Tools & Utilities

Tiny little helpers & “must haves”

Matthias Strubel - Empalis Consulting GmbH

Trademarks

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

LinuxONE	developerWorks*	z Systems
LinuxONE Emperor	z/Architecture*	IBM*
LinixONE Rockhopper		IBM (logo)*

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.

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.

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

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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, Windows Server and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Node.js is an official trademark of Joyent. IBM SDK for Node.js is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

Java, JavaScript and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

PostgreSQL, and the PostgreSQL logo, are trademarks or registered trademarks of PostgreSQL Global Development Group.

Docker and the Docker logo are trademarks or registered trademarks of Docker, Inc. in the United States and/or other countries.

"Python" is a registered trademark of the PSF. The Python logos (in several variants) are use trademarks of the PSF as well. "PyCon" is a trademark of the PSF.

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.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

IBM

Today's topics

- Automate & Managing you systems
- Monitoring & Analysis
- LinuxOne specific tools
- “Debugging” network issues
- qcow2 handling & creating backups

Disclaimer

- Use & try at your own risk.
- Not every tool / script is supported by IBM.
- Please test the tools & tweaks on a lab. system **before** deploying to production system.
- Most tips will help you on other platforms, too.

Automate & managing you systems

- Taking care of a few pets

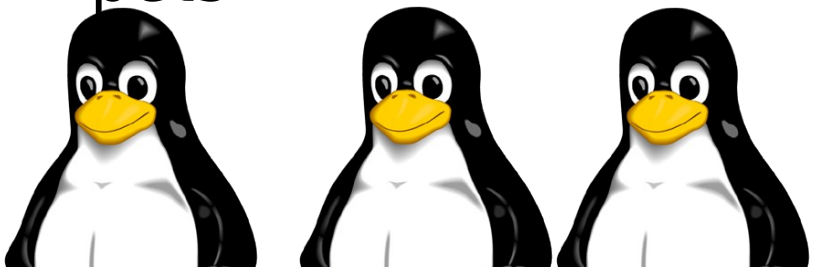

- or having a zoo full of animals

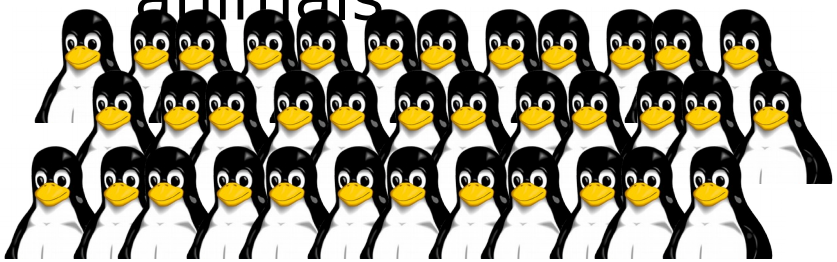

- Usually make use of configuration and “setup” tools like
  - CHEF
  - Puppet
  - Rundeck
  - OpenStack
  - .. you name it

Fabric & Cuisine

- Fabric
  - Fabric is a Python™ (2.5-2.7) library and command-line tool for streamlining the use of SSH for application deployment or systems administration tasks. http://www.fabfile.org/
- Cuisine
  - Cuisine is a small set of functions that sit on top of Fabric, to abstract common administration operations such as file/dir operations, user/group creation, package install/upgrade, making it easier to write portable administration and deployment scripts. https://github.com/sebastien/cuisine

Both libraries are Python2 based.
They don't need to be installed directly on KVM for LinuxONE or a managed guest system.

Fabric – What is ssh streamlining?

- Shell example

```
$  ssh username@mymachine sudo yum update
```

- More complex-- still pretty easy..

```
$ ssh username@mymachine sh -c “sudo yum install apache2 &&  \
                                sudo systemctl start httpd”
```

- Fabric example.
  - Create fabfile.py:

```
#!/usr/bin/python
from fabric.api import run
def setup:
 run('sudo yum update')
 run('sudo yum install apache2')
 run('sudo systemctl start httpd')
 run('sudo mkdir -p /var/www/mystuff && chown -R
httpd:httpd /var/www/mystuff')
```

  - Execute in commandline.

```
$ fab -u username -H mymachine,myothermachine setup
```

Cuisine – army knife on top of fabric

- Cuisine example inside Fabric:
  - Create `fabfile.py`:

```
#!/usr/bin/python
import cuisine

def setup:
 with mode_sudo():
 cuisine.select_package('yum')   # default is apt
 cuisine.package_ensure('apache2')
 cuisine.run('systemctl start httpd')
 cuisine.dir_ensure('/var/www/mystuff',
                    owner=httpd, group=htttpd,
                        recursive=True)
```

- Execute:

```
$ fab -u username -H mymachine,myothermachine setup
```

Cuisine – Functional overview

- Text-processing
- File & Directory handling
- Package management
- Direct shell commands
- User & Group

```
          /     /

 ___          ___    ___  ___
|    |   )| |___ | |   )|___)
|__  |__/ |  __/ | |  / |__

-- Chef-like functionality for Fabric
```

- Examples:
  - Cuisine stand-alone: https://goo.gl/yGoyN1 (pojnts to bitbucket.org)
  - As a fabfile https://goo.gl/tW4vyO (points to Github.com)
- Documentation:
  - Cuisine: https://github.com/sebastien/cuisine
  - Fabric: http://www.fabfile.org/
  - Ipython (interactive python): http://ipython.org/

Monitoring & Analysis

```
Mar  1 08:39:18 oc1163652161 kernel: input: ThinkPad Extra
Buttons as /devices/platform/thinkpad_acpi/input/input15
Mar  1 08:39:18 oc1163652161 kernel: input: HDA Intel PCH Mic
as /devices/pci0000:00/0000:00:1b.0/sound/card1/input16
Mar  1 08:39:18 oc1163652161 kernel: input: HDA Intel PCH
Headphone as
/devices/pci0000:00/0000:00:1b.0/sound/card1/input17
Mar  1 08:39:18 oc1163652161 kernel: e1000e 0000:00:19.0:
eth0: registered PHC clock
Mar  1 08:39:18 oc1163652161 kernel: e1000e 0000:00:19.0:
eth0: (PCI Express:2.5GT/s:Width x1) 54:ee:75:61:f9:87
Mar  1 08:39:18 oc1163652161 kernel: e1000e 0000:00:19.0:
eth0: Intel(R) PRO/1000 Network Connection
Mar  1 08:39:18 oc1163652161 kernel: e1000e 0000:00:19.0:
eth0: MAC: 11, PHY: 12, PBA No: 1000FF-0FF
Mar  1 08:39:18 oc1163652161 kernel: shpchp: Standard Hot
Plug PCI Controller Driver version: 0.4
Mar  1 08:39:18 oc1163652161 kernel: ACPI Warning: SystemIO
range 0x0000000000001828-0x000000000000182f conflicts
with OpRegion 0x0000000000001800-0x000000000000187f (\
_SB_.PCI0.LPC_.PMIO) (20090903/utaddress-254)
Mar  1 08:39:18 oc1163652161 kernel: ACPI: If an ACPI driver is
available for this device, you should use it instead of the native
driver
Mar  1 08:39:18 oc1163652161 kernel: ACPI Warning: SystemIO
range 0x0000000000000800-0x000000000000083f conflicts
with OpRegion 0x0000000000000800-0x000000000000087f (\
_SB_.PCI0.LPC_.LPIO) (20090903/utaddress-254)
```

