



From LPAR to Virtual Servers in Two Days: Day 2 - Linux

Michael Maclsaac - mikemac@us.ibm.com
Jin Xiong - jinxiong@us.ibm.com
Thursday, August 25th, 3:00 PM
Session 9217



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.

On the second day, you install Linux twice onto the same virtual machine. Think of it as a dual boot PC. The first install is called the golden image which will be cloned from. The second install is called the controller which will do the cloning, but also other tasks. Both of these images are configured with the option of using scripts for speed, or manually to better understand the configuration. All of the data on a PC Linux NFS server which is needed to 'bootstrap' the first two installations is copied over the zSeries controller so the PC can be retired. Then cloning is addressed and the supplied clone script is described. After 'day 2,' you should be ready to clone the golden image and configure the clones into Web, LDAP, file and print, development or other servers. Finally backup and restore issues are addressed.

Who are we?, who are you?

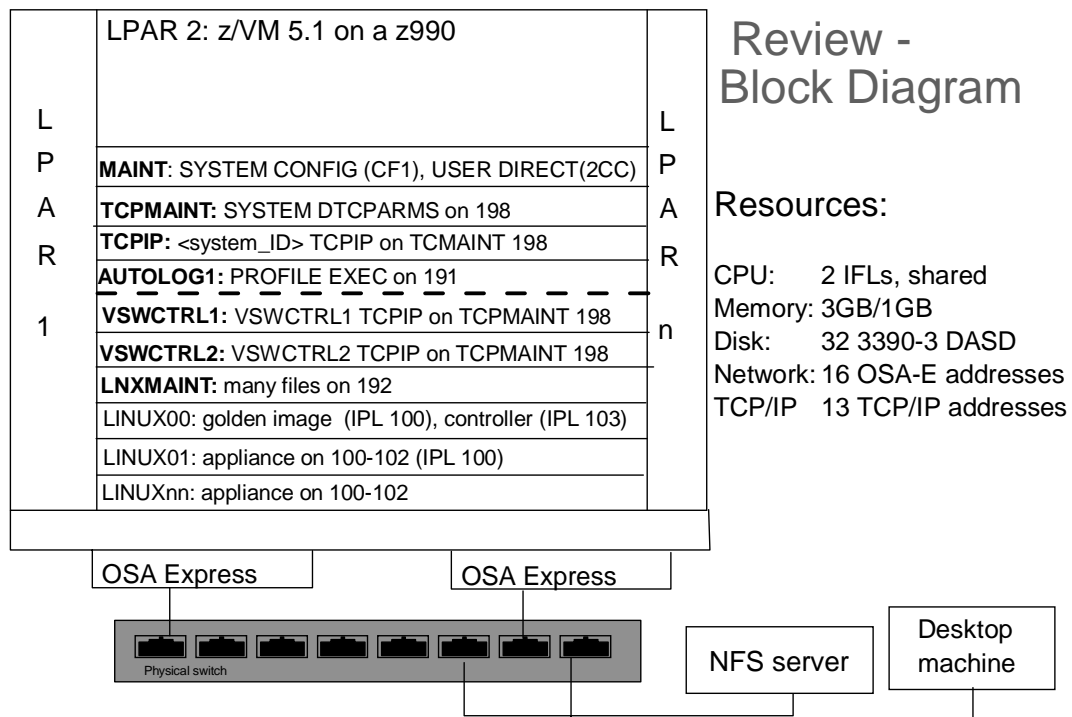
- Mike MacIsaac, mikemac@us.ibm.com
- Jin Xiong, jinxiong@us.ibm.com
- Who are you?
 - ▶ Have you tried the steps in this redbook?
 - ▶ Are you thinking about using this redbook?
 - ▶ Have you never heard about this redbook?
 - ▶ ...
 - ▶ Any other categories of attendees?

