Tips, Tools, and History of Linux on System Z: A New Mainframer's View of Linux on Z

Deric R. Abel
America First Credit Union

Thursday, March 3rd, 2011
Session 8500

Anatomy of a new Mainframer

- Deric Abel
- I've been in IT since 1997
- First installed Linux as a High School project my senior year (1999)
- Hired as a Linux Admin in 2000
- First experience with Virtualization in 2005
- Hired at America First Credit Union as a z/Linux admin in 2008
- Joined the zNextgen group and attended my first SHARE conference in 2008
- Currently serving as a Deputy Project Manger for zNextgen

Agenda

- Brief History of Linux
- Linux on System Z
  - A new mainframer's perspective
  - Differences between Distributed systems and Z
- Linux Tips and Tools
  - Native Linux Tools
  - Green Screen! Are you kidding me?
  - CP commands in Linux
  - Monitoring and Performance
  - Problem/Issue resolution
- Conclusion and Q/A

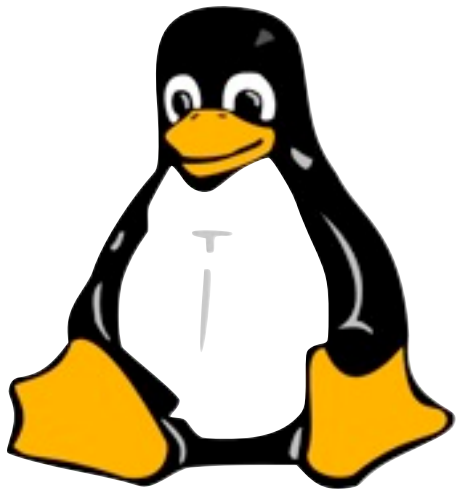

Brief History of Linux

- Linux was first introduced as a hobby project in 1991
- IBM published a collect of patches to linux on December 18th, 1999
- Linux on System Z was formally announced in 2000 along with the Integrated Facility for Linux (or IFL) engines.

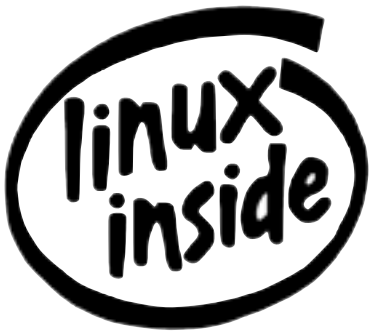

Linux on System Z

- IBM currently supports two linux Distributions:
  - Red Hat Enterprise Linux (RHEL)
  - SuSE Linux Enterprise Server (SLES)
- Three ways to run linux on Z
  - Whole mainframe
  - LPAR
  - z/VM

A new mainframer's perspective

A new mainframer's perspective

- Terminology Differences

Storage vs Memory

- Memory on the mainframe is typically referred to as storage (or more technically, real storage).
- The term memory is used as the equivalent in Distributed Systems.
- Some people in the mainframe community refer to hard disk units as “storage”, but it’s recommended we use the term storage only when referring to memory.
- Three types, or levels, of Storage: Central, Expanded, and Page Space.

Storage Cont'd

- Central Storage
  - Central storage contains the current running operating system and any processes or programs and data being used by the operating system.
- Expanded Storage
  - Expanded storage is needed to exploit certain special software facilities and also used as a faster paging device.
- Page Space
  - Similar to Swap disks on Linux

DASD vs Disk

- Disks in the mainframe environment are usually referred to as DASD (Direct Access Storage Device)
- DASD is similar to a PC hard disk, except they are external to the mainframe and they comprise many drives in a far more sophisticated arrangement.
- Two other types of disk accessible by Linux are FCP and iSCSI, these are directly attached to the Linux guest and are managed at the Linux level instead of at z/VM.

