Linux on System z Problem Determination: Real customer cases

Sven Schütz sven@de.ibm.com

Team Lead Linux on System z Service Team

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	IBM*	PowerVM	System z10	z/OS*
BladeCenter*	IBM eServer	PR/SM	WebSphere*	zSeries*
DataPower*	IBM (logo)*	Smarter Planet	z9*	z/VM*
DB2*	InfiniBand*	System x*	z10 BC	z/VSE
FICON*	Parallel Sysplex*	System z*	z10 EC	
GDPS*	POWER*	System z9*	zEnterprise	
HiperSockets	POWER7*			

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Windows Server and the Windows logo are trademarks of the Microsoft group of countries.
InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* 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.

Agenda

- Introduction
- DASD maintenance halts system
- Parallel IPL of many guests causes network problems
- IPL problems after zipl invocation
- Network connection stalls
- Lost access to SCSI storage during SVC maintenance
- Daily network recoveries
- Time-outs in HA environment
- Closing remarks

Introduction

- **Problem analysis looks straight forward on the charts but it might have taken weeks to get it done.**

  A problem does not necessarily show up on the place of origin

- **The more information is available, the sooner the problem can be solved, because gathering and submitting additional information again and again usually introduces delays.**
- **This presentation can only introduce some tools and how the tools can be used, comprehensive documentation on their capabilities is to be found in the documentation of the corresponding tool.**
- **Do not forget to update your systems**

DASD maintenance halts system (1/4)

- **Configuration:**
  - SLES11 system running under z/VM
  - DASD volumes attached
- **Problem description:**
  - For maintenance, some DASD CHPIDs were set offline
  - System halted
- **Tools used for problem determination:**
  - dbginfo.sh

DASD maintenance halts system (2/4)

lscss output:

```
Device   Subchan.  DevType CU Type Use  PIM PAM POM  CHPIDs
----------------------------------------------------------------------
0.0.0100 0.0.0000  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.0101 0.0.0001  3390/0c 3990/e9 yes  80  80  ff   e0000000 00000000
0.0.0700 0.0.0002  0000/00 3215/00 yes  80  80  ff   e0000000 00000000
0.0.0592 0.0.0003  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.000c 0.0.0015  0000/00 2540/00      80  80  ff   e0000000 00000000
0.0.000d 0.0.0016  0000/00 2540/00      80  80  ff   e0000000 00000000
0.0.000e 0.0.0017  0000/00 1403/00      80  80  ff   e0000000 00000000
0.0.0190 0.0.0018  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.019d 0.0.0019  3390/0c 3990/e9      80  80  ff   e0000000 00000000
```

What's interesting?

DASD maintenance halts system (3/4)

lscss output:

```
Device   Subchan.  DevType CU Type Use  PIM PAM POM  CHPIDs
----------------------------------------------------------------------
0.0.0100 0.0.0000  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.0101 0.0.0001  3390/0c 3990/e9 yes  80  80  ff   e0000000 00000000
0.0.0700 0.0.0002  0000/00 3215/00 yes  80  80  ff   e0000000 00000000
0.0.0592 0.0.0003  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.000c 0.0.0015  0000/00 2540/00      80  80  ff   e0000000 00000000
0.0.000d 0.0.0016  0000/00 2540/00      80  80  ff   e0000000 00000000
0.0.000e 0.0.0017  0000/00 1403/00      80  80  ff   e0000000 00000000
0.0.0190 0.0.0018  3390/0c 3990/e9      80  80  ff   e0000000 00000000
0.0.019d 0.0.0019  3390/0c 3990/e9      80  80  ff   e0000000 00000000
```

That's interesting!

