

s390-tools - In a Nutshell


The Swiss Army Knife for Linux on System z System Administration

Session 8647

Tuesday, March 1, 2011: 1:30 PM-2:30 PM
Room 203B (Anaheim Convention Center)

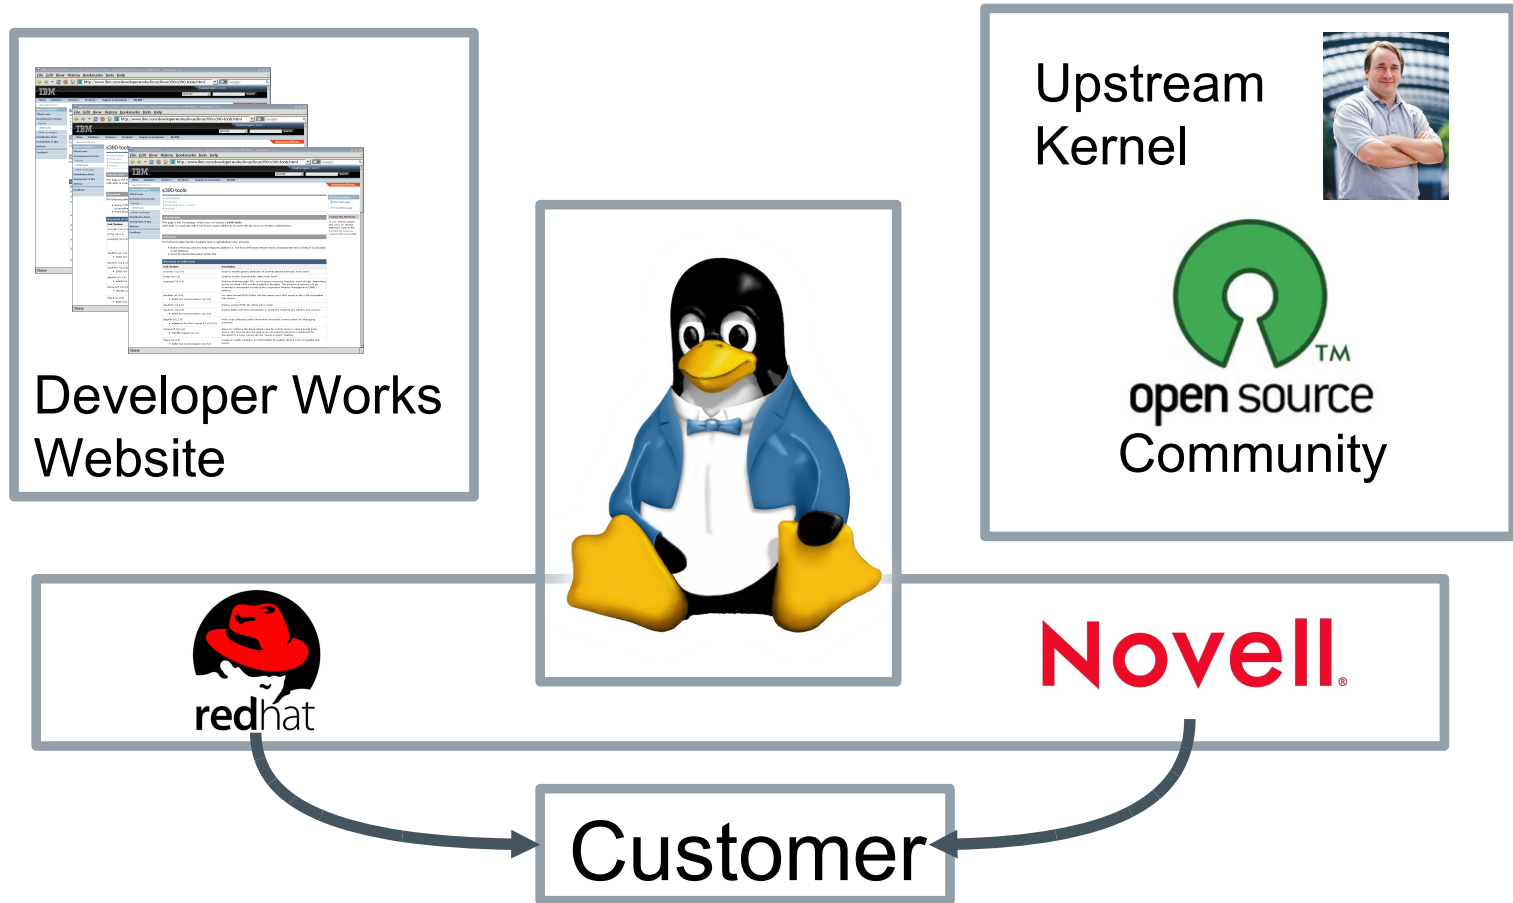


The Linux on System z Life Cycle



IBM Linux on System z Development

IBM Linux on System z Development contributes in the following areas: Kernel, s390-tools, Open Source Tools (e.g. eclipse, ooprofile), GCC, GLIBC, Binutils



...the code you use is the result of the efforts of an anonymous army of blue penguins involved in developing, testing, documenting,



What is the s390-tools package?

s390-tools is **the** essential set of Linux on System z system administration tools:

It contains everything from the boot loader to dump related tools for system crash analysis .

