



IBM Systems Group

Cloning WebSphere, DB2 and WebSphere MQ Series on Linux under z/VM

Michael MacIsaac - mikemac@us.ibm.com

Tuesday, August 15th, 1:30 PM

Session 1189/2954/9310

© 2006 IBM Corporation

IBM Systems Group



Abstract

Creating WebSphere Application Server, DB2 and WebSphere MQ Series servers typically involves installing an operating system, then installing and configuring the software. The data associated with each solution also resides on the same server. When it is time to upgrade to a new release, each server must be upgraded individually. As the number of these types of servers increases, upgrading becomes cumbersome. This presentation describes how to install the middleware's executable files on virtual Linux servers under z/VM: call these the master servers. Other virtual servers can be cloned and configured to link the master servers' disks read-only: call these the clones. In this fashion the clones access the executables read-only while the application data is read-write. **When a new version of middleware needs to be rolled out, it can be installed onto another set of disks on the masters for testing purposes. When that updated version of middleware is deemed acceptable, the clones can easily switch to the new level.**

© 2006 IBM Corporation



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?
 - ▶ =====
 - ▶ WebSphere Application Server in production on any platform?
 - ▶ DB2 in production?
 - ▶ MQ Series in production?

© 2006 IBM Corporation



Outline for the presentation

- Overview
- Roles
- Setting up a cloning infrastructure
 - ▶ z/VM
 - ▶ Linux
- WebSphere cloning
 - ▶ Installing on the "master"
 - ▶ Setting up a clone manually
 - ▶ Automating cloning
- DB2 cloning
 - ▶ Installing on the "master"
 - ▶ Setting up a clone manually
 - ▶ Automating cloning
- MQ Series cloning
 - ▶ Installing on the "master"
 - ▶ Setting up a clone manually
 - ▶ Automating cloning

© 2006 IBM Corporation



Overview: assumptions

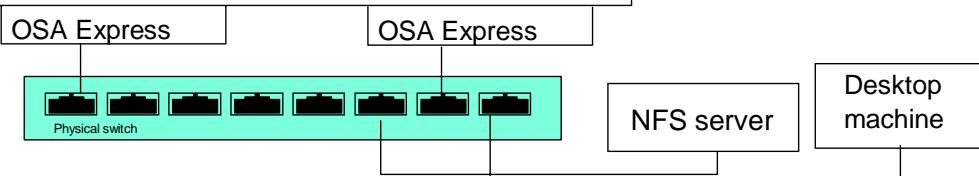
- Discusses some IBM middleware products
 - ▶ WebSphere Application Server V6.0.2 (V6.1?)
 - ▶ DB2 V8.2
 - ▶ MQ Series V6.0
- Install binaries on a master server
- Link binaries R/O and create application data R/W on clones
 - ▶ WebSphere uses profiles
 - ▶ DB2 uses instances
 - ▶ MQ Series uses queue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

