



science + computing

| A Bull Group Company



Virtualisierung im Cloud-Zeitalter

Spielt *der* Hypervisor noch eine Rolle?

LinuxTag 2013

Dipl.-Inform. (FH) Holger Gantikow

science + computing ag

IT-Dienstleistungen und Software für anspruchsvolle Rechnernetze

Tübingen | München | Berlin | Düsseldorf



Holger Gantikow

Dipl.-Inform. (FH)

System Engineer
science + computing ag STANDARD

72070 Tübingen
Deutschland

Persönliches

Ich suche

neue Kontakte

Ich biete

Abgeschlossene Diplomarbeit ("Virtualisierung im Kontext von Hoherfügbarkeit"), IT-Know-How, Erfahrung mit Linux, speziell Debian&Red Hat, Windows, Mac OS X, Solaris, *BSD, HP-UX, AIX, Netzwerkadministration, Netzwerktechnik, Hardware, Asterisk, VoIP-Systeme, Server Administration, Cluster Computing, Hochverfügbarkeit, Virtualisierung, HA, HPC, Autor von Fachartikeln zum Thema Cloud Computing, RHCT, RHSA, Python Programmierung, Neugierde, Flexibilität

Interessen

IT-spezifisch momentan: Virtualisierung (Xen, ESX, ESXi, KVM), Cluster Computing (HPC, HA), Cloud Computing (speziell: IaaS, HPCaaS), OpenNebula, OpenSolaris, ZFS, XMPP, SunRay ThinClients - ansonsten: Freie Software, Musik, Gitarre, Fotografie