This software package is contained in all major (and IBM supported) distributions which support s390

RedHat Enterprise Linux 5 (s390-tools-1.8.1 since RHEL 5.4)

RedHat Enterprise Linux 5 (s390-tools-1.8.2 since RHEL 6)

SuSE Linux Enterprise Server 10 (s390-tools-1.6.3 since SLES 10 SP2)

SuSE Linux Enterprise Server 11 (s390-tools-1.8.0)

The latest version 1.12.0 was released on 27.01.2011

Website: <http://www.ibm.com/developerworks/linux/linux390/s390-tools.html>



The Content

chccwdev chchp chreipl chshut chzcrypt chmem CHANGE	dasdfmt dasdinfo dasdview fdasd tunedasd DASD	dbginfo dumpconf zfcpdump zfcpdbf zgetdump scsi_logging_level DUMP & DEBUG
lscss lschp lsdasd lsluns lsqeth lsreipl lsshut lstape lszcrypt lszfcp lsmem DISPLAY	mon_fsstatd mon_procd ziomon MONITOR	vmconvert vmcp vmur z/VM
	ip_watcher osasnmpd qetharp qethconf NETWORK	cpuplugd iucvconn lucvtty ts-shell MISC
	tape390_crypt tape390_display TAPE	zipl BOOT



Future Linux on System z Technology

Software which has already been developed and externally published – but is **not** yet available in any Enterprise Linux Distribution



hyptop - Display hypervisor performance data

(1.12.0)

The `hyptop` command provides a dynamic real-time view of a hypervisor environment on System z.

- It works with both the z/VM and the LPAR PR/SM hypervisor.
- Depending on the available data it shows, for example, CPU and memory information about running LPARs or z/VM guest operating systems.

The following things are required to run hyptop:

- The `debugfs` file system must be mounted.
- The `hyptop` user must have read permission for the required `debugfs` files:
 - z/VM: `<debugfs mount point>/s390_hypfs/diag_2fc`
 - LPAR: `<debugfs mount point>/s390_hypfs/diag_204`
- To monitor all LPARs or z/VM guest operating systems of the hypervisor, your system must have additional permissions:
 - For z/VM: The guest must be class B.
 - For LPAR: On the HMC or SE security menu of the LPAR activation profile, select the Global performance data control checkbox.

When running Linux on System z on z10 in LPAR, the minimum required z10 code level is the Following: Driver 79 MCL N24404.008 in the SE-LPAR stream



hyptop – Displaying hypervisor performance data

Displaying performance data for the z/VM hypervisor

```
10:11:56 CPU-I: UN(16) ?=help
```

system (str)	#cpu (#)	cpu (%)	Cpu+ (hm)	online (dhm)	memuse (GiB)	memmax (GiB)	wcur (#)
T6360003	6	506.92	3404:17	44:20:53	7.99	8.00	100
T6360017	2	199.58	8:37	29:23:50	0.75	0.75	100
T6360004	6	99.84	989:37	62:00:00	1.33	2.00	100
T6360005	2	0.77	0:16	5:23:06	0.55	2.00	100
T6360015	4	0.15	9:42	18:23:04	0.34	0.75	100
T6360035	2	0.11	0:26	7:18:15	0.77	1.00	100
T6360027	2	0.07	2:53	62:21:46	0.75	0.75	100
T6360049	2	0.06	1:27	61:17:35	0.65	1.00	100
T6360010	6	0.06	5:55	61:20:56	0.83	1.00	100
T6360021	2	0.06	1:04	48:19:08	0.34	4.00	100
T6360048	2	0.04	0:27	49:00:51	0.29	1.00	100
T6360016	2	0.04	6:09	34:19:37	0.30	0.75	100
T6360008	2	0.04	3:49	47:23:10	0.35	0.75	100
T6360006	2	0.03	0:57	25:20:37	0.54	1.00	100
NSLCF1	1	0.01	0:02	62:21:46	0.03	0.25	500
VTAM	1	0.00	0:01	62:21:46	0.01	0.03	100
T6360023	2	0.00	0:04	6:21:20	0.46	0.75	100
PERFSVM	1	0.00	2:12	7:18:04	0.05	0.06	0
AUTOVM	1	0.00	0:03	62:21:46	0.00	0.03	100
FTPSEVE	1	0.00	0:00	62:21:47	0.01	0.03	100
TCPIP	1	0.00	0:01	62:21:47	0.01	0.12	3000
DATAMOVE	1	0.00	0:06	62:21:47	0.00	0.03	100
VMSERVU	1	0.00	0:00	62:21:47	0.00	0.03	1500
OPERSYMP	1	0.00	0:00	62:21:47	0.00	0.03	100

hyptop – Displaying hypervisor performance data

Displaying performance data for a single LPAR

```
10:16:59 H05LP30 CPU-I: IFL(18) CP(3) UN(2)                                     ?=help
```

cpu_id (#)	type (str)	cpu (%)	mgm (%)	visual (vis)
0	IFL	29.34	0.72	#####
1	IFL	28.17	0.70	#####
2	IFL	32.86	0.74	#####
3	IFL	31.29	0.75	#####
4	IFL	32.86	0.72	#####
5	IFL	30.94	0.68	#####
6	IFL	0.00	0.00	
7	IFL	0.00	0.00	
8	IFL	0.00	0.00	
9	IFL	0.00	0.00	
=:V:N		185.46	4.30	



