KVM Virtualization

wega Informatik AG

Aeschengraben 20
CH-4051 Basel

Phone +41 61 270 87 87
Fax +41 61 270 87 88

info@wega-informatik.ch

KVM: Why Talk About It?

- Already No. 1 supervisor under Linux
- Linux is going to be No. 1 VM host
- Free Open Source, no license costs
- Linux vanilla kernel
- Big company behind it (Red Hat)
- Well-known even outside Linux community
- Many complex features
- **Few experienced experts**

Once, There Was qemu

- http://www.qemu.org
- Emulator: dynamic code translation
- Userspace-Emulation: applications
- QEMU is a trademark (Fabrice Bellard)
- GPL, LGPL

Emulator: Host Emulates Guest

- Guest (program) executes opcode
- qemu Host emulates instructions...
- ...and executes them on virtual hardware
- But qemu can also translate opcode...
- ...and run it natively!

qemu: Host OS

- GNU/Linux
- Windows
- FreeBSD, NetBSD, OpenBSD
- OpenSolaris
- DOS
- OS X

qemu: Host Hardware

- x86
- x86_64
- PowerPC

qemu: Guest OS

- Windows NT
- DOS
- Linux
- OS X
- BeOS
- CP/M
- ReactOS
- ...!

qemu: Guest Hardware

- x86
- x86_64
- PowerPC
- ARM
- MIPS-32
- Sparc32/64
- ColdFire (m68k)
- ...

qemu: Emulated Hardware

- depends on emulated platform
- CD-ROM, Floppy
- video cards (OpenGL)
- NIC
- PCI, ISA, ATA bus
- sound card
- USB controller
- Parallel port

How Do We Store Guests?

- Disk Images in files
- Disk Images on the network
- Different formats

qemu: Image Devices

- regular files
- sparse files

qemu: Network Block Devices

- nbd daemon on block level
- access to LVM over TCP/IP
- Encryption
- RAID
- low-cost SAN with > 300 MB/s

qemu: Raw Image Format

- no fancy features
- fully covered by filesystem journal
- very efficient
- open with loopback device
- 1:1 disk images

qemu: qcow2 Image Format

- snapshots
- overlays
- non-sparse files
- encryption
- compression
- beware of power losses and SIGKILL!

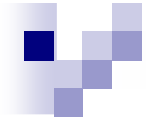

qemu: Features

- Far behind as far as UI is concerned
- Up to date as far as functionality goes
- (Quite) stable
- Ahead of its time: guest live migration

qemu: Guest Live Migration

- move guest from one CPU to another
- even over the network
- without stopping or killing the guest
- example: extend server center
- example: redistribute computation power

kqemu

- Part of qemu
- emulate x86 on x86
- fast
- not suited for production

Back to the Beginning: kvm

- x86_64 hardware layer for qemu mod
- makes use of Intel/AMD VT instruction sets
- very fast
- part of the Linux kernel since 2.6.16
- developed by Qumranet since 2006
- bought by RedHat in 2008
- makes Xen, UML, ... obsolete

kvm: Paravirtualization

- Speeds up devices to about 90-95%
- paravirtualized NIC driver (Linux, Windows)
- paravirtualized block device driver (Linux)

kvm: Memory Management

- great possibilities
- very efficient
- Ballooning, if supported by guest
- swapping
- page sharing
- live migration

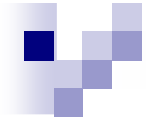

kvm: Future

RPM, RHEL, Cygwin, JBoss, Hibernate...

...kvm?

KVM Virtualization

Thank you

Questions ?

wega Informatik AG

Aeschengraben 20
CH-4051 Basel

Phone +41 61 270 87 87
Fax +41 61 270 87 88

info@wega-informatik.ch