PM
Session 9216

Abstract

An updated redbook is available: 'z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES9.' The goal of the redbook is to allow you to install and configure z/VM, install and configure Linux and be cloning Linux in two working days. Rexx EXECs and shell scripts are provided with the book to make this goal a reality. This presentation focuses on the first day planning and desktop tools, and then describes a z/VM 5.2 install from DVD. Once installed, the following topics are addressed:

* The SYSTEM CONFIG file
* TCP/IP configuration
* Configuring paging volumes
* Creating a user ID for common files
* Defining a highly available VSWITCH
* System startup and shutdown
* Servicing z/VM
* Monitoring z/VM

The task of setting up a PC NFS server is also addressed. This allows you to supply the EXECs and scripts associated with the book and a Linux install tree.

Who am I?, who are you?

- Mike MacIsaac, 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 z/VM first 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
- **Monitoring z/VM and Linux**
- Backup and restore

Overview

- This presentation is based on a book and associated files
  - *z/VM and Linux on IBM System z: The Virtual Server Cookbook*
- Started as a zBLC working group on **Linux Appliances**
- Tangible goal: From a bare LPAR (with resources):
  - Install and configure z/VM and Linux
  - Be cloning Linux *in 2 working days*
  - Be creating clones on third day
- Many usability tests conducted
  - Most achieved goal of cloning in two days

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 Director, Levanta, Rocket Provisioning Expert 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
  - Cloning and manual install hinge on CMS parameter files
    - Cloning done from Linux, no VM service machine needed

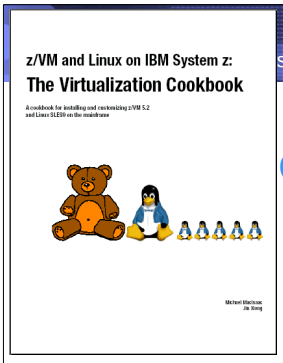

Overview: Redbook, whitebook, red, red, white

5 *The Virtualization Cookbook 2* published on linuxvm.org, 8/06

2 *The Virtualization Cookbook* published on linuxvm.org, 2/06

Project started: 11/04

2004 | 2005 | 2006

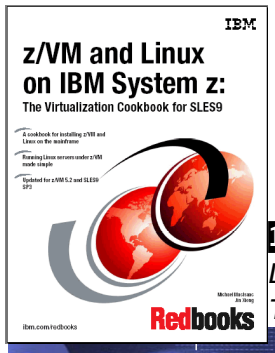

3 Redbook published *The Virtualization Cookbook for SLES9*, SG24-6695-01, 4/06

First draft: 3/05

1 Redbook published *From LPAR to Virtual Servers in Two Days*, SG24-6695-00: 6/05

4 First draft of redbook: *The Virtualization Cookbook for RHEL4*, SG24-7272-00, 7/06

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

Overview: What is new?

- First z/VM customization is IPWIZARD
  - Get off the HMC as soon as possible
- Controller (cloner) is on minidisk 200 not 103
  - New master image disks separate
- Cloning of first Linux manually
  - Less magic, better understanding
- New chapter: WebSphere, DB2 and MQ Series
  - Installation summaries
  - Read-only binaries
  - Disks for production and test
  - Manual cloning
  - Cloning via clone.sh script

Planning - bill of materials

- Hardware
  - zSeries LPAR - 2 IFLs recommended
    - 3GB central:1GB expanded storage (1.5G:512M OK)
    - 37 3390-3 DASD or more
    - 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-9 SP3 Linux distro ISO images (6 vanilla + 3 SPs)
  - Code associated with redbook
  - IBM Middleware (some demo copies available)
- Networking resources
  - TCP/IP address for z/VM
  - 12 TCP/IP addresses for Linux
  - 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

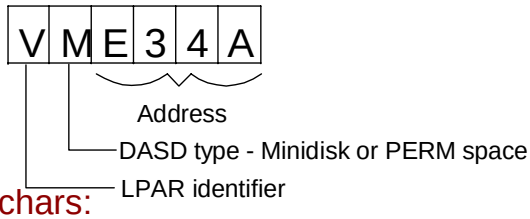

Outline

- ~~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

Configuring a Windows desktop

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

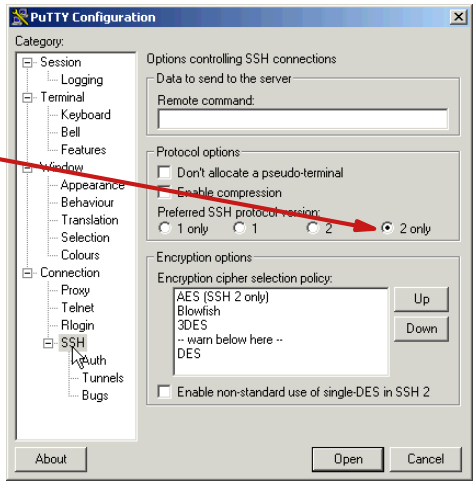

Configuring a Windows desktop (cont'd)