Outline for the two presentations

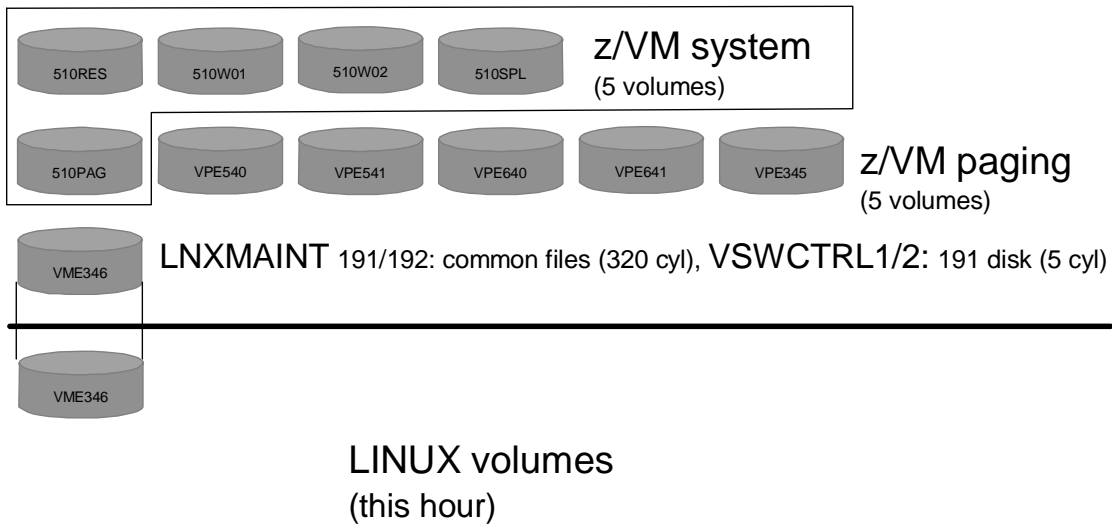
- Planning
- Configuring a Windows desktop
- Installing and configuring z/VM
- Configuring an NFS server
- ---
- **Install and configure Linux**
- **Configure NFS on controller**
- **Configure Linux for cloning - Jin Xiong**
- **Basic Linux virtual servers - Jin Xiong**
- **A virtual communications server**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

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



Review - End of 1 day - DASD view



Install and configure Linux

- Overall steps
 - ▶ Define a user ID LINUX00 - think dual boot CD
 - ▶ Prepare bootstrap files
 - ▶ Install the golden image (IPL 100)
 - Minidisk 100 - root file system - 3038 cylinders (~2GB)
 - Device 101 - VDISK swap space - 256MB
 - Minidisk 102 - swap space - 300 cylinders
 - ▶ Configure the golden image
 - ▶ Install the controller (IPL 103)
 - Minidisk 103 - root file system
 - Swap spaces at 101 and 102 can be reused
 - Minidisk 104 - /backup file system - 3003 cylinders
 - Minidisks 105-108 - logical volume for /nfs file system (~9GB)
 - ▶ Configure the controller

Define user ID LINUX00

```

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 LINUX00 LNX4VM 256M 1G BG
  INCLUDE LNXDFLT
  OPTION LNKNOPAS APPLMON
  MDISK 100 3390 0001 3038 <VME347> MR LNX4VM LNX4VM LNX4VM
  MDISK 102 3390 3039 0300 <VME347> MR LNX4VM LNX4VM LNX4VM
  MINIOPT NOMDC
  MDISK 103 3390 0001 3038 <VME348> MR LNX4VM LNX4VM LNX4VM
  MDISK 104 3390 0331 3003 <VME346> MR LNX4VM LNX4VM LNX4VM
  MDISK 105 3390 0001 3338 <VME349> MR LNX4VM LNX4VM LNX4VM
  MDISK 106 3390 0001 3338 <VME34A> MR LNX4VM LNX4VM LNX4VM
  MDISK 107 3390 0001 3338 <VME34B> MR LNX4VM LNX4VM LNX4VM
  MDISK 108 3390 0001 3338 <VME34C> MR LNX4VM LNX4VM LNX4VM