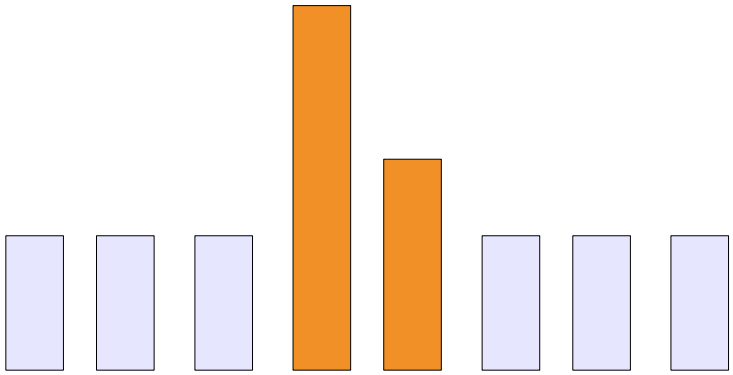

KVM for LinuxONE

Preinstalled:

- Nagios agent
- perf
- sadc
- ....

sadc / sar

- Characteristics: Very comprehensive, statistics data on device level
- Objective: Suitable for permanent system monitoring and detailed analysis
- Usage (recommended):
  - monitor /usr/lib64/sa/sadc [-S XALL] [interval in sec] [outfile]
  - View: sar -A -f [outfile]
- Package: RHEL: sysstat.s390x SLES: sysstat KVM: preinstalled
- Shows:
  - CPU utilization
  - Disk I/O overview and on device level
  - Network I/O and errors on device level
  - Memory usage/Swapping
  - ... and much more
  - Reports statistics data over time and creates average values for each item
- Hints
  - sadc parameter “-S XALL” enables the gathering of further optional data
  - Shared memory is listed under 'cache'
  - [outfile] is a binary file, which contains all values. It is formatted using sar
    - enables the creation of item specific reports, e.g. network only
    - enables the specification of a start and end time → time of interest

- Setup before problems occur
- Create regularly plain text report output
- Place output to /var/log/sar
  => will be included to while collecting support data

SAR – example: CPU utilization

Per CPU values:
watch out for
- system time (kernel)
- user (applications)
- irq/soft (kernel, interrupt handling)
- idle (nothing to do)
- iowait time (runnable but waiting for I/O)
- steal time (runnable but utilized somewhere else)

root@h42lp42

File Edit View Terminal Help

14:14:55	CPU	%user	%nice	%system	%iowait	%steal	%idle
14:15:05	all	26.64	0.00	12.03	25.92	6.24	29.16
14:15:05	0	43.81	0.00	5.49	23.25	4.99	22.46
14:15:05	1	4.30	0.00	10.19	28.67	9.89	46.95
14:15:05	2	11.81	0.00	28.03	45.15	5.01	10.01
14:15:05	3	46.61	0.00	4.49	6.79	4.99	37.13
14:15:15	all	27.19	0.00	11.93	25.11	7.75	28.01
14:15:15	0	90.60	0.00	3.70	0.00	5.70	0.00
14:15:15	1	9.24	0.00	22.49	41.57	9.24	17.47
14:15:15	2	5.98	0.00	14.64	46.71	9.06	23.61
14:15:15	3	2.90	0.00	6.99	12.09	7.09	70.93

SAR – example: Disk I/O – per device

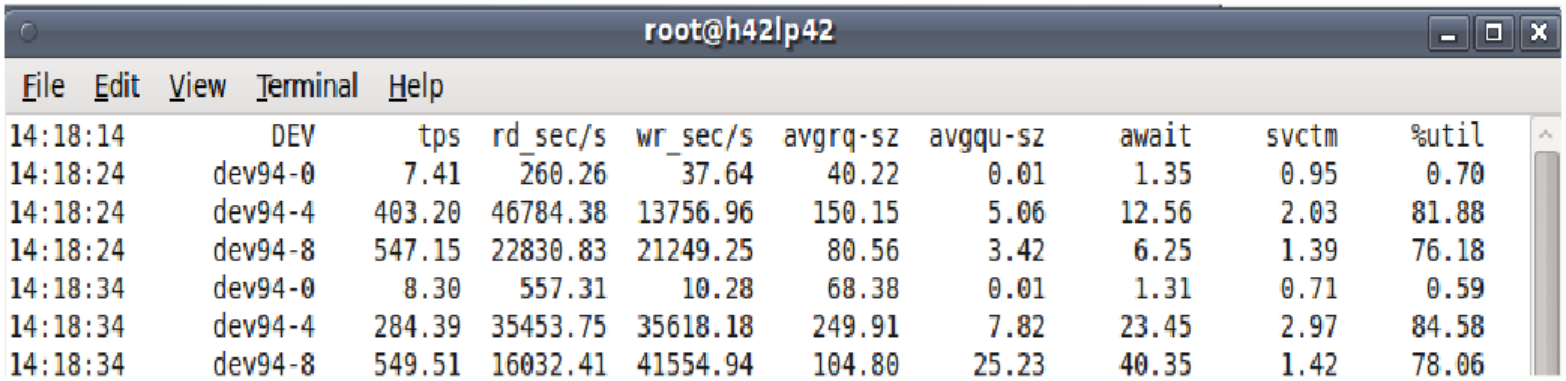

Is your I/O balanced across devices?
Imbalances can indicate issues wit a LV setup.

tps and avgrq-sz combined can be important.
Do they match your sizing assumptions?

Await shows the time the application has to wait.

dbginfo.sh - Collecting data for support purposes

- Collects debugging information and system configuration
- dbginfo.sh script is required to run before rebooting the system
- dbginfo.sh script continues to run even on issues during data collection
- dbginfo.sh script mounts debugfs/s390dbf automatically to collect LinuxOne specific trace data
- Collecting the sysfs can take some time dependent on the number of devices being attached
- Running dbginfo.sh script requires 'enough' disk space under /tmp

- Check out:
  http://www.ibm.com/developerworks/linux/linux390/s390-tools.html

dbginfo.sh – Example output

```
[root@system]# dbginfo.sh
dbginfo.sh: Debug information script version 1.15.0-0.136.3
Copyright IBM Corp. 2002, 2013

Hardware platform     = s390x
Kernel version        = 3.0.76 (3.0.76-0.7-default)
Runtime environment   = z/VM

1 of 7: Collecting command output

2 of 7: Collecting z/VM command output

3 of 7: Collecting procfs

4 of 7: Collecting sysfs

5 of 7: Collecting log files

6 of 7: Collecting config files

7 of 7: Collecting osa oat output skipped - not available

Finalizing: Creating archive with collected data

Collected data was saved to:
>>  /tmp/DBGINFO-2014-06-20-10-42-42-system-123456.tgz  <<
```

