

Using Virtualization in the Real world

Whoami ?

Kris Buytaert@inuits.be

- Senior Linux and Open Source Consultant
- „Infrastructure Architect“
- Linux since 0.98
- Linux Expert @ Newtec.eu
- Large Scale Deployment , High Availability
-
- Surviving the 10th floor test
- CoAuthor @

Virtualization with Xen(tm): Including Xenenterprise, Xenserver, and XenExpress by David E. Williams

Agenda

- Why use Virtualization
- Do you know Xen ?
- A Too Early Adopter
- An uninformed Customer
- A telco environment
 - Why Xen @ Newtec
 - Automated Xen Installations
 - Complex Xen Networking
 - Problem Situation
 - Virtual Build Instances with Xen
- Lessons Learned

Why Virtualisation Matters ?

- Consolidation
- Saving Idle Cycles
- Security
- Separating Development/Staging/Production platforms
- A Greener Environment
- Following The Hype
- Having a budget that needs to be emptied.
-

Virtualisation vs Virtualisation

- Single OS image: Ensim, Vservers, Solaris Zones
 - Group user processes into resource containers
 - Hard to get strong isolation
 - Typically used for Hosting/Virtual Hosting/ISP

Chroot's on steroids

Virtualisation vs Virtualisation

- Full virtualization: VMware, VirtualPC, Qemu
 - Run multiple unmodified guest OSes
 - Hard to efficiently virtualize x86
 - No Access to the hardware
 - No Realtime guarantees
 - Problematic for certain privileged instructions

Virtualisation vs Virtualisation

- Para-virtualization: UML, Xen,
 - Run multiple guest OSes ported to special arch
 - Arch Xen/x86 is very close to normal x86
 - Some Access to underlying hardware
 - Better Performance
 - Needs modification to the OS

Virtualisation vs Virtualisation

Hardware Supported Virtualisation

- Xen
- KVM
- Others adopted slowly
- No more changes to the software ,
- Hardware solves the problem

Let's talk about Xen

- Pioneered ParaVirtualisation
- Pioneered VT support
- Went Mainstream real fast !
- Stable and performant platform
- Scriptable
- Open Source

- Scriptable
- Targeted at Server Platforms

- 98% CPU Speed
- Open Source
- Known Developers
-
- Let's not talk about KVM (yet)