https://www.xing.com/profile/Holger_Gantikow

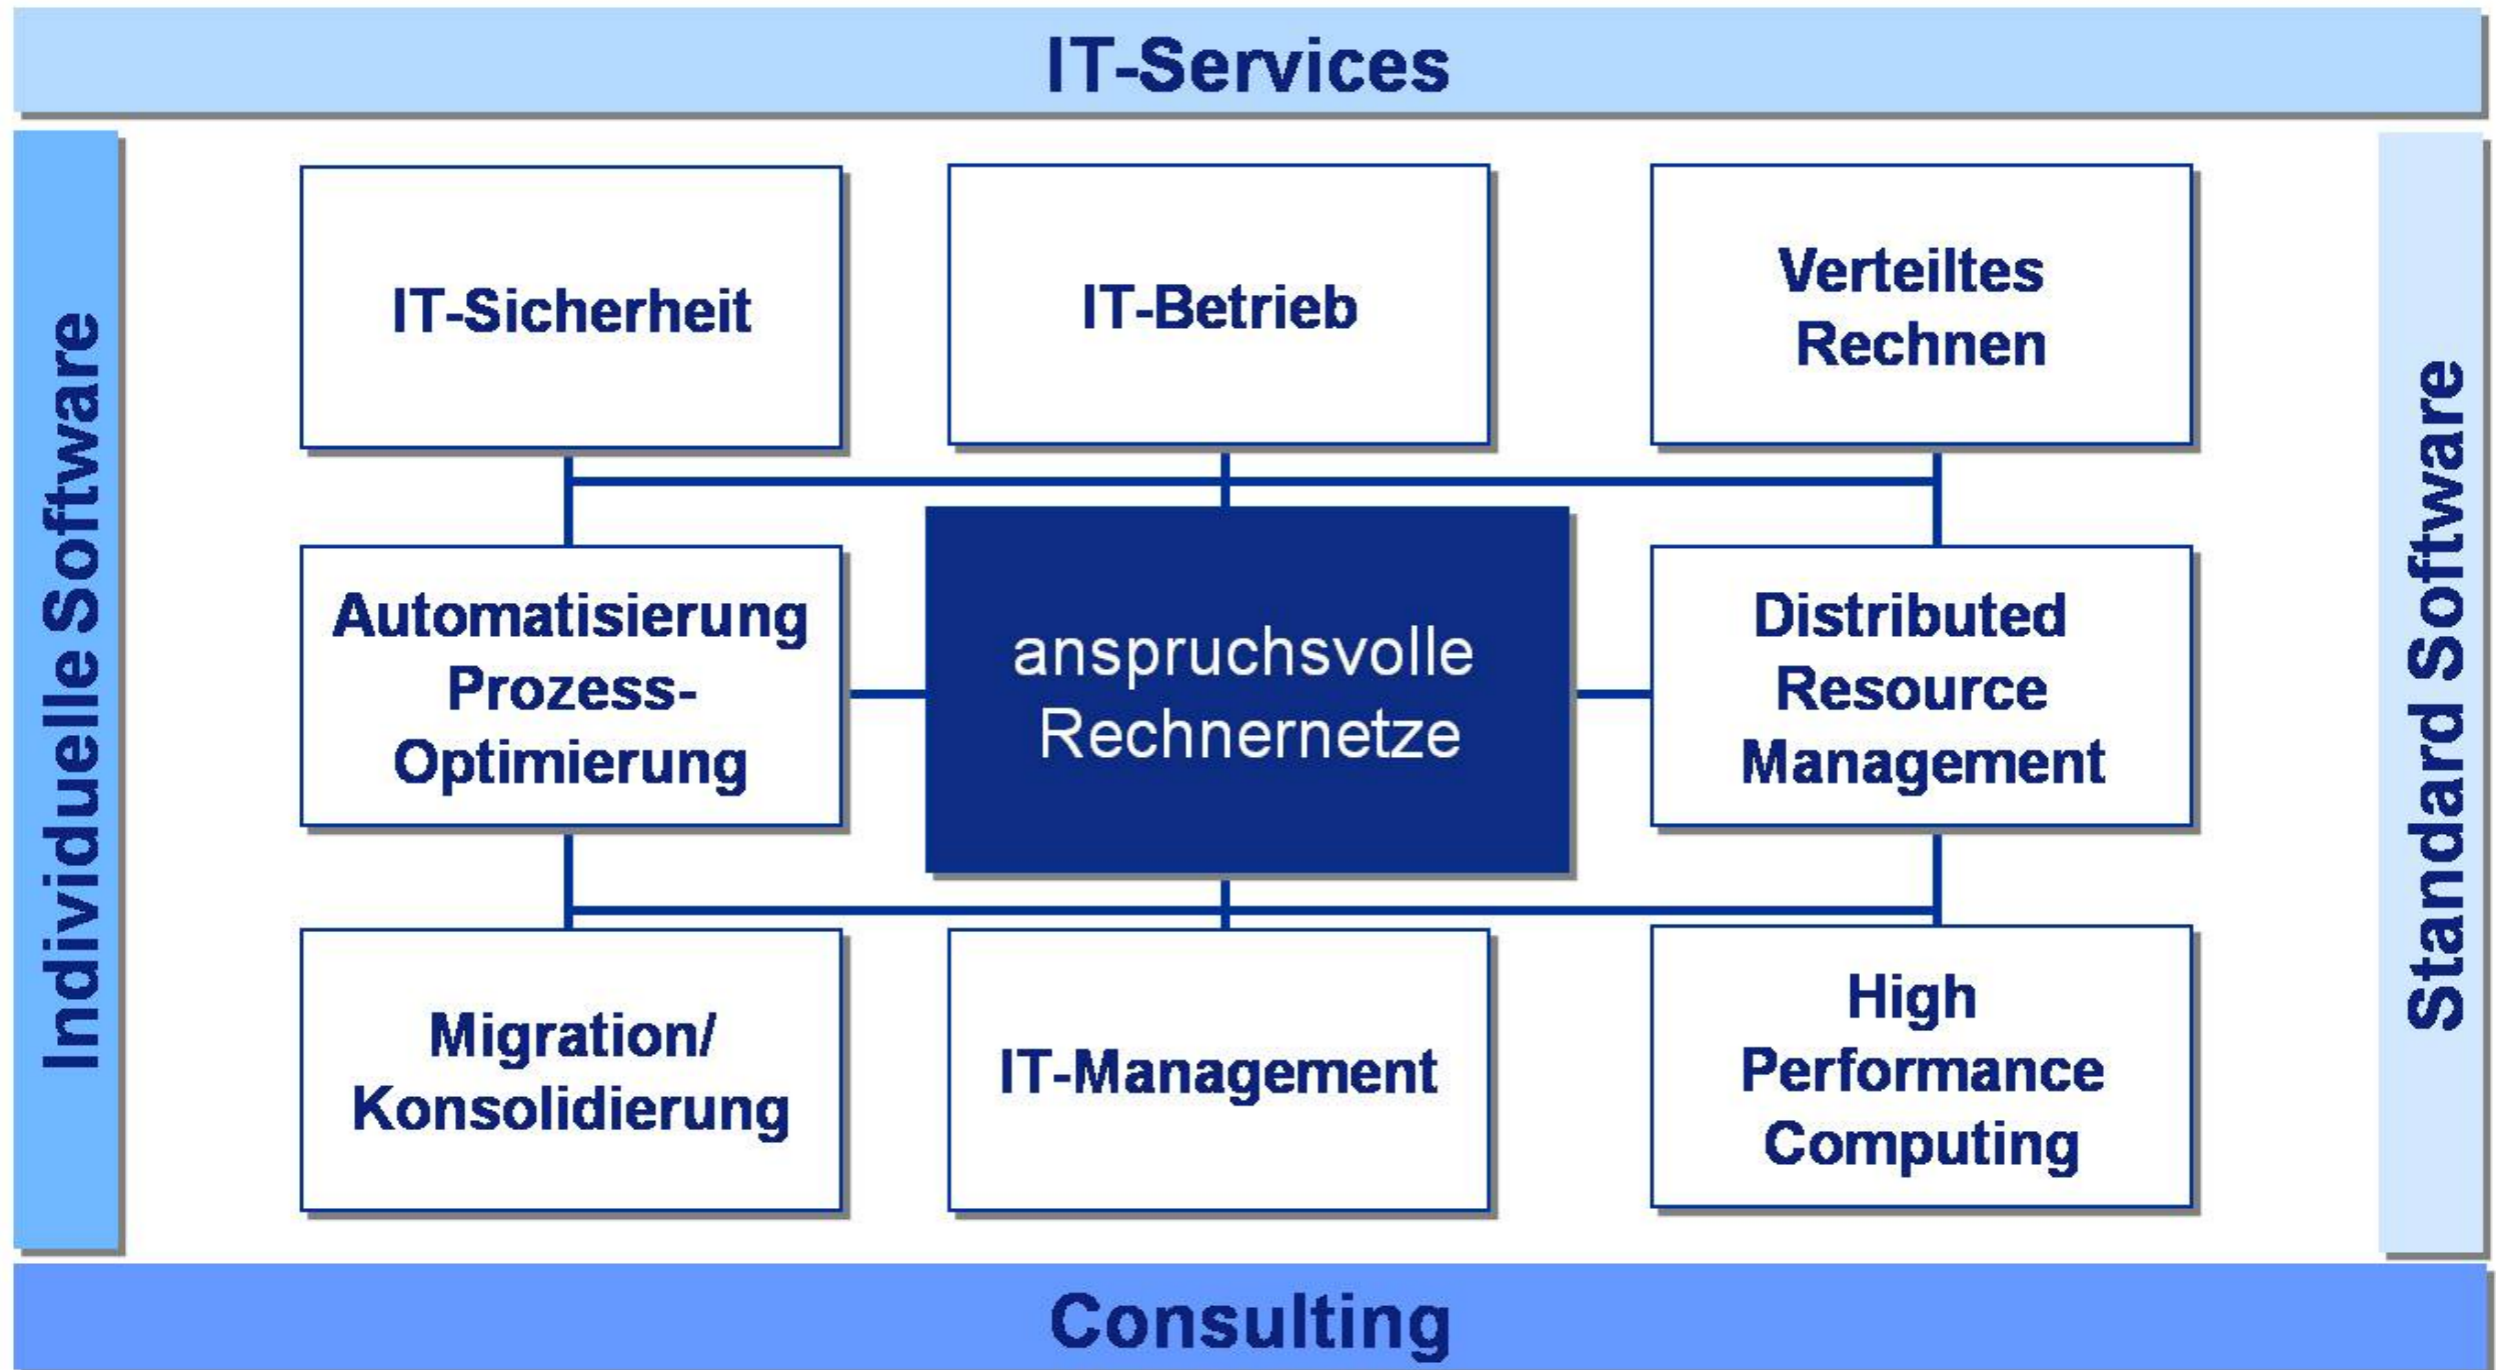
science+computing ag

Unser Fokus:
IT-Dienstleistungen und Software für
technische Berechnungsumgebungen

Gegründet	1989
Büros	Tübingen München Berlin Düsseldorf
Mitarbeiter	300+
Besitzer	Bull S.A. (100%) seit 10/2008
Jahresumsatz (10/11)	26,66 Mio. Euro



Kernkompetenzen



Teil I: Virtualisierung mit Linux

Teil II: *Relevanz-Diskussion*

Teil III: Benchmark

Teil 0: Einleitung

Wie war das noch *gestern*...?

Gartner's Top 10

Strategic Technologies for 2009

1. **Virtualization**
2. **Business Intelligence**
3. **Cloud Computing**
4. **Green IT**
5. **Unified Communications**
6. **Social Software and Social Networking**
7. **Web Oriented Architecture**
8. **Enterprise Mashups**
9. **Specialized Systems**
10. **Servers – Beyond Blades**

Quelle:

http://blogs.gartner.com/david_cearley/2008/10/14/gartner's-top-10-strategic-technologies-for-2009/

Gartner's Top 10

Strategic Technologies for 2010

1. **Cloud Computing**
2. **Advanced Analytics**
3. **Client Computing**
4. **IT for Green**
5. **Reshaping the Datacenter**
6. **Social Computing**
7. **Security - Activity Monitoring**
8. **Flash Memory**
9. **Virtualization for Availability**
10. **Mobile Applications**

Quelle:

<http://gartner.com/it/page.jsp?id=1210613>

14.10.2008

Und *heute*...?

