Xen and the Art of Consolidation

Tom Eastep
Linuxfest NW
April 29, 2007
http://www.shorewall.net/Linuxfest-2007.pdf

Outline

- Introduction
- Why I use Xen
- Xen Overview
- Xen Installation
- Xen Networking Options
  - Bridged
  - Routed
  - NAT
- Q&A

Introduction

- About me
  - Developing or Supporting Software for 38 years
  - Software Architect for HP
  - Developer of Shorewall (Linux firewall tool)
  - Disclaimers
    - I am here as an individual and not as an HP employee
    - What I tell you represents my own views and not those of HP
    - I am not affiliated with the Xen project or with XenSource
    - I am not affiliated with Novell/SuSE

Why I use Xen - Before I used Xen

- Home office crowded with 5 systems, three monitors a scanner and a printer
- Systems:
  - Firewall
  - Public shorewall.net Server (Web, FTP, Mail)
  - Private File/Print Server
  - Personal Desktop (Linux)
  - Work System (Docked Laptop running Windows™ XP)

Why I use Xen - Problems with my old setup

- Crowding
- Noise
- Heat
- High power consumption
- Lots of systems to administer
- Systems underutilized

Why I use Xen - Before I used Xen

- Why so many systems?
  - Isolate Internet-accessible systems from other systems
  - Isolate functions
  - Strong isolation preferred
  - Reproduce Shorewall user configurations
- I attended a presentation by Ian Pratt and became interested (2004 - pre-Xen 2 timeframe)

Why I use Xen - After I Adopted Xen

- Three Systems with Two Monitors and an HP AllInOne™ Printer/Copier/FAX
- Systems:
  - Combination Firewall/Public Server/Private Server/Wireless Gateway using Xen.
  - Desktop
  - Work System
- KVM switch and LCD Monitor with Dual Video Inputs (although the Xen system is essentially headless)

Why I use Xen - Benefits

- Less clutter
- Less noise
- Less power consumed
- Fewer systems to administer
- Fuller utilization of hardware assets
- There are other benefits that I don't take advantage of
  - Dynamic re-balancing to achieve SLAs
  - Avoid downtime with VM relocation

Xen Overview

- Developed at Cambridge University
- Allows multiple virtual machines within one physical machine
- Open Source version available as option with Linux Distributions
  - Text-mode tools (but that is improving)
- Now also provided by XenSource (“Commercial Open Source”)
  - Free XenServer Express (complete distribution – currently 32-bit only)
  - Nice graphical tools
  - I haven't used it yet

Xen Overview – VM Terminology

- Common Terminology
  - Host: The system hosting one or more virtual machines
  - Host OS: The operating system running on the Host
  - Guest: A virtual machine environment
  - Guest OS: The OS running on a particular guest

Xen Overview - Virtualization

- Single OS Image – Virtuozo, Vservers, OpenVZ, Zones.
  - “chroot on steroids”
  - Hard to establish protection zones
  - Confuse Administrators
- Full Virtualization – Vmware, VirtualPC, QEMU
  - Run multiple unmodified guest Oses
  - Hard to Virtualize X86 efficiently
- Para-virtualization – Xen
  - Run multiple guest Oses ported to a special architecture (Xen/X86)
  - Xen/X86 is very close to true X86
  - Also Xen/X86_64

Xen Overview - Xen and Para-virtualization

- Initially supported only p*ara*-virtualization
- Now supports full virtualization
  - Requires hardware support available in current server-class CPUs
  - AMD & Intel's *Pacifica* and *Vanderpool* extensions (think of them as ring “-1”).
- I use paravirtualization

Xen Overview – Pros/Cons of Para-virtualization

- PROS:
  - The main advantage of para-virtualization over other approaches is performance
  - Varies with benchmark
  - www.cl.cam.ac.uk/netos/papers/2003-xensosp.pdf
- CONS:
  - With para-virtualization, you can't just drop an unmodified installation DVD into the drive and install the Guest OS

Xen Overview - Domains

- Xen refers to virtual machines as *domains*
- The host OS runs in a special domain called domain-0 (Dom0)
- The guest OSes run in *User Domains (DomUs).*
- Domains have a domain number and a domain name.
  - Name is fixed; number varies (except for Dom0)
- But all Domains are guests of the *Hypervisor*