DomU vs Dom0

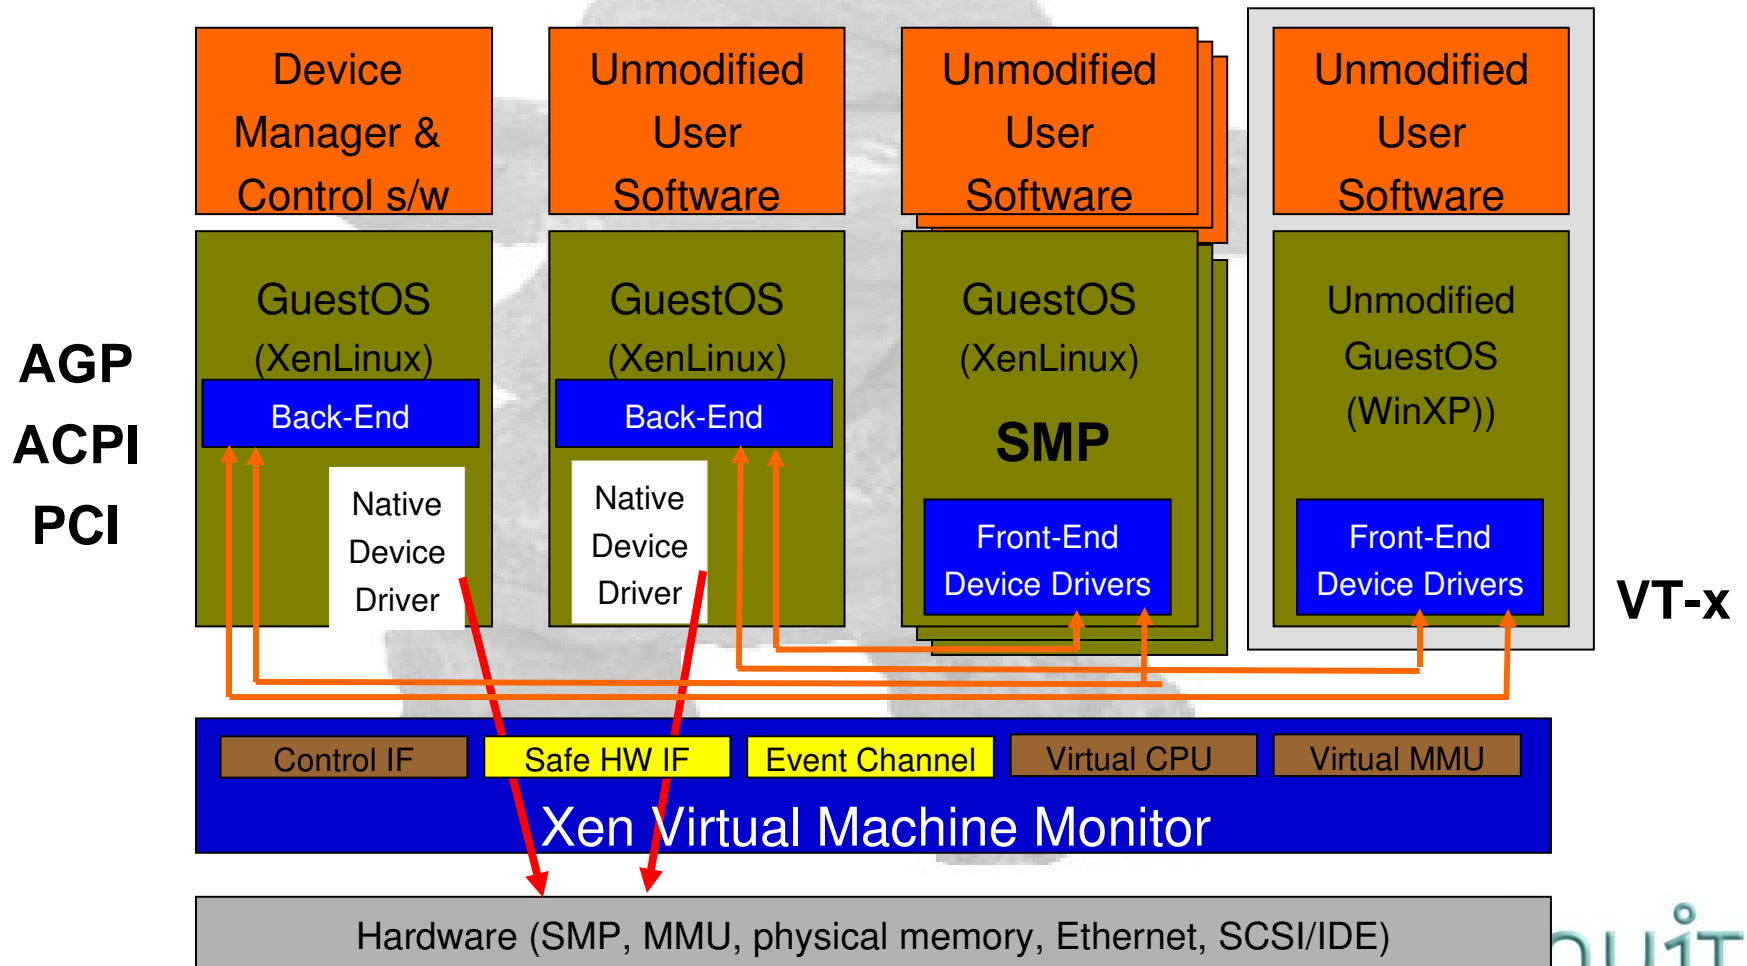
- Xen boots
- Bye bye “x86”

```
HOSTA:/etc/xen/scripts # xm list
```

Name	Id	Mem(MB)	CPU	State	Time(s)	Console
Domain-0	0	123	0	r----	41.2	

- Domain0 = management
- DomainU = virtual machines

Xen 3.0 Arch



A too early adopter

- Chose a core service first !
- Add more complexity than needed
 - Complexity is the enemy of High Availability
- Mix beta Software with more beta software
- Do not test
- Act first, then plan
- Go directly to production
- Downtime guaranteed
-

An uninformed Customer

- Major Hardware Vendor
 - Orders 2 weeks Consultancy
 - Orders software licenses
 - Does not allow direct contact with customer
- First Contact with customer
 - Just install the boxes please
 - How do you want them ?
- The Magical Migration
 - Of an uninventarized services

- Oh and they ordered licenses for the wrong architecture :)

Newtec

- Newtec is world renowned as a market leader in Satellite Communications, delivering an extensive range of highly innovative products and solutions. Since being founded in 1985, TV broadcasters, Telecom Service Providers, Integrators and Satellite Operators from all over the world have been relying on the exceptional performance and reliability of our unrivaled products.
- They build devices such as Modulators / Demodulators
- Platforms around these devices
 - DVB-RCS
 - SatMode
 - Sat3Play
 - Menos
 - ...
-

Why Open Source @ Newtec ?