LinuxOne specific tools

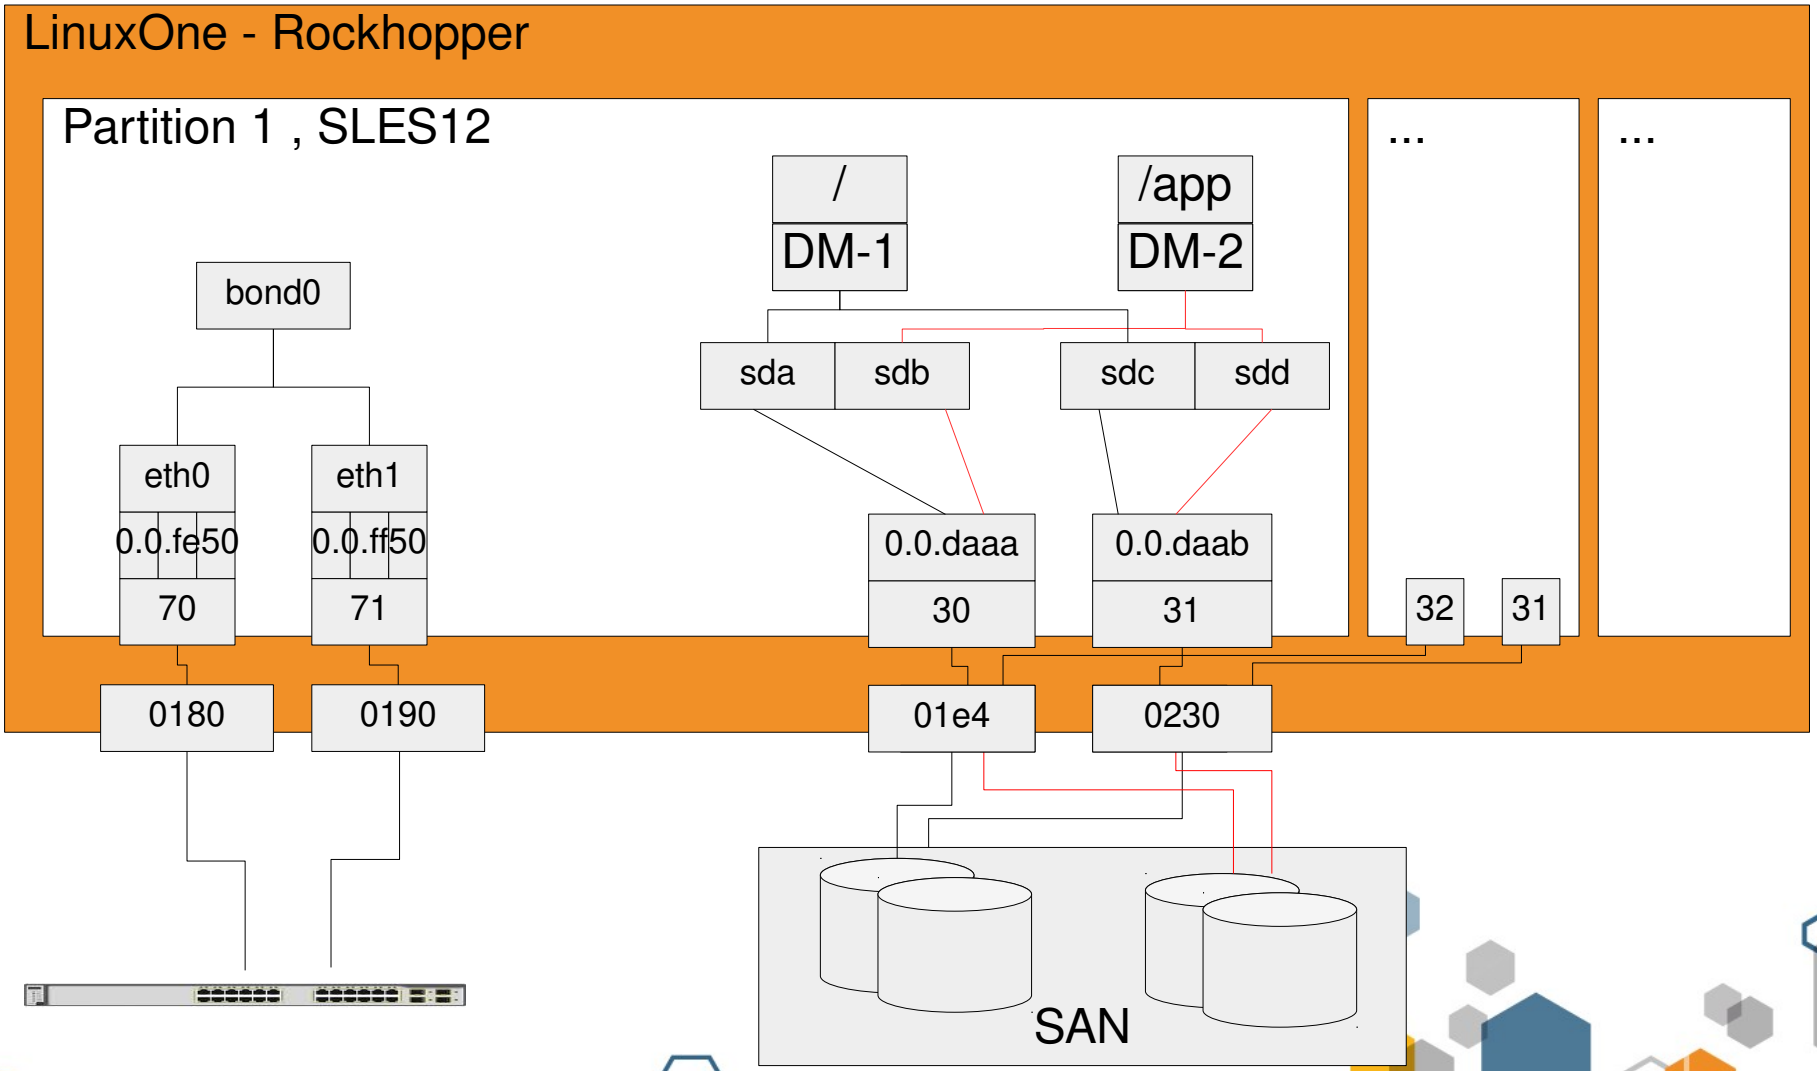

LinuxOne specific tools – Which devices are visible?

```
# lscss
Device   Subchan.  DevType CU Type Use  PIM PAM POM  CHPIDs
----------------------------------------------------------------------
0.0.fe50 0.0.0009  1732/01 1731/01 yes  80  80  ff   70000000 00000000
[..]
0.0.ff50 0.0.0015  1732/01 1731/01 yes  80  80  ff   71000000 00000000
[..]
0.0.daaa 0.0.001c  1732/03 1731/03 yes  80  80  ff   30000000 00000000
0.0.daab 0.0.001c  1732/03 1731/03 yes  80  80  ff   31000000 00000000

# lszfcp
0.0.daaa host1
0.0.daab host2
```