Xen Overview - Hypervisor

- Under Xen, the bootloader does not load the OS!
- It loads the Xen Hypervisor (Virtual Machine Monitor)
- The Hypervisor then loads the host OS
- From /boot/grub/menu.lst

```
title Kernel-2.6.18.8-0.1-xen
    root (hd0,5)
    kernel /boot/xen.gz <- Hypervisor
    module /boot/vmlinuz-2.6.18.8-0.1-xen root=/dev/sda6 vga=0x31a …
    module /boot/initrd-2.6.18.8-0.1-xen
```

Xen Overview – Architecture V3

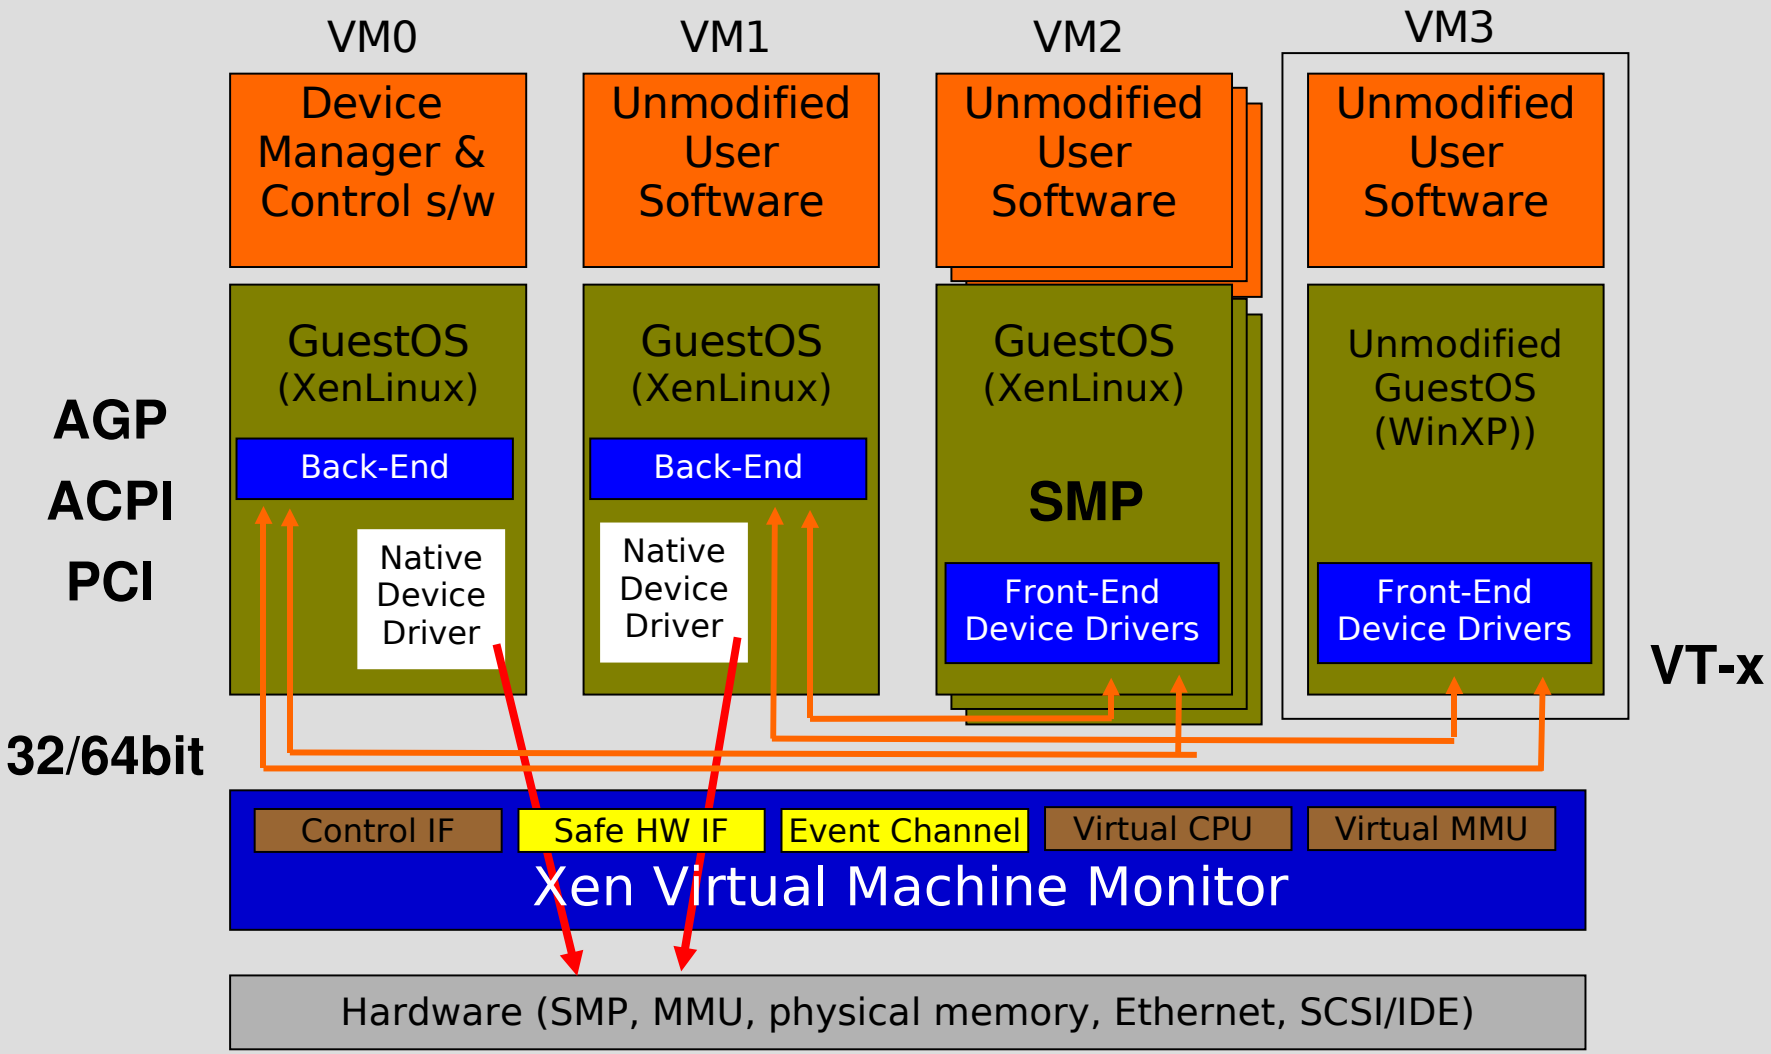

Xen Overview – Dom0 Services

- xend
  - manipulates the Dom0 environment to accommodate Xen
- xendomains
  - Auto start/stop DomU domains defined in /etc/xen/auto/
  - I personally create symbolic links from /etc/xen/auto to /etc/xen/vm (Where yast[2] creates domains)

Xen Overview - xm

- xm is a text mode tool for managing domains
- Runs in Dom0
- The 'list' command lists the domains

```
gateway:~ # xm list
Name                                        ID Mem(MiB) VCPUs State   Time(s)
Domain-0                                     0     1242     2 r-----  83315.2
server                                       2      512     1 -b----   3062.8
gateway:~ #
```

Xen Overview - xm

- Also supports a 'console' command

```
gateway:~ # xm console server
lists login: root
Password:
Last login: Sat Apr  7 07:57:58 PDT 2007 from 192.168.2.10 on pts/5
Have a lot of fun...
lists:~ # ps aux
USER       PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root         1  0.0  0.0    808    76 ?        Ss   Apr05   0:00 init [5]
root         2  0.0  0.0      0     0 ?        S    Apr05   0:00 [migration/0]
root         3  0.0  0.0      0     0 ?        SN   Apr05   0:00 [ksoftirqd/0]
```