DASD maintenance halts system (4/4)

- **Problem cause:**
  - Devices with no real physical CHPID are “attached” to a physical one
    - Puncher
    - Reader
    - Console
  - When CHPID needs to be “configured off” for maintenance, all attached devices are lost
  - When console is disconnected, system is down
- **Problem solution:**
  - Check your CHPID assignments before maintenance to prevent outages
  - z/VM working on APAR to pick free CHPID for such devices

Parallel IPL of many guests causes network problems (1/4)

- **Configuration:**
  - SLES11 SP1 system running under z/VM
  - A (relative) high number of IFLs per Linux (four in this case)
  - Regular maintenance window where all guests are IPLed at the same time
  - At least two network interfaces configured to layer 3
- **Problem description:**
  - When all servers are booted at the same time, some servers cannot bring up all networking interface correctly (one comes up in layer 3, one comes up in layer 2 ↔ conflict with z/VM VSWITCH configuration
  - Second reboot of problem-servers resolves the situation
- **Tools used for problem determination:**
  - dbginfo.sh

Parallel IPL of many guests causes network problems (2/4)

- **Looks like an easy catch:**
  - Customer running stock SLES11 SP1 kernel and misses some fixes
  - e.g. one bug "*qeth: wait for recovery finish in open function*" which seems likely
  - Customer upgraded to latest kernel
  - Next maintenance window: problem still occurring
- **Further investigation**
  - Problem only occurs on servers with low priority in z/VM
  - Resource problem, race condition?

Parallel IPL of many guests causes network problems (3/4)

Console output during IPL:

```
qeth.87067b: loading core functions
qeth_l3: module is already loaded
qeth.21a074: 0.0.1400: There is no kernel module to support
discipline 0
qeth.2c6def: register layer 3 discipline
qeth.933eb7: register layer 2 discipline
qeth.5cb8a3: 0.0.1400: The qeth device is not configured for the
OSI layer required by z/VM
qeth.3acf0c: 0.0.1400: The qeth device driver failed to recover
an error on the device
```