```

© 2005 IBM Corporation

Install and configure Linux (cont'd)

- Prepare bootstrap files
 - ▶ FTP from SLES9 install tree on Linux PC to LNXMAINT 192 disk:
 - Kernel SLES9 KERNEL
 - Initial RAMdisk SLES9 INITRD
 - ▶ Install the golden image on 100-102
 - Trickier install panels are documented well
 - ▶ Configure the golden image
 - Apply service if necessary - Y.O.U. is documented
 - Copy associate files from NFS server via "scp -r"
 - Script **config-golden** can be used to automate configuration, OR
 - All steps can be done manually
 - Add additional RPMs
 - Turn off unneeded services
 - Configure rsyncd to allow /etc/ to be copied by controller
 - Configure sitar

© 2005 IBM Corporation

Install and configure Linux (cont'd)

- All steps can be done manually (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 golden except
 - ▶ Minidisk 100, /dev/dasda, is /sles9golden (Don't format!!)
 - ▶ Mindisk 104, /dev/dasde, becomes a file system over /backup
 - ▶ Mindisks 105-108 become a logical volume mounted over /nfs
- Configure the controller
 - ▶ Apply service, if necessary
 - ▶ Copy files from NFS server
 - ▶ Automate configuration with script **config-controller**, or do manually
 - Remove some RPMs
 - Add some RPMs, including cpint by Neale Ferguson

Install and configure Linux (cont'd)

- ▶ Manual configuration (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)

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

Configure NFS on controller

- Copy files from PC NFS server"
 - # `scp -r pc.nfs.server.ip:/nfs/sles9root /nfs`
 - # `scp -r pc.nfs.server.ip:/nfs/lpar2vs /nfs`
- Configure NFS server
 - ▶ Same steps as on PC server
- Change YaST install tree location
 - ▶ Change source on the golden image
 - # `chroot /sles9golden`
 - `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

Configure Linux for cloning

- Define new user ID, LINUX01
 - ▶ Add new user to USER DIRECT
 - ▶ Add new Linux ID to AUTOLOG1 PROFILE EXEC
 - ▶ Create a parameter file
- Clone one new virtual server
- Clone six new virtual servers - LINUX02-LINUX07
- The clone script - Overview
 - ▶ Check logged off, ask are you sure
 - ▶ Reset root password
 - ▶ Copy disks via FLASHCOPY, else Linux dasdfmt/dd
 - ▶ Modify via PARMFILE settings
 - ▶ XAUTOLOG new appliance - detect on network
 - ▶ Modify SSH keys for key-based authentication