DASD Cont'd

- DASD devices show up in linux as a /dev/dasda1, where the a1 denotes the first DASD and first partition.
- An example of a Disk Free (or df) command on Linux:

```
Filesystem                  Size  Used Avail Use% Mounted on
/dev/dasdc1                 2.3G  1.9G  307M  87% /
/dev/mapper/sanvg-sanvol     99G   84G   11G  90% /data
/dev/mapper/sanvg2-home      50G   18G   30G  38% /home
```

OSA and Networking

- The LAN adapter for the mainframe is known as the Open System Adapter (OSA).
- This device shows up as a qeth device as opposed to eth
- Another networking device is known as a hipersocket, which is a memory to memory channel between two guests.
- The hipersocket shows up as a hsi device.

Channel I/O

- One of the main strengths of the mainframe is the ability to deal with a large number of simultaneous I/O operations.
- The channel subsystem (CSS) manages the flow of information between I/O devices and central memory. This relieves CPUs of the task of communicating directly with I/O devices.
- On Distributed systems, the CPU handles a large part of I/O, which is why a system slows to it's knees during high disk activity.

Linux Tools

Native Linux Tools

- uptime
- vmstat
- free
- ps
- iostat
- w
- sar
- mpstat
- netstat
- last
- du
- df

Native Linux Tools

- uptime - This command is simple. It gives you a quick snapshot of system performance and the amount of time the system has been live since the last reboot. An example of the command's output is below:

```
11:52  up 1 day, 12:06, 2 users, load averages: 0.39 0.28 0.26
```

Native Linux Tools

- vmstat - The vmstat (virtual memory statistics) command has nothing to do with virtualization but rather it has to do with the health of your system from a swap space point-of-view.

```
root@fs:~# vmstat 1
procs -----------memory---------- ---swap-- -----io---- -system-- ----cpu----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa
 3  0  65448  28504  35360 2515048    0    0   113    19    1    6  9 11 79  0
 0  0  65448  28496  35360 2514956    0    0     0     0 6287 13090  9 17 73  0
 1  0  65448  28124  35368 2515244    0    0   288    16 6159 12970 11 10 78  0
 2  0  65448  80028  35096 2463796    0    0   564     0 5604 12355 15 15 70  0
 1  0  65448  80028  35096 2463976    0    0   256     0 6275 13062 13 11 76  0
 1  0  65448  79748  35096 2464536    0    0   512     0 6274 13181 17 14 69  0
 1  0  65448  80184  35096 2464516    0    0     0     0 6219 13016 10 13 78  0
```

Native Linux Tools

- free - Free displays the amount of free physical memory (RAM) in a system, the used physical memory, free and used swap memory and buffers used by the kernel.

```
root@fs:~# free
             total       used       free     shared    buffers     cached
Mem:       4057736    3990140      67596          0      37964    2478728
-/+ buffers/cache:    1473448    2584288
Swap:      2097144      65468    2031676

root@fs:~# free -m
             total       used       free     shared    buffers     cached
Mem:          3962       3897         65          0         37       2421
-/+ buffers/cache:       1438       2523
Swap:         2047         63       1984
```

Native Linux Tools

- ps - The ps command shows you a snapshot of currently running processes. It has several possible switches (or options) but the most common is the ps -ef command. Any user may issue the ps command.

```
root@fs:~# ps af
  PID TTY      STAT   TIME COMMAND
24881 pts/0    Ss     0:00 -bash
24906 pts/0    R+     0:00  \_ ps af
 1369 tty1     Ss+    0:00 /sbin/getty -8 38400 tty1
 1025 tty6     Ss+    0:00 /sbin/getty -8 38400 tty6
 1023 tty3     Ss+    0:00 /sbin/getty -8 38400 tty3
 1022 tty2     Ss+    0:00 /sbin/getty -8 38400 tty2
 1006 tty5     Ss+    0:00 /sbin/getty -8 38400 tty5
 1002 tty4     Ss+    0:00 /sbin/getty -8 38400 tty4
```