0 = layer 3

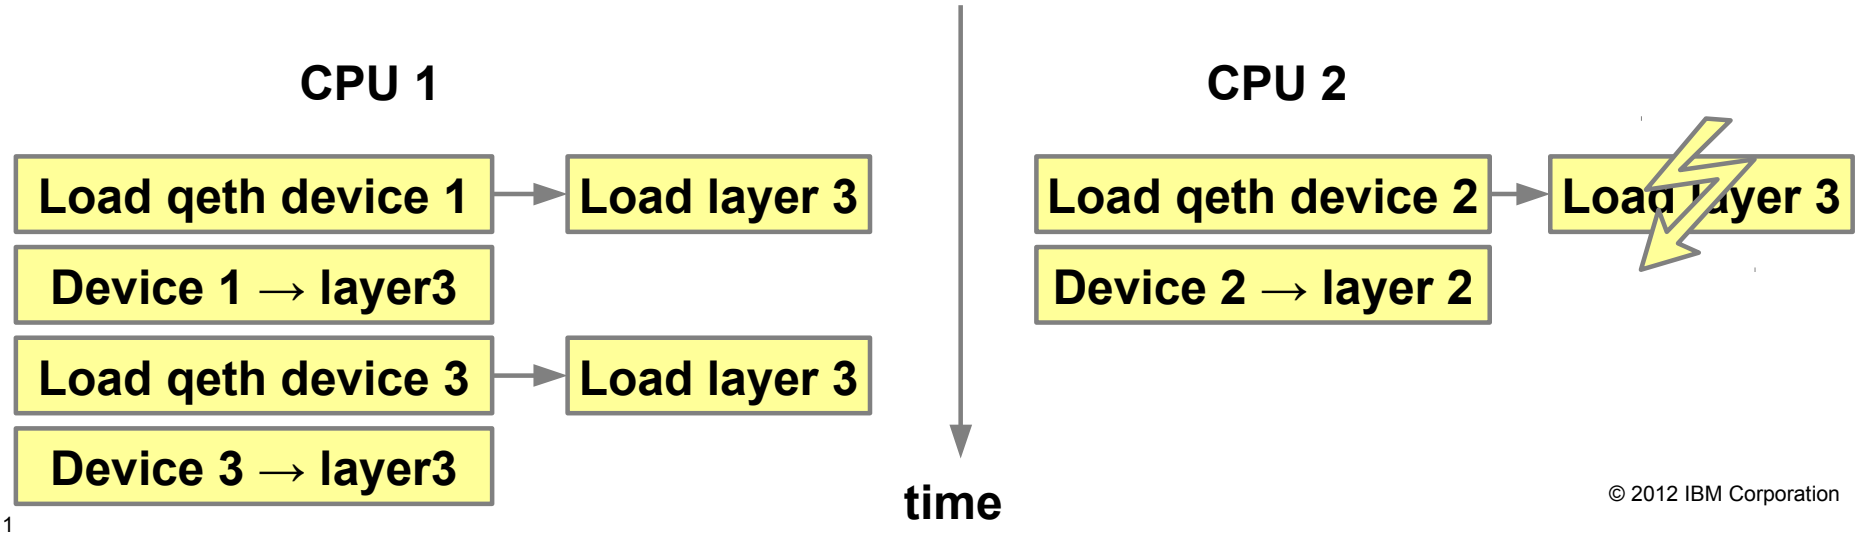

Parallel IPL of many guests causes network problems (4/4)