- Because We Love It
- T.i.n.a
- We are part of the community
- We contribute back
 - Patches
 - Documentation
 - Xen in 2004
 - Production with Xen2
 - MySQL Cluster in 2005
 - ...
 -

Linux in a DVB-RCS/S3P Platform

Hub : Original Environment :

+/-198 machines (maximum)

- 1 x Common = 2 x 4 Applications HA
- 16 x Blade
 - 6 x Cache
 - 2 x shaping
 - 2 x measurement
 - 2 x pop management
 - 2 x device management

Terminals:

Our specific reasons for Virtualisation

- Testing Large Scale System Deployment
 - How to you test bootstrapping a large environment ?
 - Buy machines for your test platform ?
 - Interrupt regular services
 - Take down a chain ?
 - How do you test upgrade and rollback procedures ?
- No Budget for full test platforms :)

Virtual Platforms

- Full Hub in 1 machine
- Testing network paths
- Testing Failover
- Prototyping
- Testing Configuration Management
-

Consolidation in the 2Way Sat Hub

- Each application had it's own machine
- Highly Available
- Over Powered

- Small identical web applications
- Each on their own identical network
- In a different VLAN

- Perfect match for Virtualisation

Deploying Virtual Machines in a Large Environment

- - Consolidation of previously deployed machines
 - Repeated work is boring and error prone
 - We automated the physical machine deployment already so ...
- ==> Automating Xen Virtual Machine Deployment

Xen in Action

- $6 + 16 \times (2 \times 3 + 6) = 198$ machines
 - Consolidated already 1 application
 - now $6 + 16 \times (2 \times 2 + 6) = 166$ machines ($6 + 2 \times (2 \times 2 + 6) = 24$)
 - we moved already 2 redundant applications to 1 of the 6 shared machines
- Less Cabling
- Less Switches
- Less Power
-

Virtual Gateway Cluster

- Different test and validation platforms
- All with identical network topologies
- Gateways (SSH) to each virtual platform
- GW-1, GW-2, GW-3 , GW-8
- ==> vGW1- vGW8
- Became highly critical
-
- GW-X & GW-Y both containing vGW-1 - vGW-14

Solving Netfilter Problems

- Building a truly transparent proxy
 - Integration of LVS and Tproxy fails
 - Required multiple machines to work
 - CONNTRACK module conflicts
- Used Xen to build this on 1 machine

Solution: Dynamic Build Environments

Use a virtual environment with linux instances:

- that are set up every time specific software needs to be build
- for every type of software

These Linux instances have their own filesystem image:

- With specific build software installed
- Can only be used to build one type of software
- As small as possible

Information about the type of software

- Is stored in release notes (versioning systems)
- Will be used to determine the linux image to use
- Can be stored together with the software

.. On one physical machine

Lessons Learned:

- Only virtualise when needed
- Test , test , test
- Put delays between machine booting
- Xen highlighted some issues with typical linux setups (bridging code) and improved it there
- Customers need to be stopped in time
- Complexity is the enemy of availability



Contact:

Kris Buytaert <Kris.Buytaert@inuïts.be>

Sint Antoniuslei 45
B - 2930 Brasschaat

T: +32 (0)3 633.26.75

Further Reading

<http://www.krisbuytaert.be/blog/>

<http://wiki.xensource.com/>

www.inuïts.be

Hybrid Deployment Goals

- Hands off (Virtual) Machine Deployment
- Minimal impact on the current infrastructure
- Fast deployments
- Easy to redeploy / reproduce
- No vendor lock in's
- We don't like doing the same stuff for different distributions

The 10th Floor Test

- Grab a random machine
(don't take a backup before)
 - Throw it out a 10th floor window
 - Can you recover it in <10 minutes ?
 - Even for Virtual Machines ?
- (Steve Traugot)