Native Linux Tools

- iostat - This command reports CPU, disk and partition (I/O) statistics. The iostat has several possible switches available to it for specific output. It is part of the sysstat package and may not be installed by default.

```
root@fs:~# iostat
Linux 2.6.32-23-server (fs)   08/02/2010          _x86_64_   (2 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           8.84    0.01   11.40    0.45    0.00   79.29

Device:            tps   Blk_read/s   Blk_wrtn/s   Blk_read   Blk_wrtn
sda               3.60       486.98        67.64 1076429627  149507752
sdb               0.84         5.44        13.43   12021232   29685832
dm-0              1.92         5.36        13.31   11850810   29417048
dm-1              0.02         0.08         0.12     167224     268704
dm-2              8.18       486.98        67.64 1076428781  149507752
```

Native Linux Tools

- w - The w (what) command is better than the *who* command for seeing who's logged in and what they are doing.

```
root@fs:~# w
 10:41:24 up 25 days, 14:07,  1 user,  load average: 0.45, 0.42, 0.50
USER     TTY      FROM              LOGIN@   IDLE   JCPU   PCPU WHAT
dabel    pts/0    192.168.12.1     10:41    0.00s  0.52s  0.11s sshd: dabel [priv]
```

Native Linux Tools

- sar - The sar (System Activity Reporter) command is part of the sysstat package.

```
$ sar
Linux 2.6.18-53.el5 (system.domain.com) 04/28/2010

12:00:01 AM CPU %user %nice %system %iowait %steal %idle
12:10:01 AM all 0.49 0.00 0.52 0.05 0.00 98.94
12:20:01 AM all 0.13 0.00 0.51 0.08 0.00 99.28
12:30:01 AM all 0.12 0.00 0.53 0.05 0.00 99.29
12:40:01 AM all 0.12 0.00 0.52 0.05 0.00 99.31
12:50:01 AM all 0.13 0.00 0.55 0.07 0.00 99.25
01:00:01 AM all 0.13 0.00 0.65 0.06 0.00 99.16
01:10:01 AM all 0.54 0.00 0.50 0.08 0.00 98.88
01:20:01 AM all 0.13 0.00 0.51 0.08 0.00 99.28
01:30:01 AM all 0.12 0.00 0.52 0.08 0.00 99.28
01:40:01 AM all 0.13 0.00 0.50 0.07 0.00 99.30
```

Native Linux Tools

- mpstat - The mpstat command provides you with Multi-processor, CPU-related statistics. It is part of the sysstat package.

```
root@fs:~# mpstat 1
Linux 2.6.32-23-server (fs)   08/02/2010          _x86_64_   (2 CPU)

10:54:45 AM  CPU    %usr   %nice    %sys %iowait    %irq   %soft  %steal  %guest   %idle
10:54:46 AM  all    3.17    0.00   12.22    0.00    0.00    1.81    0.00    3.62   79.19
10:54:47 AM  all    5.38    0.00   11.21    0.00    0.00    0.90    0.00    4.48   78.03
10:54:48 AM  all    3.59    0.00   13.45    0.45    0.00    1.35    0.00    2.69   78.48
10:54:49 AM  all    5.41    0.00    6.76    0.00    0.00    1.80    0.00    2.70   83.33
10:54:50 AM  all    5.41    0.00    8.11    0.00    0.45    1.35    0.00    3.15   81.53
10:54:51 AM  all    5.00    0.00    7.27    0.00    0.00    0.91    0.00    4.55   82.27
10:54:52 AM  all    5.41    0.00    7.66    0.45    0.00    0.45    0.00    2.70   83.33
10:54:53 AM  all    6.36    0.00    9.55    0.00    0.00    1.82    0.00    2.73   79.55
```