Upstream Updates

Version 1.10.0

- New tools :None
- Changes to existing tools:
 - chchp: Use /proc/cio_settle
 - znetconf: Add support for new CHPIDs OSX and OSM introduced with z196
- Bug Fixes

Version 1.11.0

- New tools :None
- Changes to existing tools:
 - cmsfs-fuse: Add write support With this support it is possible to add, delete, and modify CMS files under Linux
 - zipl: Add support for automatic menus When the keyword "defaultauto" is specified in the defaultboot section of a zipl.conf file, zipl will automatically build and install a boot menu including all IPL sections listed in the configuration file
- Bug Fixes



Upstream Updates: 1.12.0 beyond hyptop

Changes to existing tools

- **chreipl:** various enhancements
 - Add support to re-IPL from named saved systems (NSS)
 - Add support to specify additional kernel parameters for re-IPL
 - Add "auto target" support
 - Add support to re-IPL from device-mapper multipath devices
- **cio_ignore:** Add new option `-i / --is-ignored` to determine if a device with a given ID is on the blacklist

```
# cio_ignore --is-ignored 0.0.0190
Device 0.0.0190 is ignored
# cio_ignore -i 0.0.0009
Device 0.0.0009 is not ignored
```
- **cmsfs-fuse:** Add a configuration file for automatic translation from EBCIDC to ASCII based on the file type
- **tunedasd:** Add new option `-Q / --query_reserve` to determine the reservation status of a device
Prerequisites: kernel 2.6.37 feature "dasd: API to query DASD reservation status"
- **zgetdump:** Add `kdump` support for `--info` option
- **zfcpdump/zipl:** Disable automatic activation of LUNs
- **Bug Fixes**



Current Linux on System z Technology

Features & Functionality contained in the Novell
& Red Hat Distributions



How can you read files on a CMS disk with Linux?

About the CMS user space file system (fuse) support



11.1

- Allows to mount a z/VM minidisk to a Linux mount point
- z/VM minidisk needs to be in the enhanced disk format (EDF)
- The cmsfs fuse file system transparently integrates the files on the minidisk into the Linux VFS, no special command required

```
root@larsson:~> cmsfs-fuse /dev/dasde /mnt/cms
root@larsson:~> ls -la /mnt/cms/PROFILE.EXEC
-r--r----- 1 root root 3360 Jun 26 2009
/mnt/fuse/PROFILE.EXEC
```

- By default no conversion is performed
 - Mount with '-t' to get automatic EBCDIC to ASCII conversion

```
root@larsson:~> cmsfs-fuse -t /dev/dasde /mnt/cms
```

- Write support is work in progress, almost completed
 - use “vi” to edit PROFILE.EXEC anyone ?
- Use fusermount to unmount the file system again

```
root@larsson:~> fusermount -u /mnt/cms
```



lsmem - Show online status information about memory blocks



11.1

The lsmem command lists the ranges of available memory with their online status.

- The listed memory blocks correspond to the memory block representation in sysfs.
- The command also shows the memory block size, the device size, and the amount of memory in online and offline state.

The output of this command, shows ranges of adjacent memory blocks with similar attributes.

```
root@larsson:~> lsmem
```

```
Address range                               Size (MB) State Removable Device
```

```
=====
```

```
0x0000000000000000-0x000000000ffffffff 256      online   no       0
0x0000000010000000-0x000000002ffffffff 512      online   yes      1-2
0x0000000030000000-0x000000003ffffffff 256      online   no       3
0x0000000040000000-0x000000006ffffffff 768      online   yes      4-6
0x0000000070000000-0x00000000ffffffff 2304     offline  -       7-15
```

```
Memory device size : 256 MB
```

```
Memory block size : 256 MB
```

```
Total online memory : 1792 MB
```

```
Total offline memory: 2304 MB
```



chmem - Set memory online or offline



11.1

The chmem command sets a particular size or range of memory online or offline.

- Setting memory online can fail if the hypervisor does not have enough memory left, for example because memory was overcommitted. Setting memory offline
- can fail if Linux cannot free the memory.
- If only part of the requested memory can be set online or offline, a message tells you how much memory was set online or offline instead of the requested amount.

This command requests 1024 MB of memory to be set online.

```
root@larsson:~> chmem --enable 1024
```