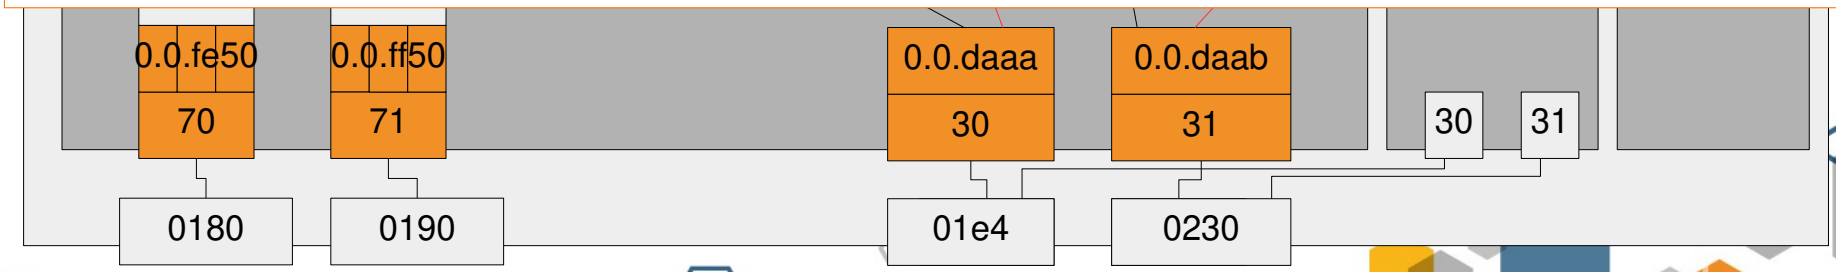

LinuxOne specific tools – Network card details

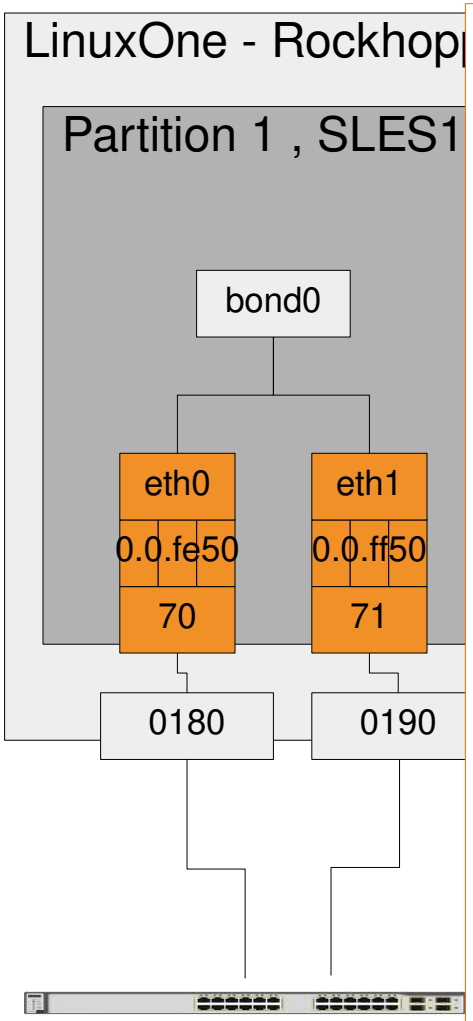


```
# lsqeth eth0

Device name                    : eth0
-------------------------------------------------------------
        card_type              : OSD_1000
        cdev0                  : 0.0.fe50
        cdev1                  : 0.0.fe51
        cdev2                  : 0.0.fe52
        chpid                  : 70
        online                 : 1
        portno                 : 0
        state                  : UP (LAN ONLINE)
        priority_queueing      : always queue 0
        buffer_count           : 64
        layer2                 : 1
        isolation              : none
```

LinuxOne specific tools – SAN configuration

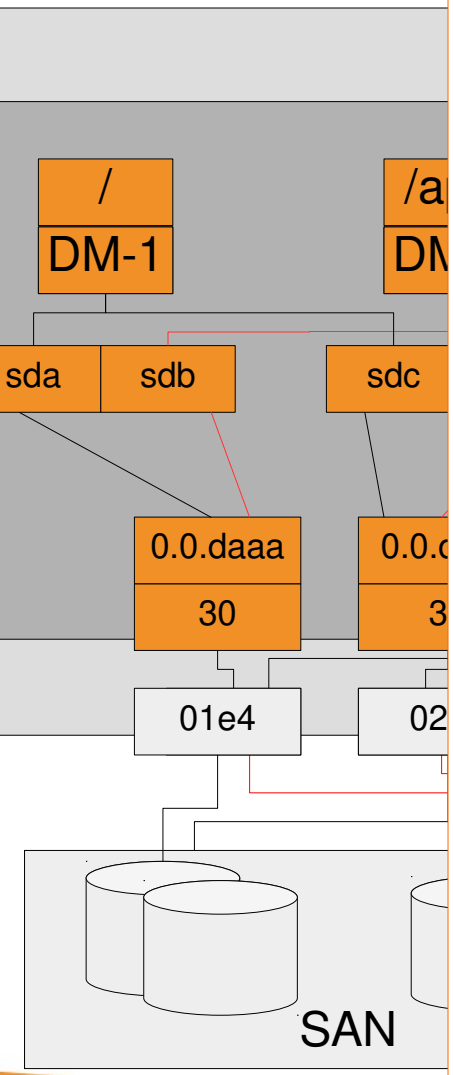


```
# multipath -l
LUN8093 (36005076305ffc1ae0000000000008093) dm-1 IBM     ,210
size=15G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=0 status=active
  |- 4:0:0:1083392128 sda 8:0  active undef running
  |- 5:0:0:1083392128 sdc 8:32 active undef running

LUN24ab (36005076305ffc1ae0000000000024ab) dm-2 IBM     ,210
size=1000G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=0 status=active
  |- 0:0:1:1084964900 sdb 8:16  active undef running
  |- 1:0:3:1084964900 sdd 8:48  active undef running

# lsluns
Scanning for LUNs on adapter 0.0.daaa
    at port 0x50050763050841ae:
        0x4080409200000000
        0x4080409300000000
        0x4081409200000000
        0x4081409300000000
Scanning for LUNs on adapter 0.0.daab
    at port 0x50050763050341ae:
        0x4080409200000000
        0x4080409300000000
        0x4081409200000000
        0x4081409300000000
```

LinuxOne specific tools – further documentation

- Tools for change hardware configuration: chchp, chcpu, cio_ignore, qethconf ….

- Find tool documentation in: “Device Drivers Features and Commands”
- http://www.ibm.com/developerworks/linux/linux390/documentation_dev.html