- **Solution: Introduce Locking in loading of layer 3 module**

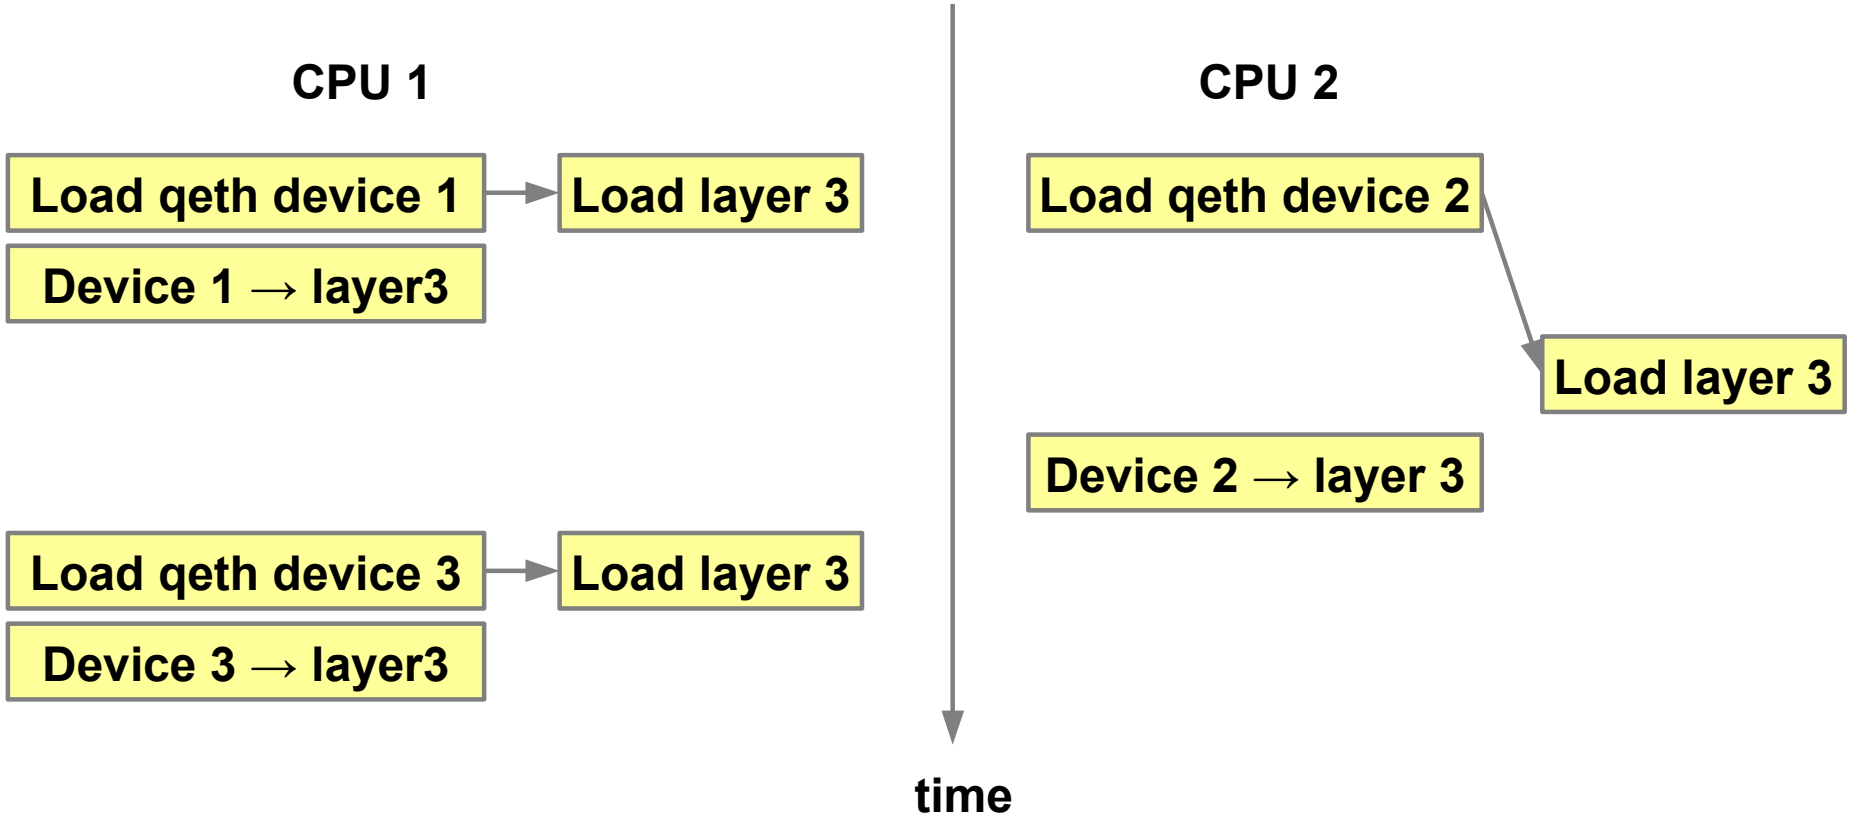

- **Kernel update available via Maintweb shortly**

IPL problems after zipl invocation (1/7)

- **Configuration:**
  - SLES11 SP1 system running under native LPAR on a z10
  - Scripted procedure to run “zipl”
- **Problem description:**
  - zipl is called by a script
  - Sometimes error messages are encountered after invocation
  - After those error messages, the system comes up with a wrong kernel or does not come up at all
- **Tools used for problem determination:**
  - Debug output from server (captured with dbginfo.sh and terminal logs)
  - C program to reproduce problem

IPL problems after zipl invocation (2/7)

1) change /proc/cio_ignore to recognize the target dasd devices.
2) set the target dasd devices online by using chccwdev command.
3) vgscan
4) set the target VG active by using vgchange command.
5) mount the root file system
6) edit system config files
(/etc/fstab, /etc/udev/rules.d/xxxx)
7) execute mkinitrd
8) edit /etc/zipl.conf
9) execute zipl command
10) umount
11) deactivate the VG
12) set the device offline by using chccwdev command

IPL problems after zipl invocation (3/7)

1) change /proc/cio_ignor
2) make the target dasd d
3) vgscan
4) mount the root file syst
5) make the target VG acti
6) edit system config files
   (/etc/fstab, /etc/udev/ru
7) execute mkinitrd
8) edit /etc/zipl.conf
9) execute zipl command
10) umount
11) deactivate the VG
12) make the device offline by using chccwdev command

```
# zipl
Using config file ' /etc/zipl.conf
Building bootmap in '/boot/zipl'
Building menu 'menu'
Adding #1: IPL section
'ipl_prod' (default)
Preparing boot device: dasdt (1048).
Done.
```