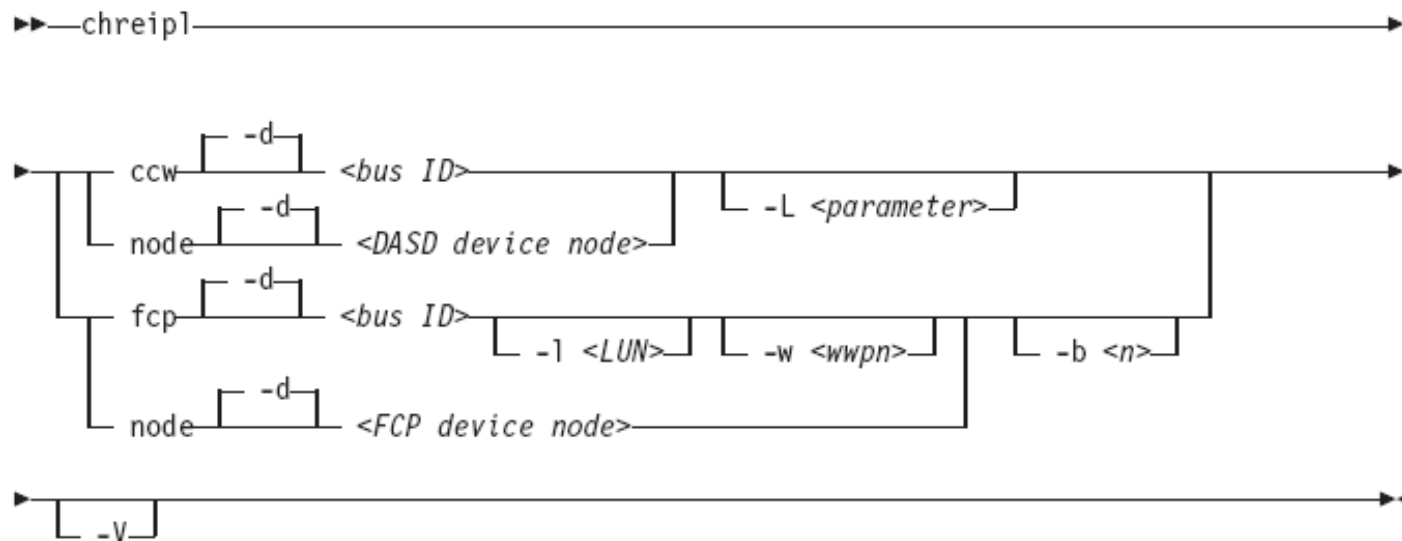
This command requests the memory range starting with 0x00000000e4000000 and ending with 0x00000000f3ffffff to be set offline.

```
root@larsson:~> chmem --disable 0x00000000e4000000 -  
0x00000000f3ffffff
```



Shutdown action tools

chreipl: Configure a disk or change a an entry in the Boot menu for the next boot cycle.



```

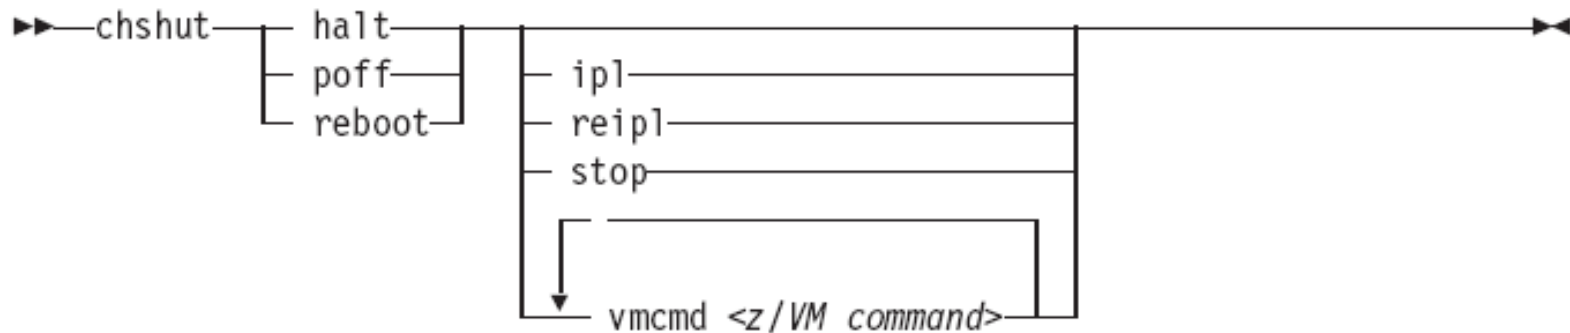
root@larsson:~> chreipl node /dev/dasda
root@larsson:~> chreipl node /dev/sda
root@larsson:~> chreipl ccw -d 0.0.7e78 -L 1
root@larsson:~> chreipl fcp --wwpn 0x500507630300c562 \
--lun 0x401040B300000000 -d 0.0.1700

```



Shutdown action tools (cont.)

chshut: Change the entries in `/sys/firmware` to configure the shutdown behaviour



```

root@larsson:~> chshut halt ipl
root@larsson:~> chshut halt vmcmd LOGOFF
root@larsson:~> chshut poff vmcmd "MSG MASTER Going
down" \ vmcmd "LOGOFF"
  
```



Shutdown action tools (cont.)

lsreipl: command to see from which device your system will boot after you issue the `reboot` command. Further you can query the system for information about the current boot device.

```
root@larsson:~> lsreipl
Re-IPL type:      ccw
Device:          0.0.4bb8
Loadparm:
root@larsson:~> lsreipl -i
```

lsshut: command to see what the system should do in one of the following states.

```
root@larsson:~> lsshut
Trigger          Action
=====
Halt             stop
Panic            stop
Power off        stop
Reboot           reipl
```



11

5.4



Isluns

Use the **Isluns** command to discover and scan LUNs in Fibre Channel Storage Area Networks (SANs).



This example shows all LUNs for port 0x500507630300c562:

```
root@larsson:~> isluns --port 0x500507630300c562  
Scanning for LUNs on adapter 0.0.5922  
at port 0x500507630300c562:  
0x4010400000000000  
0x4010400100000000  
[...]
```


This example shows all LUNs for adapter 0.0.5922:

```
root@larsson:~> isluns -c 0.0.5922  
at port 0x500507630300c562:  
0x4010400000000000  
[...]  
at port 0x500507630303c562:  
0x4010400000000000  
[...]
```



lszcrypt

Use the **lszcrypt** command to display information about cryptographic adapters managed by zcrypt and zcrypt's AP bus attributes



- To display card type and online status of all available cryptographic adapters:

```
root@larsson:~> lszcrypt -V
```

- To display card type, online status, hardware card type, hardware queue depth, and request count for cryptographic adapters 0, 1, 10, and 12

```
root@larsson:~> lszcrypt -VV 0 1 10 12
```


- To display AP bus information:

```
root@larsson:~> lszcrypt -b
```



chzcrypt

Use the **chzcrypt** command to configure cryptographic adapters managed by zcrypt and modify zcrypt's AP bus attributes.