Native Linux Tools

- netstat - The netstat command, replete with options and switches, provides you with diagnostic information about your network statistics including interface statistics, routing tables, network connections and more. A wise SA uses netstat to diagnose network problems, attacks and to see a list of services and connections. An example is shown below.

```
root@fs:~# netstat -a |grep LISTEN
tcp        0      0 *:mysql                 *:*                     LISTEN
tcp        0      0 *:50316                 *:*                     LISTEN
tcp        0      0 localhost:5900          *:*                     LISTEN
tcp        0      0 *:6543                  *:*                     LISTEN
tcp        0      0 *:sunrpc                *:*                     LISTEN
tcp        0      0 *:6544                  *:*                     LISTEN
tcp        0      0 *:www                   *:*                     LISTEN
```

Native Linux Tools

- last - The last command shows a listing of last logged in users. This will also show past reboots as well as who did it.

```
root@lnxXXXXX:PROD:~ # last
root     pts/0        x.x.x.x          Mon Aug  2 19:09   still logged in
******** pts/0        x.x.x.x          Mon Aug  2 17:35 - 17:57  (00:22)
******** pts/1        x.x.x.x          Mon Jun 21 16:17 - 16:43  (00:26)
******** pts/0        x.x.x.x          Mon Jun 21 15:57 - 22:01  (06:04)
opc_op   pts/1        x.x.x.x          Sun Jun 20 21:43 - 21:43  (00:00)
******** pts/0        x.x.x.x          Sun Jun 20 10:20 - 07:39  (21:19)
******** pts/0        x.x.x.x          Sun Jun 20 10:12 - 10:19  (00:07)
******** pts/2        x.x.x.x          Sun Jun 20 10:04 - 10:52  (00:48)
reboot   system boot  2.6.16.60-0.59.1 Sun Jun 20 10:03         (43+09:06)
******** pts/1        x.x.x.x          Sun Jun 20 09:39 - down   (00:13)
******** pts/0        x.x.x.x          Sun Jun 20 09:37 - 09:52  (00:14)
opc_op   pts/0        x.x.x.x          Sat Jun 19 21:55 - 21:55  (00:00)
```

Native Linux Tools

- du - The du command reports on disk usage. You can use it to look at all filesystems or a single one.

```
root@fs:/# du -sh data
2.0T     data
```

- du with find:

```
root@fs:/# for i in `find -maxdepth 1 -type d`; do du -sh $i; done
2.0T     .
204K     ./dev
74M      ./boot
0        ./proc
597M     ./lib
8.9M     ./home
7.8M     ./bin
2.4G     ./usr
9.5M     ./sbin
2.8G     ./var
988M     ./root
11M      ./etc
2.0T     ./data
4.0K     ./opt
```

Native Linux Tools

- df - The df command reports the amount of used vs. free space you have on your filesystems.

```
root@fs:/# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/System-root
                       20G  6.9G   12G  37% /
none                  2.0G  200K  2.0G   1% /dev
none                  2.0G  4.0K  2.0G   1% /dev/shm
none                   20G  6.9G   12G  37% /var/lib/ureadahead/debugfs
/dev/mapper/raid-data
                      2.8T  2.0T  830G  71% /data
/dev/sdb1             114M   80M   29M  74% /boot
```

Green Screen! Are you kidding me?

- What is this?