(means RC=0, successful)

IPL problems after zipl invocation (4/7)

1) change /proc/cio_ignor
2) make the target dasd d
3) vgscan
4) mount the root file syst
5) make the target VG acti
6) edit system config files
   (/etc/fstab, /etc/udev/ru
7) execute mkinitrd
8) edit /etc/zipl.conf

```
# zipl
Using config file ' /etc/zipl.conf
Building bootmap in '/boot/zipl'
Building menu 'menu'
Adding #1: IPL section
'ipl_prod' (default)
Preparing boot device: dasdt (1048).
Done.
```

(means RC=0, succesful)

But, in syslog:

```
Nov 12 17:12:50 prod_serv zipl: Boot loader written to dasdt (1048) - 5e:340
Nov 12 17:12:55 prod_serv kernel: end_request: I/O error, dev dasdt, sector 8
Nov 12 17:12:55 prod_serv kernel: Buffer I/O error on device dasdt, logical block 1
Nov 12 17:12:55 prod_serv kernel: lost page write due to I/O error on dasdt
Nov 12 17:12:55 prod_serv kernel: end_request: I/O error, dev dasdt, sector 24
Nov 12 17:12:55 prod_serv kernel: Buffer I/O error on device dasdt, logical block 3
Nov 12 17:12:55 prod_serv kernel: end_request: I/O error, dev dasdt, sector 24
```

IPL problems after zipl invocation (5/7)

- **Possible problem causes:**
  - Read only dasd
    - Customer insisted this was not happening
  - zipl does not “sync” properly
    - Checked code: zipl syncs
  - “sync” is not working correctly
    - Man page says:

      According to the standard specification (e.g., POSIX.1-2001), sync() schedules the writes, but may return before the actual writing is done. However, since version 1.3.20 Linux does actually wait. (This still does not guarantee data integrity: modern disks have large caches.)
- **Created C Program to reproduce situation**

IPL problems after zipl invocation (6/7)

