

Ready; T=0.01/0.01 20:49:10  
ipl 160 clear  
zIPL v1.8.0-44.45.2 interactive boot menu

0. default (ipl)

1. ipl

Note: VM users please use '#cp vi vmsg <number> <kernel-parameters>'

Please choose (default will boot in 10 seconds):

Booting default (ipl)...

Initializing cgroup subsys cpuset

Initializing cgroup subsys cpu

Linux version 2.6.32.36-0.5-default (geeko@buildhost) (gcc version 4.3.4 [gcc-4\_3-branch revision 152973] (SUSE Linux) ) #1 SMP 2011

-04-14 10:12:31 +0200

setup.1a06a7: Linux is running as a z/VM guest operating system in 64-bit mode

Zone PFN ranges:

DMA 0x00000000 -> 0x00080000

Normal 0x00080000 -> 0x00080000

Movable zone start PFN for each node

early\_node\_map[1] active PFN ranges

0: 0x00000000 -> 0x00020000

PERCPU: Embedded 11 pages/cpu @0000000000e92000 s12544 r8192 d24320 u65536

pcpu-alloc: s12544 r8192 d24320 u65536 alloc=16\*4096

pcpu-alloc: [0] 00 [0] 01 [0] 02 [0] 03 [0] 04 [0] 05 [0] 06 [0] 07

pcpu-alloc: [0] 08 [0] 09 [0] 10 [0] 11 [0] 12 [0] 13 [0] 14 [0] 15

pcpu-alloc: [0] 16 [0] 17 [0] 18 [0] 19 [0] 20 [0] 21 [0] 22 [0] 23

pcpu-alloc: [0] 24 [0] 25 [0] 26 [0] 27 [0] 28 [0] 29 [0] 30 [0] 31

pcpu-alloc: [0] 32 [0] 33 [0] 34 [0] 35 [0] 36 [0] 37 [0] 38 [0] 39

pcpu-alloc: [0] 40 [0] 41 [0] 42 [0] 43 [0] 44 [0] 45 [0] 46 [0] 47

pcpu-alloc: [0] 48 [0] 49 [0] 50 [0] 51 [0] 52 [0] 53 [0] 54 [0] 55

pcpu-alloc: [0] 56 [0] 57 [0] 58 [0] 59 [0] 60 [0] 61 [0] 62 [0] 63

Built 1 zonelists in Zone order, mobility grouping on. Total pages: 129280

Kernel command line: root=/dev/disk/by-path/ccw-0.0.0160-part1

TERM=dumb BOOT\_IMAGE=0

PID hash table entries: 2048 (order: 2, 16384 bytes)

Dentry cache hash table entries: 65536 (order: 7, 524288 bytes)

Inode-cache hash table entries: 32768 (order: 6, 262144 bytes)

Memory: 499540k/524288k available (4393k kernel code, 0k reserved, 2104k data, 228k init)

Write protected kernel read-only data: 0x100000 - 0x5ffffff

Hierarchical RCU implementation.