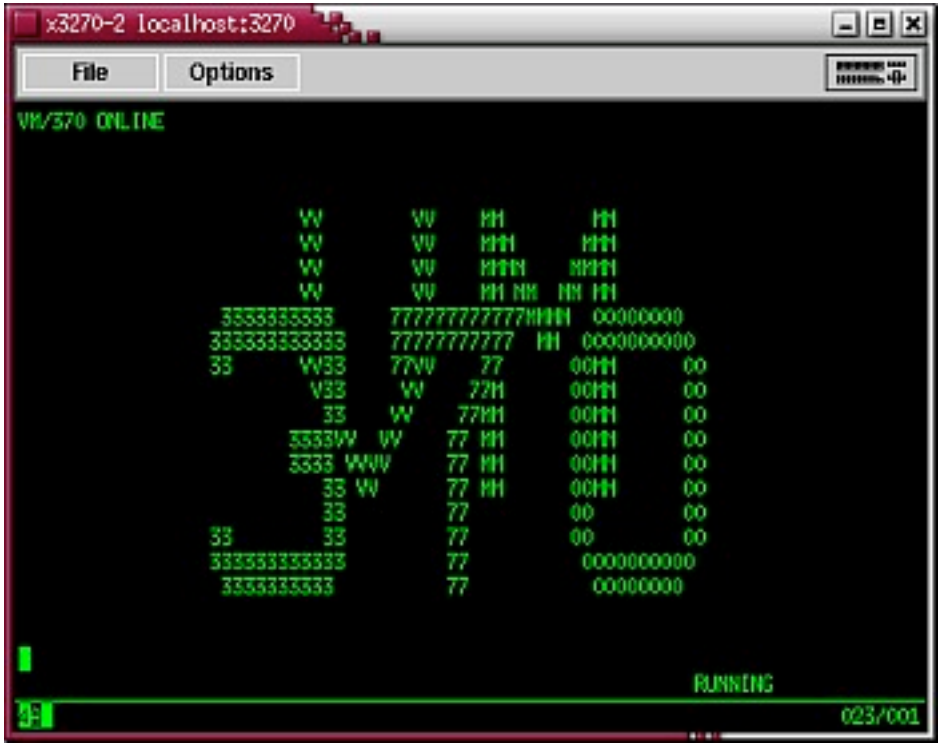

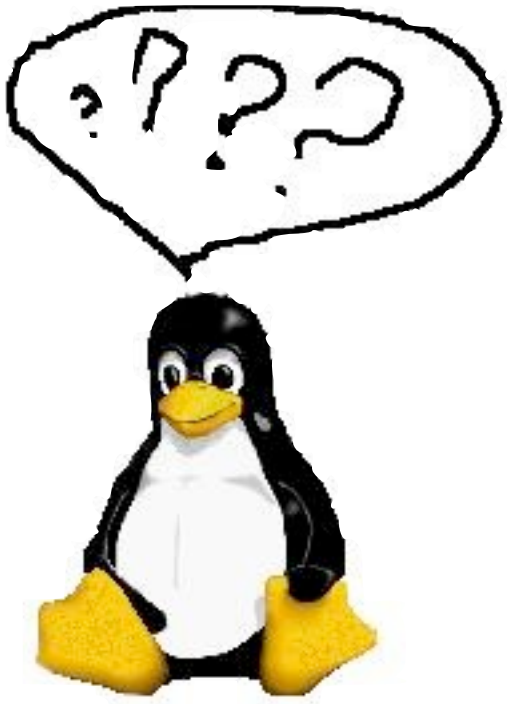

CP commands from Linux

- Two ways to access CP from linux
  - CPINT module (No longer provided with Distros)
  - VMCP module (Comes with most modern Distros)
-

Linux Tools and Tips

- Monitoring
  - Linux was originally designed on distributed hardware.
  - Most built in tools can be very expensive to run.
  - Third Party tools
    - SNMP
    - Zabbix
    - zVPS from Velocity Software

Linux Tools and Tips Cont'd

- Problem research and resolution
  - Vendor support
  - Linux390 listserv
  - Google

References and Links

- Linux-390 Listserv
  - Linux-390 focuses on Linux on System z including Linux on z/VM. To subscribe to the LINUX-390 discussion, send e-mail note to: LISTSERV@vm.marist.edu
  - In the body of the note, write only the following line: SUBSCRIBE LINUX-390 your-name-here
  - Or you can view/search the list archives: http://www.marist.edu/htbin/wlvindex?linux-390

Links

- http://www.linuxvm.org
- http://www-03.ibm.com/systems/z/os/linux/

Questions ?

SHARE
Technology · Connections · Results

Thank you for your attention!

Please do not forget to fill in evaluation forms.

Session 8500

SHARE in Anaheim 2011