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdio.h>
#include <stdlib.h>
main(int argc, char** argv)
{
        int fd, i;
        char *a;
        int count = atoi(argv[1]);
        a = malloc(2048);
        fd = open("/dev/dasdbc", O_RDWR | O_CREAT);
        for(i = 0; i < count; i++)
                write(fd, a, 2048);
        close(fd);
        printf("before sync\n");
        sync();
        printf("after sync\n");
}
```

- Online disk
- Execute
- Offline disk
- Online disk
- Execute
- Offline Disk
- ...

IPL problems after zipl invocation (7/7)

- **In one out of 30 invocations the same error messages could be seen**
- **Problem cause:**
  - Calling zipl changes data in the filesystem but also writes on the device node. The build in sync call only flushes filesystem buffers. This did not ensure that data on the device node /dev/dasdX, e.g. the pointer to the bootmap, has been written. This could cause the wrong bootmap to be used during ipl and the wrong kernel to be loaded or even an unbootable system.
- **Problem fix:**
  - Adding an fsync() for all writable opened file descriptors solves this problem.
- **Patch integrated in SLES11 SP2**

Network connection stalls (1/6)

- Configuration:
  - z10
  - High network load (TSM Server under Linux)
  - Network (qeth) parameter buffer_count = 128
- Problem Description:
  - Network connection stalls regularly and sometimes recovers after about one hour
- Tools used for problem determination:
  - tcpdump / wireshark
  - crash/lcrash
  - dbginfo.sh
  - VM and HW traces

Network connection stalls (2/6)

- Various network traces with tcpdump showed missing arp packets
  → **initially thought to be the cause of the problem**
- Research done why ARP packages are missing
- Collected performance data

	CPU	%usr	%nice	%sys	%iowait	%steal	%irq	%soft	%guest	%idle
01:31:05	all	27.83	0.00	0.30	0.42	69.27	0.11	2.08	0.00	0.00
01:31:05	0	29.87	0.00	0.43	0.83	65.22	0.16	3.49	0.00	0.00
01:31:05	1	25.79	0.00	0.14	0.02	73.34	0.04	0.67	0.00	0.00
01:32:04	all	57.50	0.00	0.31	0.02	38.13	0.20	3.86	0.00	0.00
01:32:04	0	53.82	0.00	0.33	0.00	40.38	0.30	5.16	0.00	0.00
01:32:04	1	61.13	0.00	0.30	0.03	35.86	0.12	2.55	0.00	0.00

→ performance bottle neck

- Assumption:
  HW has to discard incoming packets because Linux has no resources, missing ARP packets cause connection stalls

Network connection stalls (3/6)

- Attempt to prove with OSA traces via Single Object Operations, card specific advanced facilities:

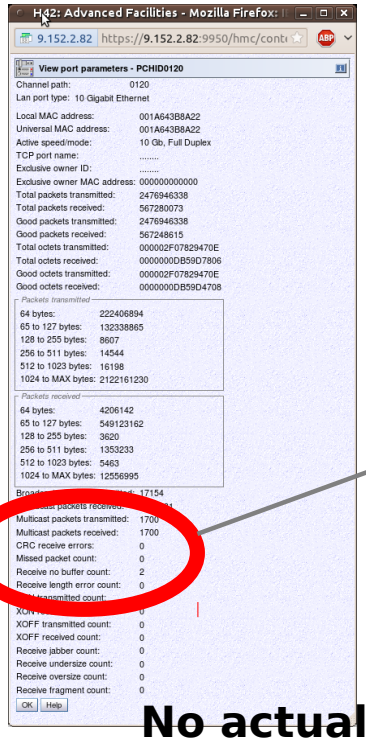

No actual customer screenshot

Network connection stalls (4/6)

- CPU resources added to VM and Linux
- Performance data looked better, steal time vanished
- **Problem persisted**
- Started further debugging with dbginfo.sh and captured a dump during the time of the stall
- Dump looked suspicious – one bit (DSCI – device state change indicator) indicating that new packets arrived was not set
- Did Linux reset it by accident? Did the HW / OSA Card or VM forget to set it?
  → beginning of extensive concurrent data collection in Linux, VM and HW

Network connection stalls (5/6)

- Long process, circumvention needed for customer
- Created script which detects stall and recovers network device:

```
Get RX packets from ifconfig
Ping gateway
Wait one second
Get RX packets from ifconfig
If RX packets did not change AND ping not sucessful
       Start network recovery:
       echo 1 > /sys/bus/ccwgroup/drivers/qeth/$devno/recover
Start from beginning
```

- Detection and recovery was fast enough that TCP connections survived

Network connection stalls: Solution (6/6)

- After numerous data collections in Linux, VM and OSA card root cause was found in OSA firmware
- Buffers 127/128 have special treatment in OSA code which contained a bug
  Temporary circumvention: use buffer_count of 125
- Fix released: `DR79 bundle 46b, MCL N24398.005`
  (System z9 also affected and patched firmware available)
- Make sure that not only Linux and z/VM is up to date, but also Firmware levels

Lost access to SCSI storage during SVC maintenance (1/5)

- **Configuration:**
  - SLES11 SP1 system running under z/VM 6.1
  - SCSI storage via a two-node San Volume Controller
  - z196
- **Problem description:**
  - Maintenance was tested:
    - Power down SVC node one, reboot Node one
    - Power down SVC node two → disks offline
- **Tools used for problem determination:**
  - Configuration file analysis

Lost access to SCSI storage during SVC maintenance (2/5)

multipath.conf:

```
defaults {
      polling_interval       30
      failback               immediate
      no_path_retry          5
      rr_min_io              100
      path_checker           tur
      user_friendly_names    yes
      }
      # SVC
           device {
                 vendor                 “IBM”
                 product                “2145”
                 path_grouping_policy   group_by_prio
                 prio_callout           “/sbin/mpath_prio_alua /dev/%n”
```

Lost access to SCSI storage during SVC maintenance (3/5)

Before node one power down:

```
# multipath -ll
mpathc (uuid...xxx) dm-4 IBM,2145
size=4.0G features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:3:2    sdj   8:144  active ready running
| |- 0:0:1:2    sdm   8:192  active ready running
| |- 1:0:1:2    sdz   65:144 active ready running
| `- 1:0:0:2    sdac  65:192 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:2:2    sdg   8:96   active ready running
  |- 0:0:0:2    sdl   8:176  active ready running
  |- 1:0:3:2    sdx   65:112 active ready running
  `- 1:0:2:2    sdaa  65:160 active ready running
```

Lost access to SCSI storage during SVC maintenance (4/5)

After node one power down:

```
# multipath -ll
mpathc (uuid...xxx) dm-4 IBM,2145
size=4.0G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
  |- 0:0:3:2    sdj   8:144  active ready running
  |- 0:0:1:2    sdm   8:192  active ready running
  |- 1:0:1:2    sdz   65:144 active ready running
  `- 1:0:0:2    sdac  65:192 active ready running
```

Lost access to SCSI storage during SVC maintenance (5/5)

- **Problem: “no_path_retry 5”**
  - Device fails, 5 retries are started
  - If device comes back during that time, it will re re-instated
  - If not, paths are removed permanently
- **Solution:**
  - **Use “no_path_retry queue”**
  - Paths will also be marked as failed, but will not be removed
  - SVC to update online documentation

Daily network recoveries (1/4)

- Configuration:
  - z10, SLES11 SP1
  - High CPU load on z/VM as well as memory shortage
  - Guestlan Type Hiper
- Problem Description:
  - Network devices recover every night
- Tools used for problem determination:
  - dbginfo.sh
  - System tap
  - crash

Daily network recoveries: Symptom (2/4)

```
Jul  5 01:20:33 server1 -- MARK --
Jul  5 01:40:33 server1 -- MARK --
Jul  5 02:11:45 server1 kernel: qeth.fd0b7c: 0.0.8000: A recovery process has been started for the device
Jul  5 02:11:45 server1 kernel: klogd 1.4.1, ---------- state change ----------
Jul  5 02:11:46 server1 kernel: qdio: 0.0.8002 HS on SC 5 using AI:1 QEBSM:0 PCI:0 TDD:1 SIGA: W AO
Jul  5 02:11:46 server1 kernel: qeth.26d434: 0.0.8000: Device is a Guest LAN Hiper card (level: V543)
Jul  5 02:11:46 server1 kernel: with link type GuestLAN Hiper.
Jul  5 02:11:46 server1 kernel: qeth.47953b: 0.0.8000: Hardware IP fragmentation not supported on hsia1
Jul  5 02:11:46 server1 kernel: qeth.066069: 0.0.8000: Inbound source MAC-address not supported on hsia1
Jul  5 02:11:46 server1 kernel: qeth.d7fdb4: 0.0.8000: VLAN enabled
Jul  5 02:11:46 server1 kernel: qeth.e90c78: 0.0.8000: Multicast enabled
Jul  5 02:11:46 server1 kernel: qeth.5a9d02: 0.0.8000: IPV6 enabled
Jul  5 02:11:46 server1 kernel: qeth.184d8a: 0.0.8000: Broadcast enabled
Jul  5 02:11:46 server1 kernel: qeth.dac2aa: 0.0.8000: Using SW checksumming on hsia1.
Jul  5 02:11:46 server1 kernel: qeth.9c4c89: 0.0.8000: Outbound TSO not supported on hsia1
Jul  5 02:11:46 server1 kernel: qeth.bad88b: 0.0.8000: Device successfully recovered!
Jul  5 02:31:46 server1 -- MARK --
Jul  5 02:51:46 server1 -- MARK --
```

- Application servers had problems if recovery was not fast enough
- Root cause needed to be found, SLES10 servers were not affected

Daily network recoveries (3/4)

- s390dbf output useless: debug area is cleared after device recovery
- Alternate method: We need to get a dump of the system *before* the recovery
- Idea: Use Systemtap
- After kenel panic: Create a dump

