

2012

**IBM System z Technical University**

Enabling the infrastructure for smarter computing

# **Problem Determination with Linux on System z Real Customer Cases**

**zLG08**

Susanne Wintenberger



## Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*, AS/400®, e business (logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

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 therefrom.

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.

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.

\* All other products may be trademarks or registered trademarks of their respective companies.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

- Remarks about customer incidents
- Customer reported incidents
  - zfcpl auto port rescan triggered frequently
  - DASD maintenance halts system
  - Network connection stalls
  - Lost access to SCSI storage during SVC maintenance
  - Daily network recoveries
  - Disk I/O bottlenecks
  - FCP disk configuration issues
  - Guest spontaneously reboots
  - Unable to mount file system after LVM changes
  - Which DASD should be considered as PAV DASDs
  - Network connection too slow
  - lcrash/crash could not read the dump

## Introductory Remarks

- The incidents reported here are real customer incidents
  - Red Hat Enterprise Linux, and Suse Linux Enterprise Server distributions
  - Linux running in LPAR and z/VM of different versions
- While problem analysis looks rather straight forward on the charts, it might have taken weeks to get it done.
- The more information is available, the sooner the problem can be solved, because gathering and submitting additional information again and again usually introduces delays.
  - See First Aid Kit at the beginning of this presentation.
- This presentation focuses on how the tools have been used, comprehensive documentation on their capabilities is in the docs of the corresponding tool.

## zfcplib auto port rescan triggered frequently

- Configuration:
  - RHEL6, but this problem also applies to SLES11 SP1
  - Shared FCP adapter card among the systems
  - FCP setup in NPIV mode with multiple Linux systems in the same zone
- Problem Description:
  - Customer complained about TSM backup problems, that its system lost access to the tape storage
- Tools used for problem determination:
  - dbginfo.sh

## zfcplib auto port rescan triggered frequently (cont'd)

### ▪ Problem Origin:

```
zfcplib.4a463f: 0.0.15d2: ERP failed for remote port 0xc05076ffe4801a01
zfcplib.4a463f: 0.0.15d2: ERP failed for remote port 0xc05076e7bd003991
zfcplib.4a463f: 0.0.15d2: ERP failed for remote port 0xc05076ffe4803011
zfcplib.4a463f: 0.0.15d2: ERP failed for remote port 0xc05076ffe4805831
zfcplib.4a463f: 0.0.15d2: ERP failed for remote port 0xc05076ffe4803931
zfcplib.4a463f: 0.0.15c2: ERP failed for remote port 0xc05076ffe4801a11
zfcplib.4a463f: 0.0.15c2: ERP failed for remote port 0xc05076ffe5002831
```

- Single fabric event (RSCN) was triggered
  - e.g. fabric login of other Linux system after setting FCP adapter online, switch reboot, storage server re-configuration
- All Linux systems in the same zone receives the event
  - This triggers the auto port rescan
  - The port discovery tries to reopen all already known ports, which fails for the other initiator ports (ports of the other Linux systems)

## zfcpc auto port rescan triggered frequently (cont'd)

- Problem Origin:

- When port discovery takes too long, an FCP adapter recovery is initiated

```
qdio: 0.0.15d0 ZFCP on SC 1 using AI:1 QEBSM:1 PCI:1 TDD:1 SIGA: W A0
```

- When all FCP adapters to one SCSI device recover at the same time, all paths to the LUN are gone and the LUN becomes unresponsive
- This leads to I/O stalls and unresponsiveness of applications, if applicable

## zfcplib auto port rescan triggered frequently (cont'd)

### ▪ Solution:

- Implement single initiator zones by means of soft zones
  - The Linux system is no longer able to recognize other initiator ports
  - This reduces the amount of fabric events being received by each individual Linux system
  - Clean up sysfs for unexpected remote ports (e.g. initiator ports)
- FCP adapter firmware upgrade (upcoming)
  - This change filters the list of initiator ports shown to Linux system
  - This prevents the Linux from logging into initiator ports
- Linux kernel upgrade with 'no\_auto\_port\_rescan' feature (upcoming)
  - This allows to disable the auto port scan that is triggered by the fabric events and FCP adapter recovery
  - The auto port scan is still active when switching an adapter online (e.g. Linux boot), therefore it might be required to clean up sysfs for unexpected remote ports (e.g. initiator ports)



## DASD maintenance halts system

- 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 (cont'd)

- Iscss output

Device	Subchan.	DevType	CU	Type	Use	PIM	PAM	POM	CHPIDs
0.0.f5f0	0.0.0000	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.f5f1	0.0.0001	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.f5f2	0.0.0002	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.e9a2	0.0.0003	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.e9a3	0.0.0004	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.e9a4	0.0.0005	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.0191	0.0.0006	3390/0c	3990/e9			f0	f0	ff	12303234 36000000
0.0.0009	0.0.0007	0000/00	3215/00		yes	80	80	ff	12000000 00000000
0.0.000c	0.0.000e	0000/00	2540/00			80	80	ff	12000000 00000000
0.0.000d	0.0.000f	0000/00	2540/00			80	80	ff	12000000 00000000
0.0.000e	0.0.0010	0000/00	1403/00			80	80	ff	12000000 00000000
0.0.0190	0.0.0011	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.019d	0.0.0012	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.019e	0.0.0013	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.0592	0.0.0014	3390/0c	3990/e9			f0	f0	ff	30323436 00000000

## DASD maintenance halts system (cont'd)

- lscss output

Device	Subchan.	DevType	CU	Type	Use	PIM	PAM	POM	CHPIDs
0.0.f5f0	0.0.0000	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.f5f1	0.0.0001	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.f5f2	0.0.0002	1732/01	1731/01		yes	80	80	ff	76000000 00000000
0.0.e9a2	0.0.0003	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.e9a3	0.0.0004	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.e9a4	0.0.0005	3390/0c	3990/e9		yes	f0	f0	ff	12303234 36000000
0.0.0191	0.0.0006	3390/0c	3990/e9			f0	f0	ff	12303234 36000000
0.0.0009	0.0.0007	0000/00	3215/00		yes	80	80	ff	12000000 00000000
0.0.000c	0.0.000e	0000/00	2540/00			80	80	ff	12000000 00000000
0.0.000d	0.0.000f	0000/00	2540/00			80	80	ff	12000000 00000000
0.0.000e	0.0.0010	0000/00	1403/00			80	80	ff	12000000 00000000
0.0.0190	0.0.0011	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.019d	0.0.0012	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.019e	0.0.0013	3390/0c	3990/e9			f0	f0	ff	30323436 00000000
0.0.0592	0.0.0014	3390/0c	3990/e9			f0	f0	ff	30323436 00000000

## DASD maintenance halts system (cont'd)

- Problem Origin:

- In z/VM fully simulated devices are assigned to a real physical CHPID
  - Puncher
  - Reader
  - Console
- When this shared CHPID needs to be “varied/configured offline” for maintenance, all attached devices are lost
- When the console is disconnected, system is down