Gartner's Top 10


Strategic Technologies for 2013

1. Mobile Device Battles
2. Mobile Applications and HTML5
3. Personal Cloud
4. Enterprise App Stores
5. The Internet of Things
6. Hybrid IT and Cloud Computing
7. Strategic Big Data
8. Actionable Analytics
9. In Memory Computing
10. Integrated Ecosystems

Quelle:

<http://www.gartner.com/newsroom/id/2209615>




23.10.2013



If you can touch it and you can see it, it's **REAL**.
If you can touch it but you can't see it, it's **TRANSPARENT**.
If you can't touch it but you can see it, it's **VIRTUAL**.
If you can't touch it and you can't see it, it's **GONE**.
(Unbekannte Quelle)

Warum brauchen wir *Virtualisierung* nochmal?

Übergang zur Virtualisierung



Geschichtliche Entwicklung

Hier noch drauf hinweisen was eine VM eigentlich ist

<http://www.vmware.com/de/technology/virtual-machine.html>

60er und 70er

- IBM CP-40, System/370 (Mainframe)

80er

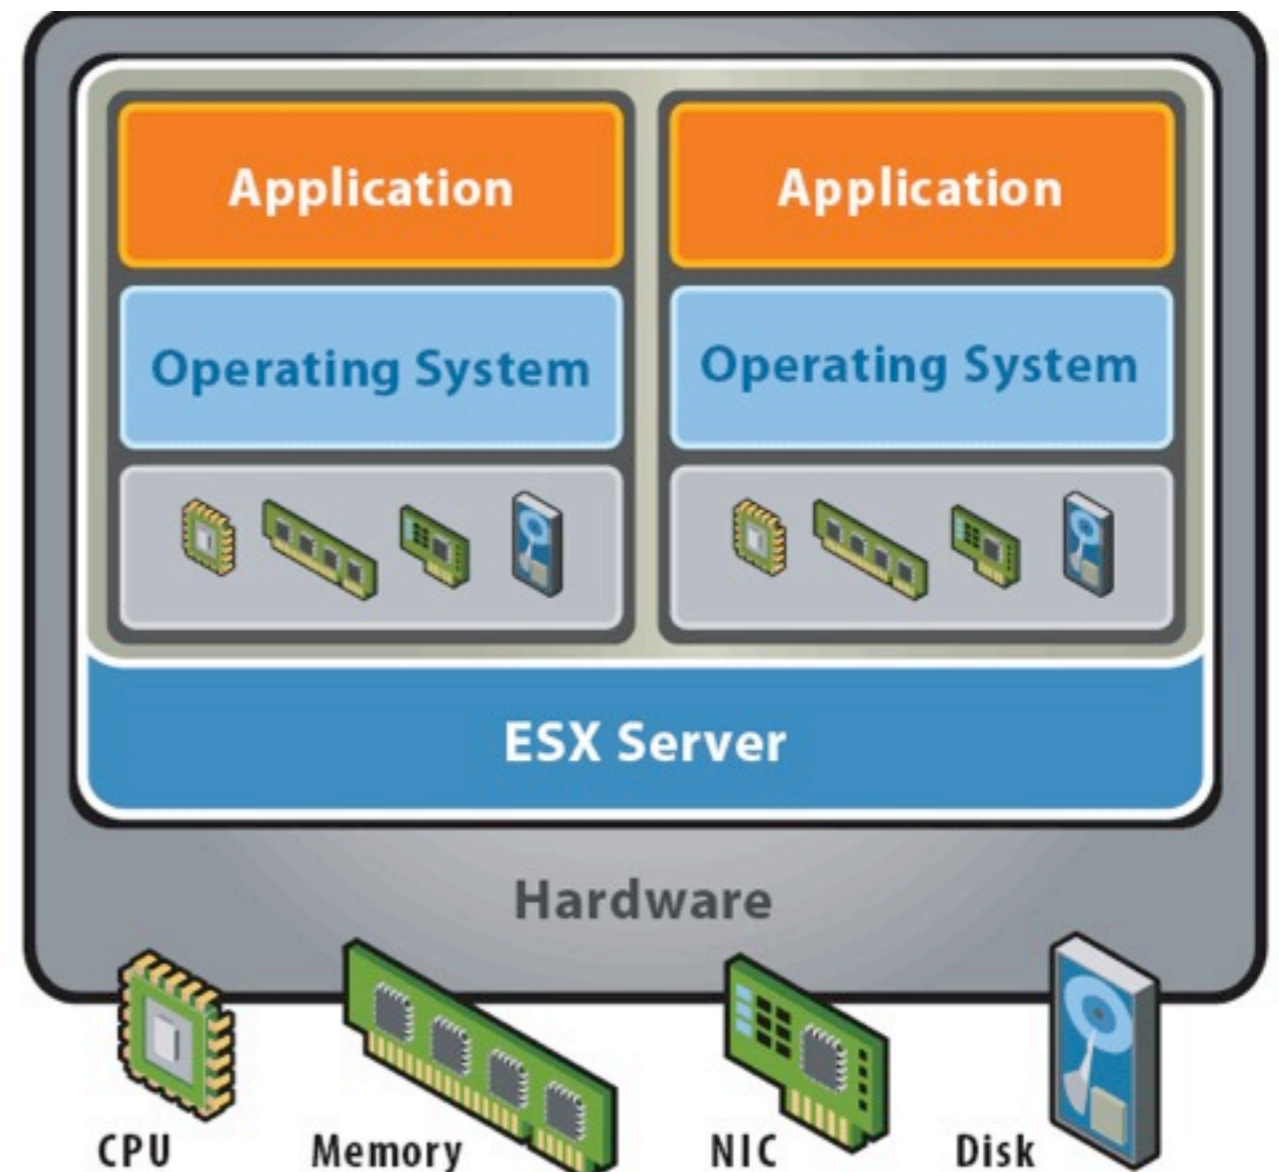
- Mach Microkernel
- Logical Partitioning (LPAR)

90er

- VMware: erstes Produkt für x86 ('99)

2000 - Gegenwart

- 2003: Xen
- 2007: Kernel-based Virtual Machine



Quelle: VMware Dokumentation

Teil I:

Virtualisierung mit Linux



**Eine Welt voller
Alternativen....**

Xen - Totgesagte...





Xen is now a Collaborative Project of The Linux Foundation

Laying the foundation and providing the infrastructure for the next evolutions in Xen's growth, adoption, and innovation.

Home Products Support Community Blog

Check out our new community site at xenproject.org!

What is the Xen Project?

The Xen Project™ – hosted by Xen.org – is the home for several virtualization related open source projects and XEN PROJECT is the new name for the XEN open source community. The Xen Project™ develops virtualization technologies powering the world's largest clouds in production and is the foundation of many commercial products. Technology developed by the Xen Project™ powers public clouds such as *Amazon Web Services, Rackspace Public Cloud* and many others. Examples of Xen based server products include *Huawei UVP, Oracle VM and XenServer*. Examples of client products and appliances include *QubesOS, XenClient and Netscaler*. The Xen Project™ is 10 years old, and the technology it has created is mature and its stability and versatility is second to none.

