



Linux Bootloaders on System z


Current & Future Implementations

John Jolly
SUSE

March 13th, 2014
Session 14802

www.SHARE.org





SHARE


Technology • Connections • Results

Agenda

- Boot Process on System Z
 - Initial Program Load (IPL)
 - zIPL and limitations
 - Grub 2 function
 - Grub 2 s390 implementation



IPL Process




SHARE

Technology • Connections • Results


Initial Program Load (IPL)

- Formalized with System 360
 - Set the dials, push the button
- Initiated from CP console for z/VM or HMC
- 1 PSW and 2 CCW loaded from
 - That's enough to load a Command Program



IPL Process for Linux on System Z

Linux Initial Program Load (zIPL)




SHARE

Technology • Connections • Results

- Linux Bootloader for System Z
- Configured with `/etc/zipl.conf`
- Bootloader written using `/sbin/zipl`

ziPL Configuration file (/etc/zipl.conf)



SHARE


Technology • Connections • Results

```
[defaultboot]
defaultmenu = menu
```

```
[SLES11_SP3]
  image = /boot/image-3.0.76-0.11-default
  target = /boot/zipl
  ramdisk = /boot/initrd-3.0.76-0.11-default,0x2000000
  parameters = "root=/dev/dasda2 hvc_iucv=8 TERM=dumb resume=/dev/dasda1"
```

```
:menu
  default = 1
  prompt = 1
  target = /boot/zipl
  timeout = 10
  1 = SLES11_SP3
```


zIPL Operation



SHARE

Technology • Connections • Results

- Multi-stage operation
 - Stage 0 – 24 bytes – 2 CCWs to load and TIC to Stage1
 - Stage 1 – 104 byte channel program to load Stage 1b
 - Stage 1b – 1k channel program to load Stage 2
 - Stage 2 – 8k maximum menu and kernel loader
 - Stage 3 – Kernel parameter parser and execution




SHARE

Technology • Connections • Results

ZIPL Limitations - Devices

- Limited device capabilities
 - Restricted to specific Channel I/O device types
 - DASD
 - SCSI
 - Tape
 - Boot from the Network? Nope.
 - Device-mapper tends to be finicky

zIPL Limitations – File Systems




SHARE

Technology • Connections • Results

- Limited file system support
 - Only ext2, ext3
 - Btrfs and XFS not supported

Grub 2 on System Z




SHARE

Technology • Connections • Results

Challenges of Grub 2 on System Z

- Large Binary Size
 - 1.3MB cannot be easily loaded via CCW
- Unique System Z Drivers
 - DASD, zFCP, QETH all should be cleanly implemented
 - Maintenance a major factor
- Unique Display Configuration
 - Both 3270 and 3215 need to be supported
-




SHARE

Technology • Connections • Results

Solution : grub2-emu

- User-mode grub2
 - Relatively little modification necessary
- Requires a linux kernel to be booted
 - All supported s390 drivers available
 - udev resolves devices
 - Still uses zIPL, but in a “Stage 4” role
- Full filesystem support
 - BTRFS snapshots
 - Unusual device-mapper configurations
- Two stage boot




SHARE


Grub 2 boot process

zIPL Stage



grub2-emu Stage






SHARE

Technology • Connections • Results

How to boot another kernel?

- kexec
 - Kernel-supported reboot
 - Safe shutdown and restart of devices
- Implementation of kexec in Grub 2
 - New feature for grub2-emu
 - Options for kernel and initrd are used and verified
 - Only available to root user



SHARE

Technology • Connections • Results

Configuration : grub2

- Special grub2 boot partition
 - Normally mounted to /boot/zipl
- Contains grub2 kernel and initrd
- Contains grub2 config file
 - Allows for dynamic configuration
 - No more need to run zipl when kernel configuration changes

Demonstration

Unpublished Work of SUSE. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE.

Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE.

Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