- Solution:

- Check your CHPID assignments of fully simulated devices before maintenance to prevent outages
- z/VM released an APAR (VM65117) to pick an unused real CHPID for fully simulated devices

## Network connection stalls

- 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
  - z/VM and OSA hardware traces

## Network connection stalls (cont'd)

- 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

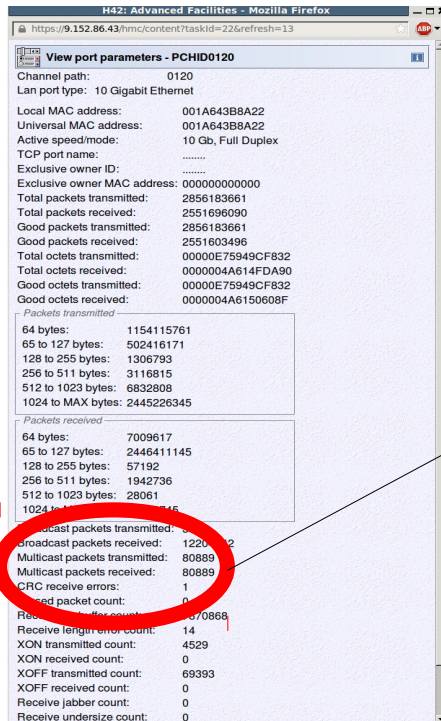
CPU	%usr	%nice	%sys	%iowait	%steal	%irq	%soft	%guest	%idle
all	27.83	0.00	0.30	0.42	69.27	0.11	2.08	0.00	0.00
0	29.87	0.00	0.43	0.83	65.22	0.16	3.49	0.00	0.00
1	25.79	0.00	0.14	0.02	73.34	0.04	0.67	0.00	0.00
all	57.50	0.00	0.31	0.02	38.13	0.20	3.86	0.00	0.00
0	53.82	0.00	0.33	0.00	40.38	0.30	5.16	0.00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

→ performance bottle neck

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

## Network connection stalls (cont'd)

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



Packets received

64 bytes:	7009617
65 to 127 bytes:	2446411145
128 to 255 bytes:	57192
256 to 511 bytes:	1942736
512 to 1023 bytes:	28061
1024 to MAX bytes:	96154745

Broadcast packets transmitted: 3442  
Broadcast packets received: 12204642  
Multicast packets transmitted: 80889  
Multicast packets received: 80889  
CRC receive errors: 1  
Missed packet count: 0  
Receive no buffer count: 7870868  
Receive length error count: 14  
XON transmitted count: 4529  
XON received count: 0  
XOFF transmitted count: 69393  
XOFF received count: 0

No actual customer screenshot

## Network connection stalls (cont'd)

- 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 had arrived was not set
  
- Did Linux reset it by accident? Did the HW / OSA Card or z/VM forget to set it?
  - beginning of extensive concurrent data collection in Linux, VM and HW



## Network connection stalls (cont'd)

- 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 successful
    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 (cont'd)

- Problem Origin:

- After numerous data collections in Linux, z/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

- Solution:

- 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

- 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 (cont'd)

- 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 (cont'd)

- 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 (cont'd)

- 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 (cont'd)

- Problem Origin:

- In `/etc/multipath.conf`: `"no_path_retry 5"`
- When a path fails, 5 retries are started
- If the path to the device comes back during this time, it will be re-instated
- If not, paths are removed permanently

- Solution:

- Use instead `"no_path_retry queue"`
- Paths will also be marked as failed, but will not be removed
- SVC to update online documentation

## Daily network recoveries

- 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 (cont'd)

```
kernel: qeth.fd0b7c: 0.0.8000: A recovery process has been started for the
device
kernel: klogd 1.4.1, ----- state change -----
kernel: qdio: 0.0.8002 HS on SC 5 using AI:1 QEBSM:0 PCI:0 TDD:1 SIGA: W A0
kernel: qeth.26d434: 0.0.8000: Device is a Guest LAN Hiper card (level: V543)
kernel: with link type GuestLAN Hiper.
kernel: qeth.47953b: 0.0.8000: Hardware IP fragmentation not supported on hsia1
kernel: qeth.066069: 0.0.8000: Inbound source MAC-address not supported on hsia1
kernel: qeth.d7fdb4: 0.0.8000: VLAN enabled
kernel: qeth.e90c78: 0.0.8000: Multicast enabled
kernel: qeth.5a9d02: 0.0.8000: IPV6 enabled
kernel: qeth.184d8a: 0.0.8000: Broadcast enabled
kernel: qeth.dac2aa: 0.0.8000: Using SW checksumming on hsia1.
kernel: qeth.9c4c89: 0.0.8000: Outbound TSO not supported on hsia1
kernel: qeth.bad88b: 0.0.8000: Device successfully recovered!
```

- Application servers had problems if recovery was not fast enough
- SLES10 servers were not affected

## Daily network recoveries (cont'd)

- s390dbf output useless:  
debug area is  
cleared after device  
recovery
- Alternate method: We  
need to get a dump  
of the system before  
the device recovery
- Idea: Use Systemtap
- After kernel 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 (cont'd)

- Problem Origin:
  - Dump showed that a specific “busy” code was being returned by our (virtual) hardware
  - If this happened again 100 $\mu$ s later, a network recovery was triggered
  - 100 $\mu$ s are too short in a shared environment
  - SLES10 had longer timeouts → no recovery in SLES10
- Solution:
  - Changed code in SLES11

## Disk I/O bottlenecks

- Configuration:
  - Customer has distributed I/O workload to multiple volumes using z/VM minidisk and LVM striping
  - This problem also applies to non-LVM and non minidisk configurations
- Problem Description:
  - Multi-disk I/O performance is worse than expected by projecting single disk benchmark to more complex solution.
- Tools used for problem determination:
  - dbginfo.sh
  - linux for system z debug feature
  - sadc/sar, iostat and DASD statistics
  - z/VM monitor data
  - storage controller DASD statistics

## Disk I/O bottlenecks (cont'd)

- Problem Origin:

- bottleneck other than the device – e.g.:

- z/VM minidisks are associated to same physical disk
    - SAN bandwidth not sufficient
    - Storage controller HBA bandwidth not sufficient
    - Multiple disks used are in the same rank of storage controller

- Solution:

- Check your disk configuration and configure for best performance

- Make sure, minidisks used in parallel are not on the same physical disk

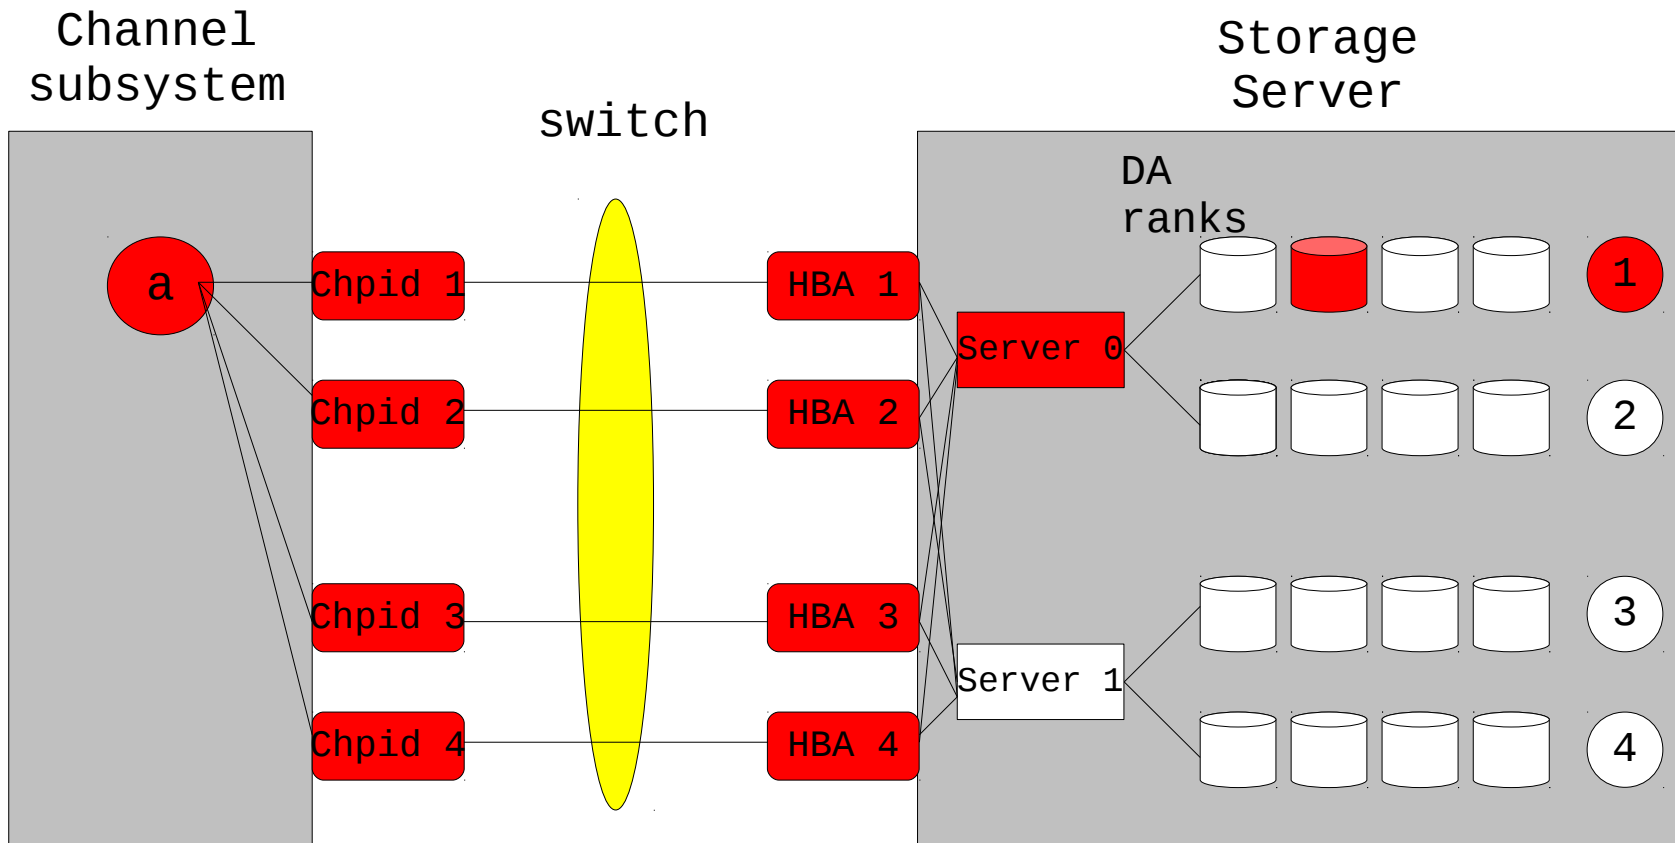
- Distribution of I/O workload (striped LVs, PAV or HyperPAV)

- For optimal disk performance configurations read and take into account

- [http://www.ibm.com/developerworks/linux/linux390/perf/tuning\\_diskio.html](http://www.ibm.com/developerworks/linux/linux390/perf/tuning_diskio.html)

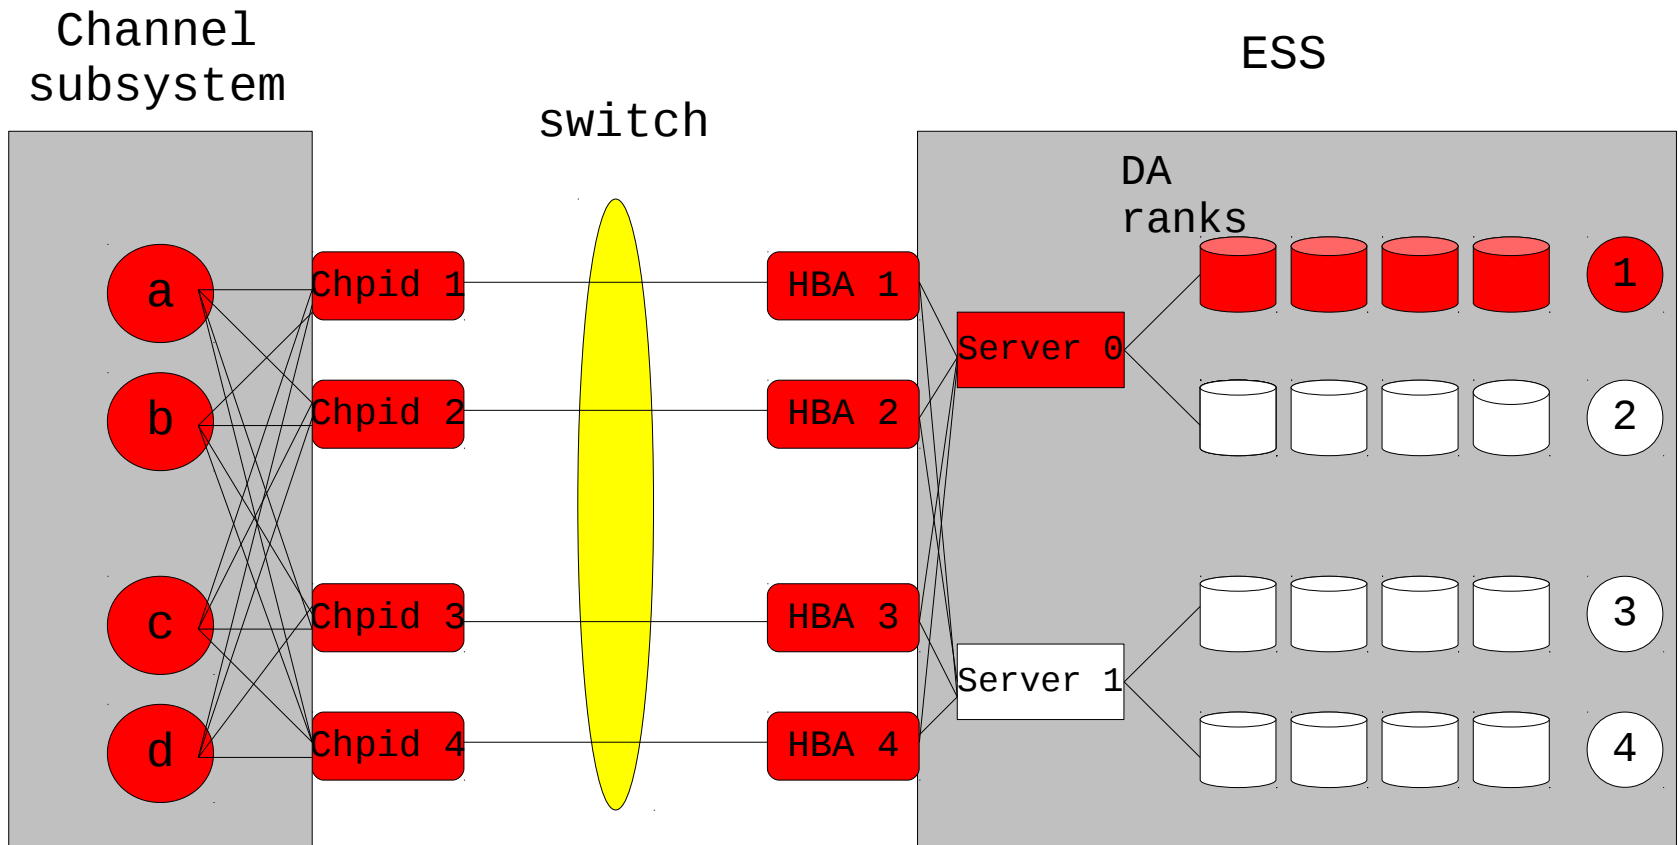
## Disk I/O bottlenecks (cont'd)