To set the cryptographic adapters 0, 1, 4, 5, and 12 online:

```
root@larsson:~> chzcrypt -e 0 1 4 5 12
```

To set all available cryptographic adapters offline:

```
root@larsson:~> chzcrypt -d -a
```

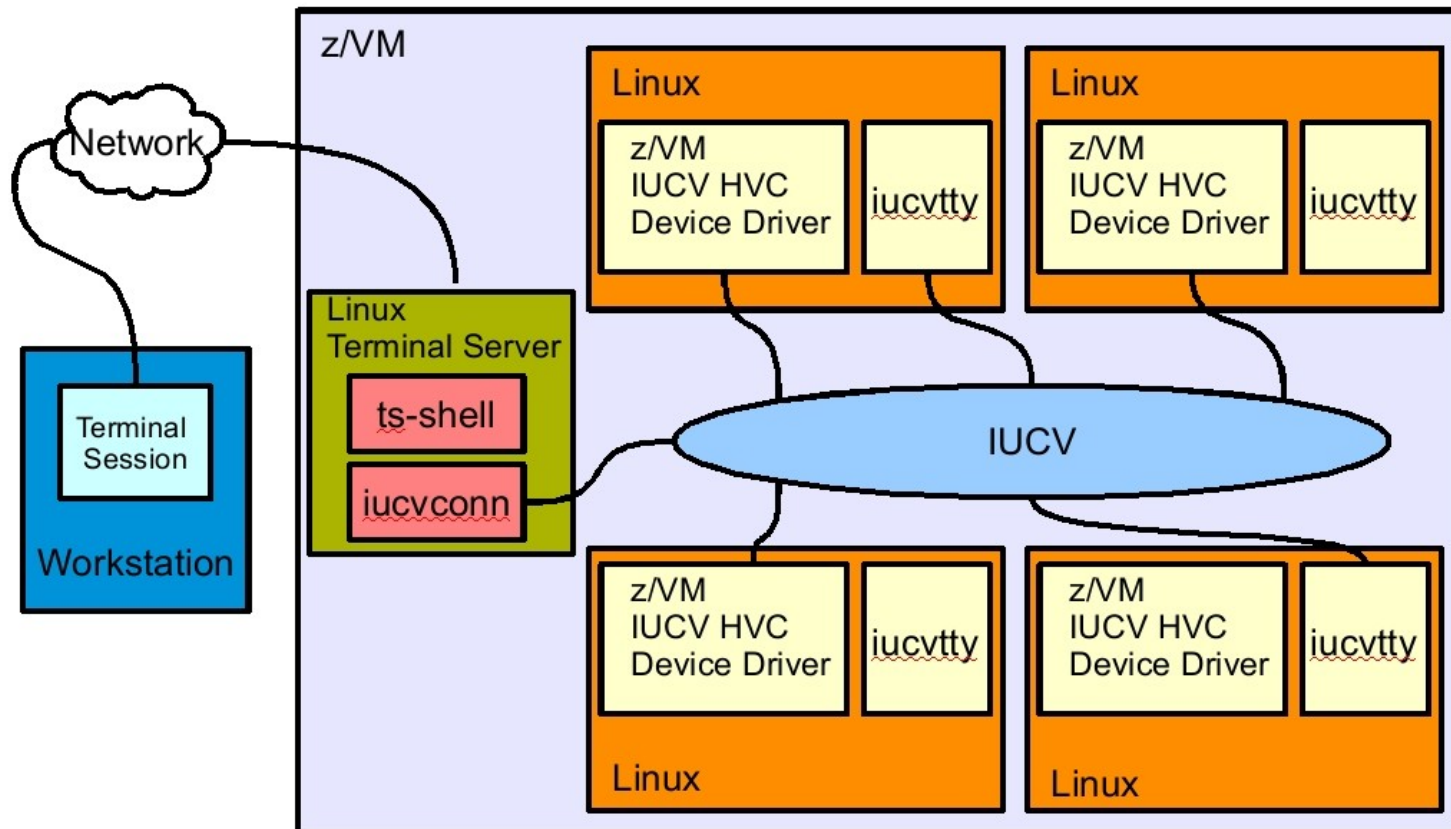
To set the configuration timer for re-scanning the AP bus to 60 seconds and disable zcrypt's poll thread:

```
root@larsson:~> chzcrypt -c 60 -n
```



IUCV Terminals

- Full-screen terminal access to Linux guest operating systems on the same z/VM
- Access Linux instances with no external network because IUCV is independent from TCP/IP



IUCV terminal applications

- The IUCV terminal applications consist of:
 - **iucvconn** – Start terminal connection over IUCV
 - **iucvtty** – Allow remote logins over IUCV
 - **ts-shell** – Login shell for terminal servers over IUCV
- Terminal access over IUCV is provided by:
 - **iucvtty**, or
 - z/VM IUCV hypervisor console device driver (Linux kernel)



