

Virtualization on CentOS

Daniël de Kok (daniel@centos.org)

14 oktober 2007



Who am I?

- ▶ Student
- ▶ Part-time "hacker"
- ▶ CentOS devteam member
- ▶ CentOS working areas: yum, virtualization SIG, backporting/fixing stuff



Heritage

- ▶ Based on sources made available by a PNAELV (Prominent North American Linux Vendor)
- ▶ We aim for 100% binary compatibility



Defining the Enterprise

What makes it Enterprise?

- ▶ Long support cycles
- ▶ Security updates for several years
- ▶ ABI/API stability – patches get backported
- ▶ Support by hardware/software vendors
- ▶ Certifications are possible



Defining the Enterprise

The cons

- ▶ No new software during a release cycle
- ▶ No cutting edge software except at release point



Finding a place for CentOS

- ▶ Each release is supported for around 7 years
- ▶ Security patches get backported
- ▶ APIs/ABIs don't change
- ▶ Your RHCE works on CentOS
- ▶ Paid support is be available



When to run an enterprise distribution

There are several reasons to run a distribution with long time support. Here are just a few of them:

- ▶ You cannot afford to update all your servers each year
- ▶ Some software you use requires a supported distribution
- ▶ Your own software needs stable interfaces
- ▶ You need stable desktop software for your business
- ▶ You need support



Project goals

These are some of the goals of the CentOS project

- ▶ Easy maintenance via yum
- ▶ Building a self hosting distribution – CentOS rpms are built on a CentOS environment
- ▶ Build up a friendly environment for users and maintainers
- ▶ Long-term support of the base distribution
- ▶ Community infrastructure



This is where we are right now

- ▶ CentOS 2.1 is in maintenance mode and will get security patches until May 31st, 2009
- ▶ CentOS 3 (3.9 is current) is in maintenance mode also and will get security patches until October 31st, 2010
- ▶ CentOS 4 (4.5 is current) will get full updates until February 29th, 2008. Maintenance mode will end on February 29th, 2012.
- ▶ CentOS 5 (5.0 is current) will get full updates until March 31, 2010. Maintenance mode will end on March 31, 2014.



Who uses CentOS?

Userbase

- ▶ Absolute numbers: we have no idea
- ▶ Guesstimation: 1- 3 million machines
- ▶ Cluster users that we know of run CentOS on in excess of 2000 nodes



Who uses CentOS?

Derived systems

- ▶ Trixbox (formerly Asterisk at home)
- ▶ Rocks cluster
- ▶ SME server
- ▶ Strongbolt Linux
- ▶ Literom ERP
- ▶ Unbreakable Linux?



What is virtualization?

- ▶ Very broad term.
- ▶ At the operating system level: allow one physical system to simultaneously run multiple operating systems.



Advantages

- ▶ Isolation
- ▶ Consolidation
- ▶ Run legacy applications on modern hardware.
- ▶ Development/testing



Common virtualization techniques

- ▶ Emulation (Bochs, qemu)
- ▶ Full virtualization (qemu, VirtualBox, VMWare)
- ▶ Hardware-assisted virtualization (Xen, KVM, VMWare, VirtualBox)
- ▶ Paravirtualization (Xen)
- ▶ OS-level virtualization (OpenVZ, Linux-VServer)



Emulation

- ▶ Emulate machine hardware
- ▶ Emulate the system CPU



Emulation

Advantages

- ▶ Support for non-native platforms (e.g. for running ARM software on a x86 machine).
- ▶ Runs practically any operating system which supports the emulated hardware.
- ▶ Nice for system-level debugging.



Emulation

Disadvantages

- ▶ Since everything, including the CPU is emulated, it is very slow.



Full virtualization

- ▶ If the virtualized system has the same architecture as the physical machine, why not let the virtual machine execute on the host CPU?
- ▶ Privileged instructions trap when executed in user mode. The virtual machine monitor can then handle these instructions
- ▶ Problem: x86 not virtualization friendly.
- ▶ Binary translation: scan for problematic privileged instructions, and add a trap.



Full virtualization

Advantages

- ▶ Satisfactory speed.
- ▶ Runs practically any operating system which supports the emulated hardware.



Full virtualization

Advantages

- ▶ System hardware still needs to be virtualized.
- ▶ x86 kinda sucks for virtualization.



Hardware-assisted virtualization

- ▶ New Intel and AMD CPUs make virtualization a lot easier (VT-X and AMD-V respectively).
- ▶ Allow the virtual machine monitor to run in another mode than virtual machines.
- ▶ Let instructions executed in guest mode trap, so the virtual machine monitor can handle it.
- ▶ Replaces binary translation.



Hardware-assisted virtualization

Advantage

- Fast.



Hardware-assisted virtualization

Disadvantage

- ▶ Hardware still needs to be emulated
- ▶ Future: support for virtualization in other hardware as well?