```
# Include the header that declares panic()
%{
#include <linux/kernel.h>
%}

# Wrap panic() in stap
function panic(msg:string) %{
        panic("%s", THIS->msg);
%}

# Tell the user what we're doing
probe begin {
        printf("panic on qeth_recover enabled\n")
}

probe end {
        printf("panic on qeth_recover disabled\n")
}

# probe the qeth_recover function
probe module("qeth").function("qeth_schedule_recovery") {
        panic("recovery triggered\n")
}
```

Daily network recoveries: Solution (4/4)

- Dump showed that a specific “busy” code was being returned by our (virtual) hardware
- If that was also the case 100µs later → trigger recovery
- 100µs too short in a shared environment
- SLES10 had longer timeouts → no recovery in SLES10
- Created patch, extended timeouts in SLES11

Timeouts in HA environment (1/3)

- **Configuration:**
  - SLES11 HA systems running under z/VM
  - Memory and CPU overcommitment in z/VM
- **Problem description:**
  - HA environment migrated from AIX to Linux on System z
  - Heartbeat timeouts
  - Cluster separation
- **Tools used for problem determination:**
  - sadc/sar
  - dbginfo.sh

Timeouts in HA environment (2/3)

- **Default timeout: 3 seconds**
- **Every few days a cluster separation could be observed**
- **Linux sadc/sar data like this:**

```
04:47:18 PM     CPU     %user     %nice   %system   %iowait    %steal     %idle
04:47:19 PM     all     46.23      0.00      0.50      0.00      4.02     49.25
04:47:20 PM     all     39.90      0.00      1.97      0.49     11.33     46.31
04:47:30 PM     all      1.41      0.00      0.10      0.05     97.38      1.07
04:47:31 PM     all     45.23      0.00      7.54      0.00     14.07     33.17
Average:        all     10.97      0.00      0.83      0.08     77.60     10.52
```

Timeouts in HA environment (2/3)

- **Default timeout: 3 seconds**
- **Every few days a cluster separation could be observed**
- **Linux sadc/sar data like this:**

```
04:47:18 PM     CPU     %user     %nice   %system   %iowait    %steal     %idle
04:47:19 PM     all     46.23      0.00      0.50      0.00      4.02     49.25
04:47:20 PM     all     39.90      0.00      1.97      0.49     11.33     46.31
04:47:30 PM     all      1.41      0.00      0.10      0.05     97.38      1.07
04:47:31 PM     all     45.23      0.00      7.54      0.00     14.07     33.17
Average:        all     10.97      0.00      0.83      0.08     77.60     10.52
```

… notice the gap

Time-outs in HA environment (3/3)

- **Due to memory and CPU constraints longer delays occurred in the guests**
- **OK for “normal” operations, but HA environments have special requirements**
- **After timer increase problems gone**
- **Think of the following:**
  - Virtualized environments are not real-time, response times are hard to predict
  - When migrating from distributed systems, timer configuration should be verified
- **Possible solutions:**
  - Increase time-out values
  - Do not over-commit too much in HA environments (easier said than done...)
  - Use z/VM tuning options such as “quick dispatch” for vital guests
  - LPAR might be an option for very large Linux guests

Closing remarks

- **Intention is not to scare you**
- **Linux on System z is a very reliable platform**
- **A very small fraction of PMRs result in a code fix / patch from us**
- **If you follow best practices (e.g. part of the “Problem determination” presentation) you are well prepared**
- **In case of problems, do not hesitate to contact me**