IUCV terminal applications – examples

Using the `iucvconn` program: To access the first z/VM IUCV HVC terminal on the Linux instance in z/VM guest LNXSYS02

```
root@larsson:~> iucvconn LNXSYS02 lnxhvc0
```

To create a transcript of the terminal session to the Linux instance in z/VM guest LNXSYS99

```
root@larsson:~> iucvconn -s ~/transcripts/lnxsys99 LNXSYS99 lnxhvc0
```

Using the `iucvtty` program: To allow remote logins using the terminal identifier „lnxterm“

```
root@larsson:~> iucvtty lnxterm
```

To access the „lnxterm“ terminal on the Linux instance in z/VM guest LNXSYS01

```
root@larsson:~> iucvconn LNXSYS01 lnxterm
```

To use `/sbin/sulogin` instead of `/bin/login` for terminal “suterm”

```
root@larsson:~> iucvtty suterm -- /sbin/sulogin
```



cpuplugd



Use the **cpuplugd** command to:

- Enable or disable CPUs based on a set of rules. This increases the performance of single threaded applications within a z/VM or LPAR environment with multiple CPUs.
- The rules can incorporate certain system load variables.
- Manage memory under z/VM.

- Configuration file: */etc/sysconfig/cpuplugd*
- Init-Script: */etc/init.d/cpuplugd {start, stop, restart}*



cpuplugd: Example Configuration

```
UPDATE="60"
```

```
CPU_MIN="2"
```

```
CPU_MAX="10"
```

```
HOTPLUG = "(loadavg > onumcpus +0.75) & (idle < 10.0)"
```

```
HOTUNPLUG = "(loadavg < onumcpus -0.25) | (idle > 50)"
```

```
CMM_MIN="0"
```

```
CMM_MAX="8192"
```

```
CMM_INC="256"
```

```
MEMPLUG = "swaprate > freemem+10 & freemem+10 < apcr"
```

```
MEMUNPLUG = "swaprate > freemem + 10000"
```



dbginfo.sh

dbginfo.sh is a script to collect various system related files, for debugging purposes.

- It generates a tar-archive which can be attached to PMRs / Bugzilla entries
- It is similar to the RedHat tools sosreport & the Novell support-utils

```
root@larsson:~> dbginfo.sh  
Create target directory /tmp/DBGINFO-2010-02-25-22-06-20-  
t6345057  
Change to target directory /tmp/DBGINFO-2010-02-25-22-06-  
20-t6345057  
Get procfs entries  
Saving runtime information into runtime.out  
Get file list of /sys  
Get entries of /sys  
[...]
```

Please use the data from this tool if you open a Bugzilla (Novell/RedHat) or a PMR



Multi Volume Dump

zipl can now dump to multiple DASDs. It is now possible to dump system images, which are larger than a single DASD.



10.3




5.4

You can specify up to 32 ECKD DASD partitions for a multi-volume dump

What are dumps good for?

- Full snapshot of system state taken at any point in time (e.g. after a system has crashed, or on a running system)
- Can be used to analyse system state beyond messages written to the syslog
- Internal data structures not exported to anywhere



Obtain message, which have not been written to the syslog due to a crash



Multi Volume Dump

How to prepare a set of ECKD DASD devices for a multi-volume dump? (32-bit systems only).



10.3

- We use two DASDs in this example:



5.4

```
root@larsson:~> dasdfmt -f /dev/dasdc -b 4096  
root@larsson:~> dasdfmt -f /dev/dasdd -b 4096
```

- Create the partitions with `fdasd`. The sum of the partition sizes must be sufficiently large (the memory size + 10 MB):

```
root@larsson:~> fdasd /dev/dasdc  
root@larsson:~> fdasd /dev/dasdd
```

- Create a file called `sample_dump_conf` containing the device nodes (e.g. `/dev/dasda1`) of the two partitions, separated by one or more line feed characters
- Prepare the volumes using the `zipl` command.

```
root@larsson:~> zipl -M sample_dump_conf  
[...]
```



How to obtain a dump

To obtain a dump with the multi-volume DASD dump tool, perform the following steps:

- Stop all CPUs, Store status on the IPL CPU.
- IPL the dump tool using one of the prepared volumes, either 4711 or 4712.
- After the dump tool is IPLed, you'll see a messages that indicates the progress of the dump. Then you can IPL Linux again

```
==> cp cpu all stop  
==> cp cpu 0 store status  
==> cp ipl 4711
```

- Copying a multi-volume dump to a file
- Use zgetdump command without any option to copy the dump parts to a file:

```
root@larsson:~> zgetdump /dev/dasdc > mv_dump_file
```



How to obtain information about a multi volume dumps

Display information on the involved volumes:

```
root@larsson:~> zgetdump -d /dev/dasdc
'/dev/dasdc' is part of Version 1 multi-volume dump, which is
spread along the following DASD volumes:
0.0.4711 (online, valid)
0.0.4712 (online, valid)
[...]
```

Display information about the dump itself:

```
root@larsson:~> zgetdump -i /dev/dasdc
Dump device: /dev/dasdc
>>> Dump header information <<<
Dump created on: Thu Feb 25 15:12:41 2010
[...]
Multi-volume dump: Disk 1 (of 2)
Reading dump contents from
0.0.4711.....
Dump ended on: Thu Feb 25 15:12:52 2010
Dump End Marker found: this dump is valid.
```



Handling large dumps

Compress the dump and split it into parts of 1 GB

```
root@larsson:~> zgetdump /dev/dasdc1 | gzip | split -b 1G
```

Several compressed files such as xaa, xab, xac, are created

Create md5 sums of the compressed files

```
root@larsson:~> md5sum xa* > dump.md5
```

Upload all parts together with the md5 information.