- FICON/ECKD dasd I/O to a single disk



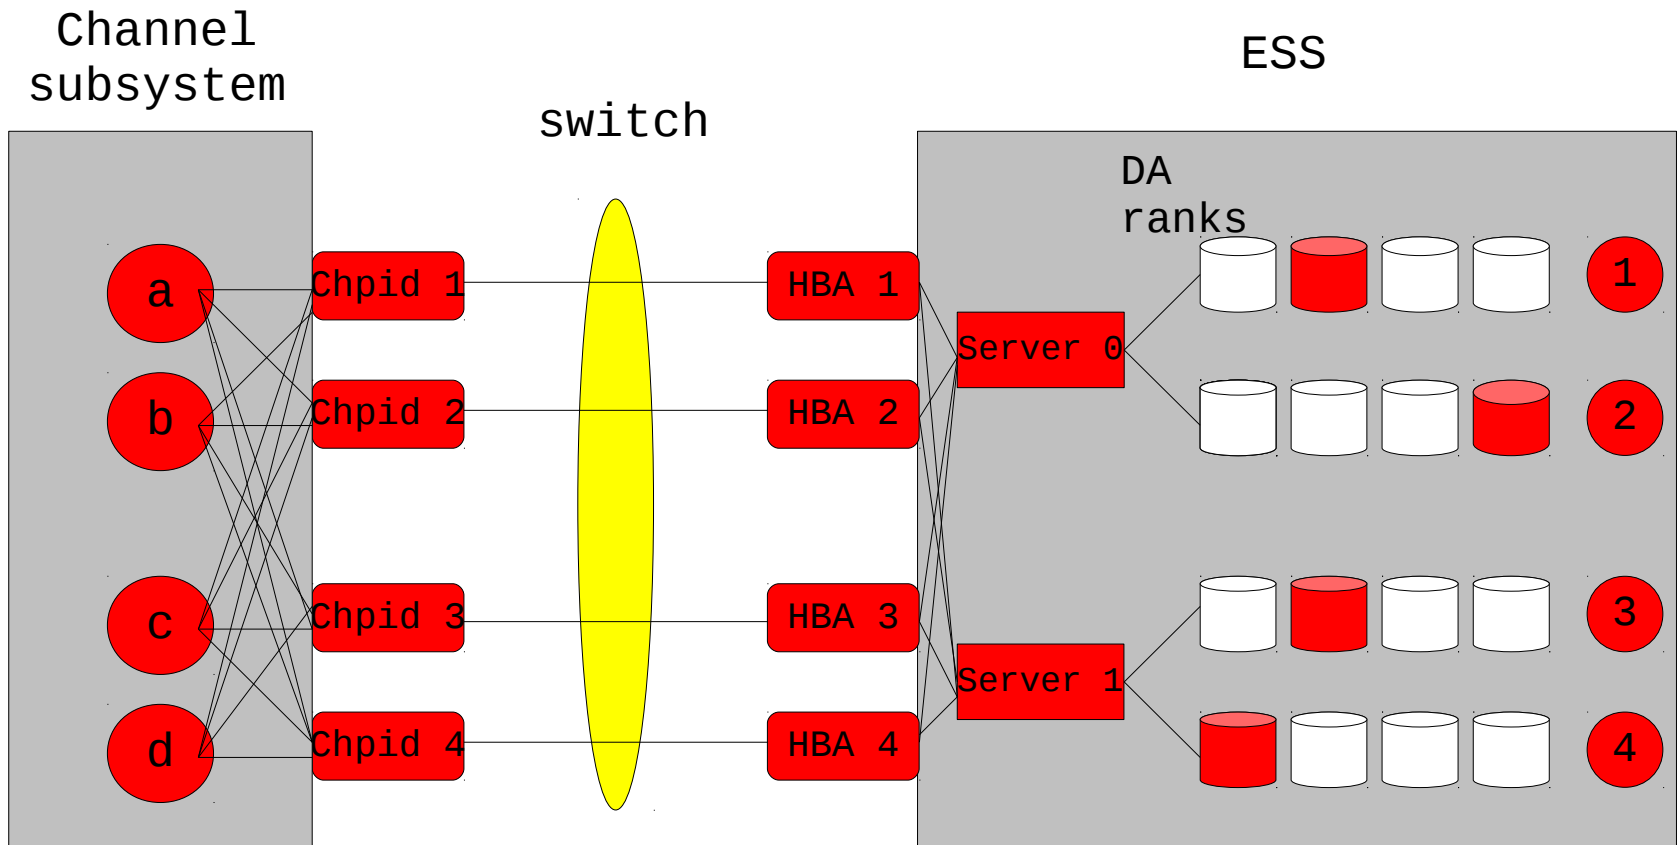
## Disk I/O bottlenecks (cont'd)