Xen Projects

The Xen Project™ has created a number of open source technologies for specific purposes:

- The **Xen® Hypervisor** is the open source standard for hardware virtualization and the foundation for all solutions hosted on Xen.org. — [Learn More](#)
- The **Xen Cloud Platform** is an enterprise-ready virtualization platform based on the Xen® Hypervisor, adding further functionality and many integrations for cloud, storage and networking solutions. — [Learn More](#)
- **Xen ARM** is a Xen based Hypervisor that targets embedded and mobile devices on the ARM architecture. — [Learn More](#)

If you are unsure, which Xen project is best for you, you may want to read the [Xen Overview](#) document which introduces the main components and some of the key choices you will have to make.

Open Source

Quelle:
<http://www.xen.org/>

```
[root@centos64-host ~]# cat /boot/grub/grub.conf
[...]
```

```
title Xen with CentOS (3.9.2-1.el6xen.x86_64)
    root (hd0,0)
    kernel /xen.gz dom0_mem=1024M
    module /vmlinuz-3.9.2-1.el6xen.x86_64 [...]
    module /initramfs-3.9.2-1.el6xen.x86_64.img
```



```
title CentOS (2.6.32-358.6.1.el6.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 [...]
    initrd /initramfs-2.6.32-358.6.1.el6.x86_64.img
```

Status - KVM

Status - KVM



Go Search

MAIN PAGE

STATUS

GUEST SUPPORT STATUS

KVM-AUTOTEST

HOWTO

MIGRATION

LISTS, IRC

DOCUMENTS

DOWNLOADS

BUGS

CODE

TODO

FAQ



Status

KVM is included in the mainline linux kernel since 2.6.20 and is stable and fast for most workloads.

The userspace tools are part of mainline QEMU since 1.3.

It is also available as a patch for recent Linux kernel versions and as an external module that can be used with your favorite distro- provided kernel going back up to 2.6.16, therefore including all latest versions for Enterprise Linux Distributions.

Working:

- Intel-based hosts (requires VT capable processors)
- AMD-based hosts (requires SVM capable processors)
- Windows/Linux/Unix guests (32-bit and 64-bit)
- SMP hosts
- SMP guests (as of kvm-61, max 16 cpu supported)
- Live [Migration](#) of guests from one host to another (32-bit and 64-bit)
- See the [Guest Support Status](#) page for a list of guest operating systems known to work
- See the [Host Support Status](#) page for information on host hardware.
- Guest swapping
- [Paravirtualized networking](#)
- [Paravirtualized block device](#)
- [PCI-Express passthrough](#)

In progress:

- [PowerPC port](#)
- IA64 port
- xenner (<http://kraxel.fedorapeople.org/xenner>), a project to run x86 xen guest (domU) kernels
- [ARM port](#)
- [VGA device assignment](#)

Related

- [FLOSS weekly KVM interview by Avi Kivity & D...](#)

Contents [hide]

- 1 Status
 - 1.1 Working:
 - 1.2 In progress:
 - 1.3 Related

Quelle:
<http://www.linux-kvm.org/page/Status>



VirtualBox

[Login](#) [Preferences](#)

Welcome to VirtualBox.org!

VirtualBox is a powerful x86 and AMD64/Intel64 [virtualization](#) product for enterprise as well as home use. Not only is VirtualBox an extremely feature rich, high performance product for enterprise customers, it is also the only professional solution that is freely available as Open Source Software under the terms of the GNU General Public License (GPL) version 2. See "[About VirtualBox](#)" for an introduction.

Presently, VirtualBox runs on Windows, Linux, Macintosh, and Solaris hosts and supports a large number of [guest operating systems](#) including but not limited to Windows (NT 4.0, 2000, XP, Server 2003, Vista, Windows 7), DOS/Windows 3.x, Linux (2.4 and 2.6), Solaris and OpenSolaris, OS/2, and OpenBSD.

VirtualBox is being actively developed with frequent releases and has an ever growing list of features, supported guest operating systems and platforms it runs on. VirtualBox is a community effort backed by a dedicated company: everyone is encouraged to contribute while Oracle ensures the product always meets professional quality criteria.

Hot picks:

- Pre-built virtual machines for developers over at [Oracle Tech Network](#)
- **phpVirtualBox** AJAX web interface [project site](#)
- **IQEmu** automated Windows VM creation, application integration [project site](#)

News Flash

- **New** **April 12th, 2013**
VirtualBox 4.2.12 released!
Oracle today released VirtualBox 4.2.12, a maintenance release of VirtualBox 4.2 which improves stability and fixes regressions. See the [ChangeLog](#) for details. VirtualBox 4.1.26, a maintenance release of VirtualBox 4.1 is available as well. See the [Changelog](#) for details.
- **New** **September 13th, 2012**
VirtualBox 4.2 released!
Read the official [press release](#) for more details.
- **Attention** **January 20th, 2012**
This site has switched to Oracle single-sign-on authentication.
Get your account at [myprofile.oracle.com](#) and on first login on this site associate it with a new or existing nick name.

[More information...](#)
[Contact](#) - [Privacy policy](#) - [Terms of Use](#)

Quelle:
<https://www.virtualbox.org/>

VMware ESXi 5.1.0 (VMKernel Release Build 1065491)

innotek GmbH VirtualBox

Intel(R) Core(TM) i5-3427U CPU @ 1.80GHz
2 GiB Memory

VMware vSphere Hypervisor (ESXi)

Loading module tcpip3 ...

Quelle:
VMware ESXi Installer

Teil II:
Relevanz-Diskussion?

Spielt der *Hypervisor* noch eine Rolle?

Äussere Faktoren

Tools!

[Home](#)[News](#)[Applications](#)[Downloads](#)[Documentation](#)[Wiki](#)[FAQ](#)[Bug reports](#)[Contact](#)[Test suites](#)[Related Links](#)[Sitemap](#)

The virtualization API

libvirt is:

- A toolkit to interact with the virtualization capabilities of recent versions of Linux (and other OSes), see our [project goals](#) for details.
- Free software available under the [GNU Lesser General Public License](#).
- A long term stable C API
- A set of bindings for common languages
- A [CIM provider](#) for the DMTF virtualization schema
- A [QMF agent](#) for the AMQP/QPid messaging system

libvirt supports:

- The [KVM/QEMU](#) Linux hypervisor
- The [Xen](#) hypervisor on Linux and Solaris hosts.
- The [LXC](#) Linux container system
- The [OpenVZ](#) Linux container system
- The [User Mode Linux](#) paravirtualized kernel
- The [VirtualBox](#) hypervisor
- The [VMware ESX and GSX](#) hypervisors
- The [VMware Workstation and Player](#) hypervisors
- The [Microsoft Hyper-V](#) hypervisor
- The [IBM PowerVM](#) hypervisor
- The [Parallels](#) hypervisor
- Virtual networks using bridging, NAT, VEPA and VN-LINK.
- Storage on IDE/SCSI/USB disks, FibreChannel, LVM, iSCSI, NFS and filesystems

libvirt provides:

- Remote management using TLS encryption and x509 certificates
- Remote management authenticating with Kerberos and SASL
- Local access control using PolicyKit
- Zero-conf discovery using Avahi multicast-DNS
- Management of virtual machines, virtual networks and storage

Quelle:
<http://libvirt.org/>

```
holgrrr — root@lt-centos:~ — ssh — 80x24
root@lt-centos:~
VIRSH(1)          Virtualization Support          VIRSH(1)

NAME
  virsh - management user interface

SYNOPSIS
  virsh [OPTION]... [COMMAND STRING]

  virsh [OPTION]... COMMAND [ARG]...

DESCRIPTION
  The virsh program is the main interface for managing virsh guest domains. The program can be used to create, pause, and shutdown domains. It can also be used to list current domains. Libvirt is a C toolkit to interact with the virtualization capabilities of recent versions of Linux (and other OSes). It is free software available under the GNU Lesser General Public License. Virtualization of the Linux Operating System means the ability to run multiple instances of Operating Systems concurrently on a single hardware system where the basic resources are driven by a Linux instance. The library aims at providing a long term stable C API. It currently supports Xen, QEmu, KVM, LXC, OpenVZ, VirtualBox and VMware ESX.
```

Quelle:
man virsh ;)

```
ssh
root 6042 1.9 0.3 924728 12476 ? Ssl 00:40 0:15 /usr
rtld --listen
qemu 6146 6.3 9.1 3392144 365476 ? Sl 00:40 0:48 /usr
vm -name opensuse12-kvm-64bit-SA -S -M pc-1.3 -enable-kvm -m 2048 -sm
=1,cores=1,threads=1 -uuid 41a5aef0-e337-a060-b980-a77001ccca84 -no-u
-ndefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu
-kvm-64bit-SA.monitor,server,nowait -mon chardev=charmonitor,id=moni
trol -rtc base=utc -no-shutdown -device piix3-usb-uhci,id=usb,bus=pci
.0x2 -drive file=/var/lib/kvm/images/opensuse12-kvm-64bit-SA/disk0.ra
d=drive-virtio-disk0,format=raw -device virtio-blk-pci,scsi=off,bus=p
x5,drive=drive-virtio-disk0,id=virtio-disk0,bootindex=1 -drive file=/
E-12.3-NET-x86_64.iso,if=none,id=drive-ide0-0-0,readonly=on,format=ra
de-cd,bus=ide.0,unit=0,drive=drive-ide0-0-0,id=ide0-0-0 -netdev user,
-device virtio-net-pci,netdev=hostnet0,id=net0,mac=52:54:00:34:89:4e
addr=0x3 -vnc 127.0.0.1:0 -vga cirrus -device AC97,id=sound0,bus=pci.
-device virtio-balloon-pci,id=balloon0,bus=pci.0,addr=0x6
root 7143 4.2 1.3 637372 54776 ? Ss 00:52 0:03 pyth
re/virt-manager/virt-manager.py
root 7226 0.0 0.0 10372 952 pts/2 S+ 00:53 0:00 grep
to virt
opensuse123-host:~ # lsmod | grep kvm
kvm_intel 136915 3
kvm 475491 1 kvm_intel
opensuse123-host:~ #
```

```
Handle 0x2000, DMI type 32, 11 bytes
System Boot Information
Status: No errors detected
Handle 0x7F00, DMI type 127, 4 bytes
End Of Table
opensuse123-kvm:~ # cat /proc/cpuinfo
processor : 0
vendor_id : GenuineIntel
cpu family : 6
model : 2
model name : QEMU Virtual CPU version 1.3.0
stepping : 3
microcode : 0x1
cpu MHz : 1861.999
cache size : 4096 KB
cpu : yes
fpu_exception : yes
cpuid level : 4
wp : yes
flags : fpu de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pse36 clflush mmx f
_lm
bogomips : 3723.99
clflush size : 64
cache_alignment : 64
address sizes : 40 bits physical, 48 bits virtual
power management:
opensuse123-kvm:~ #
```

Virtual Machine Manager

Datei Bearbeiten Anzeigen Hilfe

Neu Öffnen Ausführen Pause Herunterfahren

Name	CPU-Verbr
localhost (QEMU)	
opensuse12 Ausgeschaltet	
opensuse12-1 Ausgeschaltet	
opensuse12-kvm-64bit-SA Wird ausgeführt	
opensuse12-kvm-x86_64 Ausgeschaltet	
localhost (xen) - Nicht verbunden	

```
opensuse123-kvm:~ # cat /proc/cpuinfo
processor : 0
vendor_id : GenuineIntel
cpu family : 6
model : 2
model name : QEMU Virtual CPU version 1.3.0
stepping : 3
microcode : 0x1
cpu MHz : 1861.999
cache size : 4096 KB
cpu : yes
fpu_exception : yes
cpuid level : 4
wp : yes
flags : fpu de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pse36 clflush mmx f
_lm
bogomips : 3723.99
clflush size : 64
cache_alignment : 64
address sizes : 40 bits physical, 48 bits virtual
power management:
opensuse123-kvm:~ #
```

virt-manager



ONE 4.0 is out!!

Enterprise Cloud and Datacenter Virtualization

An enterprise-ready open-source platform to manage cloud data centers

Open

Fully open-source (not open-core) software released under the Apache license

Adaptable

Highly customizable to fit into your data center and leverage existing IT investments

Proven

Rigorously tested, mature, widely used, and many massive scale deployments

Powerful

Innovative functionality for enterprise clouds and data center virtualization

Interoperable

Choice of interfaces, from open cloud, like OCCI, to de-facto standards, like AWS

No Lock-in

Broad infrastructure support and platform independent on VMware, KVM and Xen

Very Light

Highly efficient, easy to install and update, with packages for the main Linux distributions

Enterprise-ready

Single installing and updating process, and community and commercial support

Why OpenNebula? ▶



Our Users are the Industry & Research L

Some of our featured users building the most innovative and largest clouds in the

Quelle:
<http://opennebula.org/>



FOREMAN



oVirt

Open**Nebula**.org

PROXMOX

Undundundundundund?



EUCALYPTUS



Tipp:
http://www.linux-kvm.org/page/Management_Tools

Deshalb...