© 2006 IBM Corporation

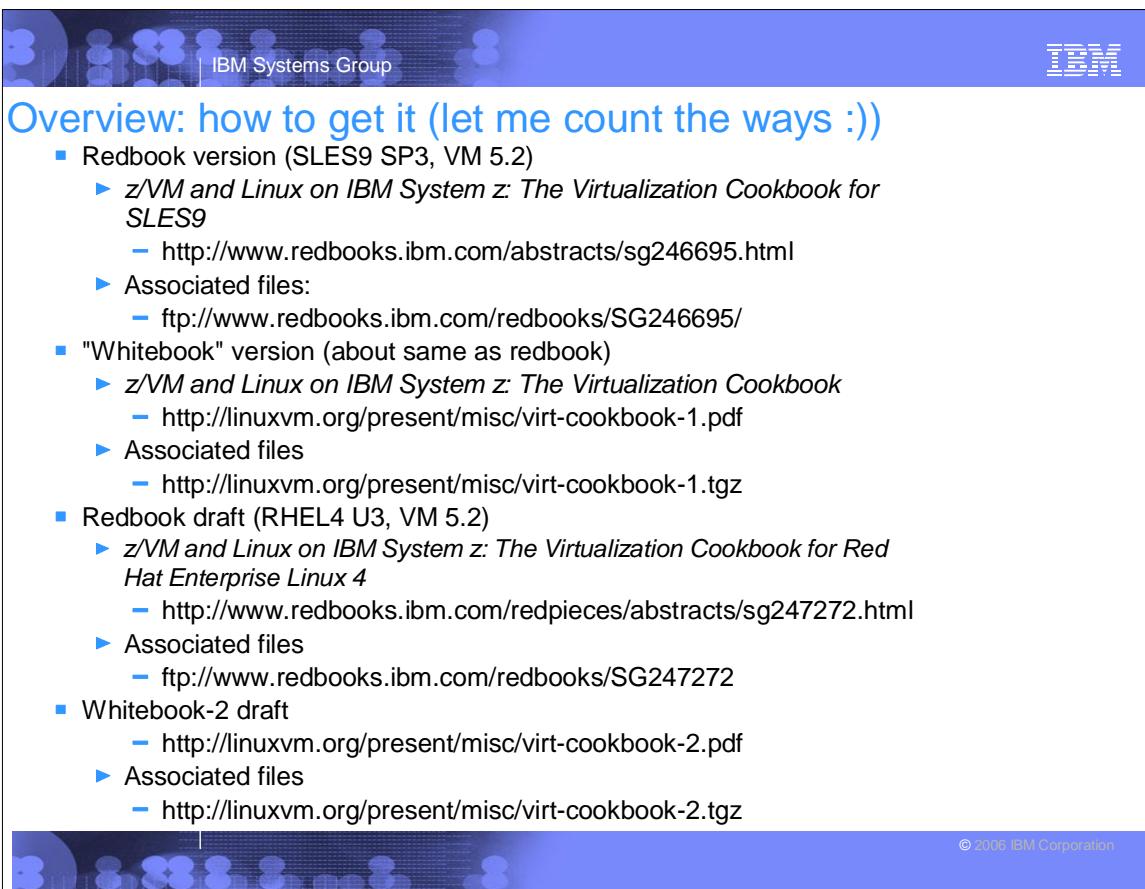
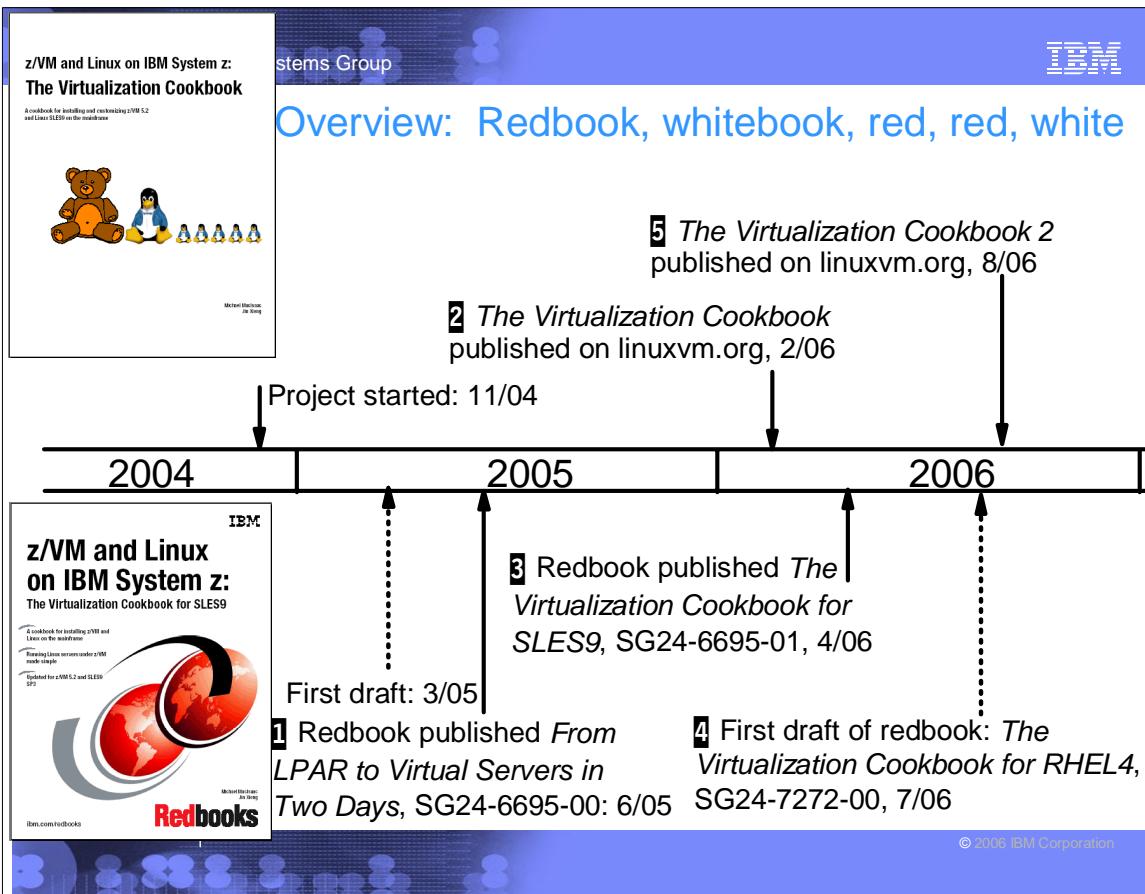
Overview: Infrastructure Block Diagram

