KVM on IBM System z Customer Experience

SHARE Anaheim, Session 14764
Tuesday, March 11, 2014: 4:30 PM

Tobias Dörkes, Finanz Informatik

finanz informatik

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

The company serves a large part of German retail banking market

Finanz Informatik – Company	
Revenue (in mill. €) (2012)	1,425
with saving banks	972
with state banks	202
Employees (full-time equivalents)	4,992

Customers	
Savings banks	417
State banks + DekaBank	9
State home loan banks	10
Accumulated balance sheet of supported savings banks (in bill. €) (2012)	1,063

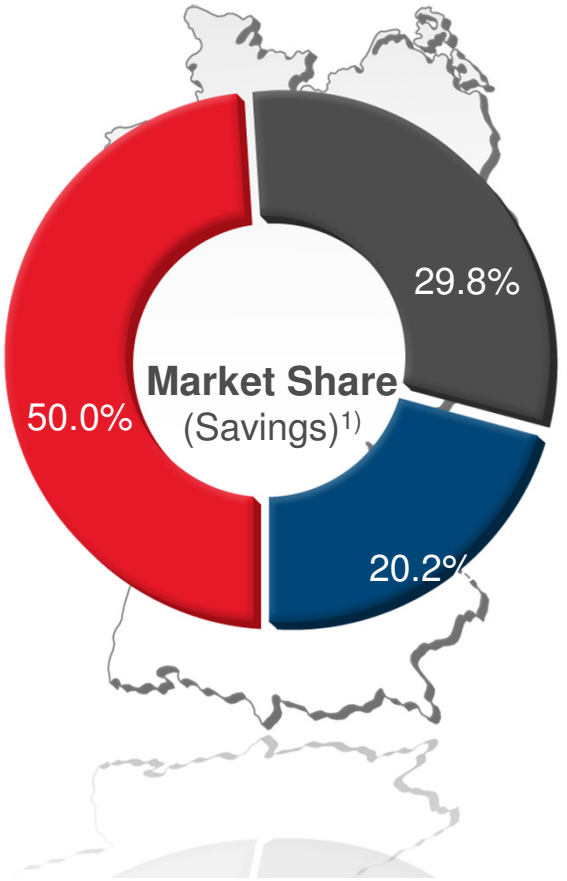

Savings Banks Financial Group Credit Unions Private Banks, other

December 31st, 2013

1) Sources: DSGV (12/31/2012); German Federal Bank; Others.

finanz informatik

Significant scale can be achieved through bundling volume IT services

Supported financial institutions	
Branches of supported savings banks (2012)	15,097
Bank-specific employees of supported savings banks (2012)	190,095

Processing volumes	
Supported accounts (in mill.)	125

End devices	
ATMs	24,325
Statement printers	14,890
Other self-service terminals	14,055
Booked entries per annum (in bill.)	11

December 31st, 2013

finanz informatik

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Virtualization basics

What does hardware virtualization mean ?

- Virtualization means serving multiple operating systems in parallel on one physical hardware.
- Virtualized operating systems are running in so called »virtual machines« (VM)
- A very easy naming convention is to call a VM a »guest« and the serving system providing the virtualization capabilities a »host«.

Virtualization basics

Different types of hardware virtualization

- **Full virtualization:** Almost complete simulation of the actual hardware. This allows guests to run unmodified (e.g. QEMU, VMware workstation, z/VM)
- **Partial virtualization:** Not all of the actual hardware is simulated, only an address space is created. So the guest may need to be modified to run in this environment (eg. z/OS, Linux)
- **Paravirtualization:** No simulation at all. Guests are executed in their own domains. The VM is similar to the underlying hardware, but not the same - paravirtualized drivers are needed. (eg XEN, KVM)
- **Hardware-assisted virtualization:** The Host-CPU improves hardware virtualization efficiency by owning virtualization functions. (e.g. Intel-VT, AMD-V, System z)

Virtualization basics

Different types of hypervisors

A hypervisor or virtual machine manager creates and runs VMs. There are two different types of hypervisor implementation

- **Type 1 (Native, bare metal):**
  The hypervisor runs directly on the host‘s hardware. No need for an additional OS to control and manage the VMs. (eg. PR/SM, z/VM, XenServer, VMware ESX, Hyper-V)

- **Type 2 (Hosted):**
  The hypervisor runs as an application on top of a normal operating system. (eg. VMware workstation, VirtualBox)

And KVM? Personally, i would say it is a Type 1 hypervisor as it runs in the kernel and not on some second software level. But there are other opinions.

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

What is KVM and how does it fit with QEMU and Libvirt ?

KVM means »Kernel-based Virtual Machine«

KVM is no emulator itself.
KVM just provides an interface /dev/kvm to set up VMs.

KVM is a Linux kernel module that allows a user space program access to the hardware virtualization features of various processors, with which QEMU is able to offer virtualization for x86, PowerPC, and S/390 guests. When the target architecture is the same as the host architecture, QEMU can make use of KVM particular features, such as acceleration.

Wikipedia at http://en.wikipedia.org/wiki/Qemu

What is KVM and how does it fit with QEMU and Libvirt ?

QEMU stands for »**Quick EMUlator**« and is a processor emulator that relies on dynamic binary translation to achieve a reasonable speed while being easy to port to new host CPU architectures.

Wikipedia at http://en.wikipedia.org/wiki/Qemu

QEMU emulates:

- CPUs, even for different architectures.
- various hardware components needed to create a VM (network card, storage, …)

QEMU does I/O, KVM does CPU, memory and interrupt controller.
QEMU uses KVM as an accelerator to access hardware features.

What is KVM and how does it fit with QEMU and Libvirt ?

libvirt is an open source API, daemon and management tool for managing platform virtualization. It can be used to manage Linux KVM, Xen, VMware ESX, QEMU and other virtualization technologies. These APIs are widely used in Orchestration Layer for Hypervisors in the development of a cloud based solution.

Wikipedia at http://en.wikipedia.org/wiki/Libvirt

libvirt provides:

- A directory for configuration data and operational state of VMs
- The **libvirtd daemon** is the server side daemon component of the libvirt virtualization management system. (man page)
  It runs on host servers and provides remote management services.
- The **virsh command shell** is the main interface for managing guest domains. It is an interactive shell and batch scriptable tool. (man page)

What is KVM and how does it fit with QEMU and Libvirt ?