**Spielt der *Hypervisor* noch eine
Rolle?**

Teil III: Benchmark

Motivation

openSoftware-Pakete

```
a2ps acroread ant apache2 argouml autoconf automake bash bison bonnie++  
cpan-5.10.1 cpan-5.8.4 cpio curl cvs ddd dia diffutils ding eclipse emacs  
findutils firefox flex fping freemind gawk gcc-3.4.6 gcc-4.1.1 gcc-4.3.2  
gcc-4.3.3 gcc-4.4.0 gcc-4.4.1 gdb ggrep ghostscript ghostview gm4 gmake  
gnupg gnuplot gqview grace gsed gtar gv gzip hpijs hping imagemagick java  
jedit jxplorer keepass kerberos less lsof lynx lyx mc mpeg_play mtools  
mutt ncd u nedit netbeans netcat netpbm ntp nttcp octave openSoftwareLib  
openoffice openssh patch perl-5.10.1 perl-5.8.4 pidgin pine postgresql  
procmail pstree psutils pvm python-2.6 python qt rcs rdesktop recode  
recordmydesktop rsync samba screen sox subversion sudo tcltk tcsh teTeX  
texinfo tgif thunderbird tkcvs tkdiff tomcat valgrind vim vnc wget x11vnc  
x3270 xemacs xev xfce4 xfig xftp xpdf xv Xaw3d Xvfb apache berkeleydb  
coreutils ctags cvs-local cyrus-sasl dejagnu dmalloc doxygen dvd+rw-tools  
expat f2c fftw freetype gcc-4.2.1 gdbm gettext ghostpcl gimp2 gmp gnuplot  
gpa gridengine groff html2ps iperf irfanview jpilot kerberos lapack libIDL  
libart libbzip2 libgd libiconv libjpeg libpng libungif libxml2 lpd monitor  
mpfr mysql ncftp ncurses netscape nmon nvi nx openldap openscenegraph  
openssl pcre pdflib perl-ldap perl-spezial pgp php pilot-link plan pts pv  
python qt rcs readline rootsh samba-2.2 samba sgzip smartmontools squid  
tcpdump test tiff tkX-7.3.1 top xaff xemacs-large xforms xmon xscope  
xspread zip zlib
```

Drum prüfe wer sich bindet...

openSoftware-Pakete

```
a2ps acroread ant apache2 argouml autoconf automake bash bison bonnie++  
cpan-5.10.1 cpan-5.8.4 cpio curl cvs ddd dia diffutils ding eclipse emacs  
findutils firefox flex fping freemind gawk gcc-3.4.6 gcc-4.1.1 gcc-4.3.2  
gcc-4.3.3 gcc-4.4.0 gcc-4.4.1 gdb ggrep ghostscript ghostview gm4 gmake  
gnupg gnuplot gqview grace gsed gtar gv gzip hpijs hping imagemagick java  
jedit jxplorer keepass kerberos less lsof lynx lyx mc mpeg_play mtools  
mutt ncd u nedit netbeans netcat netpbm ntp nttcp octave openSoftwareLib  
openoffice openssh patch perl-5.10.1 perl-5.8.4 pidgin pine postgresql  
procmail pstree psutils pvm python-2.6 python qt rcs rdesktop recode  
recordmydesktop rsync samba screen sox subversion sudo tcltk tcsh teTeX  
texinfo tgif thunderbird tkcvs tkdiff tomcat valgrind vim vnc wget x11vnc  
x3270 xemacs xev xfce4 xfig xftp xpdf xv Xaw3d Xvfb apache berkeleydb  
coreutils ctags cvs-local cyrus-sasl dejagnu dmalloc doxygen dvd+rw-tools  
expat f2c fftw freetype gcc-4.2.1 gdbm gettext ghostpcl gimp2 gmp gnuplot  
gpa gridengine groff html2ps iperf irfanview jpilot kerberos lapack libIDL  
libart libbzip2 libgd libiconv libjpeg libpng libungif libxml2 lpd monitor  
mpfr mysql ncftp ncurses netscape nmon nvi nx openldap openscenegraph  
openssl pcre pdflib perl-ldap perl-spezial pgp php pilot-link plan pts pv  
python qt rcs readline rootsh samba-2.2 samba sgzip smartmontools squid  
tcpdump test tiff tkX-7.3.1 top xaff xemacs-large xforms xmon xscope  
xspread zip zlib
```

Testumgebung

Hardware

```
holgrrr — ssh — 82x24
ssh
opensuse123-host:~ # dmidecode | egrep '(Manufacturer|Product Name)' | head -n 2
    Manufacturer: Dell Inc.
    Product Name: OptiPlex 745
opensuse123-host:~ #
opensuse123-host:~ #
opensuse123-host:~ # egrep '(model name)' /proc/cpuinfo
model name      : Intel(R) Core(TM)2 CPU          6300  @ 1.86GHz
model name      : Intel(R) Core(TM)2 CPU          6300  @ 1.86GHz
opensuse123-host:~ #
opensuse123-host:~ #
opensuse123-host:~ # egrep '^flags.*(vmx|svm)' /proc/cpuinfo | head -n 1
flags           : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pa
t pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx lm constant_tsc
arch_perfmon pebs bts nopl aperfmperf pni dtes64 monitor ds_cpl vmx est tm2 ssse3
cx16 xtpr pdcm lahf_lm dtherm tpr_shadow
opensuse123-host:~ #
opensuse123-host:~ #
opensuse123-host:~ # free -m
              total          used          free          shared          buffers          cached
Mem:           3895          2604          1291             0             70            2007
-/+ buffers/cache:           526          3369
Swap:           2053             0          2053
opensuse123-host:~ #
```

Informationen für Paket kvm:

Repository: openSUSE-12.3-1.7

Name: kvm

Version: 1.3.0-3.3.2

Arch: x86_64

Anbieter: openSUSE

Installiert: Ja

Status: aktuell

Installierte Größe: 10,9 MiB

Zusammenfassung: Kernel-based Virtual Machine

Beschreibung:

KVM (Kernel-based Virtual Machine) is virtualization software for Linux. It is based on the hardware virtualization extensions provided by Intel VT and AMD-V technologies. KVM kernel modules provide a control interface at /dev/kvm which the qemu-kvm user-space program uses to provide a hybrid emulated and actual hardware environment sufficient to run various PC operating systems (guests) in unmodified form, including Linux, Windows, and Mac OS X.

Since qemu-kvm is derived from the qemu processor emulator sources it also is able to run guests using processor emulation mode, but with the expected performance impact. Conversely, hardware virtualization features outside the processor such as iommu and sr-iovt are used by KVM allowing for improved performance.

The seabios, vgabios and ipxe open source projects are also pulled from to provide the firmware components included.

To increase performance over emulated hardware devices virtio drivers are supported, and in the case of Windows, included.

KVM is compatible with various VM management solutions, including libvirt, virt-manager and vm-install.

Informationen für Paket gcc:

Repository: openSUSE-12.3-1.7

Name: gcc

Version: 4.7-7.1.1

Arch: x86_64

Anbieter: openSUSE

Installiert: Ja

Status: aktuell

Installierte Größe: 0 B

Zusammenfassung: The system GNU C Compiler

Beschreibung:

The system GNU C Compiler.

opensuse123-host:~ # uname -a

Linux opensuse123-host 3.7.10-1.1-desktop #1 SMP PREEMPT Thu Feb 28 15:06:29 UTC 2013 (82d3f21) x86_64 x86_64 x86_64 GNU/Linux

opensuse123-host:~ #

Software

Hosts & Hypervisors & Guests

OpenSUSE 12.3



Welcome

Bienvenue

Vítejte

Benvenuto

Willkommen

Bienvenido

Добро
пожаловать

Quelle:
OpenSUSE 12.3 Installer

Welcome to CentOS 6.4!