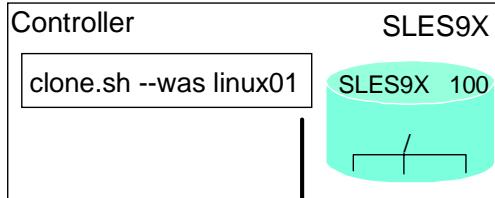
L	LPAR 2: z/VM 5.2 on a z990	L
P	MAINT: SYSTEM CONFIG (CF1), USER DIRECT(2CC)	P
A	TCPMAINT: SYSTEM DTCPARMS on 198	A
R	TCPIP: <system_ID> TCPIP on TCMAINT 198	R
	AUTOLOG1: PROFILE EXEC on 191	
1	DTCVSW1 and DTCVSW2: VSWITCH controllers	
	LNXMAINT: files on 192	n
	SLES9: master image (IPL 100), controller (IPL 200)	
	SLES9X: master image (IPL 100), controller (IPL 200)	
	LINUX0n: virtual server on 100-102 (IPL 100)	
	WAS, DB2, MQS: IPL 100, 30x: production, 40x: test	

For details see:
sessions 9216-7
Thurs@3:00 & 4:30
Room 321



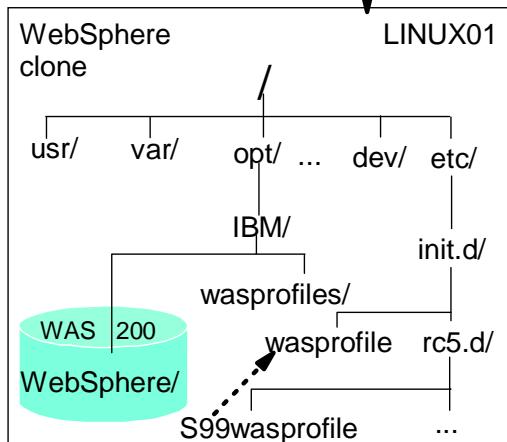
© 2006 IBM Corporation



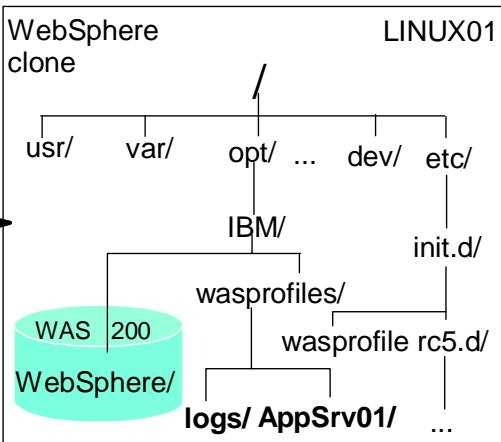


Overview: Cloning process block diagram

Before first IPL:



After first IPL:



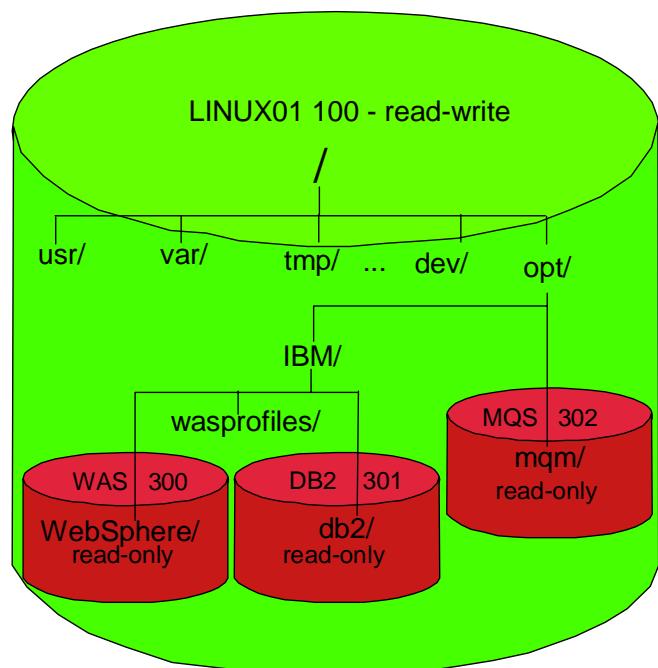
© 2006 IBM Corporation

Overview: User IDs, disks and mount points

Software	z/VM User ID of "master"	Important minidisks (production and test)	Mount point on clones
Linux	SLES9 or SLES9X	100-102: master image 200: controller	/sles9xmaster /
WebSphere	WAS	300 and 400	/opt/IBM/WebSphere/
DB2	DB2	301 and 401	/opt/IBM/db2/
MQ Series	MQS	302 and 402	/opt/mqm/

© 2006 IBM Corporation

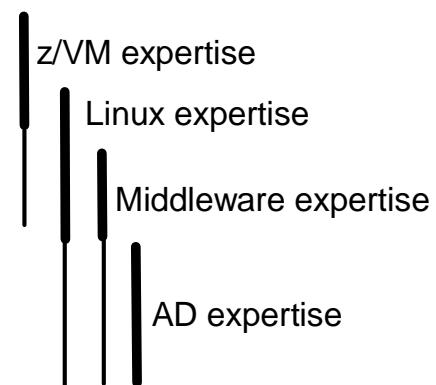
Overview: Cloned system block diagram



© 2006 IBM Corporation

Roles

- Many possibilities
 - ▶ z/VM sysadmin (MAINT, etc)
 - ▶ z/VM Network admin (TCPMAINT)
 - ▶ Linux sysadmin (LNXMAINT, SLES9X, SLES9)
 - ▶ Middleware admins
 - WebSphere admin (WAS)
 - DB2 admin (DB2)
 - MQ Series admin (MQS)
 - ▶ Application owners
 - WebSphere (LINUXnn)
 - DB2 (LINUXnn)
 - MQ Series (LINUXnn)