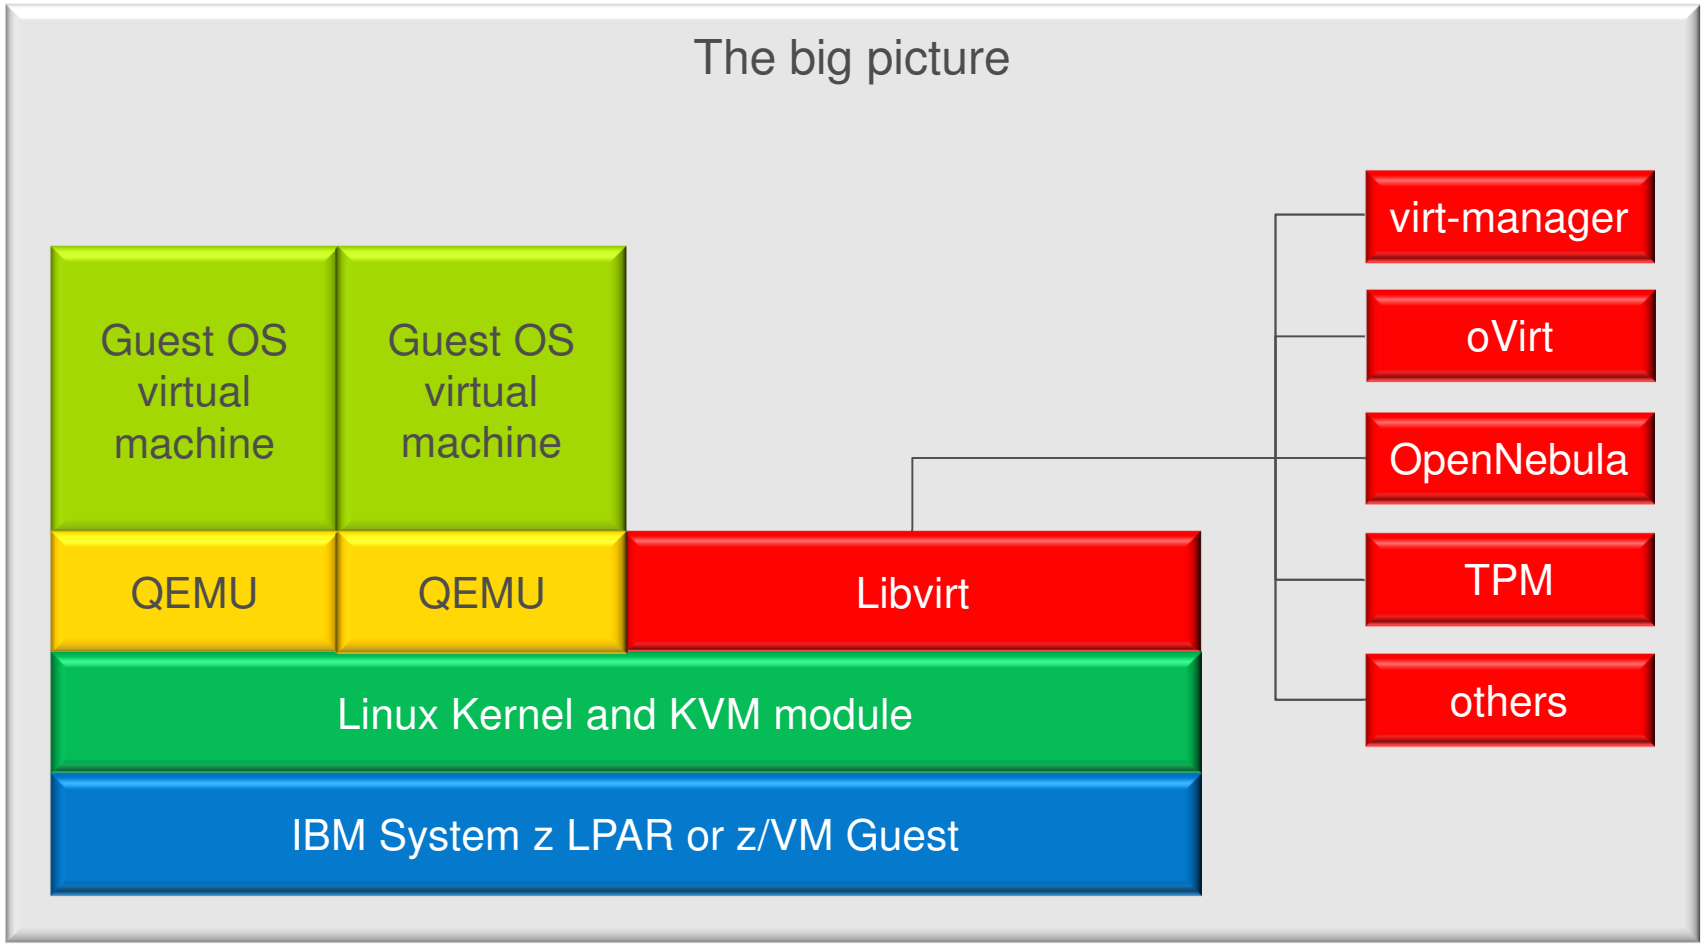

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Prerequisites for running KVM on IBM System z (s390x)

There are just a few requirements for running KVM on System z

- IBM System z9 to get hardware support for KVM

Linux distributions:

- **SLES 11 SP3** has a KVM s390x technical preview built in.
- **RHEL 6.3** has the KVM module available, but the userspace tools need to be compiled manually