console [ttyS0] enabled

allocated 5242880 bytes of page\_cgroup

please try 'cgroup\_disable=memory' option if you don't want memory cgroups

pid\_max: default: 65536 minimum: 512

Security Framework initialized

AppArmor: AppArmor initialized

Mount-cache hash table entries: 256

Initializing cgroup subsys ns

Initializing cgroup subsys cpuacct

Initializing cgroup subsys memory

Initializing cgroup subsys devices

Initializing cgroup subsys freezer

Initializing cgroup subsys net\_cls

cpu.33a262: 1 configured CPUs, 0 standby CPUs

cpu.17772b: Processor 0 started, address 0, identification 018942

```
Brought up 1 CPUs
devtmpfs: initialized
NET: Registered protocol family 16
bio: create slab <bio-0> at 0
NetLabel: Initializing
NetLabel: domain hash size = 128
NetLabel: protocols = UNLABELED CIPSOv4
NetLabel: unlabeled traffic allowed by default
AppArmor: AppArmor Filesystem Enabled
NET: Registered protocol family 2
IP route cache hash table entries: 4096 (order: 3, 32768 bytes)
TCP established hash table entries: 16384 (order: 6, 262144 bytes)
TCP bind hash table entries: 16384 (order: 6, 262144 bytes)
TCP: Hash tables configured (established 16384 bind 16384)
TCP reno registered
NET: Registered protocol family 1
Unpacking initramfs...
Freeing initrd memory: 5647k freed
audit: initializing netlink socket (disabled)
type=2000 audit(1311986968.620:1): initialized
HugeTLB registered 1 MB page size, pre-allocated 0 pages
VFS: Disk quotas dquot_6.5.2
Dquot-cache hash table entries: 512 (order 0, 4096 bytes)
msgmni has been set to 246
alg: No test for stdrng (krng)
Block layer SCSI generic (bsg) driver version 0.4 loaded (major 254)
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered (default)
io scheduler cfq registered
cio.b5d5f6: Channel measurement facility initialized using format extended (mode
autodetected)
TCP cubic registered
registered taskstats version 1
Freeing unused kernel memory: 228k freed
doing fast boot
Creating device nodes with udev
udev version 128 started
dasd-eckd.90fb0d: 0.0.0160: New DASD 3390/0C (CU 3990/01) with
10017 cylinders, 15 heads, 224 sectors
dasd-eckd.412b53: 0.0.0160: DASD with 4 KB/block, 7212240 KB
total size, 48 KB/track, compatible disk layout
  dasda:VOL1/ 0X0160: dasda1 dasda2 dasda3
mount: devpts already mounted or /dev/pts busy
mount: according to mtab, devpts is already mounted on /dev/pts
Boot logging started on /dev/ttyS0 (/dev/console) at Fri Jul 29 20:49:29 2011
kjournald starting. Commit interval 15 seconds
EXT3 FS on dasda1, internal journal
EXT3-fs: mounted filesystem with ordered data mode.
resume device not found (ignoring)
Waiting for device /dev/disk/by-path/ccw-0.0.0160-part1 to
appear: ok
fsck from util-linux-ng 2.16
[/sbin/fsck.ext3 (1) -- /] fsck.ext3 -a /dev/dasda1
/dev/dasda1: clean, 4185/102400 files, 65982/102396 blocks
fsck succeeded. Mounting root device read-write.
Mounting root /dev/disk/by-path/ccw-0.0.0160-part1
```

```
mount -o rw,acl,user_xattr -t ext3 /dev/disk/by-path/ccw-
0.0.0160-part1 /root
INIT: version 2.86 booting
System Boot Control: Running /etc/init.d/boot
Mounting sysfs at /sys..done
Mounting debugfs at /sys/kernel/debug..done
Copying static /dev content..done
Mounting devpts at /dev/pts..done
Boot logging started on /dev/ttyS0(/dev/console) at Fri Jul 29 20:49:32 2011
udev version 128 started
SCSI subsystem initialized
vmur.c793e5: z/VM virtual unit record device driver loaded.
lcs.4fa72b: Loading LCS driver
qeth.87067b: loading core functions
lcs.0aabaa: LCS device eth0 without IPv6 support
lcs.86dbe7: LCS device eth0 with Multicast support
Adding 524872k swap on /dev/dasda2. Priority:-1 extents:1 across:524872k
device-mapper: uevent: version 1.0.3
device-mapper: ioctl: 4.15.0-ioctl (2009-04-01) initialised: dm-devel@redhat.com
Starting udevd: ..done
Loading drivers, configuring devices: ..done
Activating swap-devices in /etc/fstab...
..done
Setting up the System Clock..done
..done
Activating device mapper...
..done
Loading required kernel modules
..done
Starting MD Raid ln: creating symbolic link `/var/run/mdadm': No such file or
directory
..unused
Waiting for udev to settle...
Scanning for LVM volume groups...
  Reading all physical volumes. This may take a while...
  Found volume group "vg1" using metadata type lvm2
Activating LVM volume groups...
  6 logical volume(s) in volume group "vg1" now active
..done
blogd: no message logging because /var file system is not accessible
Waiting for /dev/vg1/usrx . no more events
Checking file systems...
fsck from util-linux-ng 2.16
Checking all file systems.
[/sbin/fsck.ext3 (1) -- /home] fsck.ext3 -a /dev/mapper/vg1-home
/dev/mapper/vg1-home: clean, 30/25600 files, 1863/25600 blocks
[/sbin/fsck.ext3 (1) -- /opt] fsck.ext3 -a /dev/mapper/vg1-opt
/dev/mapper/vg1-opt: clean, 11/19456 files, 1647/19456 blocks
[/sbin/fsck.ext2 (1) -- /repository] fsck.ext2 -a /dev/mapper/vg1-vmarc
/dev/mapper/vg1-vmarc: clean, 3477/131072 files, 14677/524288 blocks
[/sbin/fsck.ext3 (1) -- /tmp] fsck.ext3 -a /dev/mapper/vg1-tmp
/dev/mapper/vg1-tmp: clean, 303/76800 files, 13287/76800 blocks
[/sbin/fsck.ext3 (1) -- /usr] fsck.ext3 -a /dev/vg1/usrx
fsck.ext3: No such file or directory while trying to open
/dev/vg1/usrx
/dev/vg1/usrx:
```