- VNC client or X server
  - RealVNC client is recommended
  - Must set option *Use Only protocol version 3.3* in recent clients

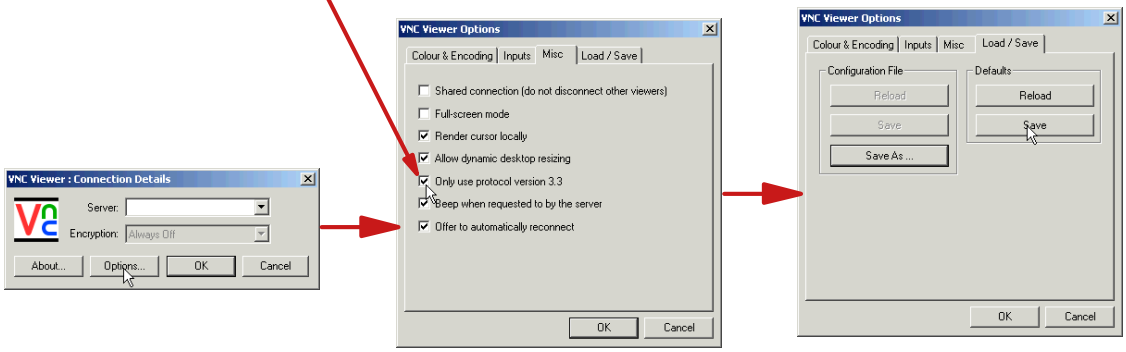

Outline

- ~~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

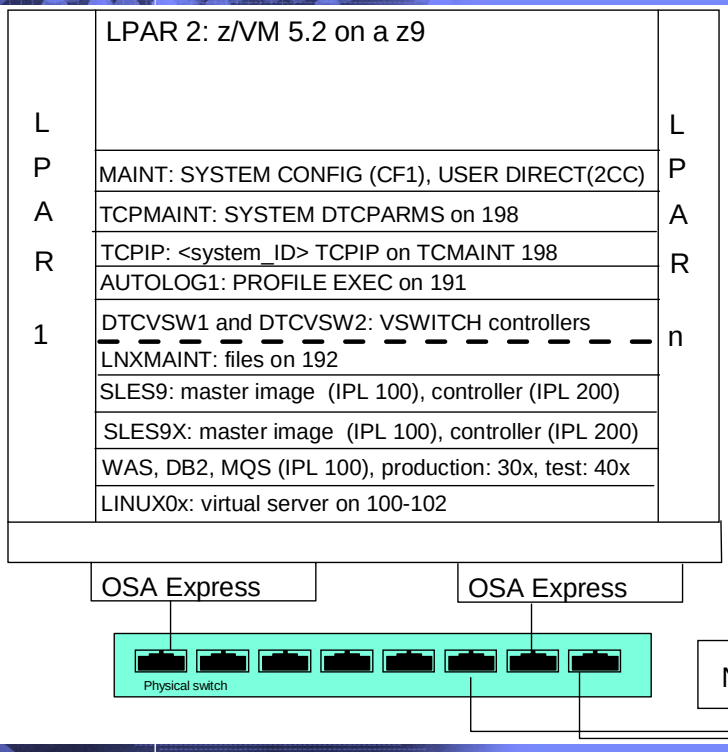

Block Diagram

Resources:

CPU: 2 IFLs, shared
Memory: 3GB/1GB
Disk: 37 3390-3 DASD
Network: 12 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
- Configure z/VM - Overview
  - Customize TCPIP with IPWIZARD
  - Customize PROFILE XEDIT
  - Customize SYSTEM CONFIG
  - Add paging volumes
  - Create a common CMS user ID - LNXMAINT
  - Customize startup and shutdown
  - Address security issues
- Backup system to tape
- Modify system labels
- Restore system

Configure z/VM

- 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, reIPL
- Customize PROFILE XEDIT
  - F12 as Retrieve key not "FILE"
- Customize SYSTEM CONFIG
  - Change system name
  - Allow users to create VDISKs
  - Define a highly available VSWITCH
  - Run CPSYNTAX
- Add paging volumes - 5 new volumes are recommended
  - Format the volumes - CPFORMAT EXEC is written
  - Update SYSTEM CONFIG w/paging volumes and reIPL