Cloning details

```

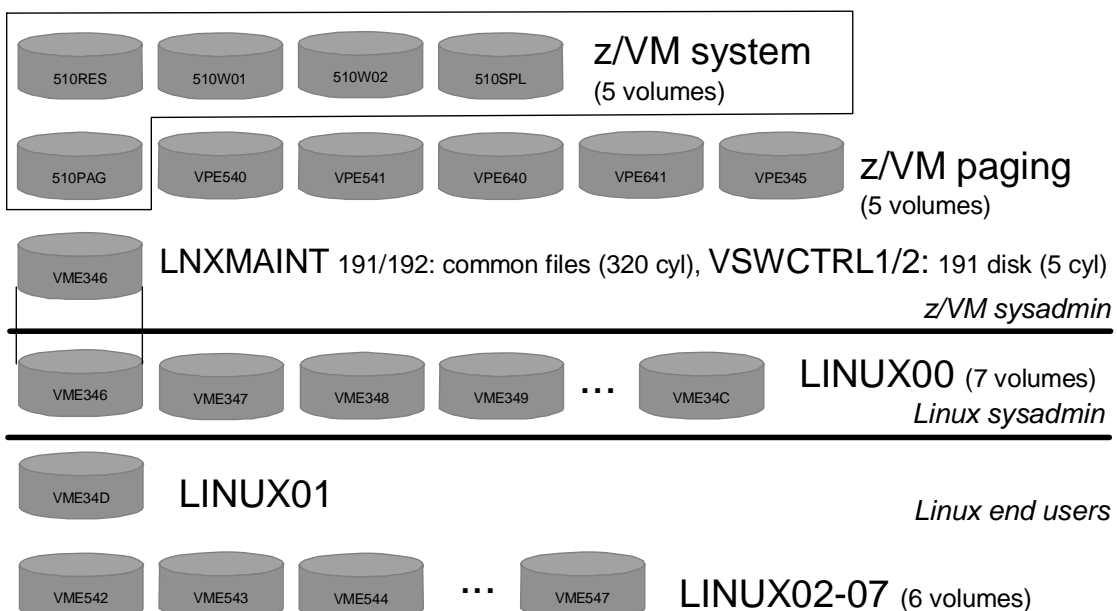
# main()
gold_mnt_pt="/sles9golden"
backup_dir="/backup/linux"
if [ $# != 1 ]; then
    echo "Error: incorrect number of arguments"
    help
fi
new_linux_id=`echo $1 | tr 'abcdefghijklmnopqrstuvwxyz' 'ABCDEFGHIJKLMNOPQRSTUVWXYZ'`
check_logged_off
check_for_parmfile
get_target_info
get_source_info
link_minidisk $new_linux_id 100 1100
link_minidisk $new_linux_id 102 1102
ask_are_you_sure

# modify networking info, SSH keys, root password under /sles9golden
modify_golden_image
# copy the 102 disk then the 100 sleeping in between if FLASHCOPY succeeds
copy_102_100
# XAUTOLOG the new appliance and wait for a ping to answer, then undo changes under /sles9golden
wait_for_sign_of_life
on_network=$?
restore_golden_image
# if the new appliance is up, make a backup directory and add new SSH key to known_hosts file
if [ $on_network = 1 ]; then
    make_backup_dir
    update_known_host
fi

```

© 2005 IBM Corporation

Review - DASD and role view



© 2005 IBM Corporation

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

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

A virtual Web server

- Clone a server and install Apache2 RPMs
 - ▶ use the script or do it manually
- 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
- After install, retrieve scripts and run `config-golden`
- 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
- Integrated development environment (IDE)
 - ▶ The Eclipse application development platform
 - Very popular, is the basis for many other apps
 - Can be used for virtually any type of development, supports many programming languages with plug-ins
 - ▶ Installing Eclipse for Java and C/C++ development
 - Eclipse is not officially supported on Linux for zSeries
 - Build from source works on 31-bit platform only


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

A virtual communications server

- Install IBM Communications Server for Linux (CS Linux)
 - ▶ SLES9 SP1 is required
 - ▶ Many kernel-source type RPMs are required
 - ▶ Linux Streams module must be built
 - ▶ Then CS Linux can be built
- Configure CS Linux
 - ▶ Set some environment variables
 - ▶ Start SNA service
- CS Linux example - an EE link
 - ▶ Graphical config tool xsnaadmin is briefly described
- Documentation can be found starting at:
 - ▶ <http://www-306.ibm.com/software/network/commserver/library/>

CS Linux Examples



- In the first example, CS Linux is installed on a separate z800 system to offload the main z/OS system. An Enterprise Extender link using IP, or an LLC2 link, is used to connect the two systems.
- In the second example, CS Linux is installed on one or more VMs or LPARs in the main z/OS system. Although CS Linux and z/OS Communications Server are on the same mainframe, they are two separate SNA nodes, and so an Enterprise Extender link using HyperSockets IP or an LLC2 link is still required between them.