© 2006 IBM Corporation

Setting up: z/VM

- Define 3 user IDs for masters

```

USER WAS      WAS      512M 1G G
INCLUDE LNXDFLT
OPTION APPLMON
MDISK 100 3390 0001 3038 VMA78B MR LNX4VM LNX4VM LNX4VM
MDISK 102 3390 3039 0300 VMA78B MR LNX4VM LNX4VM LNX4VM
MINIOPT NOMDC
MDISK 300 3390 0001 3338 VMA78C MR LNX4VM LNX4VM LNX4VM
MDISK 400 3390 0001 3338 VMA78D MR LNX4VM LNX4VM LNX4VM
USER DB2      DB2      512M 1G G
INCLUDE LNXDFLT
OPTION APPLMON
MDISK 100 3390 0001 3038 VMA78E MR LNX4VM LNX4VM LNX4VM
MDISK 102 3390 3039 0300 VMA78E MR LNX4VM LNX4VM LNX4VM
MINIOPT NOMDC
MDISK 301 3390 0001 1669 VMA78F MR LNX4VM LNX4VM LNX4VM
MDISK 401 3390 1670 1669 VMA78F MR LNX4VM LNX4VM LNX4VM
USER MQS      MQS      512M 1G G
INCLUDE LNXDFLT
OPTION APPLMON
MDISK 100 3390 0001 3038 VMA790 MR LNX4VM LNX4VM LNX4VM
MDISK 102 3390 3039 0300 VMA790 MR LNX4VM LNX4VM LNX4VM
MINIOPT NOMDC
MDISK 302 3390 0001 1669 VMA791 MR LNX4VM LNX4VM LNX4VM
MDISK 402 3390 1670 1669 VMA791 MR LNX4VM LNX4VM LNX4VM

```

© 2006 IBM Corporation

Setting up: z/VM (cont'd)

- Define user IDs for clones

```

USER LINUX05 LNX4VM 512M 1G EG
INCLUDE LNXDFLT
OPTION APPLMON
MDISK 100 3390 0001 3038 VMA788 MR LNX4VM LNX4VM LNX4VM
MDISK 102 3390 3039 0300 VMA788 MR LNX4VM LNX4VM LNX4VM
MINIOPT NOMDC
LINK WAS 300 300 RR
LINK WAS 400 400 RR
LINK DB2 301 301 RR
LINK DB2 401 401 RR
LINK MQS 302 302 RR
LINK MQS 402 402 RR
*
USER LINUX06 LNX4VM 512M 1G G
...
USER LINUX07 LNX4VM 512M 1G G
...

```

- Bring changes online
 - DISKMAP
 - DIRECTXA

© 2006 IBM Corporation

Setting up: Linux master image

- Set new clones to have Java, compat libraries and empty mount points
- Modify the master image with **chroot**:
 - ▶ Add RPMs


```
# chroot /sles9master/
# yast -i IBMJava2-JRE IBMJava2-SDK compat compat-32bit
...
```
 - ▶ Create mount points


```
# mkdir -p /opt/IBM/WebSphere /opt/IBM/db2 /opt/mqm
# exit
```
 - ▶ Could also be done *after* cloning
 - To keep all clones as lean as possible
- Clone a new Linux


```
# clone.sh WAS
```

© 2006 IBM Corporation

WebSphere cloning - installing the master

- Install WebSphere Network Deployment V6.0.2 on master image
 - ▶ SSH into new Linux running on WAS user ID
 - ▶ Activate minidisk 300, format, partition and make into file system:


```
# chccwdev -e 0.0.0300
# dasdfmt -b 4096 -y -f /dev/dasdd
# fdasd -a /dev/dasdd
```
 - ▶ Make into file system with label "was-prod"


```
# mke2fs -j -L was-prod /dev/dasdd1
```
 - ▶ Modify /etc/zipl.conf - add two new minidisks - run zipl


```
[ip1]
target = /boot/zipl
image = /boot/image
ramdisk = /boot/initrd,0x1000000
parameters = "root=/dev/dasd1 dasd=100-102,300,400 selinux=0 TERM=dumb elevator=cfq"
# zipl
```
 - ▶ Modify /etc/fstab - mount one minidisk (production)


```
/dev/dasd1      /
/dev/dasdb1     swap      swap      pri=42      0  0
/dev/dasdc1     swap      swap      pri=42      0  0
LABEL=was-prod  /opt/IBM/WebSphere ext3      acl,user_xattr  0  0
devpts         /dev/pts   devpts   mode=0620,gid=5  0  0
...
```

© 2006 IBM Corporation

Installing the master WebSphere (cont'd)