Verification of the parts for a receiver

```
root@larsson:~> md5sum -c dump.md5  
xaa: OK  
[....]
```

Merge the parts and uncompress the dump

```
root@larsson:~> cat xa* | gunzip -c > dump
```



Transferring dumps

Transferring single volume dumps with ssh

```
root@larsson:~> zgetdump /dev/dasdc1 | ssh user@host "cat >
dump_file_on_target_host"
```

Transferring multi-volume dumps with ssh

```
root@larsson:~> zgetdump /dev/dasdc | ssh user@host "cat >
multi_volume_dump_file_on_target_host"
```

Transferring a dump with ftp.

Establish an ftp session with the target host, login and set the transfer mode to Binary. Send the dump to the host

```
root@larsson:~> ftp> put |"zgetdump /dev/dasdc1"
<dump_file_on_target_host>
```



dumpconf

The dumpconf tool configures a dump device that is used for automatic dump in case of a kernel panic.



- The command can be installed as service script under `/etc/init.d/dumpconf` or can be called manually.
- Start service: `service dumpconf start`
- It reads the configuration file `/etc/sysconfig/dumpconf`.
- Example configuration for CCW dump device (DASD) and reipl after dump:

```
ON_PANIC=dump_reipl
DUMP_TYPE=ccw
DEVICE=0.0.4711
```



dumpconf (cont.)

Example configuration for FCP dump device (SCSI disk):

```
ON_PANIC=dump
DUMP_TYPE=fcp
DEVICE=0.0.4714
WWPN=0x5005076303004712
LUN=0x4047401300000000
BOOTPROG=0
BR_LBA=0
```

Example configuration for re-IPL without taking a dump, if a kernel panic occurs:

```
ON_PANIC=reipl
```

Example of executing a CP command, and rebooting from device 4711 if a kernel panic occurs (MASTER is a VM Guest Name):

```
ON_PANIC=vmcmd
VMCMD_1="MSG MASTER Starting VMDUMP"
VMCMD_2="VMDUMP"
VMCMD_3="IPL 4711"
```



vmcp

Using the z/VM CP interface device driver (vmcp), you can send control program (CP) commands to the VM hypervisor and display VM's response.

```
root@larsson:~> modprobe vmcp
root@larsson:~> vmcp "q dasd"|grep T6345057
DASD 4DE0 ATTACHED TO T6345057 4DE0 R/W 0X4DE0
DASD 4DE1 ATTACHED TO T6345057 4DE1 R/W 0X4DE1
DASD 4DE2 ATTACHED TO T6345057 4DE2 R/W 0X4DE
DASD 4DE3 ATTACHED TO T6345057 4DE3 R/W 0X4DE3
```

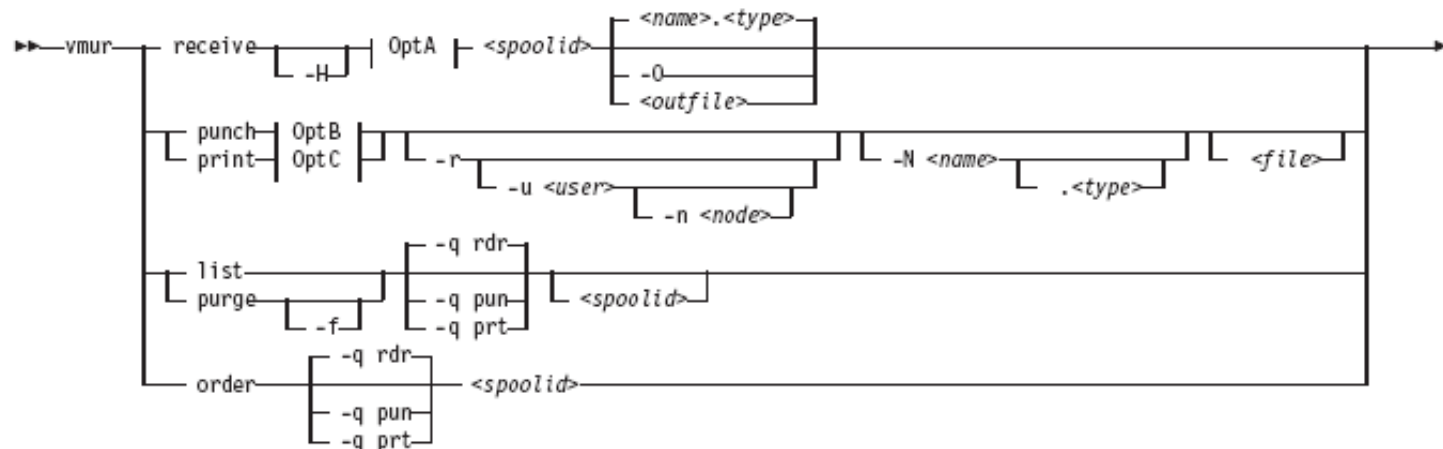


vmur

The **vmur** command provides all functions required to work with z/VM spool file queues:



- Receive: Read data from the z/VM reader file queue
- Punch or print: Write data to the z/VM punch or printer file queue and transfer it to another user's virtual reader, optionally on a remote z/VM node.
- List: Display detailed information about one or all files on the specified spool file queue.
- Purge: Remove one or all files on the specified spool file queue.
- Order: Position a file at the top of the specified spool file queue.



vmur: Log and read Linux guest machine console

Begin console spooling:

```
root@larsson:~> vmcp sp cons start
```

Produce output to VM console (for example, with CP TRACE).

