

Making your life easier with automation

Michael MacIsaac

MVMUA

Oct 18, 2016

Abstract

- As the number of Linux instances under z/VM grows in your organization, so too does the workload of the z/VM and Linux systems administrators. Automation of processes is a must to not only keep the size of the workload reasonable, but also to avoid having to perform repetitive tasks. The speaker will present on how to make your life easier using automation.

Agenda



Agenda

- Overview
- Integrating with other teams
- Cloning/provisioning
- Patching
- Monitoring
- Linux Automation
- Self-service
- Live demo
- Questions

Overview



Integrating with other teams



Integrating with other teams

Two rules:

1. Ya gotta do it ...
2. When you can't do it, see rule #1

Cloning/provisioning
or
Guest building ...



Cloning/building

- Build vs. Clone
 - Building is best, ideally
 - Cloning is usually a necessary evil
- Copying disks
 - COPYDISK EXEC
 - Try FLASHCOPY, fail back to DDR
 - <https://sites.google.com/site/mike99mac/copydisk.exec>

Patching



Patching

- z/VM
 - Standard SERVICE ALL and PUT2PROD
- Linux
 - Other teams do the patching
 - Sometimes we have to clean up
 - Blade logic controls
 - Bash scripts, Python, Perl, etc.

Monitoring



Monitoring (Ops Mgr)

- Availability
 - Log on & off
- Proactive
 - Paging / Spool usage
 - Multipath failures
 - OSA Port failures
 - OOM & add swap (see next charts)
- Reactive
 - Oops
 - Zadmin (monitoring the monitor)

Ops Mgr and automation

- Reactive
 - E-mails of the form:
To: Mike MacIsaac
From: OPMGRS8@VM.MF.ADP.COM
Subj: TJSIA052 OOM_Kill_Process

TJSIA052 has had an OOM message.

Ops Mgr and automation

- Proactive

- E-mails of the form:

```
To:   Mike MacIsaac
From: OPMGRS8@VM.MF.ADP.COM
Subj: TJSIA052 OOM_Kill_Process
```

```
TJSIA052 has had an OOM message.  Disk 0163 has
been added.
```

How swap disks are added

- OpsMgr traps Linux OOM messages

```
DEFRULE NAME (LINUXOOMC),MATCH(*oom_kill_process*),ACTION (LINUXOOMC), +
  PARM (OOM_Kill_Process),LIMIT (1,1800),SYSGRP (CDL)
DEFACTN NAME (LINUXOOMC),COMMAND (EXEC OOMCDL &U &P &T),ENV (SVM)
```

- Calls **OOMCDL EXEC**
- Which calls **GETASWAP EXEC**

OOMCDL EXEC

```

/**/
TRACE o
parse arg userid parmtxt textline

'pipe',
  'cp query userid',
  '| specs words 3 1',
  '| var sysname'

/* Get the directory entry */
makebuf
queue 'fn oomuser'
queue 'ft file'
queue 'ffile'
'DIRME 'userid
if rc > 0 then do
  say 'OOMCDL could not obtain directory entry for'userid
  textline = sysname userid "had an OOM, but DIRME failed."
  "MAILIT TO(DEANNA) NAMES(OPMGR)",
    "TITLE("sysname userid " had OOM, recovery failed.)",
    "TEXT "textline
  EXIT
end
dropbuf

```

OOMCDL EXEC (cont'd)

```

/* Search for highest 016* disk */
found = 1
do cnt = 2 to 15 while found
  hcnt = d2x(cnt)
  disknumber = '016'|||hcnt
  'pipe < OOMUSER FILE A',
    '| locate /MDISK 'disknumber'/' ,
    '| var mdisknumb'
  if mdisknumb = 'MDISKNUMB'
    then found = 0
end

/* Add new Swap mdisk */
if found /* All mdisks from 0162-016F are in use */
then do
  text1 = userid 'has had an OOM message. However, there are'
  text1 = text1 'no more disks in range 0162-016F to add'
  "MAILIT TO ("KMGROUP") NAMES(OPMGR) TITLE("userid parmtxt")
  TEXT "text1
end
else do /* There is an available mdisk in 0162-016F range */
  say 'Valid disknumber to add is 'disknumber
  'dirm for 'userid' amd 'disknumber' xxx autog 2900 CDLSWAP'
  'EXEC GETASWAP' userid
  text1 = userid 'has had an OOM message. Disk 'disknumber
  text1 = text1 'has been added.'
  "MAILIT TO ("KMGROUP") NAMES(OPMGR) TITLE("userid parmtxt")
  TEXT "text1
end
end
exit

