

Managing Linux under z/VM

Guide SHARE Europe, Warwick

30 Jan 2003



Malcolm Beattie

Linux Technical Consultant
IBM EMEA Enterprise Server Group
beattiem@uk.ibm.com






e-business




WWW.



IBM

Introduction

- Remote control
 - ▶ Consoles, CP and early boot configuration
- Linux dynamic DASD support
- Conclusions
- Questions



Consoles and early boot config


- Automating Linux guest operations
 - ▶ Authorising remote commands with PROP
 - ▶ Remote/automatic startup and shutdown
 - ▶ Distributing configurations via PROP



e-business



WWW.



IBM

What is a console?

- How VM sees a guest console
 - ▶ A 3270 or 3215 device
 - ▶ input of #CP FOO passes command to CP
 - ▶ DIAG X'08' also passes command to CP
- How a Linux guest sees its console
 - ▶ /dev/console is the device node
 - ▶ only supports basic line-mode I/O
 - ▶ kludges for ^c but still no raw/canon mode



e-business



WWW.



IBM

CP access from a Linux guest

- Kernel driver (cpint) mediates access to CP via DIAG X'08'

- ▶ Low level

- write CP command to /dev/cpcmd
- read back status/response

- ▶ High level utility is **hcp**

```
# hcp query 2801
```

```
DASD 2801 3390 VMLX5B R/W          1669
```


```
  CYL ON DASD  1001 SUBCHANNEL = 000C
```



e-business



WWW.



IBM

CP access from another guest

- VM allows authorised guests to
 - ▶ force input to another guest's console
 - **SEND TARGET** This is console input
 - ▶ invoke CP commands in a different guest
 - **SEND CP TARGET DETACH 2345**
 - causes guest TARGET to execute CP command DETACH 2345
- SEND itself is a CP command
- thus can be called from CMS, Linux, ...



e-business



WWW.

Authorising CP SEND access

- VM allows guest MGRGUEST to **SEND (CP) TARGET ...**
if either of the following are true:
 - ▶ MGRGUEST has privilege class C
 - to be exact: a privilege class assigned to SEND by the local VM installation, the default being C
 - ▶ MGRGUEST is “SECUSER” for TARGET *and TARGET is running disconnected*
 - This can be done in TARGET's user directory entry or via CP SET SECUSER TARGET MGRGUEST

The IBM logo, consisting of the letters 'IBM' in a white, bold, sans-serif font on a dark blue background.

IBM



Using **SEND** to trigger Linux processes

- Sometimes you want to trigger an action in a Linux guest without resorting to TCP/IP
- The CP command `EXT 1234` causes a virtual machine to see an external interrupt
- Such an interrupt can be triggered from a different guest via

`SEND CP TARGET EXT 1234`

- The Linux kernel module `ext_int` (q.v.) can translate the “hardware” interrupt into a signal, thus notifying your process



e-business



WWW.



IBM

The ext_int module

- A kernel module written by me with no warranty, no support, blah blah blah
- Introduced in 9.8 of redbook SG24-6824 (contains reference to download site)

```
# insmod ext_int code=0x0d1e pid=1 sig=2
```
- When EXT 0D1E is triggered, ext_int arranges for signal SIGINT (signal 2) to be sent to the init process (PID 1)
- init can then do a clean shutdown or reboot



Automating/Authorising SEND

- You want to automate CP SEND...
- ...or allow some users limited access to sending specific commands to specific users...
- ...and you want to authorise based on which user sends the command and what commands they're sending
- What guest/code do you trust to do that?

PROP!



e-business

PROP

- PROP is the *PR*ogrammable *OP*erator
- Part of VM
- Runs in its own little virtual machine
- Table driven by simple pattern matches in a flat CMS file
- Acts on messages sent from guests with MSG or SMSG (or console input, or...)
- who? what? where? action!

The IBM logo, consisting of the letters 'IBM' in a white, striped font on a dark blue background.


IBM



e-business



WWW.



IBM

PROP examples

- Two examples in redbook SG24-6824
- Set up a privilege class C virtual machine
 - ▶ If guests are always disconnected, we could get by with only a SECUSER entry instead of privilege class C
- Prepare PROP RTABLE file
- Ensure guest is autologged at system IPL
- PROP acts on received SMSG messages



PROP example 1


- Remote startup and shutdown of guests
 - ▶ SMSG PROP XAUTOLOG FOO
 - ▶ SMSG PROP SHUTDOWN FOO
- Message sent from CMS command, Linux command or web front-end
- PROP receives SMSG, authorises sender and takes appropriate action
 - ▶ XAUTOLOG FOO
 - ▶ SEND CP FOO EXT 0D1E
- (Recent VM and Linux support SIGNAL SHUTDOWN too)



e-business



WWW.