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Prepare the KVM host system

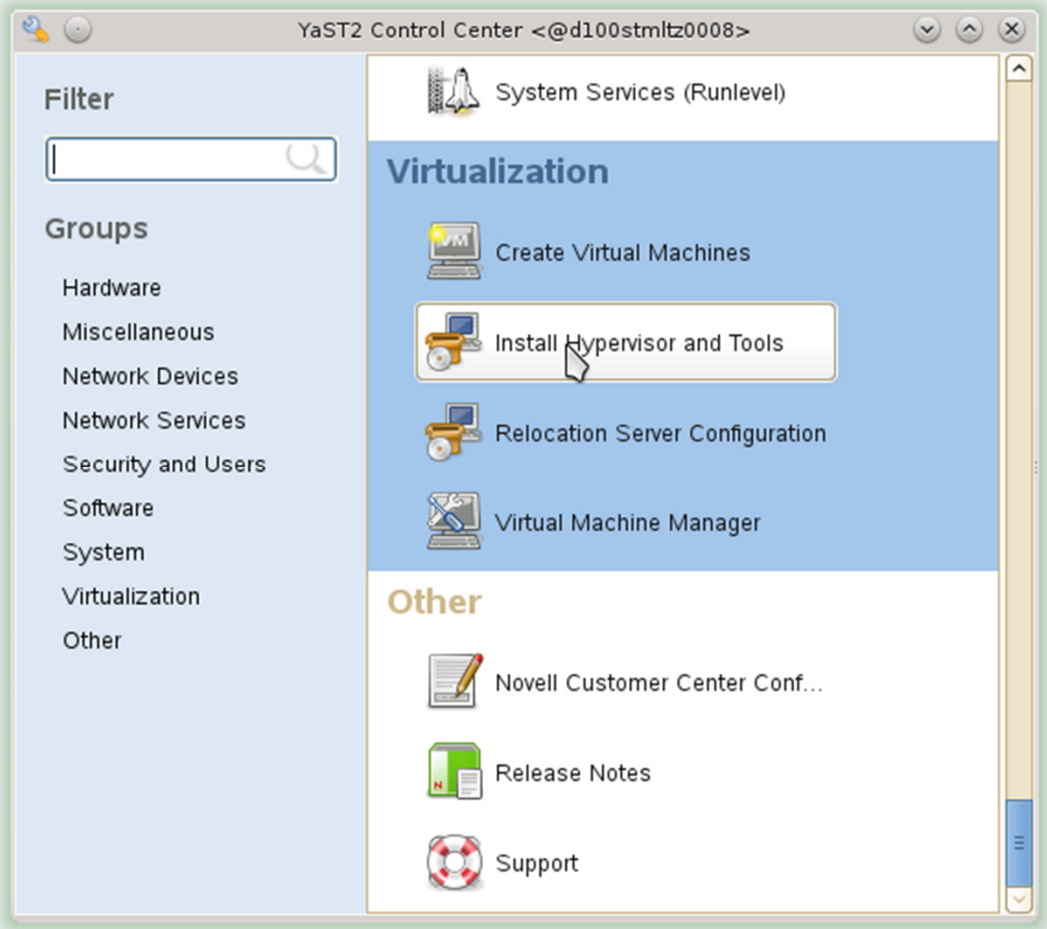

Necessary steps:

- Make sure rpm yast-vm is installed
- Make sure system is current on maintenance
- Select Virtualisation > Install Hypervisor and Tools

Prepare the KVM host system

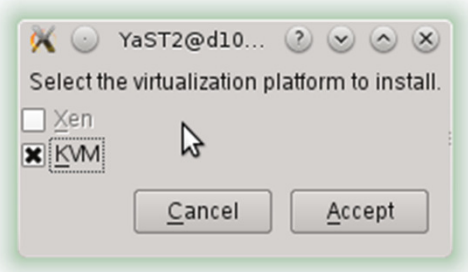

Necessary steps:

- Select KVM
- Reboot server to activate configuration changes, because switch_amode is added to kernel parmline
- kvm is mot loaded by default. Should be added to INITRD_MODULES

Prepare the KVM host system

Installation verification

1. KVM module loaded

```
root@kvmhost:~> lsmod | grep kvm
kvm                   114285  0
```

2. KVM support in QEMU enabled

```
root@kvmhost:~> qemu-system-s390x -monitor stdio -machine type=s390-
virtio,accel=kvm
VNC server running on `127.0.0.1:5900'
QEMU 1.2.0 monitor - type 'help' for more information
(qemu) info kvm
kvm support: enabled
```

3. libvirt information

```
root@kvmhost:~> virsh version
Compiled against library: libvirt 1.0.0
Using library: libvirt 1.0.0
Using API: QEMU 1.0.0
```

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Install a KVM guest system

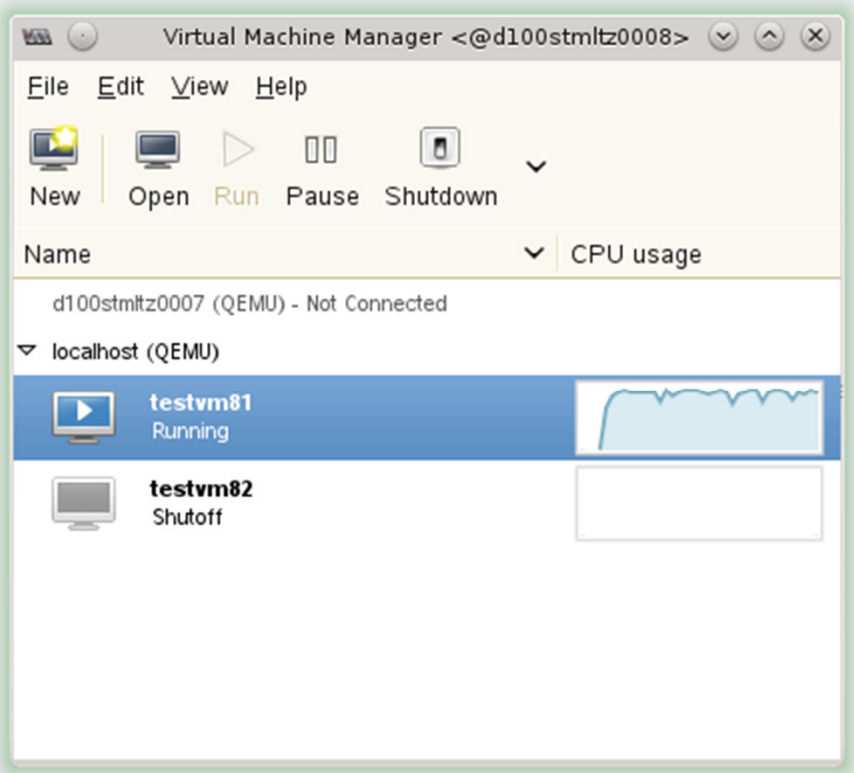

Open Virtual Machine Manager

- GUI to operate local and remote KVM guests
- Simple performance metrics
- KVM guest console interface
- Life-cycle management of KVM guests