- FICON/ECKD dasd I/O to a linear or striped logical volume in one rank



## Disk I/O bottlenecks (cont'd)

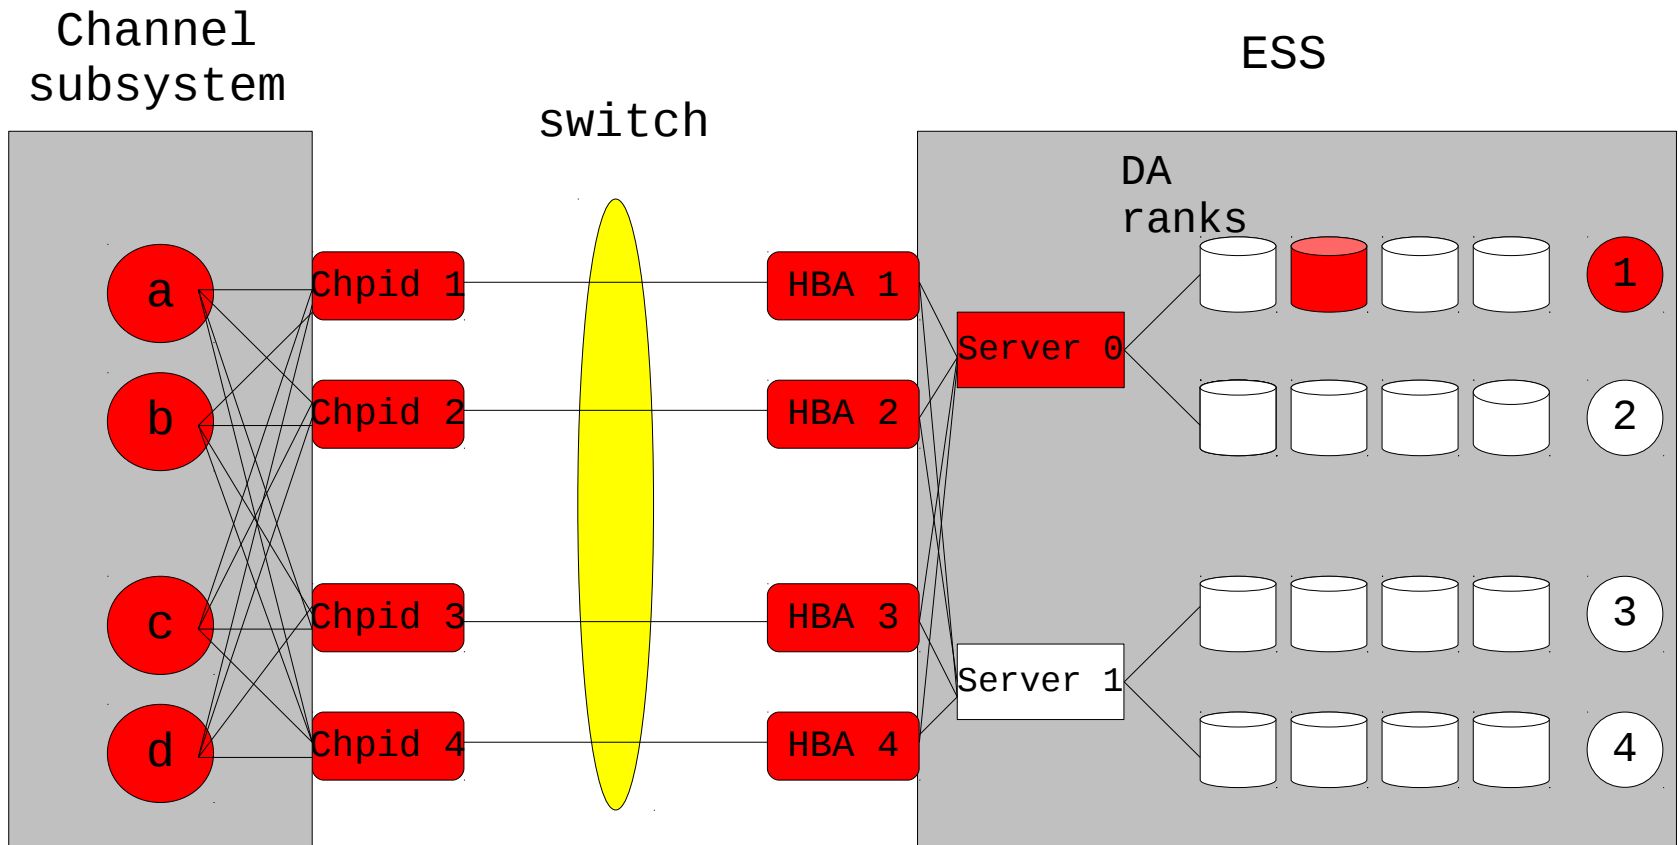
- FICON/ECKD dasd I/O to a linear or striped logical volume





## Disk I/O bottlenecks (cont'd)


- FICON/ECKD dasd I/O to a single disk with (hyper) PAV




## FCP disk: multipath configuration

- Configuration:
  - Customer is running Samba server on Linux with FCP attached disk managed by Linux LVM.
  - This problem also applies to any configuration with FCP attached disk storage
- Problem Description:
  - Accessing some files through samba causes the system to hang while accessing other files works fine
  - Local access to the same file cause a hanging shell as well
    - Indicates: this is not a network problem!
- Tools used for problem determination:
  - dbginfo.sh
- Problem Indicators:
  - Intermittent outages of disk connectivity

# FCP disk: multipath configuration (cont'd)



# FCP disk: multipath configuration (cont'd)



## FCP disk: multipath configuration (cont'd)

- Solutions:

- Configure multipathing correctly:

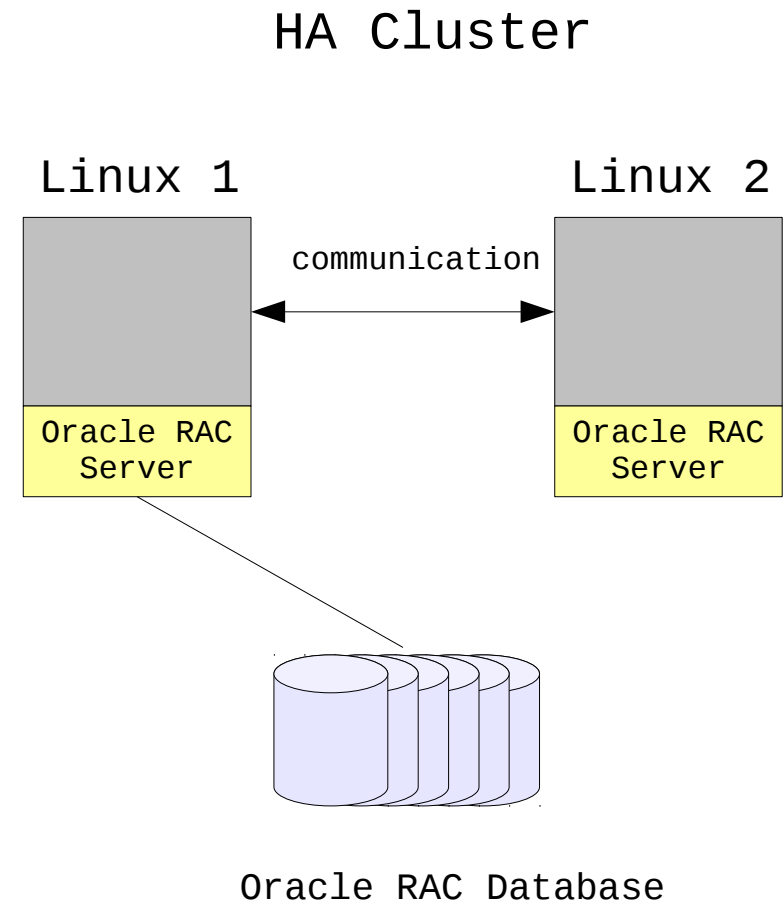
- Establish independent paths to each volume
    - Group the paths using the device-mapper-multipath package
    - Base LVM configuration on top of mpath devices instead of sd<#>

- For a more detailed description how to use FCP attached storage appropriately with Linux on System z, see

<http://public.dhe.ibm.com/software/dw/linux390/docu/lk33ts04.pdf>

## Guest spontaneously reboots

- Configuration:
  - Oracle RAC server or other HA solution under z/VM
- Problem Description:
  - Occasionally guests spontaneously reboot without any notification or console message
- Tools used for problem determination:
  - cp instruction trace of (re)IPL code
  - Crash dump taken after trace was hit



## Guest spontaneously reboots - Steps to find root cause

- Step 1
  - Check syslog
    - No message printed prior to reboot
    - Messages were found for previously reboots triggered by the operator
    - → 'shutdown -r' or 'reboot' was not used
- What mechanism could have caused the reboot?

## Guest spontaneously reboots - Steps to find root cause (cont'd)

### ▪ Step 2

– Check source code:

- The reboot code is called from  
    `machine_restart()` in `arch/s390/kernel/setup.c`
- This function is called from:  
    `drivers/char/watchdog/softdog.c`  
    `drivers/char/hangcheck-timer.c`  
    `drivers/char/sysrq.c`  
    `kernel/power/disk.c`  
    `kernel/power/sys.c`  
    `kernel/sys.c` (3 instances)  
    `kernel/panic.c`

– Some kernel functions print to syslog first

- But the message might not get its way to syslog before reboot
- Also, some functions don't print to syslog.

### ▪ Which instance is doing the reboot here? How to find out?



## Guest spontaneously reboots - Steps to find root cause (cont'd)

- We need to dump the system 'caught-in-the-act'
  - Right when calling `machine_restart()`
  - Otherwise the call stack will be lost already
  - Also the kernel log buffer is reinitialized after reboot
- Step 3
  - Find out address of (re)ipl code in the system map
  - Use this address to set instruction trace

```
cd /boot
grep machine_restart System.map-2.6.16.60-0.54.5-default
0000000000010c364 T machine_restart
000000000001171c8 t do_machine_restart
00000000000603200 D _machine_restart
```

## Guest spontaneously reboots - Steps to find root cause (cont'd)

### ▪ Step 4

- Set CP instruction trace on the reboot address
- System is halted at that address, when a reboot is triggered
- In LPAR a trace can be set up in PEDEBUG mode at the SE, too

```
CP CPU ALL TR IN R 10C364.4
```

```
HCPTRI1027I An active trace set has turned RUN off
```

```
CP Q TR
```

```
NAME INITIAL (ACTIVE)
```

```
1 INSTR PSWA 0010C364-0010C367
```

```
TERM NOPRINT NORUN SIM
```

```
SKIP 00000 PASS 00000 STOP 00000 STEP 00000
```

```
CMD NONE
```

```
-> 000000000010C364' STMF EBCFF0780024 >> 000000003A557D48
```

```
CC 2
```

## Guest spontaneously reboots - Steps to find root cause (cont'd)

- Step 5

- Take a dump, when the (re)ipl code is hit

```
cp cpu all stop
cp store status
Store complete.
```

```
cp i 4fc6
```

```
Tracing active at IPL
```

```
HCPGSP2630I The virtual machine is placed in CP mode due to a
SOGP stop and store status from CPU 00.
```

```
zIPL v1.6.3-0.24.5 dump tool (64bit)
```

```
Dumping 64 bit OS
```

```
00000128 / 00001024 MB
```

```
.....
```

```
00001024 / 00001024 MB
```

```
Dump successful
```

```
HCPIR450W CP entered, disabled wait PSW 00020000 80000000
```

```
00000000 00000000
```

## Guest spontaneously reboots - Steps to find root cause (cont'd)

- Step 6

- Save dump in a file

```
zgetdump /dev/dasdb1 > dump_file
```

```
Dump device: /dev/dasdb1
```

```
>>> Dump header information <<<
```

```
Dump created on: Wed Oct 27 12:00:40 2010
```

```
Magic number: 0xa8190173618f23fd
```

```
Version number: 4
```

```
Header size: 4096
```

```
Page size: 4096
```

```
Dumped memory: 1073741824
```

```
Dumped pages: 262144
```

```
Real memory: 1073741824
```

```
cpu id: 0xff00012320948000
```

```
System Arch: s390x (ESAME)
```

```
Build Arch: s390x (ESAME)
```

```
>>> End of Dump header <<<
```

```
Reading dump content .....
```

```
Dump ended on: Wed Oct 27 12:00:52 2010
```

```
Dump End Marker found: this dump is valid.
```

## Guest spontaneously reboots - Steps to find root cause (cont'd)

### ▪ Step 7

– Use (l)crash, to find out, which process has triggered the reboot

```

STACK:
0 start_kernel+950 [0x6a690e]
1 _stext+32 [0x100020]
=====
TASK HAS CPU (1): 0x3f720650 (oprocd.bin):
LOWCORE INFO:
- psw      : 0x0704200180000000 0x000000000010c36a
- function : machine_restart_smp+6
- prefix   : 0x3f438000
- cpu timer: 0x7fffffff 0xff9e6c00
- clock cmp: 0x00c6ca69 0x22337760
- general registers:

<snip>                                     /var/opt/oracle/product/crs/bin/oprocd.bin

STACK:
0 machine_restart+94 [0x109ad6]
1 hangcheck_fire+140 [0x8033e90f8]
2 run_timer_softirq+384 [0x13fac4]