Close the console file and transfer it to the reader queue, find the spool ID behind the FILE keyword in the corresponding CP message.

```
root@larsson:~> vmcp sp cons clo \* rdr  
RDR FILE 0398 SENT FROM T6360025 CON WAS 0398 RECS 1872  
CPY 001 T NOHOLD NOKEEP
```

Read the guest machine console file into a file on the Linux file system in the current working directory:

```
root@larsson:~> chccwdev -e 000c  
root@larsson:~> vmur re -t 398 linux_cons
```



cio_ignore

When a Linux on System z instance boots, it senses and analyses all available devices. You can use the `cio_ignore` kernel parameter to specify a list of devices that are to be ignored.

The following applies to ignored devices:

- Ignored devices are not sensed and analyzed. The device cannot be used unless it has been analyzed.
- Ignored devices are not represented in `sysfs`.
- Ignored devices do not occupy storage in the kernel.
- The subchannel to which an ignored device is attached is treated as if no device were attached.
- `cio_ignore` might hide essential devices such as the console under z/VM. The console is typically device number 0.0.0009.
-

This example specifies that all devices in the range 0.0.b100 through 0.0.b1ff, and the device 0.0.a100 are to be ignored.

```
cio_ignore=0.0.b100-0.0.b1ff,0.0.a100
```



cio_ignore (cont.)

Display ignored devices:

```
root@larsson:~> cat /proc/cio_ignore
0.0.0000-0.0.78ff
0.0.f503-0.0.ffff
```

Free a individual device from the ignore list

```
root@larsson:~> echo free 0.0.4711 >/proc/cio_ignore
```

Free all devices from the ignore list

```
root@larsson:~> echo free all >/proc/cio_ignore
```

Use cio_ignore tool to manage the I/O device exclusion list

```
root@larsson:~> cio_ignore -l
Ignored devices:
=====
0.0.0000-0.0.0008
0.0.000a-0.0.6365
[...]
```



cio_ignore (cont'd)

Use the -L option to display the devices which are accessible

```
root@larsson:~> cio_ignore -L  
Accessible devices:  
=====  
0.0.0009  
0.0.6366  
0.0.f5f0-0.0.f5f2
```

Use the -r option to remove devices from the exclusion list

```
root@larsson:~> cio_ignore -r 6366
```

The the -R option is used to free all devices

Use the -a option to add devices to the exclusion list

```
root@larsson:~> cio_ignore -a 4000-5fff
```

Use the -k option to create the kernel parameter list string

```
root@larsson:~> cio_ignore -k  
cio_ignore=all,!0009,!6366,!f5f0-f5f2
```




More Information

<http://www.ibm.com/developerworks/linux/linux390/>

Linux on System z



How to use Execute-in-Place Technology with Linux on z/VM
March, 2010



Linux on System z



How to use FC-attached SCSI devices with Linux on System z

Development stream (Kernel 2633)

Using the Dump Tools

Development stream (Kernel 2633)



Linux on System z



How to Set up a Terminal Server Environment on z/VM
June 2009

Linux Kernel 26 - Development stream

Linux on System z



Kernel Messages

Development stream (Kernel 2633)

Linux on System z



Device Drivers, Features, and Commands

Development stream (Kernel 2633)

New: Distribution specific Documentation

SC94-2584-01

SC93-8413-04



More Information




z/VM and Linux on IBM System z The Virtualization Cookbook for Red Hat Enterprise Linux 6.0

Hands-on instructions for installing z/VM and Linux on the mainframe

Updated information for z/VM V6.1 and Red Hat Enterprise Linux 6.0

New, more versatile file system layout



Brad
Michael M

Redbook

ibm.com/redbooks



z/VM and Linux on IBM System z The Virtualization Cookbook for SLES 11 SP1

Hands-on instructions for installing z/VM and Linux on the mainframe

Updated information for z/VM 6.1 and Linux SLES 11 SP1

A new, more versatile file system layout



Your Linux on System z Requirements?

Are you missing a certain feature, functionality or tool? **We'd love to hear from you!**

We will evaluate each request and (hopefully) develop the additional functionality you need.


Send your input to hans@de.ibm.com



Questions?



Hans-Joachim Picht
Linux Technology Center



*IBM Deutschland Research
& Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany*

*Phone +49 (0)7031-16-1810
Mobile +49 (0)175 - 1629201
hans@de.ibm.com*



Trademarks & Disclaimer

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DB2, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, System Storage, System z9, VM/ESA, VSE/ESA, WebSphere, xSeries, z/OS, zSeries, z/VM.

The following are trademarks or registered trademarks of other companies

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries. LINUX is a registered trademark of Linux Torvalds in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC. Intel is a registered trademark of Intel Corporation. * All other products may be trademarks or registered trademarks of their respective companies.

NOTES: Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography. References in this document to IBM products or services do not imply that IBM intends to make them available in every country. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