- ▶ Reboot - new production disk should be mounted over /opt/IBM/WebSphere/


```
# reboot
# exit
... (new SSH session a minute later)
# mount | grep opt
/dev/dasdd1 on /opt/IBM/WebSphere type ext3 (rw,acl,user_xattr)
```
- ▶ Start VNC server for a graphical environment


```
# vncserver
```
- ▶ Install WebSphere V6 but don't create a profile (see book for details)
- ▶ Install RefreshPack 2 to get to V6.0.2
- ▶ Modify new wasprofile.properties file with r/w profiles directory


```
# cd /opt/IBM/WebSphere/AppServer/properties
# vi wasprofile.properties
...
WS_CMT_LOG_HOME=/opt/IBM/wasprofiles/logs
...
WS_PROFILE_REGISTRY=/opt/IBM/wasprofiles/properties
...
```
- ▶ Shutdown and logoff the master WebSphere system!
 - Don't write to file system that is R/O to clones
 - Save resources - especially memory

© 2006 IBM Corporation

Cloning WebSphere manually

- ▶ From the controller, clone a *vanilla* Linux


```
# clone.sh linux05
...
# chccwdev -e 0.0.0300
```
- ▶ SSH into new server and activate WebSphere production disk (300)


```
# chccwdev -e 0.0.0300
```
- ▶ Modify /etc/zipl.conf to access disks read-only and run zipl


```
[zipl]
target = /boot/zipl
image = /boot/image
ramdisk = /boot/initrd,0x1000000
parameters = "root=/dev/dasda1 dasd=100-102,300(ro),400(ro) selinux=0 TERM=dumb
...
# zipl
```
- ▶ Modify /etc/fstab to mount disks read-only

/dev/dasda1	/	ext3	acl,user_xattr	1 1
/dev/dasdb1	swap	swap	pri=42	0 0
/dev/dasdcl	swap	swap	pri=42	0 0
LABEL=was-prod	/opt/IBM/WebSphere	ext2	ro,acl,user_xattr	0 0
devpts	/dev/pts	devpts	mode=0620,gid=5	0 0
proc	/proc	proc	defaults	0 0
sysfs	/sys	sysfs	noauto	0 0

© 2006 IBM Corporation

Cloning WebSphere manually (cont'd)

- ▶ Mount new r/o file system


```
# mount /opt/IBM/WebSphere
# mount | grep opt
/dev/dasd01 on /opt/IBM/WebSphere type ext2 (ro,acl,user_xattr)
```
- ▶ Reboot to test changes


```
# reboot
```
- ▶ Copy the wasprofile.properties file to the r/w /opt/IBM/wasprofiles/ directory


```
# cd /opt/IBM/wasprofiles
# mkdir properties
# cd properties
# cp /opt/IBM/WebSphere/AppServer/properties/wasprofile.properties .
```
- ▶ Add the WebSphere bin/ directory to the PATH


```
# cat /root/.bash_profile
export PATH=$PATH:/opt/IBM/wasprofiles/AppSrv01/bin
```

© 2006 IBM Corporation

Cloning WebSphere manually (cont'd)

- ▶ Create a script that is a wrapper around the wasprofile.sh command


```
# cat /usr/local/sbin/wasprofile
#!/bin/bash
HOST=`hostname -s`
/opt/IBM/WebSphere/AppServer/bin/wasprofile.sh \
-create \
-profileName AppSrv01 \
-profilePath /opt/IBM/wasprofiles/AppSrv01 \
-templatePath /opt/IBM/WebSphere/AppServer/profileTemplates/default \
-nodeName "$HOST"Node01 \
-cellName "$HOST"Node01Cell \
-hostName "$HOST".pbm.ihost.com
```
- ▶ Run the wasprofile script and view the resulting directory


```
# wasprofile
INSTCONFSUCCESS: Success: The profile now exists
# ls /opt/IBM/wasprofiles/AppSrv01
bin/      etc/      installableApps/  installedConnectors/  properties/  wstemp/
config/   firststeps/  installedApps/  logs/          temp/
```
- ▶ Start the server


```
# startServer.sh server1
...
ADMU3000I: Server server1 open for e-business; process id is 1340
```

© 2006 IBM Corporation

Test WebSphere

- http://129.40.178.133:9080/snoop

Snoop Servlet - Request/Client Information

Requested URL:

Servlet Name:

Request Information:

Request method	GET
Request URI	/snoop
Request protocol	HTTP/1.1
Servlet path	/snoop
Path info	<none>
Path translated	<none>
Character encoding	<none>
Query string	<none>
Content length	<none>
Content type	<none>
Server name	129.40.178.133