```

## Guest spontaneously reboots (cont'd)

- Problem Origin:
  - HA component erroneously detected a system hang
    - hangcheck\_timer module did not receive timer IRQ
    - z/VM 'time bomb' switch
    - TSA monitor
- z/VM cannot guarantee 'real-time' behavior if overloaded
  - Longest 'hang' observed: 37 seconds(!)
- Solution:
  - Offload HA workload from overloaded z/VM
    - e.g. use separate z/VM
    - or: run large Oracle RAC guests in LPAR
  - Critical factors:
    - Low memory overcommitment
    - Sufficient CPU capacity

## Unable to mount file system after LVM changes

- Configuration:
  - Linux HA cluster with two nodes
  - Accessing same DASDs which are exported via ocfs2
- Problem Description:
  - Added one node to cluster, brought Logical Volume online
  - Unable to mount the filesystem from any node after that
- Tools used for problem determination:
  - dbginfo.sh



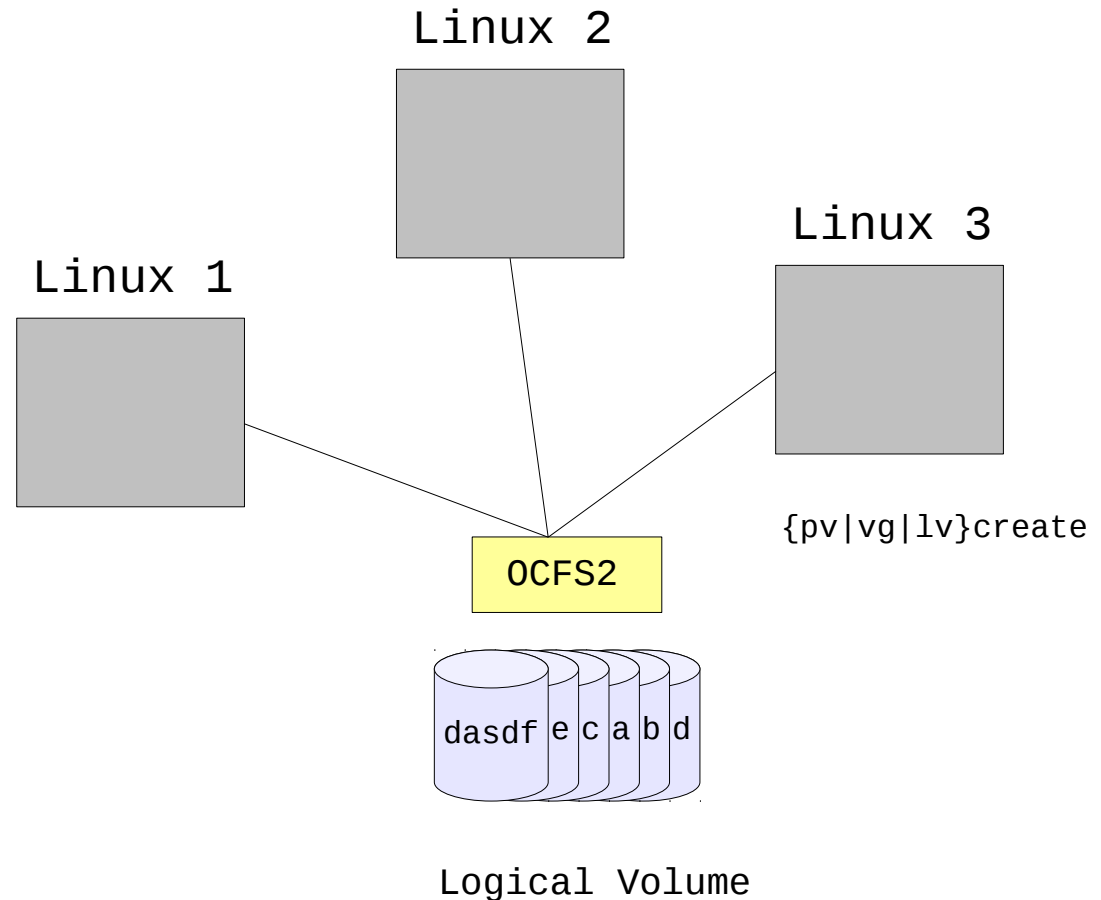
## Unable to mount file system after LVM changes (cont'd)

- Problem Origin:

- LVM meta data was overwritten when adding 3rd node
- e.g. superblock not found

- Solution:

- Extract meta data from running node (/etc/lvm/backup) and write to disk again





## Which DASDs should be considered as PAV devices

- Configuration:
  - guest under z/VM, Oracle RAC Server
  - 14 IFLs, 12 GB of storage configured,
  - 4 FICON channels, LVM managed (linear LVs), PAV partially enabled
- Problem Description:
  - customer complained about bad performance and asked to determine which DASDs should be considered for PAV
- Tools used for problem determination:
  - dbginfo.sh
  - iostat
  - z/VM MONWRITE data

## Which DASDs should be considered as PAV devices (cont'd)

- iostat data for dasdav – RDEV EC46

Device:	rrqm/s	wrqm/s	r/s	w/s	rkB/s	wkB/s	avgrq-sz	avgqu-sz	await	svctm	%util
dasdav	0.00	174.70	0.02	1.63	0.07	705.33	855.03	0.10	61.72	33.74	5.57
dasdav	0.00	357.68	0.05	3.67	0.20	1445.40	777.90	0.31	82.78	42.56	15.82
dasdav	0.00	904.50	0.08	9.53	0.33	3656.14	760.32	0.61	63.50	28.68	27.59
dasdav	0.00	394.27	0.03	3.80	0.13	1592.27	830.82	0.23	59.43	31.57	12.10
dasdav	0.00	541.58	0.05	5.60	0.20	2188.70	774.96	0.35	61.98	31.24	17.65
dasdav	0.00	318.44	0.03	3.28	0.13	1269.94	765.75	0.20	59.45	30.90	10.25
dasdav	0.00	402.53	0.05	4.18	0.20	1643.80	776.69	0.26	62.32	32.05	13.57

avgqu-sz: average length of queue, how many i/o requests are not dispatched

await: average time (ms) for i/o requests issued to the device to be serviced  
(includes the time spent by the requests in queue and the time spent servicing them).

svctm: average service time (ms) for i/o requests that were issued to the device.  
(time spent outside linux)

util: percentage of CPU time during which I/O requests were issued to the device  
(bandwidth utilization for the device)

## Which DASDs should be considered as PAV devices (cont'd)

- z/VM performance monitor data for RDEV EC46
  - FCX168 - General I/O Device Performance Log

Interval	<-Rate/s->	<----- Time (msec) ---->				<Percent>			
End Time	I/O	Avoid	Pend	Disc	Conn	Serv	Resp	Busy	READ
15:48:07	1.9	.0	.2	16.6	10.9	27.7	27.7	5	1
15:49:07	5.8	.0	.3	16.8	12.0	29.1	29.1	17	1
15:50:07	9.3	.0	.3	15.5	11.0	26.8	26.8	25	1
15:51:07	4.1	.0	.3	15.5	10.8	26.6	26.6	11	1
15:52:07	3.7	.0	.3	18.7	11.8	30.8	30.8	11	1
15:53:07	5.1	.0	.2	17.2	11.7	29.1	29.1	15	1