```

GETASWAP EXEC

```

/* This is the GETASWAP EXEC */
/* Its purpose is to get pull the addswap.sh script
   from the 191 mdisk and run it
   with the input parameter of the guest to run it on */
parse arg guest
/* start of the code */

say guest
Address "COMMAND"
cp send guest "chccwdev -d 191 > /dev/null 2>&1"
cp send guest "chccwdev -e 191 > /dev/null 2>&1"
sleep 1
cp send guest "userid=$(vmcp q userid | awk '{print$1}')"
cp send guest "echo $userid"
cp send guest "dev191='/dev/$(lsdasd -a | grep 0.0.0191 | awk '{print $3}')"
cp send guest "echo $dev191"
cp send guest "cmsfscat -a -d $dev191 ADDSWAP.SH > addswap.sh"
cp send guest "chmod +x addswap.sh"
cp send guest "mv addswap.sh /usr/local/sbin/"
cp send guest "which addswap.sh"

cp send guest "addswap.sh"

EXIT

```

Linux automation



Examples of Linux automation

- On 'first boot'
 - IP configuration: static (work) or DHCP/DDNS (better)
 - Application configuration (Oracle)
- On Restart
 - Change settings
 - systemctl
 - Get new configuration files
 - Multipath.conf
- Scheduled automation
 - Oracle DBA assistance
 - ASM (Automatic Storage Management) usage
 - Instance counts
 - SGA settings

Static IP vs. DHCP

- Static IP address are "so 2000s"
- Requirements for DHCP
 - Subnet with DHCP/DDNS working
 - VSWITCH
 - Layer 2
 - Authorize at virtual machine logon
 - Attach with correct VLAN ID
 - Unique MACID for each virtual machine
 - Prefix in SYSTEM CONFIG
 - Suffix in user directory entry
 - LDAP registration (if applicable)

Virtual machine definition

```

PROFILE LINXENG
  ACCOUNT ENG
  CPU 00 BASE
  IPL CMS PARM AUTOOCR
  MACHINE ESA 8
  ...

USER LINUX001 XXXXXXXX 1G 8G G
  INCLUDE LINXENG
  COMMAND DEFINE STORAGE 1G STANDBY 7G
  COMMAND SET VSWITCH VSW1 GRANT LINUX001 VLAN 2345
  COMMAND COUPLE 350 TO SYSTEM VSW1
  NICDEF 0350 TYPE QDIO LAN SYSTEM ESNETNG MACID FFFFC
  MDISK 0151 3390 16380 16379 MYVOL1 MR READ WRITE MULTI
  ...

```

Virtual machine definition

- Add the virtual machine without NICDEF nor disks


```

==> dirm add linux001
...

```
- Add a NICDEF with user exit invoked


```

==> dirm for linux001 NICDEF 0350 TYPE QDIO LAN SYSTEM VSW1
...
DVHXRA8888E MACID FFFFC assigned to LINUX001 0350

```
- Add disks


```

==> dirm for linux001 amd 151 x autog 16379 mod27
...
==> dirm for linux001 amd ffff x autog 1 myvol2
...