The superblock could not be read or does not describe a correct ext2 filesystem. If the device is valid and it really contains an ext2 filesystem (and not swap or ufs or something else), then the superblock is corrupt, and you might try running e2fsck with an alternate superblock:  
e2fsck -b 8193 <device>

```
[/sbin/fsck.ext3 (1) -- /var] fsck.ext3 -a /dev/mapper/vg1-var  
/dev/mapper/vg1-var: clean, 1887/103808 files, 79400/103424  
blocks  
..failed  
fsck failed for at least one filesystem (not /).  
Please repair manually and reboot.  
The root file system is is already mounted read-write.
```

Attention: Only CONTROL-D will reboot the system in this maintenance mode. shutdown or reboot will not work.

Give root password for login:  
passwordforroot

```
(repair filesystem) #  
df  
Filesystem          1K-blocks      Used Available Use% Mounted on  
/dev/dasda1         396728        251072    125180   67% /  
(repair filesystem) #  
mount  
/dev/dasda1 on / type ext3 (rw,acl,user_xattr)  
(repair filesystem) #  
cd /etc/  
(repair filesystem) #  
cat fstab  
/dev/disk/by-path/ccw-0.0.0160-part2 swap          defaults  
0 0  
/dev/disk/by-path/ccw-0.0.0160-part1 /              ext3          acl,user_xattr  
1 1  
/dev/vg1/home       /home          ext3          acl,user_xattr    1 2  
/dev/vg1/opt        /opt           ext3          acl,user_xattr    1 2  
/dev/vg1/vmarc     /repository    ext2          acl,user_xattr    1 2  
/dev/vg1/tmp        /tmp           ext3          acl,user_xattr    1 2  
/dev/vg1/usrx       /usr           ext3  
acl,user_xattr      1 2  
/dev/vg1/var        /var           ext3          acl,user_xattr    1 2  
proc                /proc          proc          defaults          0 0  
sysfs               /sys           sysfs         noauto            0 0  
debugfs             /sys/kernel/debug debugfs        noauto            0 0  
devpts              /dev/pts       devpts        mode=0620,gid=5   0 0  
(repair filesystem) #
```

```
sed -e 's/usrx/usr/' fstab
```

```
/dev/disk/by-path/ccw-0.0.0160-part2 swap swap defaults
0 0
/dev/disk/by-path/ccw-0.0.0160-part1 / ext3 acl,user_xattr
1 1
/dev/vg1/home /home ext3 acl,user_xattr 1 2
/dev/vg1/opt /opt ext3 acl,user_xattr 1 2
/dev/vg1/vmarc /repository ext2 acl,user_xattr 1 2
/dev/vg1/tmp /tmp ext3 acl,user_xattr 1 2
/dev/vg1/usr /usr ext3
acl,user_xattr 1 2
/dev/vg1/var /var ext3 acl,user_xattr 1 2
proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
(repair filesystem) #
```

```
sed -i -e 's/usrx/usr/' fstab
```

```
(repair filesystem) #
# mount -a
#
exit
Unmounting file systems (ignore error messages)
/dev/dasda1 umounted
md: stopping all md devices.
zIPL v1.8.0-44.45.2 interactive boot menu
```

```
0. default (ipl)
```

```
1. ipl
```

Note: VM users please use '#cp vi vmsg <number> <kernel-parameters>'

Please choose (default will boot in 10 seconds):