Install a KVM guest system

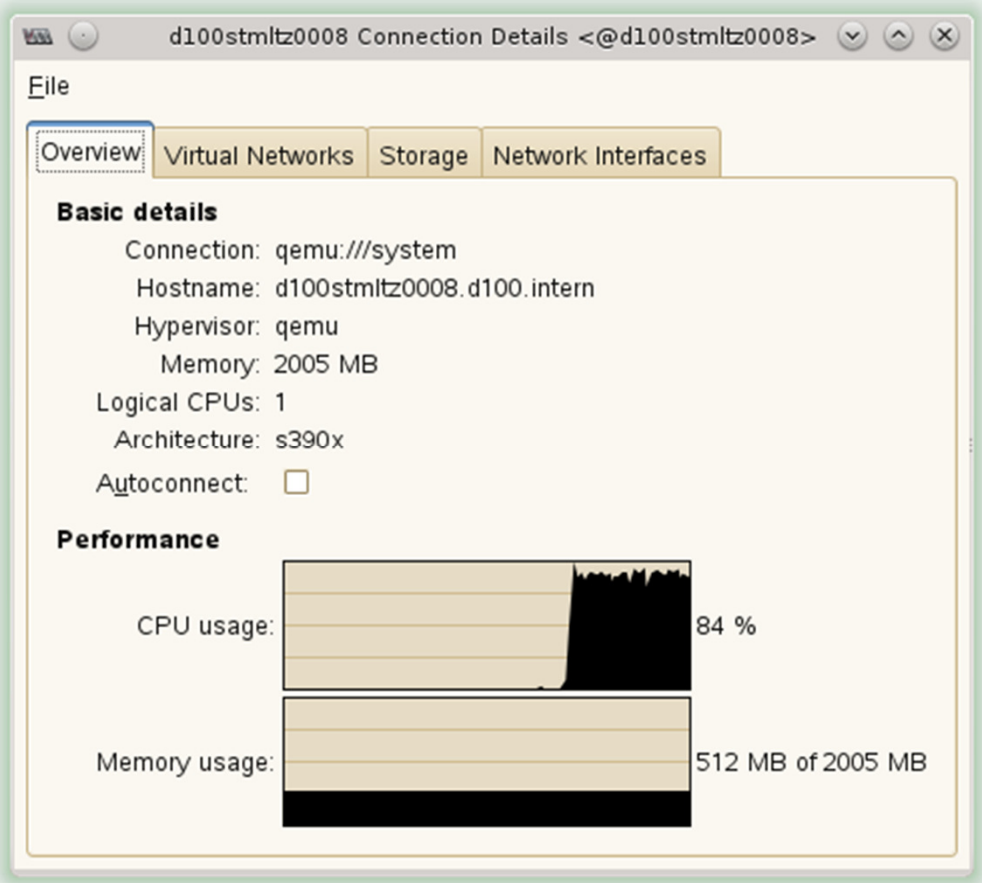

Open Connection Details

- Simple KVM host configuration interface
- Configuration options for virtual network, network interfaces and storage