```

Application automation

- Oracle servers – first boot processing

```

if [ "$configOracle" = "ora" ]; then      # Oracle standalone server
  say "*****"
  say "$this: preparing for Oracle standalone installation"
  say "*****"
  setOraGrid                             # settings common to grid and standalone
  setKernelParms ora                     # configure kernel parameters
  setFCP ora                              # configure FCP/SCSI disks
  mkPartitions ora                       # make partitions of the FCP/SCSI LUNs
  mountPartitions ora                   # mount the three new partitions
  mkMountPoints ora                     # make mount points for NFS file systems
  append2fstab ora                      # add entries to the /etc/fstab file
  append2sudoers ora                    # add entries to the /etc/sudoers file
elif [ "$configOracle" = "grid" ]; then  # Oracle CRS cluster server

```

Application automation (cont'd)

```

elif [ "$configOracle" = "grid" ]; then  # Oracle CRS cluster server
  say "*****"
  say "$this: preparing for Oracle grid installation"
  say "*****"
  setOraGrid                             # settings common to grid and standalone
  setFCP grid                            # configure FCP/SCSI disks
  mkPartitions grid                      # make partitions of the FCP/SCSI LUNs
  mountPartitions grid                  # mount the three new partitions
  mkMountPoints grid                   # make mount points for NFS file systems
  append2fstab grid                    # add entries to the /etc/fstab file
  append2sudoers grid                  # add entries to the /etc/sudoers file
elif [ "$configOracle" = "ev5app" ]; then # Oracle app server - only limited
  say "*****"
  say "$this: preparing for Oracle app server installation"
  say "*****"
  createGroupsUsers                     # define users and groups
  modifyProfileLocal                    # modify the /etc/profile.local file
  setPrism                              # configure Prism
  mkMountPoints ev5app                  # make mount points for NFS file systems
  append2fstab ev5app                   # add entries to the /etc/fstab file
  setKernelParms ev5app                 # add entries to /etc/sysctl.conf
  append2sudoers ev5app                 # add entries to the /etc/sudoers file
else                                     # not an Oracle system (should never
called)
  say "$this: 'ora', 'grid' nor 'ev5app' keyword not found - should not be c

```

Self-service



Self-service

- Solutions that were not 100% successful:
 - IBM Wave
 - IBM CSAM
 - Cloud Solution Acronym of the Month
 - Velocity zPro
 - Zoom
 - HPOO
 - Others (probably)
- Main problem: “Greenfield” vs. “Brownfield”

Greenfield vs. brownfield

- From “techopedia”

“Brownfield refers to the implementation of new systems to resolve IT problem areas while accounting for established systems. New software architecture must account for existing and running software.

A commonly used IT term, Brownfield was borrowed from the building industry, where brownfield land describes a geographical location where new buildings may be constructed after considering the area's established structures and services.

Brownfield development enhances traditional software engineering practices. Greenfield, or a blank slate, is considered the traditional target environment during software development”

...

Roll-your-own self-service

- We are looking at ‘roll-your-own’ to utilize
 - z/VM more than Linux
 - Internal ‘Guest finder’
 - IBM Operations Manager
 - Existing REXX EXECs

Live demo



Questions



Back up slides



Self-service portal

10.1.231.89/cgi-bin/ops/selfserv

Self-service prototype

You user name: **macisaam** (uid: 4506)
 Your group name: **vmlinux** (gid: 1041)

Resource Quotas		
Resource	Quota	Amount Used
CPUs	768	0
memory	1024	0

Operations	
Clone a zLinux server	Destroy a zLinux server
Power on server(s)	Power off server(s)
Start (BEGIN CPU) server(s)	Stop (STOP CPU) server(s)
Reboot (soft)	Reboot (hard)
Add CPU(s)	Remove CPU(s)
Add memory	Remove memory

Self-service clone Linux

10.1.231.89/cgi-bin/ops/cloneinput

You user name: **macisaam** (uid: 4506)
 Your group name: **vmLinux** (gid: 1041)

Resource Quotas		
Resource	Quota	Amount Used
CPUs	768	0
memory	1024	0

Clone a zLinux system

Linux distribution: <input type="radio"/> SLES 11 <input checked="" type="radio"/> SLES 12 <input type="radio"/> Ubuntu Xenial	Set no-op/verbosity: <input type="checkbox"/> No operation <input checked="" type="checkbox"/> Normal verbosity <input type="checkbox"/> Verbose
Amount of memory (GB): <input type="text" value="1"/>	Number of CPUs: <input type="text" value="2"/>

Start cloning → **Constructed clone command:**

```
ADP/MAUT WED-JUN-22-07:41:40-EDT-2016 ZCLONE SLES12 1 2 vmLinux NORMAL
```

Selfserv code

1