CPFORMAT EXEC - 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-------'
```

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

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)
    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 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
- 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**
    - Give Linux IDs access to Vswitch:
      - **set vswitch vsw1 grant <id>**
    - **XAUTOLOG** each Linux ID (but only when needed!)

Configure z/VM (cont'd)

- Security issues
  - CHPW52 XEDIT is written to change USER DIRECT passwords
  - Other considerations (no recipe)
- Backup system to tape (no recipe)
- Modify 520xxx labels?
  - Two issues with duplicate labels
    - Pause z/OS IPL
    - Multiple z/VM systems using same volumes
  - 4/5 volumes can be relabeled during z/VM install
- Restore system from tape (no recipe)

Relabel z/VM system volumes

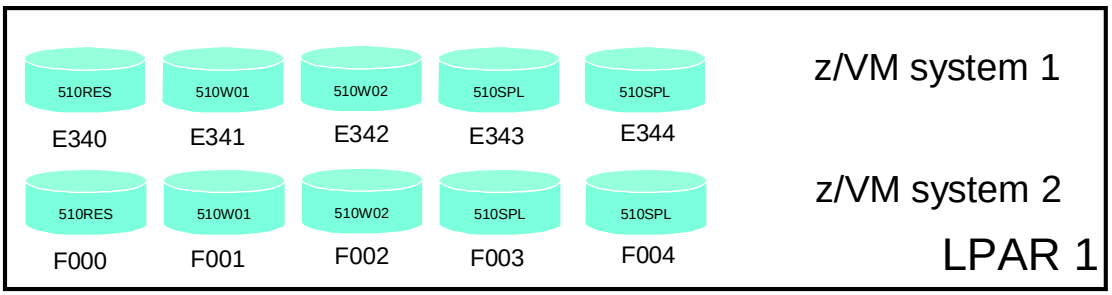

- What happens with "IPL E340"?
- What happens with "IPL F000"?

- 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
- LABEL520 EXEC and LABEL520 XEDIT written to help prevent mistakes

IBM Systems Group

Relabel z/VM system volumes

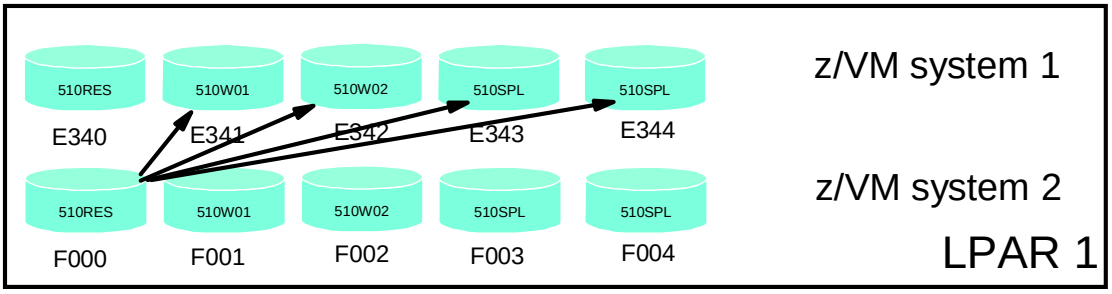

- What happens with "IPL E340"?
- What happens with "IPL F000"?

- 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
- LABEL520 EXEC and LABEL520 XEDIT written to help prevent mistakes

End of 1 day - DASD view

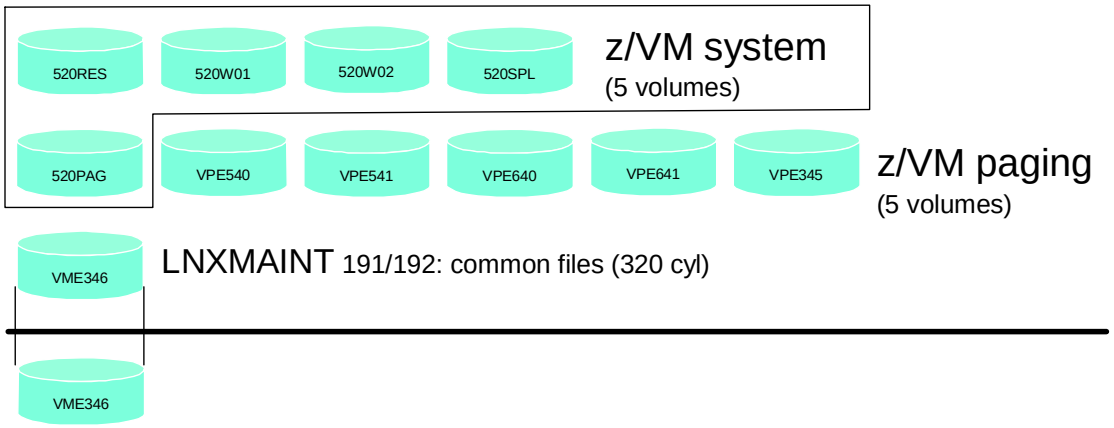

LINUX volumes
(next hour)

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- ~~Installing and configuring z/VM~~
- **Servicing z/VM**
- Configuring an NFS seerver
- 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

Servicing z/VM

- Step 1 - back up your system to tape
- Step 2? - test restoring your system from tape
- z/VM service can be one of two flavors:
  - CORrective in the form of a Programming Temporary Fix (PTF)
    - Think bug fix
  - Recommended Service Upgrade (RSU)
    - Think Service Pack
- z/VM service can be applied in one of two ways:
  - The old way - many steps for each component
    - More laborious, more granular, easier to back out
  - The new way - SERVICE and PUT2PROD
    - Much simpler, more magic, difficult to back out of
- Cookbook steps for:
  - Applying a PTF obtained from the Internet
  - Applying an RSU obtained from the Internet

Outline

- ~~Planning~~
- ~~Configuring a Windows desktop~~
- ~~Installing and configuring z/VM~~
- ~~Servicing and configuring 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

Configure an NFS server

- Installing Linux on zSeries is a *chicken and egg* problem
- Recommendation: install Linux on an Intel PC
  - Server is a temporary NFS server
  - Contents later moved to controller
- Copy files associated with this book to this NFS server
- Copy SLES-9 .iso images + Service Pack .iso images
- Create a SLES9 install tree
  - mksles9root script - looks for 6 SLES9 ISO and 2 SP1 ISO images