Find: vien Find Next Find Previous Highlight all Match case Done

IBM Systems Group

IBM

Automating cloning of WebSphere

- Create new script in master image's services directory


```
# cd /sles9master/etc/init.d/
# cat wasprofile
#!/bin/bash
echo "$0 - Creating a WAS profile ..." | tee /tmp/wasprofile.log
host=`hostname -s`
domain=`hostname -d`
/opt/IBM/WebSphere/AppServer/bin/wasprofile.sh -create \
-profileName AppSrv01 \
-profilePath /opt/IBM/wasprofiles/AppSrv01 \
-templatePath /opt/IBM/WebSphere/AppServer/profileTemplates/default \
-nodeName ${host}Node01 \
-cellName ${host}Node01Cell \
-hostName ${host}.${domain}
echo "Return code from wasprofile.sh = $" | tee -a /tmp/wasprofile.log
symlink="/etc/init.d/rc5.d/S99wasprofile"
if [ -h $symlink ]; then # remove symlink so script runs once
  rm $symlink
  echo "Removed symbolic link $symlink" | tee -a /tmp/wasprofile.log
fi
/opt/IBM/wasprofiles/AppSrv01/bin/startServer.sh server1 | \
tee -a /tmp/wasprofile.log
...

```
- Create symbolic link


```
# cd rc5.d
# ln -s ../../wasprofile S99wasprofile
```

IBM Systems Group

© 2006 IBM Corporation



IBM

Automating cloning of WebSphere (cont'd)

- ▶ Add a function to the clone.sh script

```
#+-----+
function clone_was()
#   Arg 1: Mount point of the newly cloned server
# Clone a WebSphere Application Server
#+-----+
{
    echo "Cloning WebSphere ..."
    echo "Modifying zipl.conf and running zipl ..."
    cd $1/etc
    cp zipl.conf zipl.conf.orig
    sed -i -e 's:dasd=100-102:dasd=100-102,300(ro),400(ro):g' zipl.conf
    chroot $1 zipl
    echo "Modifying fstab ..."
    cp fstab fstab.orig
    sed -i -e '4a LABEL=was-prod      /opt/IBM/WebSphere    ext2
    ro,acl,user_xattr      0 0' fstab
    echo "making symlink to /etc/init.d/wasprofile ..."
    cd $1/etc/init.d/rc5.d
    ln -s ../wasprofile S99wasprofile
    cd
}
```

© 2006 IBM Corporation



IBM

Clone a WebSphere server with automation

- ▶ Use the **clone.sh** script with the **--was** flag

```
# clone.sh --was linux05
...
WARNING!!: this will copy 100 and 102 disks to LINUX01 100 and 102
New host name will be: lat135.pbm.ihost.com
New TCP/IP address will be: 129.40.178.135
Other network data is retrieved from LINUX01 PARMFILE on 191 disk
A WebSphere system will be cloned
Are you sure you want to overwrite these disks (y/n): y
...
▶ Jump to the 3270 console after clone is XAUTOLOGed
...
/etc/init.d/rc5.d/S99wasprofile - Creating a WAS profile ...
Jun 22 12:33:13 lat133 kernel: eth0: no IPv6 routers present
INSTCONFSUCCESS: Success: The profile now exists.
Return code from wasprofile.sh = 0
Removed symbolic link /etc/init.d/rc5.d/S99wasprofile
ADMU0116I: Tool information is being logged in file
        /opt/IBM/wasprofiles/AppSrv01/logs/server1/startServer.log
ADMU0128I: Starting tool with the AppSrv01 profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 2118
...
▶ Test WebSphere
```

© 2006 IBM Corporation



DB2 cloning

- Set up the DB2 master from the controller
 - ▶ Clone a vanilla server to the DB2 master user ID


```
# clone.sh db2
...
#
```
- SSH into the master
 - ▶ Configure minidisks 301 and 401 read/write
 - Enable - **chccwdev**
 - Format - **dasdfmt**
 - Partition - **fdasd**
 - Make a file system - **mke2fs -j -L <label>**
 - Labels are **db2-prod** and **db2-test**
 - ▶ Configure /etc/zipl.conf and /etc/fstab to mount db2-prod (301) disk over /opt/IBM/db2


```
# mount | grep opt
/dev/dasdd1 on /opt/IBM/db2 type ext3 (ro,acl,user_xattr)
```
 - ▶ Reboot to test changes
- Shut down the DB2 master

© 2006 IBM Corporation

Clone a DB2 server manually

- Clone a vanilla server to the target clone user ID


```
# clone.sh linux06
...
#
```
- SSH into new clone
 - ▶ Enable disk **db2-prod** read-only
 - Modify /etc/zipl.conf, /etc/zipl.conf, run zipl, reboot to test
 - ▶ Create DB2 users and groups
 - For the instance **db2inst1**

```
# groupadd db2grp1
# mkdir /tmp/noskel
# useradd -g db2grp1 -m -k /tmp/noskel db2inst1
# passwd db2inst1
```
 - For the DB2 fenced user **db2fenc1**
 - For the DB2 administration user **db2admn1**
 - ▶ Create an instance


```
# touch /home/db2inst1/.profile
# /opt/IBM/db2/V8.1/instance/db2icrt -u db2fenc1 db2inst1
DBI1070I Program db2icrt completed successfully.
```
 - ▶ Create an administrative user


```
# touch /home/db2admn1/.profile
# /opt/IBM/db2/V8.1/instance/dascrt db2admn1
...
```