IBM

PROP example 2

- Distribute early boot-time configuration
 - ▶ IP address, IP device(s), role, ...
- Useful for cloned guests or central configuration before network or LDAP infrastructure is available
 - ▶ SMSG PROP GETMYCONF
 - ▶ Magic happens
 - ▶ Info appears in `/etc/sysconfig/vmconfig`
 - ▶ Continue boot knowing who we are



e-business



WWW.

How PROP gets our config

- Magic happens?
 - ▶ PROP knows who sent the SMSG (FOO, say) and looks up config info in a flat CMS file
 - ▶ How to get the info where FOO can see it?
 - ▶ FOO defines a virtual printer: `DEFINE PRT 001E`
 - ▶ PROP tags it with the information
 - `SEND CP FOO TAG 001E IPADDR=1.2.3.4`
 - ▶ FOO loops until tag appears, writes it to file

```
# hcp TAG QUERY 001E
PRT 001E TAG:
IPADDR=1.2.3.4
```

IBM



e-business



WWW.

Linux dynamic DASD support

- Linux supports hot add/remove of DASD
- Control is via echoing commands to `/proc/dasd/devices`
 - ▶ But redbook SG24-6824 has a convenient script named **dasd**

```
# dasd list
```

```
2800(ECKD) at ( 94: 0) is dasda      : active  
    at blocksize: 4096, 108000 blocks, 421 MB
```

```
2801(ECKD) at ( 94: 4) is dasdb      : active  
    at blocksize: 4096, 300420 blocks, 1173 MB
```

```
...
```

The IBM logo, consisting of the letters 'IBM' in a bold, blue, sans-serif font.



e-business



WWW.

Hot addition of DASD

- Add new physical DASD and connect up
- Handle IOCP (unless MVS wants to)
 - ▶ DEFINE CHPID/CU/DEVICE ...
- Vary online to VM
 - ▶ VARY ONLINE CHPID, VARY ONLINE rdev
- Carve up disk space for Linux guests
 - ▶ Dedicate volumes: ATTACH ... TO GUEST1
 - ▶ Or divide into minidisks then: LINK ...

The IBM logo, consisting of the letters 'IBM' in a white, blocky font on a dark blue background.

IBM



e-business



WWW.

Hot addition of DASD (contd)

- Bring online to Linux guest
 - # `dasd add 1234-1236`
 - # `dasd list`
 - ▶ Find newly allocated device names/numbers
- Format and partition
 - ▶ # `dasdfmt -f /dev/dasdx`
 - ▶ # `fdasd /dev/dasdx`
- Disk space `/dev/dasdx1` is now available

The IBM logo, consisting of the letters 'IBM' in a white, bold, sans-serif font, set against a dark blue background.

IBM



e-business



WWW.

Hot addition of DASD (contd)

- You could now make a new filesystem
 - # `mke2fs /dev/dasdx1`
 - # `mount /dev/dasdx1 /usr/local/someproject`
- **Or**, if you are using LVM
 - ▶ Add it into your logical volume group
 - # `pvcreate /dev/dasdx1`
 - # `vgextend myvg /dev/dasdx1`
 - ▶ The disk space is now part of your “free space pool”
- Question: How much hotter could it be?

The IBM logo, consisting of the letters 'IBM' in a white, bold, sans-serif font, is positioned at the bottom of the vertical banner.

IBM



Hot addition of DASD (contd)

- Answer: a little bit
 - ▶ What if *one filesystem* needs more room?
- Extend underlying logical volume
 - # `lvextend -L +2G /dev/myvg/foo`
- Can you schedule time to unmount it?
 - ▶ Yes? Unmount it then
 - # `resize2fs ...` Or `resize_reiserfs ...`
 - ▶ No? Maybe try `ext2online` (*unsupported, risky, ext2-specific*)



e-business



WWW.

Hot addition of DASD (contd)

- A few notes on filesystems
 - ▶ ext2resize and resize_reiserfs both support shrinking as well as enlarging filesystems
 - ▶ Methods for ext2 also apply to ext3
 - Easy to add/remove journal (hot or not)
 - Quantify exposure of short-term lack of journal
 - ▶ Other filesystems may or may not resize

The IBM logo, consisting of the letters 'IBM' in a white, sans-serif font, is positioned at the bottom of the vertical banner.

IBM



e-business



WWW.



IBM

Conclusions

- Access to consoles beats “real” hardware
- Linux takes advantage of S/390 DASD management capabilities
- Questions?

Thank you!

Malcolm Beattie

beattiem@uk.ibm.com