Install a KVM guest system

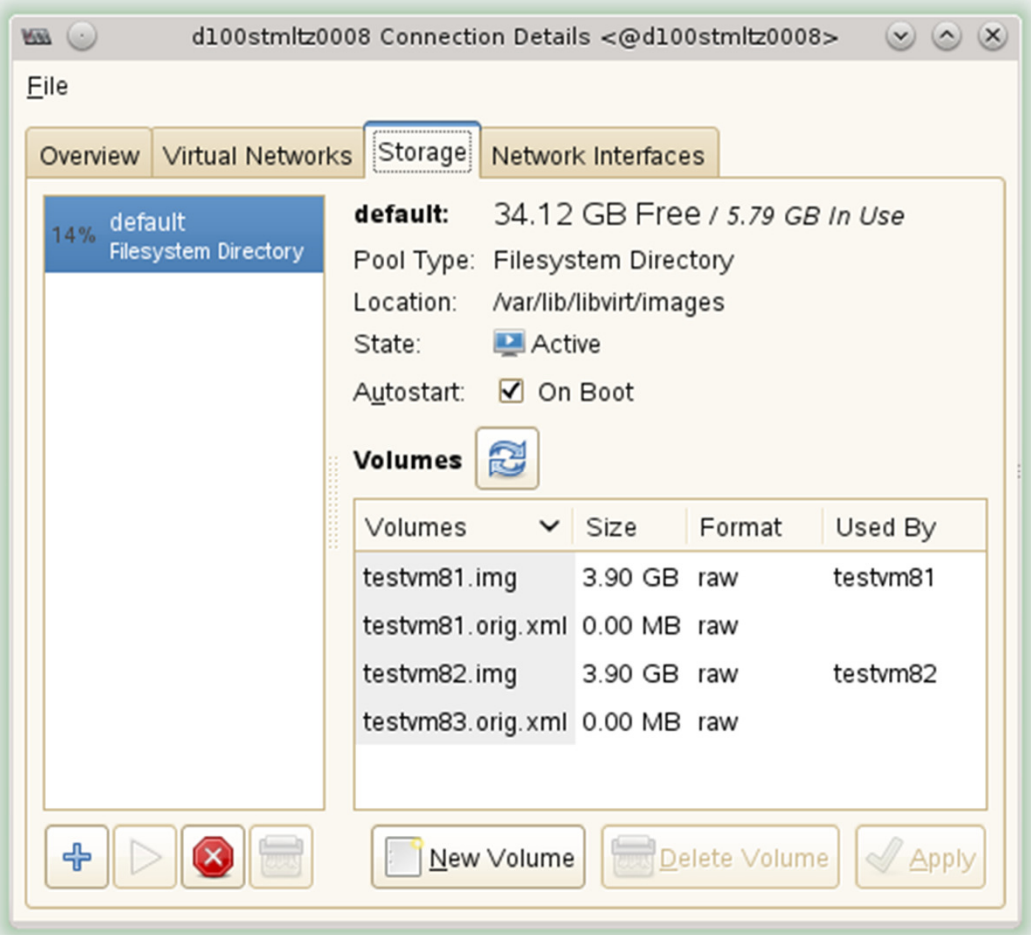

Open Storage Tab

- Storage pool overview
- Press »New volume« button to create a new dasd image

Install a KVM guest system

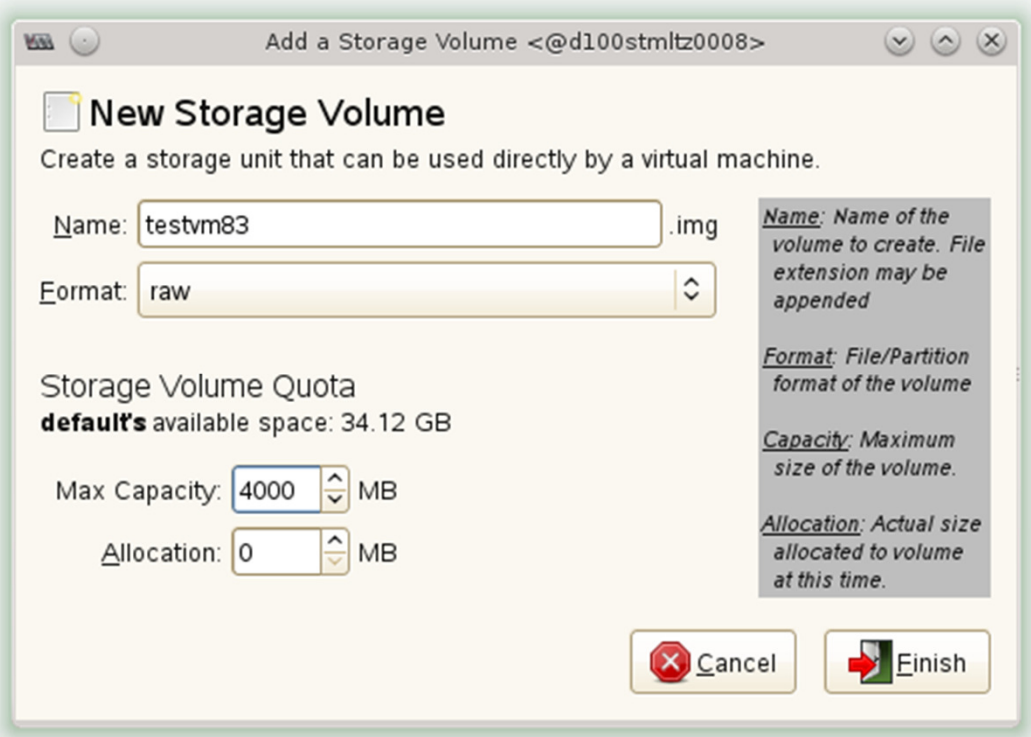

Allocate a new storage device

- Name
- Select format (raw, cow,qcow, qcow2, vmdk)
- Enter maximum capacity

Install a KVM guest system

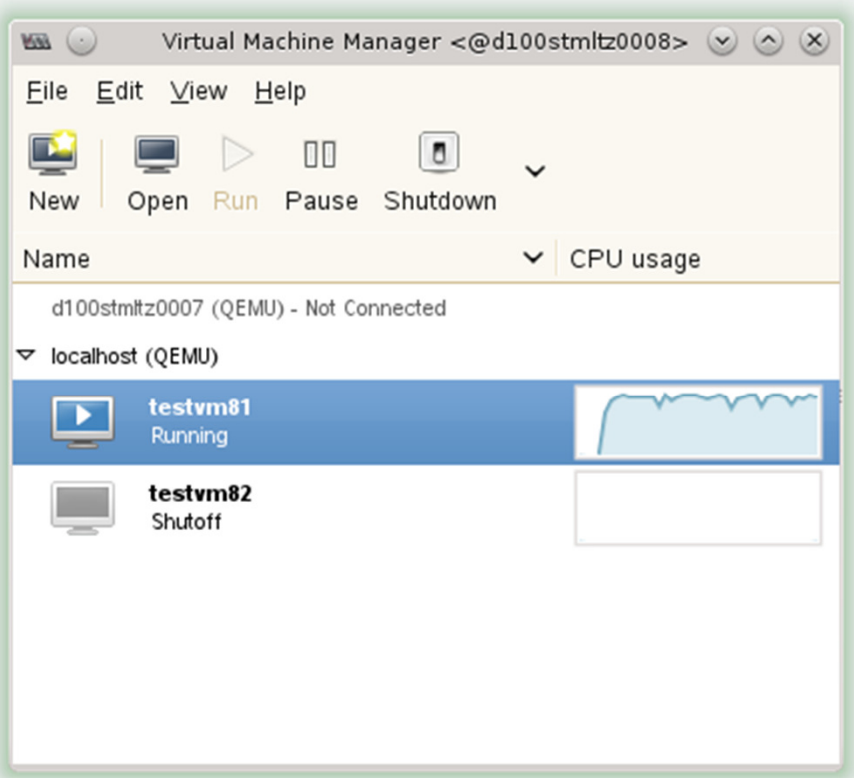