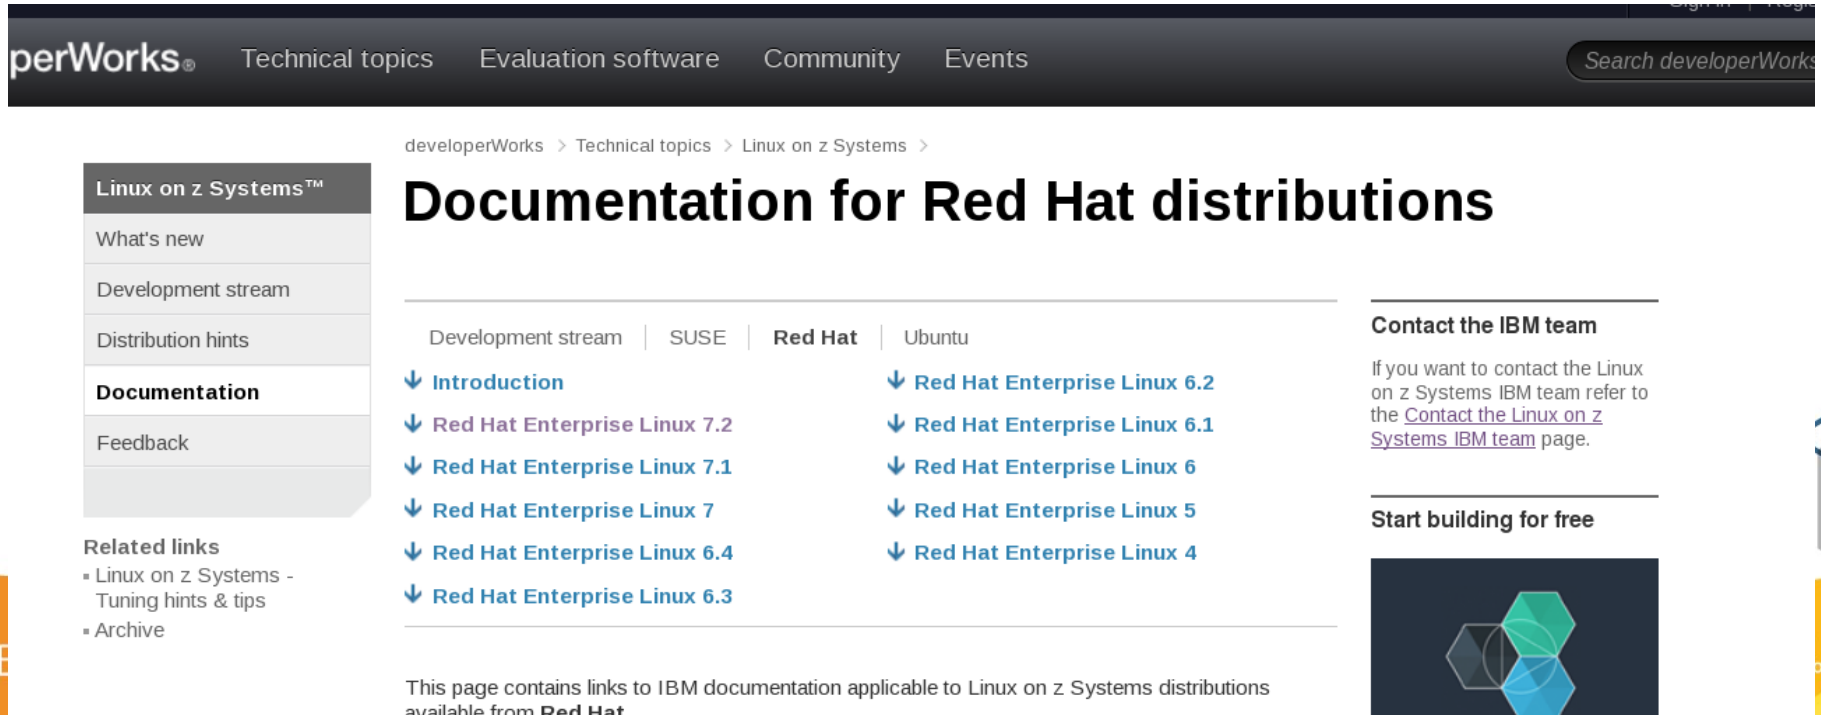

“Debugging” network issues

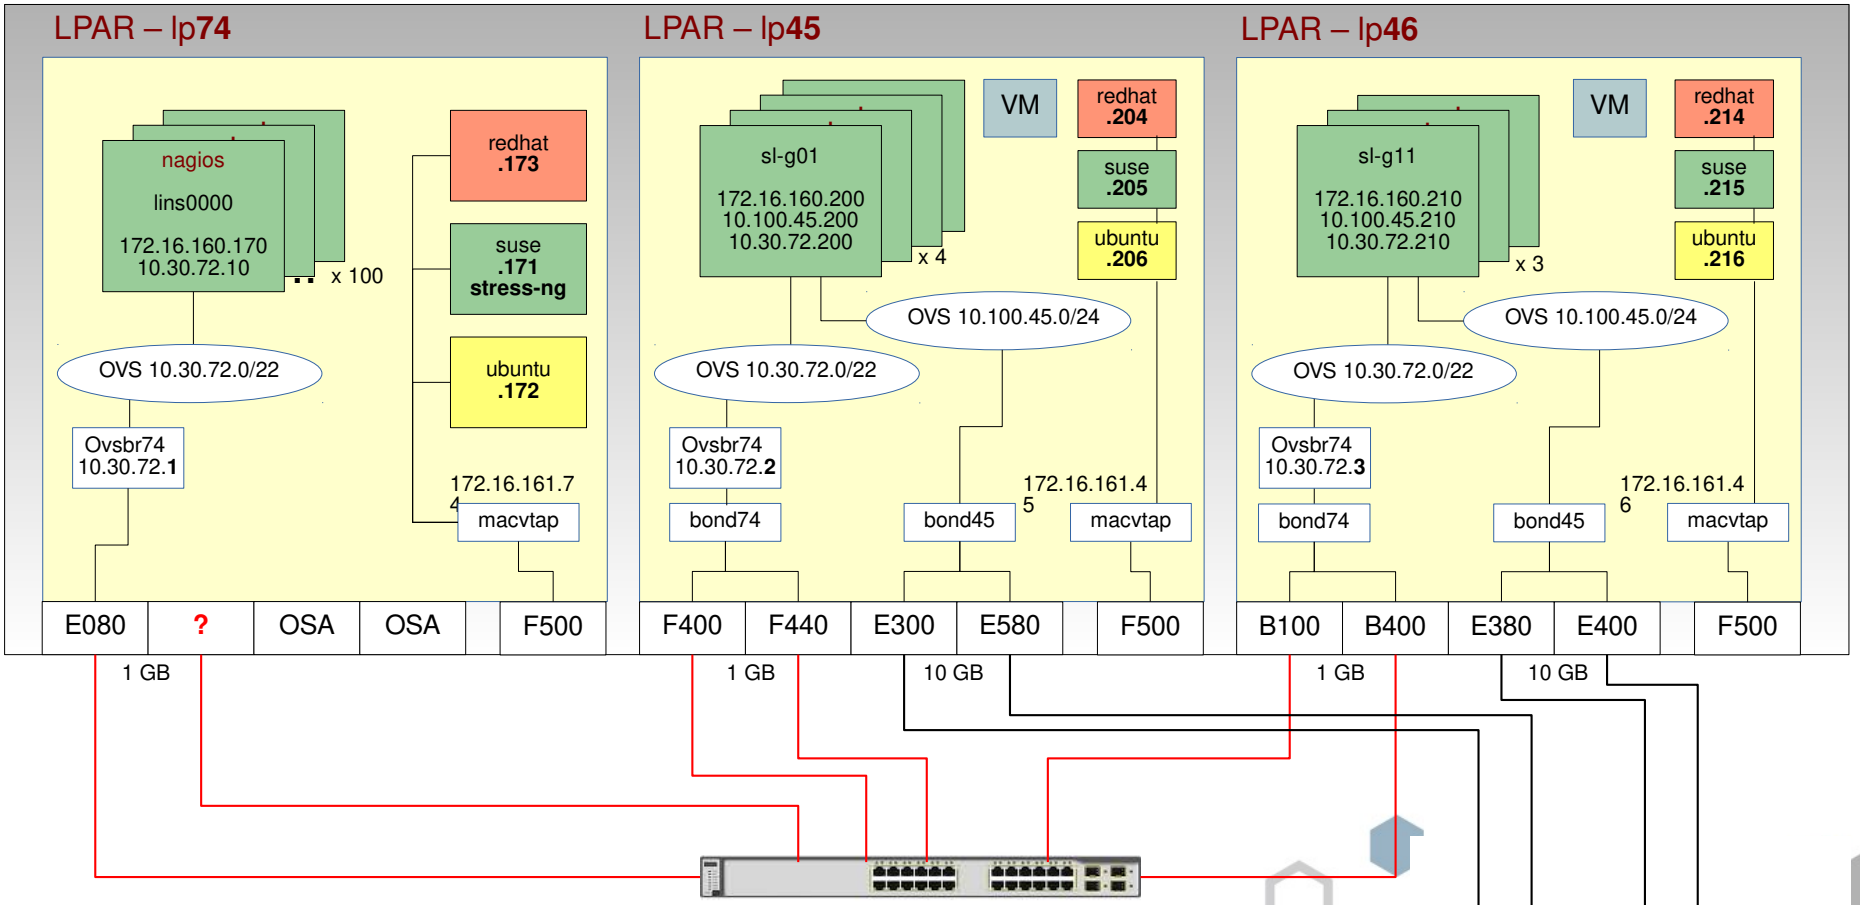

How to find out guest related interfaces...

```
$ virsh dumpxml guest01
```

```
      <alias name='virtio-disk0'/>
      <address type='ccw' cssid='0xfe' ssid='0x0' devno='0x0001'/>
    </disk>
    <interface type='network'>
      <mac address='52:54:00:07:c9:1a'/>
      <source network='default' bridge='virbr0'/>
      <target dev='vnet0'/>
      <model type='virtio'/>
      <alias name='net0'/>
      <address type='ccw' cssid='0xfe' ssid='0x0' devno='0x0000'/>
    </interface>
```

```
$ virsh domiflist guest01
   Interface  Type       Source     Model       MAC
   -------------------------------------------------------
   vnet0      network    default    virtio      52:54:00:07:c9:1a
```

… and everything is wired up correctly.

IBM

bridge-utils

```
$ brctl show
bridge name     bridge id               STP enabled     interfaces
pan0            8000.000000000000       no
virbr0          8000.525400000452       yes             virbr0-nic
                                                        vnet0
```

open-vswitch

```
# List available ovs bridges
$ ovs-vsctl list-br
    ovsbr45
    ovsbr74

# List connected interfaces
$ ovs-vsctl list-ports ovsbr45
    bond45
    vnet0
    vnet2
    vnet3
    vnet5
    vnet6
    vnet8
    vnet9
```

TCPDump or how to wiretap your network traffic!

- Characteristics: dumps network traffic to console/file
- Objective: analyze packets of applications manually
- Usage: “tcpdump ...”