© 2006 IBM Corporation

Test the DB2 server

- Change user to db2inst1, create a sample database and query it

```
# su - db2inst1
db2inst1> db2sampl
Connect to sample database query a table:
db2inst1> db2
...
db2 => connect to sample
Database server      = DB2/LINUXZ64 8.2.0
SQL authorization ID = DB2INST1
Local database alias = SAMPLE

db2 => select * from staff where dept = 20
```

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
10	Sanders	20	Mgr	7	18357.50	-
20	Pernal	20	Sales	8	18171.25	612.45
80	James	20	Clerk	-	13504.60	128.20
190	Sneider	20	Clerk	8	14252.75	126.50

© 2006 IBM Corporation

Automating cloning of DB2

- Create a script in master image's services directory

```
# cat /sles9master/etc/init.d/db2instance
#!/bin/bash
echo "Creating a DB2 instance ..." | tee /tmp/db2instance.log
mkdir /tmp/noskel          # create a DB2 instance user and group
groupadd db2grp1
useradd -g db2grp1 -m -k /tmp/noskel db2inst1
touch /home/db2inst1/.profile
groupadd db2fgrp1           # create a DB2 fenced user and group
useradd -g db2fgrp1 -m -k /tmp/noskel db2fenc1
groupadd db2agrp1           # create a DB2 administrative user and group
useradd -g db2agrp1 -m -k /tmp/noskel db2admn1
touch /home/db2admn1/.profile
/opt/IBM/db2/V8.1/instance/db2icrt -a SERVER -s ese -u db2fenc1 -p db2c_db2inst1 db2inst1
echo "Return code from db2icrt = $" | tee -a /tmp/db2instance.log
/opt/IBM/db2/V8.1/instance/dasrcrt -u db2admn1
echo "Return code from dasrcrt = $" | tee -a /tmp/db2instance.log
cp /opt/IBM/db2/V8.1/das/dasprofile /home/db2admn1/das
su - db2inst1 -c "db2iauto -on db2inst1"
echo "Return code from db2iauto = $" | tee -a /tmp/db2instance.log
symlink="/etc/init.d/rc5.d/S98db2instance"
if [ -h $symlink ]; then # remove symlink so script runs once
    rm $symlink
    echo "Removed symbolic link $symlink" | tee -a /tmp/db2instance.log
fi
su - db2inst1 -c "db2start"
echo "Return code from db2start = $" | tee -a /tmp/db2instance.log
```

© 2006 IBM Corporation

Automating cloning of DB2 (cont'd)

- ▶ Add a function to the clone.sh script

```
#+-----+
function clone_db2()
#   Arg 1: Mount point of the newly cloned server
# Clone a DB2 Server
#+-----+
{
  echo "Cloning DB2 ..."
  echo "Modifying zipl.conf and running zipl ..."
  cd $1/etc
  cp zipl.conf zipl.conf.orig
  sed -i -e 's:dasd=100-102:dasd=100-102,301(ro),401(ro):g' zipl.conf
  chroot $1 zipl
  echo "Modifying fstab ..."
  cp fstab fstab.orig
  sed -i -e '4a LABEL=db2-prod      /opt/IBM/db2      ext2
    ro,acl,user_xattr  0 0' fstab
  echo "Modifying /etc/inittab"
  echo "fmc:2345:respawn:/opt/IBM/db2/V8.1/bin/db2fmcd #DB2 Fault Monitor Coordinator" \
    >> $1/etc/inittab
  echo "Making symlink to /etc/init.d/db2instance ..."
  cd $1/etc/init.d/rc5.d
  ln -s ../db2instance S98db2instance
  cd
}
```

© 2006 IBM Corporation

Test DB2 automated cloning

- Clone a vanilla server to the target clone user ID

```
# # clone.sh --db2 linux06
...
WARNING!!: this will copy 100 and 102 disks to LINUX02 100 and 102
New host name will be: lat136.pbm.ihost.com
New TCP/IP address will be: 129.40.178.136
Other network data is retrieved from LINUX02 PARMFILE on 191 disk
A DB2 system will be cloned
Are you sure you want to overwrite these disks (y/n): y
...
```

- SSH into new clone

- ▶ View the log file

```
# cat /tmp/db2instance.log
Creating a DB2 instance ...
Return code from db2icrt = 0
Return code from dascrt = 0
Return code from db2iauto = 0
Removed symbolic link /etc/init.d/rc5.d/S98db2instance
Return code from db2start = 0
```