Open Virtual Machine Manager

- Select localhost
- Hit »New« button to finally create a new virtual machine

Install a KVM guest system

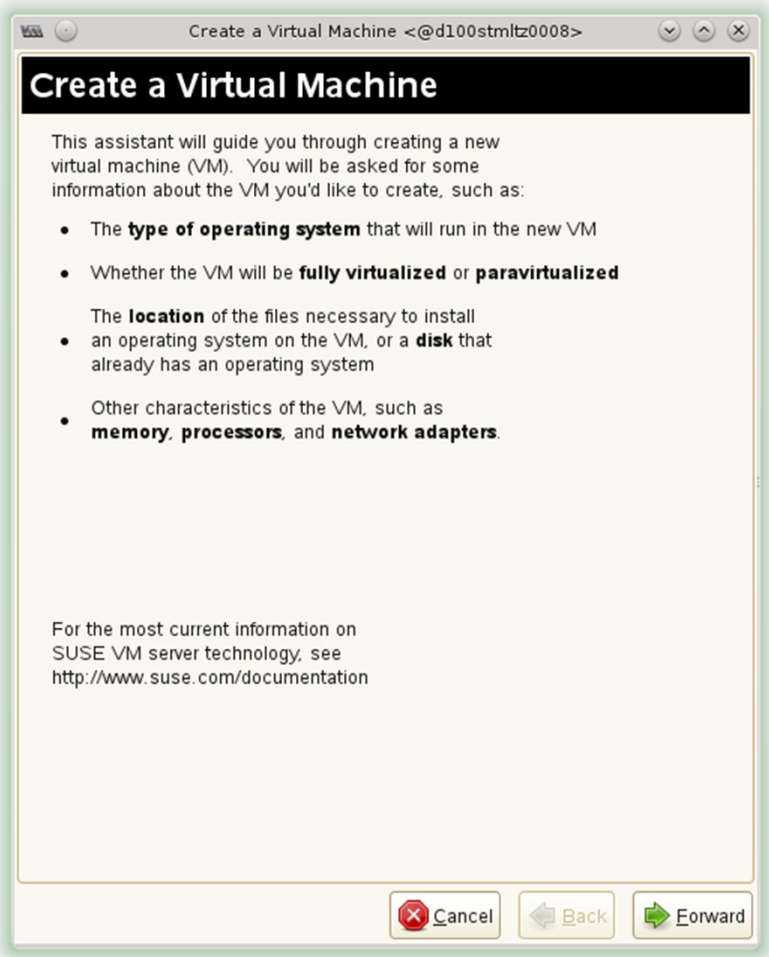

Create a Virtual Machine

- 4 Step process

Install a KVM guest system

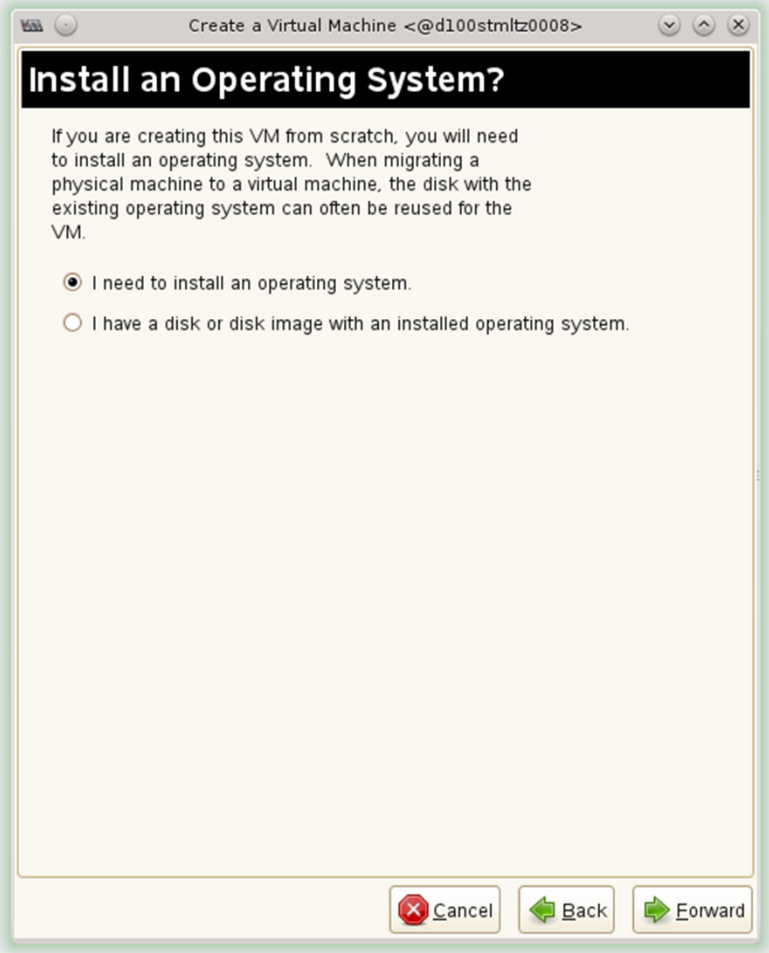

Choose wether you want to install a new operating system

- At this point you could simply use a cloned image for personalisation