  - -i <interface> to limit to a particular network inte

```
# Capture all packets on vnet0
$ tcpdump -i vnet0
```

- Package: RHEL: tcpdump SLES: tcpdump KVM: preinstalled

```
tcpdump host pserver1
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
13:30:00.326581 IP pserver1.boeblingen.de.ibm.com.38620 > p10lp35.boeblingen.de.ibm.com.ssh: Flags [.], ack 3142, win
102, options [nop,nop,TS val 972996696 ecr 346994], length 0
13:30:00.338239 IP p10lp35.boeblingen.de.ibm.com.ssh > pserver1.boeblingen.de.ibm.com.38620: Flags [P.], seq 3142:3222,
ack 2262, win 2790, options [nop,nop,TS val 346996 ecr 972996696], length 80
13:30:00.375491 IP pserver1.boeblingen.de.ibm.com.38620 > p10lp35.boeblingen.de.ibm.com.ssh: Flags [.], ack 3222, win
102, options [nop,nop,TS val 972996709 ecr 346996], length 0
[...]
^C
31 packets captured
31 packets received by filter
0 packets dropped by kernel
```

- Not all devices support dumping packets in older distribution releases
  - Also often no promiscuous mode
- Check flags or even content if your expectations are met
- -w flag exports captured unparsed data to a file for later analysis in libpcap format
  - Also supported by wireshark
- Usually you have to know what you want to look for

Simple AdHoc Webserver

- Small and simple way to test TCP based communication.
- Quickly exchanging files, or create an “adhoc repository” for testing purposes.
- !! Do not use for production !!
- !! Do not use for performance tests !!

- `$ python2 -m SimpleHTTPServer <port>`

  ```
  python2  -m SimpleHTTPServer 8080
  ```