© 2005 IBM Corporation

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

© 2005 IBM Corporation

A virtual communications controller server

- Install IBM Communication Controller for Linux (CC Linux)
 - ▶ SLES9 SP1 is required
 - ▶ Copy product code to virtual server
 - ▶ Install kernel source and other RPMs
 - ▶ Run CC Linux install program - **setuptools390.bin**
- Configure OSA card(s)
 - ▶ Modify original NCP module
 - ▶ Build the CC Linux load module
 - ▶ Transfer the load module to Linux
 - ▶ Start the CC Linux engine - **cclengine** command
- Documentation can be found starting at:
 - ▶ <http://www-306.ibm.com/software/network/ccl/>
 - ▶ <http://www.ibm.com/support/docview.wss?rs=2192&uid=swg27005786>

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~~

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 (on 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)

- ▶ These files can be backed up nightly via cron - one method (cont'd)

- A simple script, backup_vm, is written:

```
# tail -7 /etc/cron.daily/backup_vm
cd /backup/vm
getVMInfo > zVMInfo.txt
ftp ftp://maint:<lnx4vm>@<129.40.178.36> < /etc/cron.daily/MAINT.FTPcommands
ftp ftp://tcpmaint:<lnx4vm>@<129.40.178.36> < /etc/cron.daily/TCPMAINT.FTPcommands
ftp ftp://autolog1:<lnx4vm>@<129.40.178.36> < /etc/cron.daily/AUTOLOG1.FTPcommands
cd LNXMAINT
ftp ftp://lnxmaint:<lnx4vm>@<129.40.178.36> < /etc/cron.daily/LNXMAINT.FTPcommands
```

- A sample "FTP commands" file: (cont'd)

```
# cat /etc/cron.daily/MAINT.FTPcommands
ascii
get USER.DISKMAP
cd maint.2cc
get USER.DIRECT
cd maint.cf1
get SYSTEM.CONFIG
quit
```

- Incremental backup of Linux
 - ▶ Recall that a directory /backup/LINUX-on-<tcp.ip.address> directory was created for each Linux system cloned
 - ▶ Recall that key-based SSH authentication was set up to the clones

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

```
# cat /etc/cron.daily/run-sitar
#!/bin/bash
sitar --format=html --outfile=/etc/sitar.html
```

- So the output of sitar is also backed up

- Disaster recovery

- ▶ Follow good s/390 procedures
- ▶ Consider that three volumes (VME346-8 in book) contain
 - LNXMAINT 192 (Common CMS files)
 - LINUX00 100 (Golden), 103 (Controller), 104 (/backup)
 - Along with the 5 vanilla z/VM volumes, a system could be rebuilt

Directories and files in tar file

```
doc/ - Important VM, Linux documents
linux-controller/ - Files for Linux controller
|-- etc/Bastille/config - Bastille config file
| |-- cron.daily/ - Files for backing up nightly
| `-- ntp.conf.template - Template NTP config file
`-- usr/local/sbin/clone - The clone script
    `-- config-controller - Controller quick configuration
linux-golden/ - Files for Linux golden image
|-- etc/Bastille/config - Bastille config file
| |-- cron.daily/ - Files to run nightly
| `-- rsyncd.conf - To allow controller rsync access
`-- usr/local/sbin/ - Various configuration files
    `-- config-golden - Golden image quick configuration
nfs-server/ - MD5SUM files, script
`-- mksles9root - Make install tree script
vm/ - z/VM files
|-- CHANGEPW.XEDIT - Change passwords in USER DIRECT
|-- CPFORMAT.EXEC - Format many disks with CPFMTXA
|-- CPLABEL.EXEC - Label many disks with CPFMTXA
|-- DISKMAPP.EXEC/XEDIT - Quick check diskmap command
|-- LABEL510.EXEC/XEDIT - Relabel a 5.1 system
|-- LINUX00.PARMPFILE - Template parameter file
|-- PROFILE.EXEC - Appliance's PROFILE EXEC
|-- PROFILE.EXECAUT1 - AUTOLOG1 PROFILE EXEC
|-- SLES9.EXEC - SLES9 install EXEC
| `-- Other z/VM files
```

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 - ???