```
# ./mksles9root.sh s390x
Making a SLES9 install tree ...
SP3 ISO images found ...
The tree named sles9xsp3root/ will be SLES9 + SP3 ...
Making the directory structure ...
Copying SLES9 ISO images ...
  Mounting and copying SLES-9-s390x-RC5a-CD1.iso ...
```


Configure an NFS server

- Creates a SLES9 install tree (cont'd)

```
# cd /nfs/sles9x/sles9xsp3root
# ls -l
lrwxrwxrwx  1 root root   14 Jan 12 11:35 boot -> sles9/CD1/boot
lrwxrwxrwx  1 root root   17 Jan 12 11:35 content -> sles9/CD1/content
lrwxrwxrwx  1 root root   21 Jan 12 11:35 control.xml -> sles9/CD1/control.xml
drwxr-xr-x  7 root root 4096 Jan 12 11:16 core9
lrwxrwxrwx  1 root root   22 Jan 12 11:35 driverupdate -> sp3-9/CD1/driverupdate
lrwxrwxrwx  1 root root   15 Jan 12 11:35 linux -> sp3-9/CD1/linux
lrwxrwxrwx  1 root root   17 Jan 12 11:35 media.1 -> sles9/CD1/media.1
dr-xr-xr-x  3 root root 4096 Dec 20 07:21 s390x
drwxr-xr-x  3 root root 4096 Jan 12 11:16 sles9
drwxr-xr-x  5 root root 4096 Jan 12 11:16 sp3-9
drwxr-xr-x  2 root root 4096 Jan 12 11:35 yast
# cat yast/instorder
/sp3-9/CD1 /sp3-9/CD1
/sles9/CD1 /sles9/CD1
/core9/CD1 /core9/CD1
```

- Configure NFS
  - Modify /etc/exports

```
# tail -2 /etc/exports
/nfs/virt-cookbook-2            *(ro,sync)
/nfs/sles9x/sles9xsp3root       *(ro,sync)
```

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

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

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 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

Resources

- 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
- Red Hat Enterprise Linux 4 evaluation
  - http://www.redhat.com/rhel/details/eval/
- z/VM publications
  - http://www.vm.ibm.com/pubs/
- z/VM performance tips
  - http://www.vm.ibm.com/perf/tips/

Questions - ???

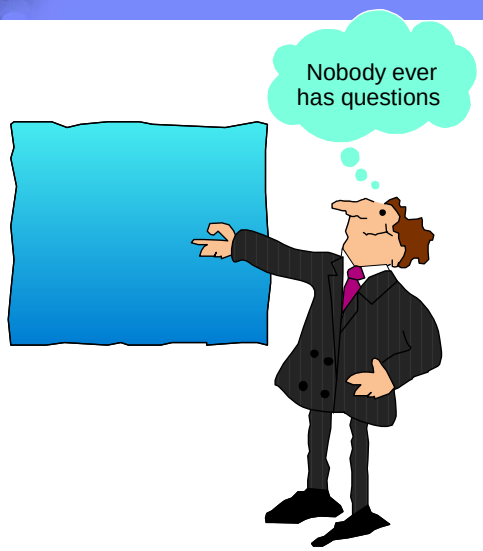

- Mike's email: mikemac@us.ibm.com
- Whitebook: http://linuxvm.org/present/misc/virt-cookbook-2.pdf
- Associated files: http://linuxvm.org/present/misc/virt-cookbook-2.tgz
- This session: 9216