- `$ python3 -m http.server <port>`

  ```
  python3 -m http.server 8080
  ```

qcow2 handling & Creating backups

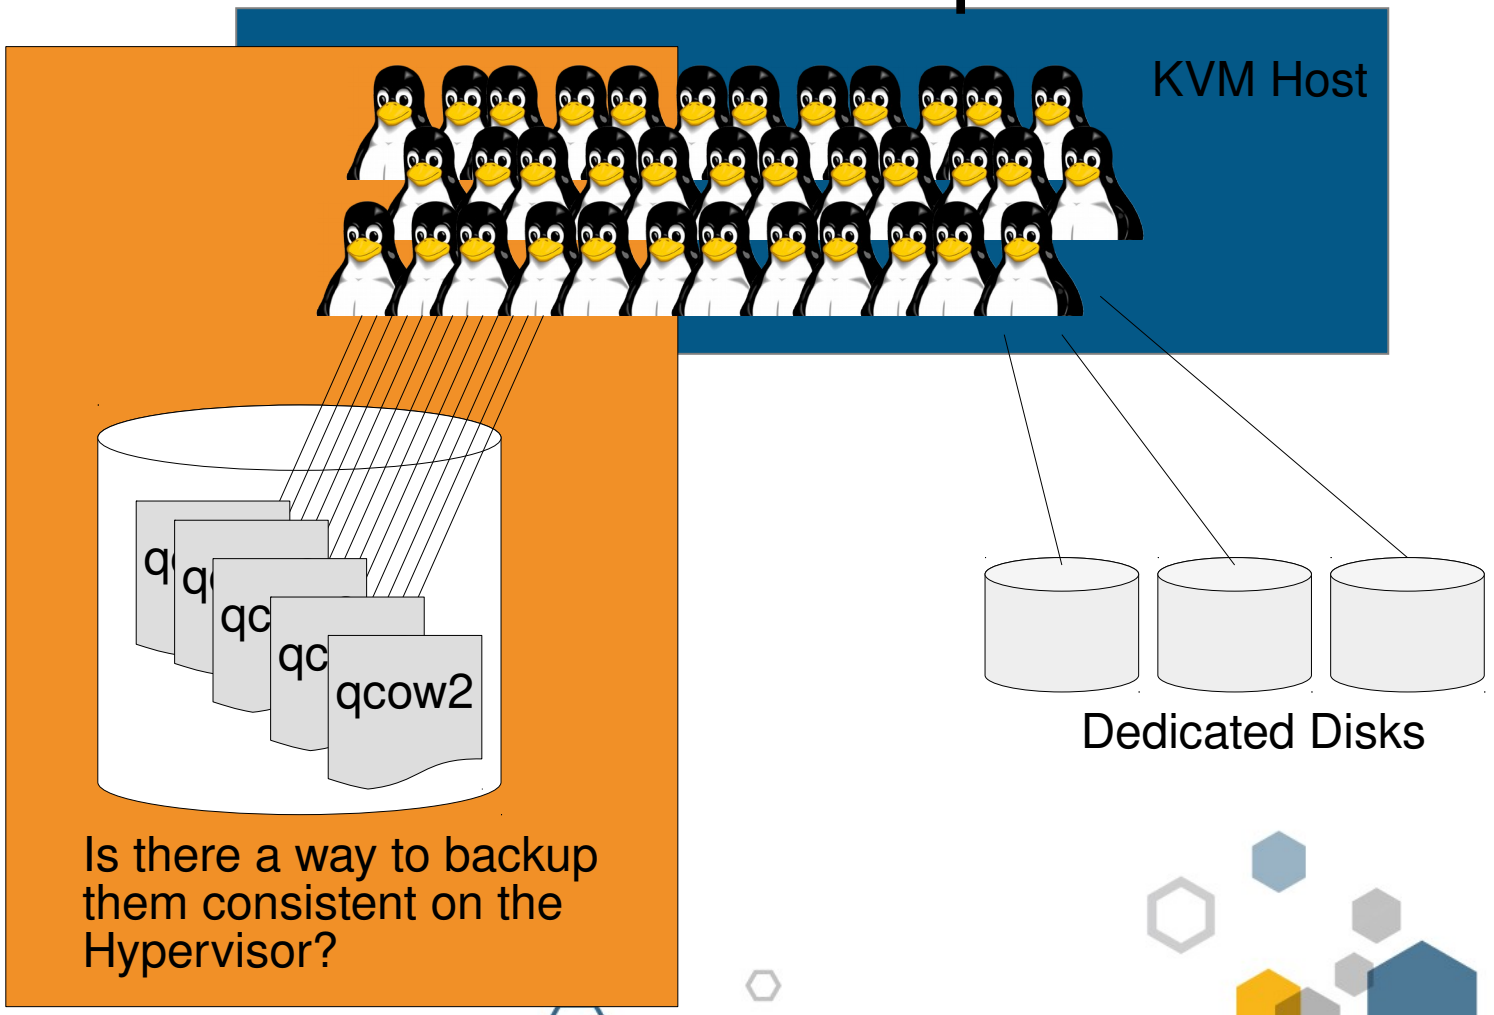

qcow2 caveats - sparse parts in image-files

- Sometimes qcow2 stores empty areas, which are zero-filled.
- These sparse areas do not occupy physical disk space, but still show the maximum virtual filesize
- This can be displayed using additional option flags.
  - `$ ls -s -lh`
    `2.9G -rw-r--r--. 1 qemu qemu 7.9G Mar 14 17:16 qcow2.img`

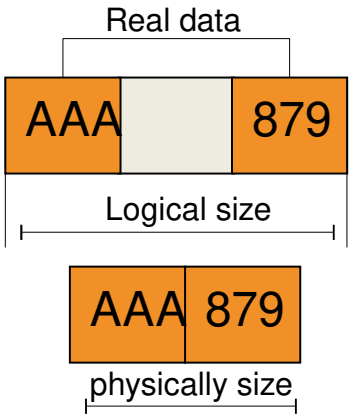

- One reason for this behavior can be
  - qcow2 create with: `-o preallocation=metadata`

- CLI tools need additional flags to work efficiently with that areas. If the tool encounters spare areas without the option set, it will copy the complete logical size instead of the used parts only.
  - `cp --sparse=auto`
  - `rsync -S` `rsync --sparse`
  - `tar --spare`

Shrink qcow2 image-files again

- Inside the guest-OS: if files are deleted, space is freed.
- On KVM, a qcow2 image can be reduced in filesize, when running a utility after the guest shutdown.

- The following example is a qcow2 image with one single filesystem & a ~500MB file, which was deleted after creation:

```
$ qemu-img convert   -O qcow2    big.qcow2      small.qcow2
$ ls -lsh
550M -rw-r--r--. 1 root root 550M Mar 14 18:32 big.qcow2
1.4M -rw-r--r--. 1 root root 1.4M Mar 14 18:33 small.qcow2
```

Backing up your VM with qcow2 images using Online Forward Incremental Backup

- Utility using shell, via libvirt/virsh & qemu-image
- OpenSource script: https://github.com/dguerri/LibVirtKvm-scripts
- Process is:
  - Halt guest via guest-agent quickly, or does a “dump domain state”
  - Create an incremental snapshot and change libvirt config
  - Resume guest, which proceed working on snapshot file.

- Also function to merge incremental together to a full image again.
- Helpful for backup-strategy based on “last modified” or
  to create recovery points during patch-days.