Disc: Average time that the device remained disc. from channel while executing I/O request (High values = short on cache/NVS, overloaded path and/or long SEEK)

Conn: Average time that the storage server transfers data on the chpid

Serv: Sum of Pending (Pend), Connected (Conn) and Disconnect (Disc) time

Resp: Service time plus the time during which an I/O req. was waiting to be started

Good response time values are a couple of ms

## Which DASDs should be considered as PAV devices (cont'd)

- Problem Origin:
  - Further z/VM performance analysis:
    - Almost all DASDs are in one LCU (= rank)
    - That SSID is the hottest LCU with 42% writes and write hits on cache only 67% of the time.
    - This is a severe bottleneck and it explains high disconnect time
  - Further dbginfo analysis showed:
    - Linear logical volumes

## Which DASDs should be considered as PAV devices (cont'd)

- Solution/Recommendation:

- The storage server is short on cache and NVS (Non Volatile Storage) for that LCU. NVS is needed in regard to write requests and too little NVS results in high service time for the write requests.
- The devices in the LCUs should be rearranged.
  - The devices with the highest service times should be moved to different LCUs, which do not have any workload problems.
  - A move of the devices with the highest service time to a PAV rank that is also in trouble would also not help, because if the storage server is already in trouble, making it do more work concurrently won't help matters. In that case it could make things worse.
  - The devices should be distributed over various storage ranks/servers
  - Striped logical volumes should be used instead of linear LVs, so that I/O workload is spread over the whole storage server.
- As soon as the storage server is no longer the problem AND there is still a high average utilization AND a device queue is forming, then PAV can be considered.

## Network connection is too slow

- Configuration:
  - z/VSE running CICS, connecting to DB2 in zLinux
  - HiperSocket connection from zLinux to z/VSE
  - But also applies to hipersocket connections between zLinux and z/OS
- Problem Description:
  - When CICS transaction were monitored, some transactions take a couple of seconds instead of milliseconds
- Tools used for problem determination:
  - dbginfo.sh
  - s390 debug feature
  - sadc/sar
  - CICS transaction monitor

## Network connection is too slow (cont'd)

- s390 debug feature

- Check for qeth errors:

```
cat /sys/kernel/debug/s390dbf/qeth_qerr
00 01282632346:099575 2 - 00 0000000180b20218 71 6f 75 74 65 72 72 00 | qouterr.
00 01282632346:099575 2 - 00 0000000180b20298 20 46 31 35 3d 31 30 00 | F15=10.
00 01282632346:099576 2 - 00 0000000180b20318 20 46 31 34 3d 30 30 00 | F14=00.
00 01282632346:099576 2 - 00 0000000180b20390 20 71 65 72 72 3d 41 46 | qerr=AF
00 01282632346:099576 2 - 00 0000000180b20408 20 73 65 72 72 3d 32 00 | serr=2.
```

- dbginfo file

- Check for buffer count:

```
cat /sys/devices/qeth/0.0.1e00/buffer_count
16
```

- Problem Origin:

- Too less inbound buffers

## Network connection is too slow (cont'd)

### ▪ Solution:

- Increase inbound buffer count (default: 16, max 128)
- Check actual buffer count with 'lsqeth -p'
- Set the inbound buffer count in the appropriate config file:
  - SUSE SLES10:
    - in /etc/sysconfig/hardware/hwcfg-qeth-bus-ccw-0.0.F200
    - add QETH\_OPTIONS="buffer\_count=128"
  - SUSE SLES11:
    - in /etc/udev/rules.d/51-qeth-0.0.f200.rules add ACTION=="add",
    - SUBSYSTEM=="ccwgroup", KERNEL=="0.0.f200",
    - ATTR{buffer\_count}="128"
  - Red Hat:
    - in /etc/sysconfig/network-scripts/ifcfg-eth0
    - add OPTIONS="buffer\_count=128"



## lcrash/crash couldn't read the dump

- Configuration:
  - linux guest under z/VM
  - Problem is independent of any distribution or any product
- Problem Description:
  - Created dump could not be read by
    - lcrash: bad map file
    - crash: WARNING: could not find MAGIC\_START!  
WARNING: cannot read linux\_banner string
  - But 'zgetdump -i' showed:

```
>>> Dump header information <<<
Dump created on: Thu Sep  2 21:04:35 2010
.....
Dump ended on:      Thu Sep  2 21:08:49 2010
Dump End Marker found: this dump is valid.
```

- Tools used for problem determination:
  - lcrash/crash, any hex editor

## Icrash/crash couldn't read the dump (cont'd)

### ■ Problem Origin:

```
>> dump 0x85425000 20
0x0000000085425000: 726f6f743d2f6465 762f646173646131 : root=/dev/dasda1
0x0000000085425010: 20646173643d3161 302d3161322c3230 : dasd=1a0-1a2,20
0x0000000085425020: 302d3230312c3430 302d343036207365 : 0-201,400-406 se
0x0000000085425030: 6c696e75783d3020 5445524d3d64756d : linux=0 TERM=dum
0x0000000085425040: 6220656c65766174 6f723d6366712073 : b elevator=cfq s
0x0000000085425050: 6176657379733d64 62326d6173747220 : avesys=db2mastr
0x0000000085425060: 424f4f545f494d41 47453d3020000000 : BOOT_IMAGE=0 ...
```

- kernel was loaded into Named Saved Segment (NSS)
- Permanent restriction: NSS or DCSS can only be dumped with VMDUMP

### ■ Solutions:

- VMDUMP
  - Takes a very long time (average 15 min per 1 GB)
  - Works noninterruptive
- To be able to use the zlinux dump function, change the kernel parameter line in `/etc/zipl.conf` to eliminate NSS or DCSS (e.g. boot from DASD)
- zLinux dump is preferred over VMDUMP

## References

- Linux on System z project at IBM DeveloperWorks:  
<http://www.ibm.com/developerworks/linux/linux390/>
- Linux on System z: Tuning Hints & Tips  
<http://www.ibm.com/developerworks/linux/linux390/perf>
- Optimize disk configuration for performance:  
[http://www.ibm.com/developerworks/linux/linux390/perf/tuning\\_rec\\_dasd\\_optimize](http://www.ibm.com/developerworks/linux/linux390/perf/tuning_rec_dasd_optimize)
- Linux-VM Performance Website:  
<http://www.vm.ibm.com/perf/tips/linuxper.html>
- IBM Redbooks:  
<http://www.redbooks.ibm.com/>
- IBM Techdocs:  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/Techdocs>

# Questions?



***Susanne Wintenberger***

*Schönaicher Strasse 220  
71032 Böblingen, Germany*

*Certified IT Specialist*

*Phone +49 (0)7031-16-3514*

*Linux on System z*

*swinten@de.ibm.com*