Paravirtualization

- ▶ Modify the guest operating system kernels to work with the virtual machine monitor
- ▶ The guest system calls the hypervisor (hypercalls) for operations it can not perform from the unprivileged mode it runs in
- ▶ The virtual machine monitor provides virtualized devices (network, block), for which the guest has drivers



Paravirtualization

Advantages

- ▶ Very fast
- ▶ A lot of potential for management



Paravirtualization

Disadvantages

- ▶ Requires a modified kernel



Virtualization in CentOS

We currently provide the following technologies in CentOS 5:

- ▶ Xen (base system)
- ▶ KVM (CentOS extras)





Xen in CentOS

- ▶ Support integrated in CentOS 5
- ▶ Can run CentOS 4.5, and 5.x paravirtualized guests
- ▶ If your CPU has VT-X or AMD-V, you can run unmodified operating systems as well (e.g. CentOS 2.1, CentOS 3, Windows)



One (point and) click Xen install




The default installation of CentOS includes a set of software applicable for general internet usage. What additional tasks would you like your system to include support for?

- ☐ Server - GUI
- ☐ Virtualization
- ☐ Clustering
- ☐ Storage Clustering


Please select any additional repositories that you want to use for software installation.



- ☐ Packages from CentOS Extras

 Add additional software repositories

You can further customize the software selection now, or after install via the software management application.

☒ Customize later ☐ Customize now

 Release Notes

 Back  Next



Post-install

After the installation:

- ▶ The Xen hypervisor/virtual machine monitor is installed
- ▶ Your fresh CentOS installation is running as a so-called administrative domain (dom0) under the hypervisor.
- ▶ You will have a set of console and graphical tools for installing and managing unprivileged domains (domUs)



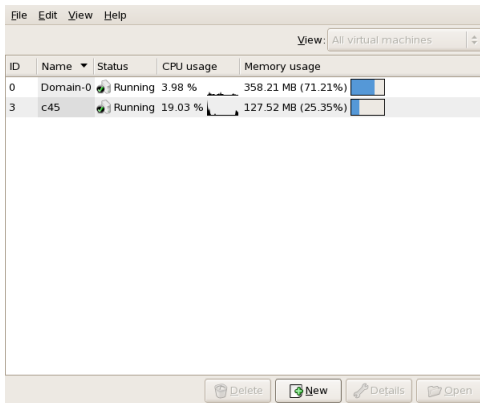
domU management

- ▶ Graphical: virt-manager
- ▶ Shell: xm, virsh



Tools

virt-manager



The screenshot shows the virt-manager application window. At the top is a menu bar with 'File', 'Edit', 'View', and 'Help'. Below the menu bar is a 'View:' dropdown menu set to 'All virtual machines'. The main area contains a table with the following columns: 'ID', 'Name', 'Status', 'CPU usage', and 'Memory usage'. There are two rows of data. The first row has ID '0', Name 'Domain-0', Status 'Running' with a green icon, CPU usage '3.98 %' with a small bar chart, and Memory usage '358.21 MB (71.21%)' with a blue progress bar. The second row has ID '3', Name 'c45', Status 'Running' with a green icon, CPU usage '19.03 %' with a small bar chart, and Memory usage '127.52 MB (25.35%)' with a blue progress bar. At the bottom of the window are four buttons: 'Delete' (trash icon), 'New' (plus icon), 'Details' (key icon), and 'Open' (folder icon).

ID	Name	Status	CPU usage	Memory usage
0	Domain-0	Running	3.98 %	358.21 MB (71.21%)
3	c45	Running	19.03 %	127.52 MB (25.35%)



Finishing touches

- ▶ Each CentOS domU has a (fake-ish) bootloader, so the domU can be fully managed
- ▶ Want to install domains automatically? Pass the *ks* parameter to the installation kernel for a kickstart install
- ▶ If the virtual framebuffer is enabled for a domain, you can connect to it with VNC



KVM in CentOS

- ▶ Is Xen too much work?
- ▶ Do you have a CPU that supports VT-X or AMD-V?
- ▶ Do you use CentOS 5?
- ▶ KVM may be a viable alternative!



Kernel-based Virtual Machine: A short introduction

- ▶ Hardware device emulation through a modified qemu
- ▶ A virtual machine monitor implemented as a kernel module that uses VT-X or AMD-V
- ▶ Runs nearly all modern x86/x86_64 systems
- ▶ Included in the CentOS 5 extras repository (*yum install kmod-kvm* does the trick)




KVM quickstart

Creating an image and booting a virtual machine with two commands:


```
$ qemu-img create -f qcow2 centos5-inst.img 4G  
$ qemu-kvm -hda centos5-inst.img -cdrom boot.iso -boot d
```



KVM in action



Dazzled? Confused? Want more?



Virtualization SIG

- ▶ CentOS virtualization Special Interest Group since September 2007
- ▶ Maintenance and testing of virtualization technologies in CentOS
- ▶ User support and discussion through the *centos-virt* list. You can subscribe at: <http://lists.centos.org/>
- ▶ Maintenance of virtualization related documentation in the CentOS wiki: <http://wiki.centos.org/>



Wrapup

CentOS is ideal for your virtual infrastructure:

- ▶ Xen is supported in the base system: one click install
- ▶ KVM is provided through CentOS-Extras: single command install
- ▶ Each version of CentOS is supported for seven years
- ▶ No license per VM, it's free as in beer and freedom
- ▶ Run legacy applications with CentOS 2.1, 3.9, or 4.5 in a virtual machine, with current security updates for the operating system
- ▶ Run your virtual machines with SELinux in operation.



Questions?

Questions?



Questions?

Answers!