Backing up your VM with qcow2 images using Online Forward Incremental Backup

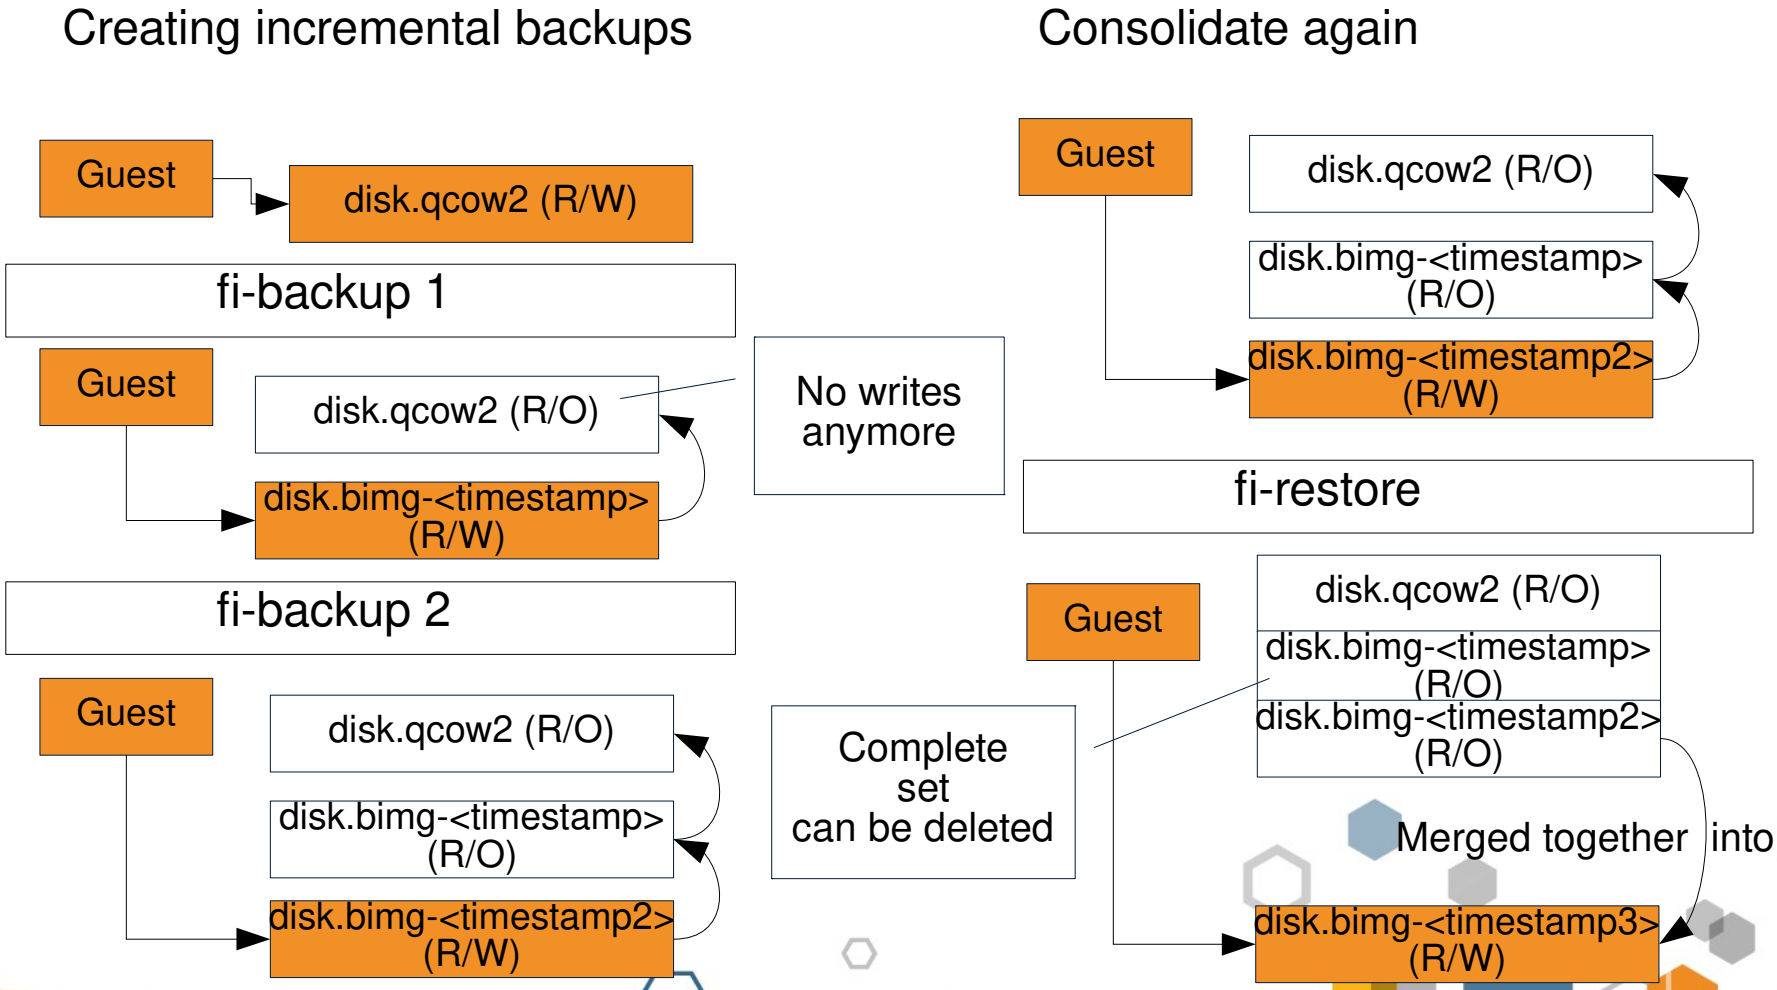

Matthias Strubel
(Business Partner)

Service & Support

Linux on z Systems
&
KVM for z Systems

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

Office:
mstrubel@de.ibm.com

Questions?

Empalis Consulting GmbH
Wankelstraße 14
70563 Stuttgart

Tel. +49 711 46 928 260

Notices and Disclaimers

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

Notices and Disclaimers (cont'd)

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM EXPRESSLY DISCLAIMS ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Bluemix, Blueworks Live, CICS, Clearcase, DOORS®, Enterprise Document Management SystemTM,Global Business Services ®, Global Technology Services ®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®,Netcool®, OMEGAMON, OpenPower, PureAnalyticsTM, PureApplication®, pureClusterTM, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, SoDA, SPSS, StoredIQ,
Tivoli®, Trusteer®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml .

Backup

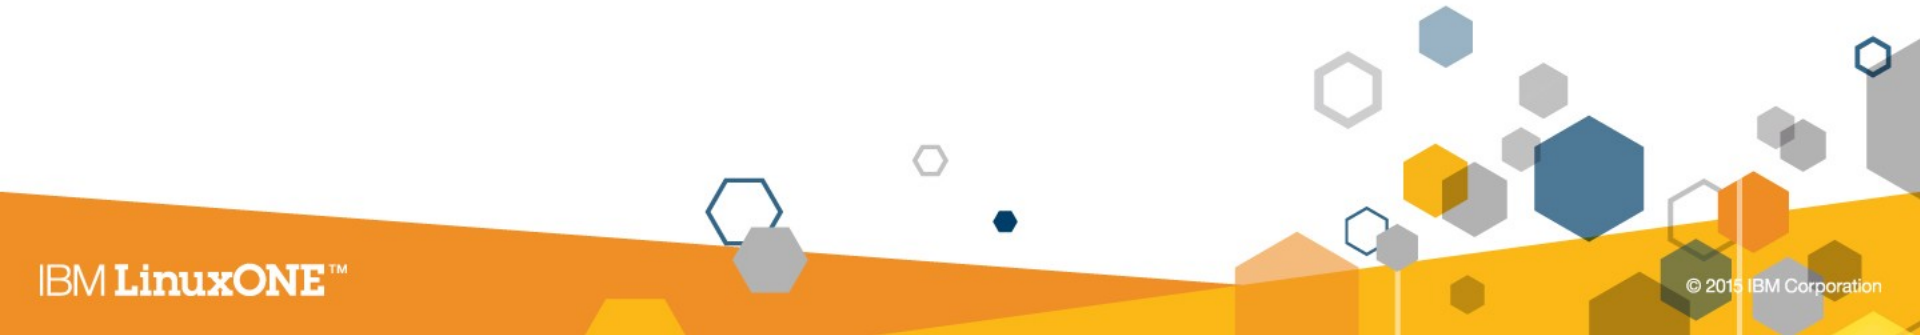

Backing up your qcow2 images II

- Creates complete running guest backup and keeps n copies.
- More consistent then using “hot” LVM snapshots.
- Needed qcow2 files are identified using the virsh guest definition.
- Shell script & virsh

- Process:
  1) Freezes VM for a couple of seconds & saves VM state.
  2) Uses “blockcopy” feature of libvirt/qemu to create a full backup of qcow2 disks.
  3) Resume VM

- Article about technical background:
  - http://soliton74.blogspot.de/2013/08/about-kvm-qcow2-live-backup.html
- Script location: https://goo.gl/mNZ1X6 (points to gist.github.com)

(unsupported by ccw mode) Decrease qcow2 imagesize again

- https://pve.proxmox.com/wiki/Shrink_Qcow2_Disk_Files
- https://chrisirwin.ca/posts/discard-with-kvm/
- Libivirtd configuration option needs to be added
  - <driver name='qemu' type='qcow2' cache='writeback' discard='unmap'/>
- Configure your VMs themselves to discard unused data
- Manually run an fstrim to discard all the currently unused crufty storage you've collected on all applicable filesystems:
  - sudo fstrim -a
- Going forward, you can either add 'discard' to the mount options in fstab, or use fstrim periodically. I opted for fstrim, as it has a systemd timer unit that can be scheduled:
  - sudo systemctl enable fstrim.timer
  - sudo systemctl start fstrim.timer
-

IBM LinuxONE™

Color palette