Xen Overview - xm

- xm top

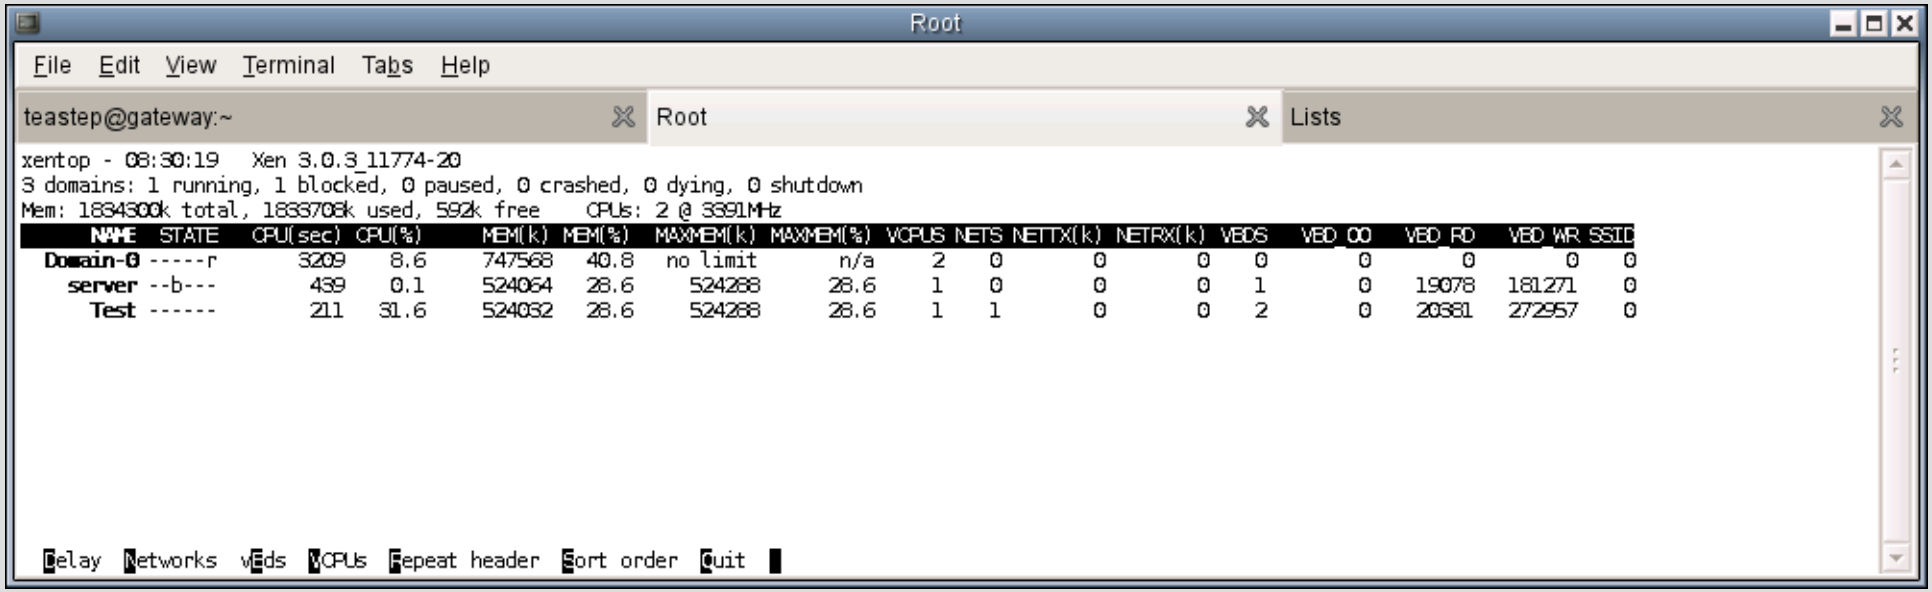

DomU Installation

- Use an Installer
  - OpenSuSE (my personal choice)
  - Fedora
  - Other Linux distributions
- Use an Image
  - http://jailtime.org/
- Roll your own

DomU Installation - Yast2

- Yast2 xen
- Select 'Add'