Install a KVM guest system

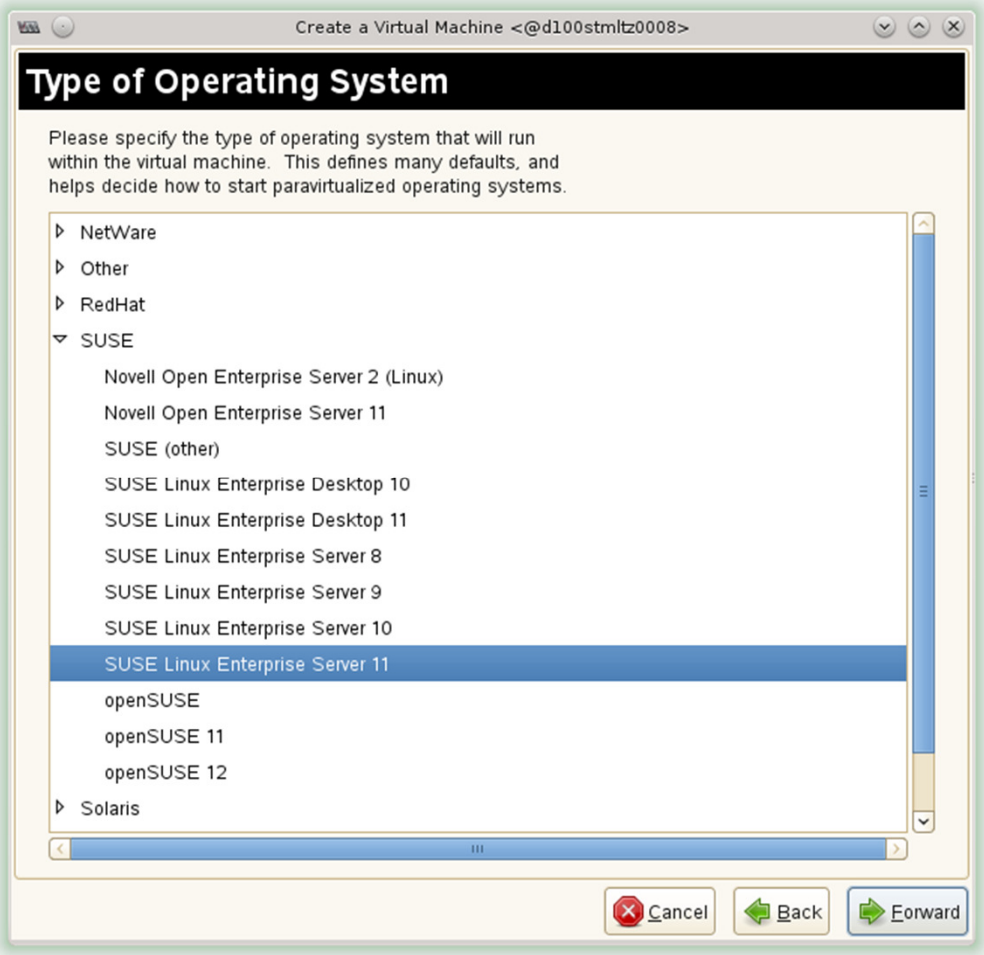

Select Operating System of your choise

- Be careful: most of the entries are x86-only

Install a KVM guest system

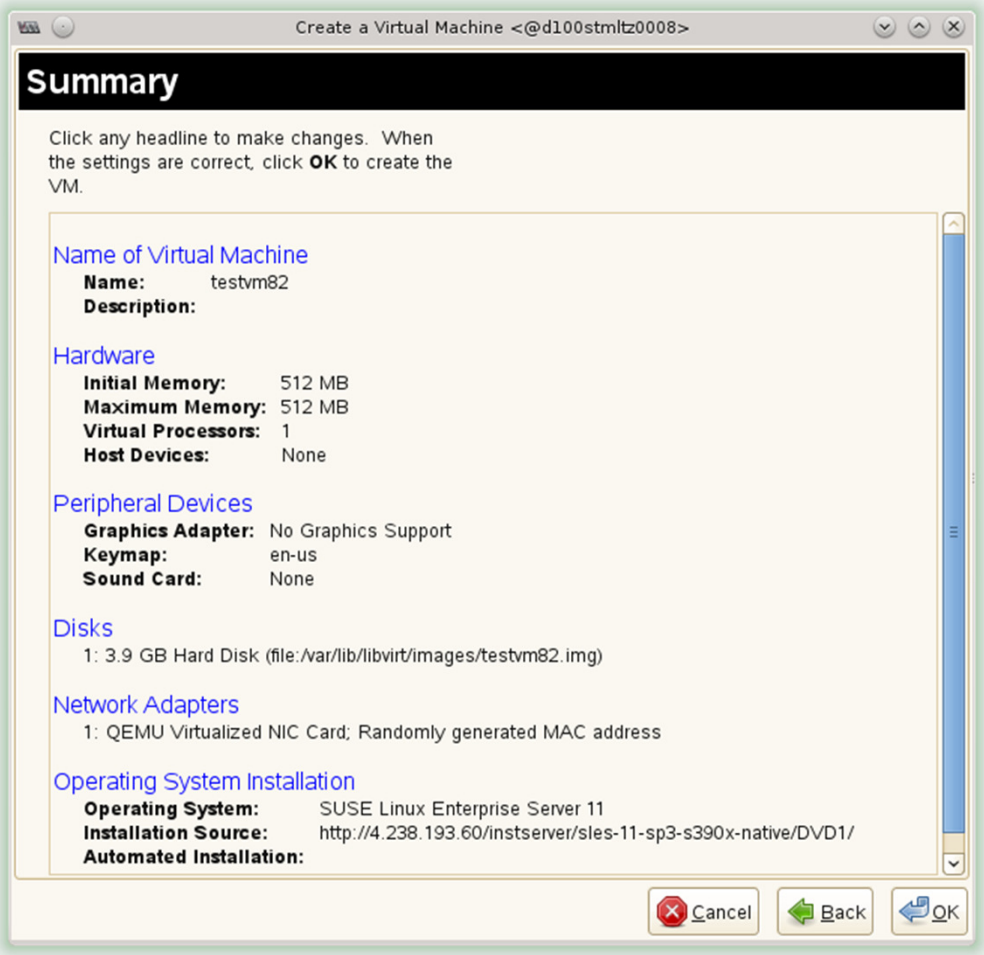

Adjust values for

- Change name
- Adjust memory configuration
- Choose the pre-allocated volume
- Specifiy MAC address to get static ip address from dhcpd