Install or upgrade an existing system
Install system with basic video driver
Rescue installed system
Boot from local drive
Memory test

CentOS 6.4

Press [Tab] to edit options

Automatic boot in 58 seconds...

CentOS 6
Community ENTERprise Operating System



Quelle:
CentOS 6.4 Installer

VMware ESXi 5.1.0 (VMKernel Release Build 1065491)

innotek GmbH VirtualBox

Intel(R) Core(TM) i5-3427U CPU @ 1.80GHz
2 GiB Memory

Welcome to xen.org, the home of the Xen® Project which dev

Welcome to xen.org, the home o...

LINUX FOUNDATION COLLABORATIVE PROJECTS

Home Wiki Blog Lists User Voice Downloads Xen Planet



LIN
CO
Xen is
Laying
Xen's g

Hypervisorputtelchen

Loading module...

Oracle VM VirtualBox

Status - KVM



KVM



Virtual

Welcome to VirtualBox

VirtualBox is a powerful x86 and AMD64/Intel64 architecture virtual machine monitor for enterprise as well as home use. Not only is it a high performance product for enterprise customer solution that is freely available as Open Source software under the GNU General Public License (GPL) version 2.0, but it also has a detailed introduction.

- About
- Screenshots
- Downloads
- Documentation

- MAIN PAGE
- STATUS
- GUEST SUPPORT STATUS
- KVM-AUTOTEST
- HOWTO
- MIGRATION

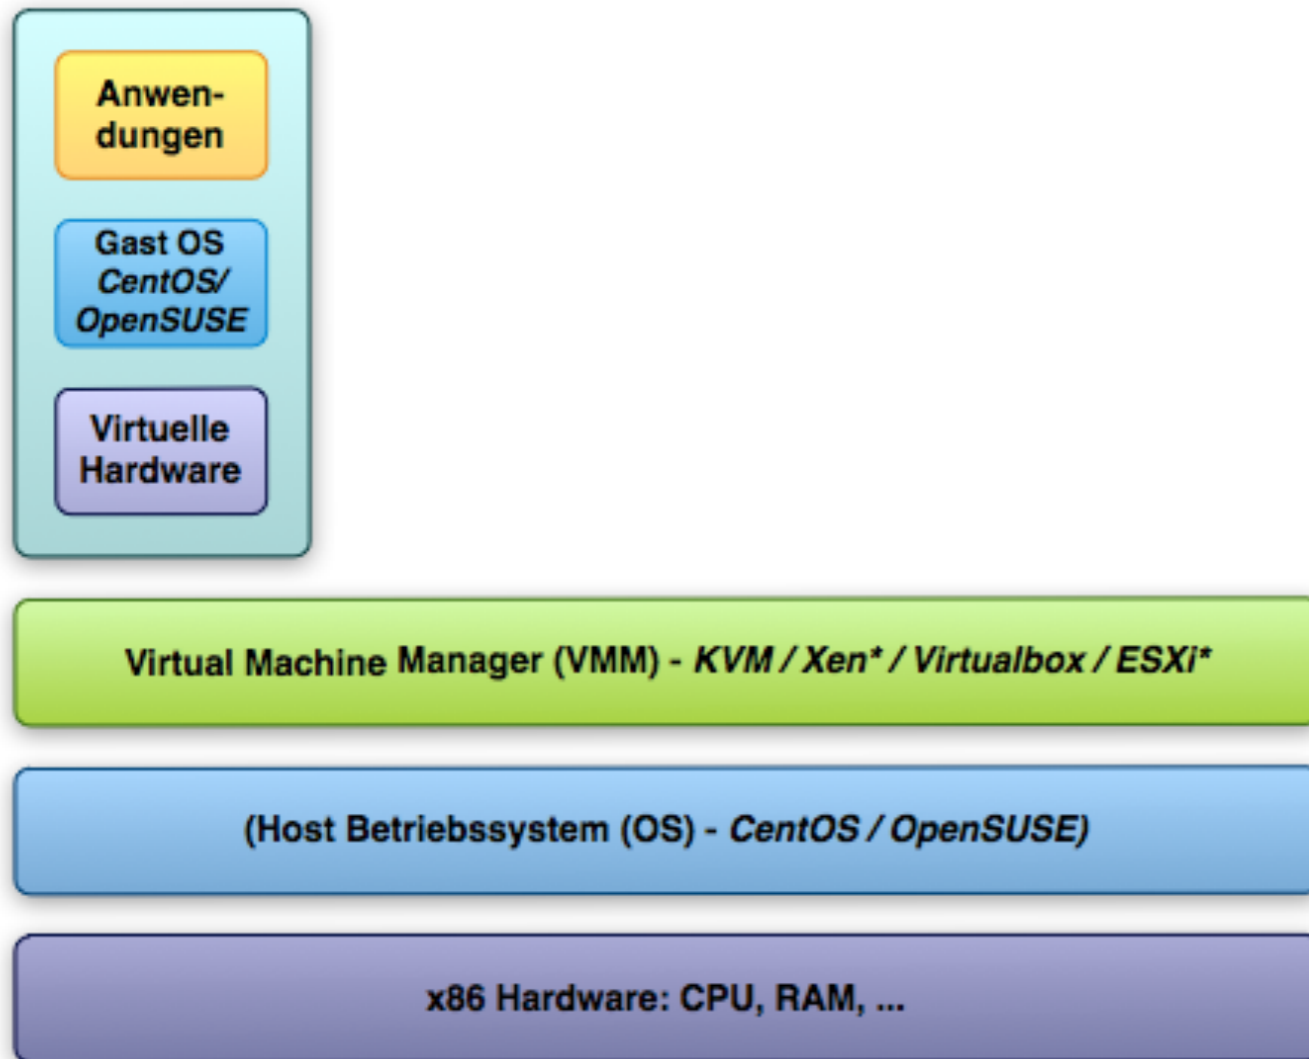
Status

KVM is included in the mainline of most workloads.

The userspace tools are part of the mainline of most workloads.

It is also available as a patch for most of your favorite distro- provided Enterprise Linux Distributions

Aufbau



Hinweise:
Diagramm bei Lösungen mit * nicht korrekt

Hinweise zur Ausstattung VM

- * 1 vCPU x86_64
- * 2 GB RAM
- * 8GB vDisk

Entspricht was wir Infrastruktur-VMs vergeben.

CentOS / OpenSUSE - Gast

- auf CentOS / OpenSUSE-Host
 - mit Xen, KVM, VirtualBox als Hypervisor
- oder auf VMware ESXi

Immer nur 1VM aktiv!