- Test DB2

- Test a reboot

© 2006 IBM Corporation

MQ Series cloning

- Set up the DB2 master from the controller
 - ▶ Clone a vanilla server to the MQ Series master user ID


```
# clone.sh mqs
...
#
```
 - SSH into the master
 - ▶ Configure minidisks 302 and 402 read/write
 - Enable - **chccwdev**
 - Format - **dasdfmt**
 - Partition - **fdasd**
 - Make a file system - **mke2fs -j -L <label>**
 - Labels are **mqs-prod** and **mqs-test**
 - ▶ Configure /etc/zipl.conf and /etc/fstab to mount mqs-prod (302) disk over /opt/mqm
 - ▶ Reboot to test changes


```
# mount | grep opt
/dev/dasdd1 on /opt/mqm type ext3 (ro,acl,user_xattr)
```
 - ▶ Install MQ Series from RPMs


```
# rpm -ivh MQSeriesRuntime-6.0.0-0.s390.rpm MQSeriesSDK-6.0.0-0.s390.rpm
MQSeriesServer-6.0.0-0.s390.rpm MQSeriesClient-6.0.0-0.s390.rpm
MQSeriesSamples-6.0.0-0.s390.rpm MQSeriesJava-6.0.0-0.s390.rpm
...
#
```
 - Shut down the MQ Series master

© 2006 IBM Corporation

Clone a MQ Series server manually

- Clone a vanilla server to the target clone user ID


```
# clone.sh linux07
...
#
```
- SSH into new clone
 - ▶ Enable disk **mqs-prod** read-only
 - Modify /etc/zipl.conf, run zipl, reboot to test
 - ▶ Create an MQ Series user and group


```
# groupadd mqm
# useradd -g mqm -d /var/mqm mqm
```
 - ▶ Create a .bash_profile in the mqm home


```
# cd /var/mqm
# vi .bash_profile
export PATH=$PATH:/opt/mqm/bin
```
- Test by creating an MQ Series queue manager


```
# su - mqm
$ crtqm -q test.queue.manager
```

© 2006 IBM Corporation



IBM

Automating cloning of MQ Series (cont'd)

- ▶ Create a script in master image's services directory

```
# cat /sles9master/etc/init.d/mqmuser
#!/bin/bash
echo "Creating an MQ Series user ..." | tee /tmp/mqmuser.log
groupadd mqm
echo "Return code from groupadd = $" | tee -a /tmp/mqmuser.log
useradd -g mqm -d /var/mqm mqm
echo "Return code from useradd = $" | tee -a /tmp/mqmuser.log
chown -R mqm.mqm /var/mqm
echo "Return code from chown = $" | tee -a /tmp/mqmuser.log
symlink="/etc/init.d/rc5.d/S97mqmuser"
if [ -h $symlink ]; then # symlink found - remove it so script runs once
    rm $symlink
    echo "Removed symbolic link $symlink" | tee -a /tmp/mqmuser.log
fi
```

© 2006 IBM Corporation



IBM

Automating cloning of MQ Series (cont'd)

- ▶ Add a function to the clone.sh script

```
#+-----+
function clone_mqs()
# Clone a MQ Series Server
#+-----+
{
    echo "Cloning MQ Series ..."
    echo "Modifying zipl.conf and running zipl ..."
    cd $1/etc
    cp zipl.conf zipl.conf.orig
    sed -i -e 's:dasd=100-102:dasd=100-102,302(ro),402(ro):g' zipl.conf
    chroot $1 zipl
    echo "Modifying fstab ..."
    cp fstab fstab.orig
    sed -i -e '4a LABEL=mqs-prod      /opt/mqm      ext2      ro,acl,user_xattr      0 0' fstab
    echo "Making symlink to /etc/init.d/mqmuser ..."
    cd $1/etc/init.d/rc5.d
    ln -s ../mqmuser S97mqmuser
    cd
}
```

© 2006 IBM Corporation



Test MQ Series automated cloning

- Clone a vanilla server to the target clone user ID

```
# clone.sh --mqseries linux07  
...  
An MQ Series system will be cloned  
Are you sure you want to overwrite these disks (y/n): y  
...  
Cloning MQ Series ...  
...
```

- SSH into new clone

- View the log file

```
# cat /tmp/mqmuser.log  
Creating an MQ Series user ...  
Return code from groupadd = 0  
Return code from useradd = 0  
Return code from chown = 0  
Removed symbolic link /etc/init.d/rc5.d/S97mqmuser
```

- Test MQ Series

- Test a reboot

© 2006 IBM Corporation

Live Demo!

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



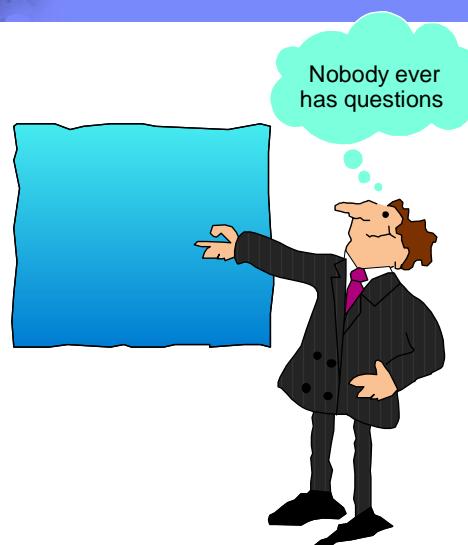
© 2006 IBM Corporation

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/linuxenterpriserver/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/>

© 2006 IBM Corporation

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: 1189/2954/9310

© 2006 IBM Corporation