Install a KVM guest system

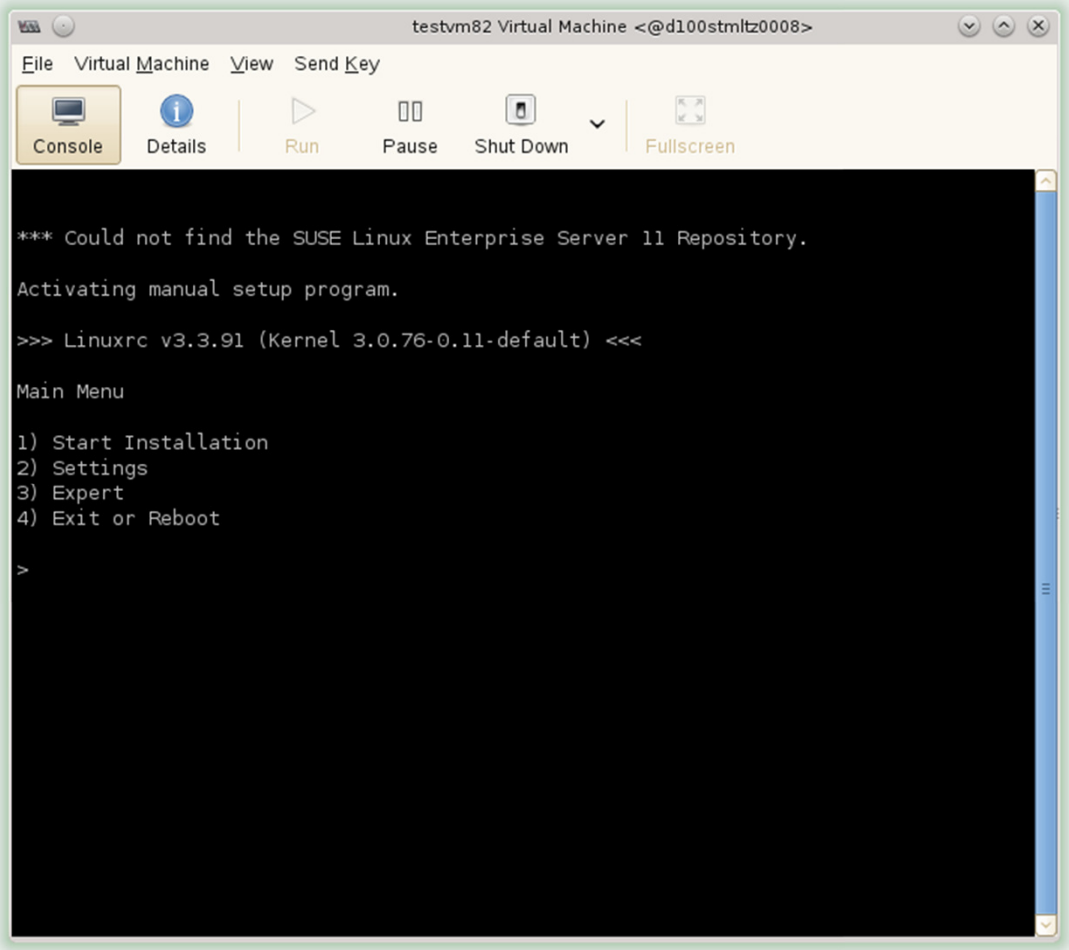

Virtual Machine Console

- At main panel choose the newly created VM and hit the »Open« button
- You will end up on the console to start the well known installation process.

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Additional tools and features
libguestfs

libguestfs is a scriptable and secure library to access images without root privileges. But it is not included in the SLES11 SP3 Tech Preview.

Libguestfs-Tools

virt-cat, virt-install, virt-rescue, virt-clone, virt-list-filesystems, virt-resize, virt-convert, virt-list-partitions, virt-tar, virt-df, virt-ls, virt-top, virt-edit, virt-make-fs, virt-viewer, virt-image, virt-manager, virt-win-reg, virt-inspector, virt-pki-validate, virt-xml-validate

```
root@kvmhost:~> virt-df -h ~/disk.img
Filesystem                                  Size    Used Available Use%
/home/tobias/disk.img:/dev/vda1           193.7M 21.6M    162.1M  12%
/home/tobias/disk.img:/dev/vg_f12x32/lv_root   5.2G  2.3G      2.6G  45%
```

Additional tools and features
Snapshot capabilities

In the KVM arena there are different ways to create a backup via snapshot depending on the configuration of host and guest

- virsh snapshot-create when using QCOW2 image
- LVM snapshots when using logical volumes as image or placing image files in a logical volume

```
# Simple example using QCOW2 and virsh snapshot-create-as
# Backup
virsh snapshot-create-as testvm82 mysnapshot1
cp /var/lib/libvirt/images/testvm82.img /backup
cp /etc/libvirt/qemu/testvm82.xml /backup
cp /var/lib/libvirt/qemu/snapshot/testvm82/mysnapshot1.xml /backup
# Restore
cp /backup/testvm82.img /var/lib/libvirt/images
virsh define /backup/testvm82.xml
virsh snapshot-create testvm82 /backup/mysnapshot1.xml --redefine
virsh snapshot-revert testvm82 mysnapshot1 --running
```

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Lessons learned

The KVM technical preview is well implemented. But there a some things you need to pay attention for:

- The virt tools are Gnome tools. So you should install a Gnome desktop instead of KDE
- You should perform an online update before installing the »KVM hypervisor and tools«. Otherwise we received errors during guest install.
- The KVM module is not added to initrd. So you should add »modprobe kvm« to /etc/init.d/boot.local or to extend initrd
- There is no security concept like PRIVCLASSES or USER for the different guests. When using libvirtd – which is running as root – all VMs will as run as root.
- In z/VM you need basic automation to start the guests after ipling the hypervisor. In libvirt there is a AUTOSTART parameter.

Lessons learned

The KVM technical preview is well implemented. But there a some things you need to pay attention for:

- Currently no unattended install. Networking interfaces are not recognized during start of the installer system to retrieve the installation media. Manual interaction is necessary

- As hypervisor and guest operating system are of the same type, there is a suite of tools (libguestfs) to get insight form the host into the guest.

- In z/VM we have a lot of unused DASD storage due to separate filesystems for /tmp, /var and /scratch. In the KVM world these are image files using thin provisioning.

- Using snapshots virtual machines can be reverted to a previous saved state. Very useful for maintenance. In case of success, just merge the delta into the snapshot concurrently.

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Captured movie

finanz informatik

Agenda

1. The IT Service Provider Finanz Informatik
2. Virtualization basics
3. What is KVM and how does it fit with QEMU and Libvirt ?
4. Prerequisites for running KVM on IBM System z (s390x)
5. Prepare the KVM host system
6. Install a KVM guest system
7. Additional tools and features
8. Lessons learned
9. Captured movie
10. Summary

Zwei Sachverhalte und ein zusammengefasstes Fazit

Pros

- Open Source alternative to z/VM for a linux-only environment.
- Easy to install
- You can use any GUI that supports libvirt to administrate a z/penguin farm.
- DASD storage overcommitment possible

Cons

- Currently no official support by any vendor or distributor
- KVM has to prove that it can handle CPU and memory constraint environments.
- Lack of performance measurement tools with mainframe scope.

Our own opinion

- Useful and cheap alternative to z/VM in a linux only environment on IBM System z.
- It will take time to get KVM and QEMU ready for a production environment, but rome wasn‘t built in a day.
- The TechPreview KVM in SLES11 SP3 is really cool stuff and worth to start playing! But libguestfs tools are still missing.

Thank you all for your attention

Thank you !!!
Karl-Heinz Doppelfeld (FI), Jeremy Koch (FI)