Versions-Stände

	OpenSUSE 12.3 64bit	CentOS (6.4) 64bit
Kernel	3.7.10-1.1-desktop	2.6.32-358.6.1.el6.x86_64 2.6.32-358.6.2.el6.x86_64 3.9.2-1.el6xen.x86_64
KVM	1.3.0-3.3.2	1.0.3
XEN	4.2.1_12-1.8.1	4.2.2
VirtualBox	4.2.6-3.1.8	4.2.12_84980_el6
GCC	4.7-7.1.1	4.4.7

Und
VMware ESXi 5.1.u1

Prozedur

Benchmark-Skript

```
[root@centos64-host 3.6.15]# cat bench.sh  
#!/bin/bash  
[...]
```

```
for run in 1 2 3 4 5  
do
```

```
    echo run $run
```

```
        (time bash -c 'tar xf samba-3.6.15.tar.gz ;  
cd samba-3.6.15/source3 ; time ./configure ;  
time make' ) >& logfile-$OS-$VIRTTYPE-$CPU-  
$run
```

```
    rm -rf samba-3.6.15
```

```
done
```

Samba:
<http://ftp.samba.org/pub/samba/stable/samba-3.6.15.tar.gz>

Benchmark-Skript

Hinweise

- * Werte der Einzeldisziplinen nur wenig abweichend
- * Deshalb für Auswertung nur Endergebnis betrachtet

```
[root@centos64-host run1]# grep ^real logfile-CENTOS-XEN-  
PARA-x86_64-*
```

logfile-CENTOS-XEN-PARA-x86_64-1:real	1m41.054s
logfile-CENTOS-XEN-PARA-x86_64-1:real	6m12.247s
logfile-CENTOS-XEN-PARA-x86_64-1:real	7m55.886s
logfile-CENTOS-XEN-PARA-x86_64-2:real	1m39.449s
logfile-CENTOS-XEN-PARA-x86_64-2:real	6m11.782s
logfile-CENTOS-XEN-PARA-x86_64-2:real	7m53.842s
logfile-CENTOS-XEN-PARA-x86_64-3:real	1m40.055s
logfile-CENTOS-XEN-PARA-x86_64-3:real	6m11.735s
logfile-CENTOS-XEN-PARA-x86_64-3:real	7m54.394s
logfile-CENTOS-XEN-PARA-x86_64-4:real	1m39.619s
logfile-CENTOS-XEN-PARA-x86_64-4:real	6m10.562s
logfile-CENTOS-XEN-PARA-x86_64-4:real	7m52.780s
logfile-CENTOS-XEN-PARA-x86_64-5:real	1m39.542s
logfile-CENTOS-XEN-PARA-x86_64-5:real	6m10.701s
logfile-CENTOS-XEN-PARA-x86_64-5:real	7m52.857s

```
(time bash -c 'tar xf samba-3.6.15.tar.gz ; cd samba-3.6.15/source3 ;  
time ./configure ; time make' ) >& logfile-$OS-$VIRTTYPE-$CPU-$run
```



Ergebnisse

Ergebnisse

	OpenSUSE 12.3				CentOS 6.4			
	Mittel	Min	Max	%	Mittel	Min	Max	%
ohne Virtualisierung	7:24	7:24	7:25	100	6:09	6:09	6:09	100
Xen Paravirtualisierung	8:33	8:28	8:40	87 115	7:54	7:53	7:56	78 128
Xen Vollvirtualisierung	12:02	12:01	12:04	61 163	10:31	10:30	10:32	59 171
KVM (X86)	11:31 (17:05)	11:29 (17:04)	11:34 (17:08)	64 156 (43 231)	9:40	9:37	9:42	64 157
VirtualBox	25:31	22:49	27:56	29 345	18:39	18:35	18:44	33 303
ESXi (mit VMware Tools)	10:06	10:04	10:11	73 137	8:28 (8:29)	8:26 (8:28)	8:32 (8:29)	73 138 (73 138)

Ergebnisse - abgespeckt

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Genau drauf geschaut...



Wahl des Betriebssystem

- * Bereits ohne Virtualisierung Unterschiede!
- * Ursache wohl hier der Compiler
- * anderer Benchmark = andere Ergebnisse

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Wahl des Betriebssystem

* Wahl des Betriebssystems (relativ) egal
 * Prozentual gesehen ist Virt.-Verlust ähnlich

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Xen Para- vs. Vollvirtualisierung

- * Paravirtualisierung nach wie vor *schnell*
- * KVM nur paravirtualisierte IO
- * kam hier nicht zum Tragen

© 2011 Red Hat Group, Inc.

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Vollvirtualisierung - Xen vs

* bei vollständiger Virtualisierung KVM besser
 * kommt hier paravirt. IO zu tragen?

© 2011 Red Hat Group, Inc.

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

VMware vs. Xen vs. KVM

* gute Performance bei VMware
 * schlechter als Xen Para., besser als KVM

© 2011 Red Hat Group, Inc.

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Sonstiges

* VirtualBox für Szenario ungeeignet

© 2011 Red Hat Group, Inc.

	OpenSUSE 12.3		CentOS 6.4	
	Mittel	%	Mittel	%
ohne Virtualisierung	7:24	100	6:09	100
Xen Paravirtualisierung	8:33	87	7:54	78
Xen Vollvirtualisierung	12:02	61	10:31	59
KVM	11:31	64	9:40	64
VirtualBox	25:31	29	18:39	33
ESXi	10:06	73	8:28	73

Siegererhrung



- 1. Xen Paravirtualisierung**
- 2. VMware vSphere Hypervisor**
- 3. KVM**
- 4. Xen Vollvirtualisierung**
- 5. VirtualBox**

Fazit?

Spielt der *Hypervisor* noch eine Rolle?

Aber sicher!

Teil IV: Zusammenfassung

Virtualisierung Grundlage für viele aktuelle Themen der IT

**Linux bietet eine *Vielzahl* von
performanten Virtualisierungs-Lösungen
und Tools**

**Wahl des *Hypervisors* weiterhin
wichtig!**



science + computing

| A Bull Group Company



Vielen Dank für Ihre Aufmerksamkeit.

Holger Gantikow

science + computing ag
www.science-computing.de

Telefon: 07071 9457 - 503
E-Mail: h.gantikow@science-computing.de

Frage?

Antwort!





science + computing

| A Bull Group Company



<http://www.science-computing.de>
<http://www.hpc-wissen.de>
<https://www.science-computing.de/jobs>

Logos

- <http://libvirt.org/libvirtLogo.png>
- <http://www.eucalyptus.com/sites/all/img/logos/eucalyptus-logo-349x83.png>
- http://pve.proxmox.com/mediawiki/skins//common/images/proxmox_logo.png
- <http://www.convirture.com/images/logo.png>
- <http://www.ovirt.org/skins/common/images/oVirt-logo.png>
- http://opennebula.org/_media/about:opennebula_logo_clean.png
- http://theforeman.org/static/images/foreman_large.png
- http://www.openqrm-enterprise.com/uploads/pics/logoOpenQrm_07.png
- <http://www.openstack.org/blog/wp-content/uploads/2011/07/openstack-cloud-software-logo.png>