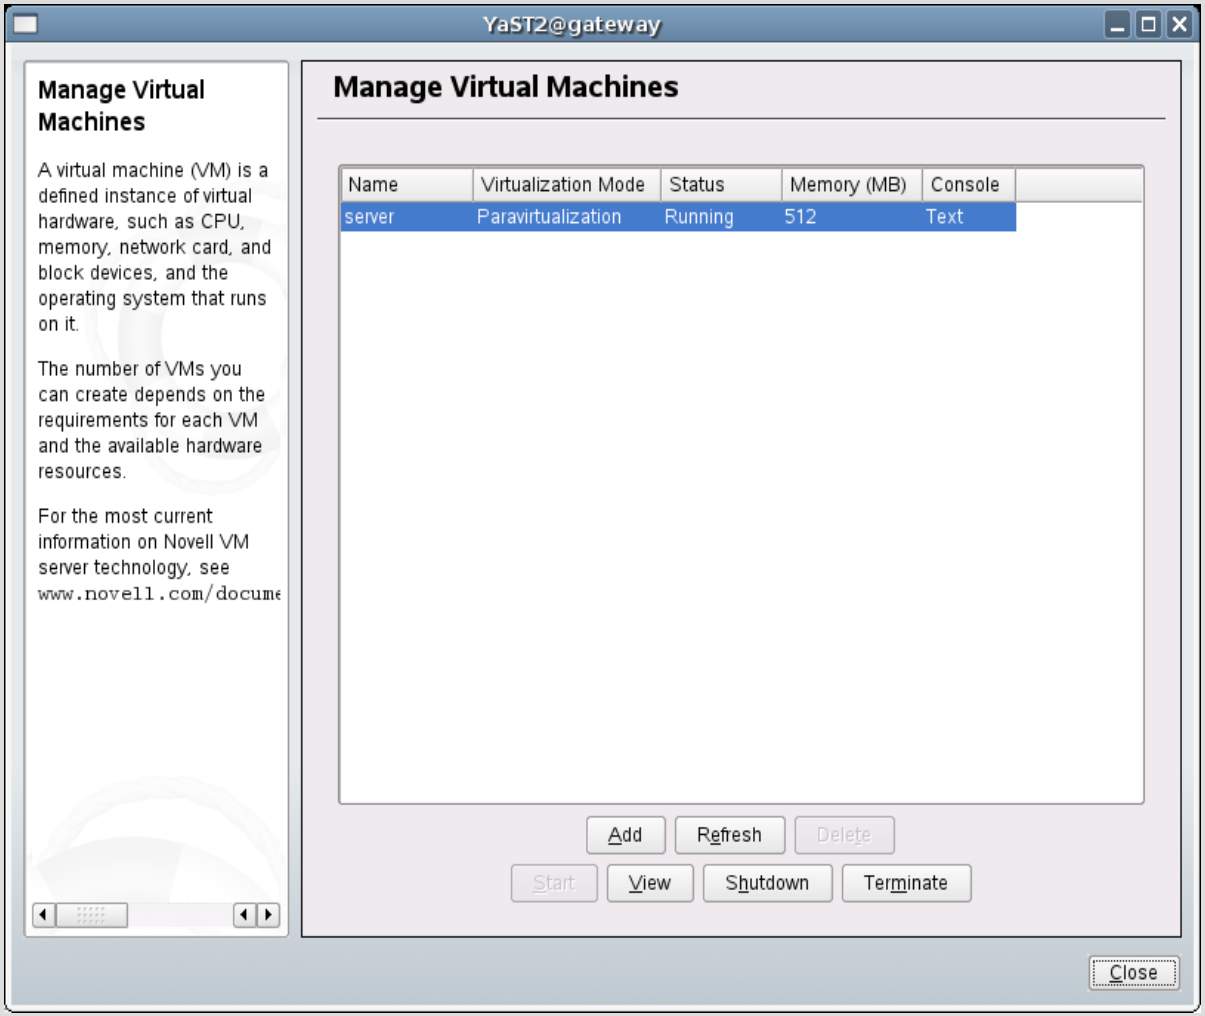

DomU Installation - Yast2

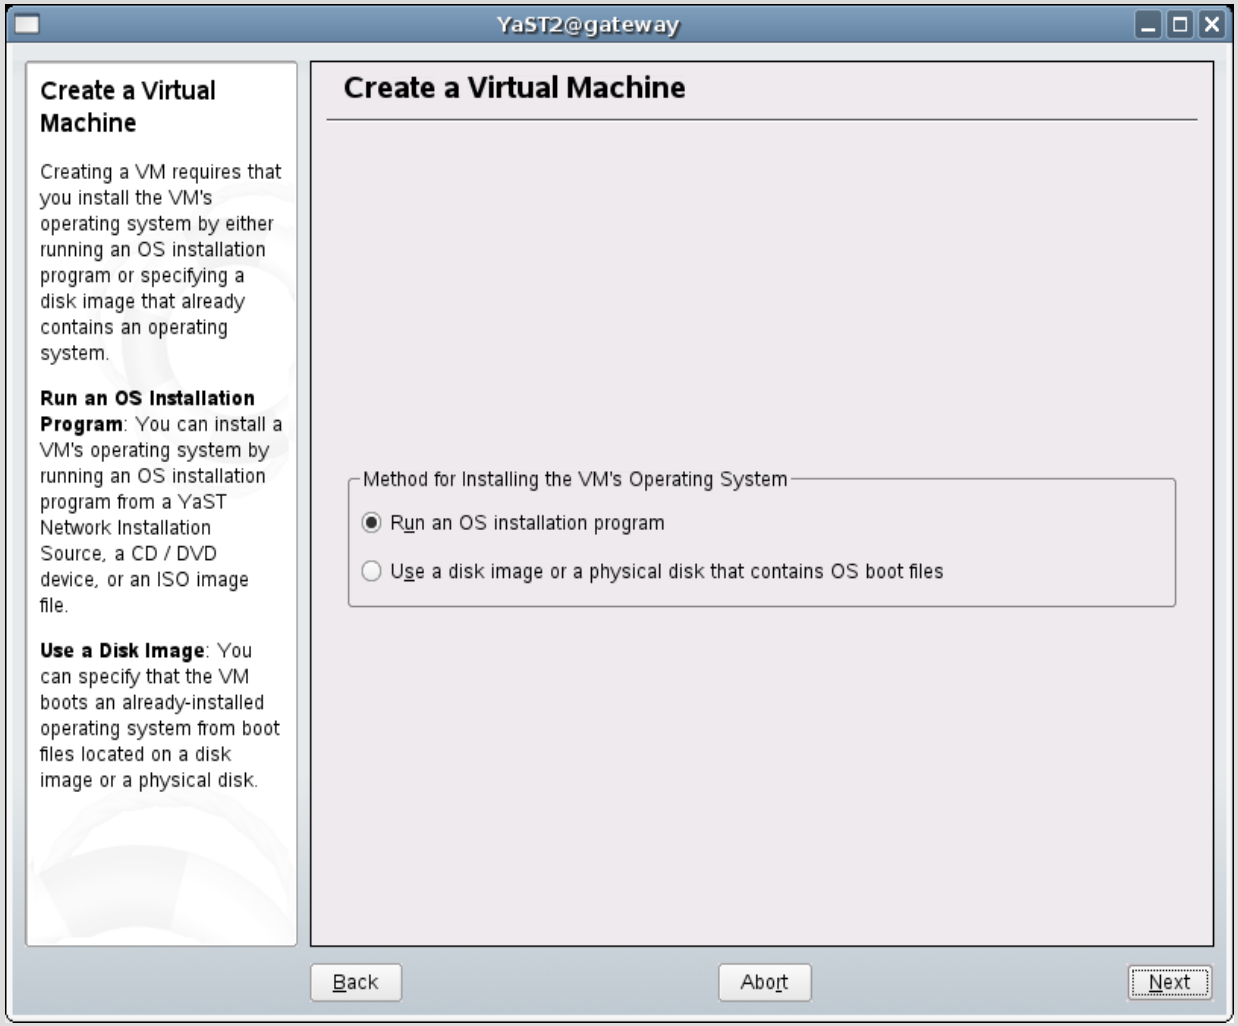

DomU Installation - Yast2

- Click “Run an OS Installation Program” and you get a choice of doing a SuSE Install or another distro's install (I've only done SuSE Installs).
  - OpenSuSE install is text mode in an XTERM after the initial configuration screens.
- Click “Use a disk image or a physical disk that contains OS boot files” to install an image.

DomU Installation - Yast2

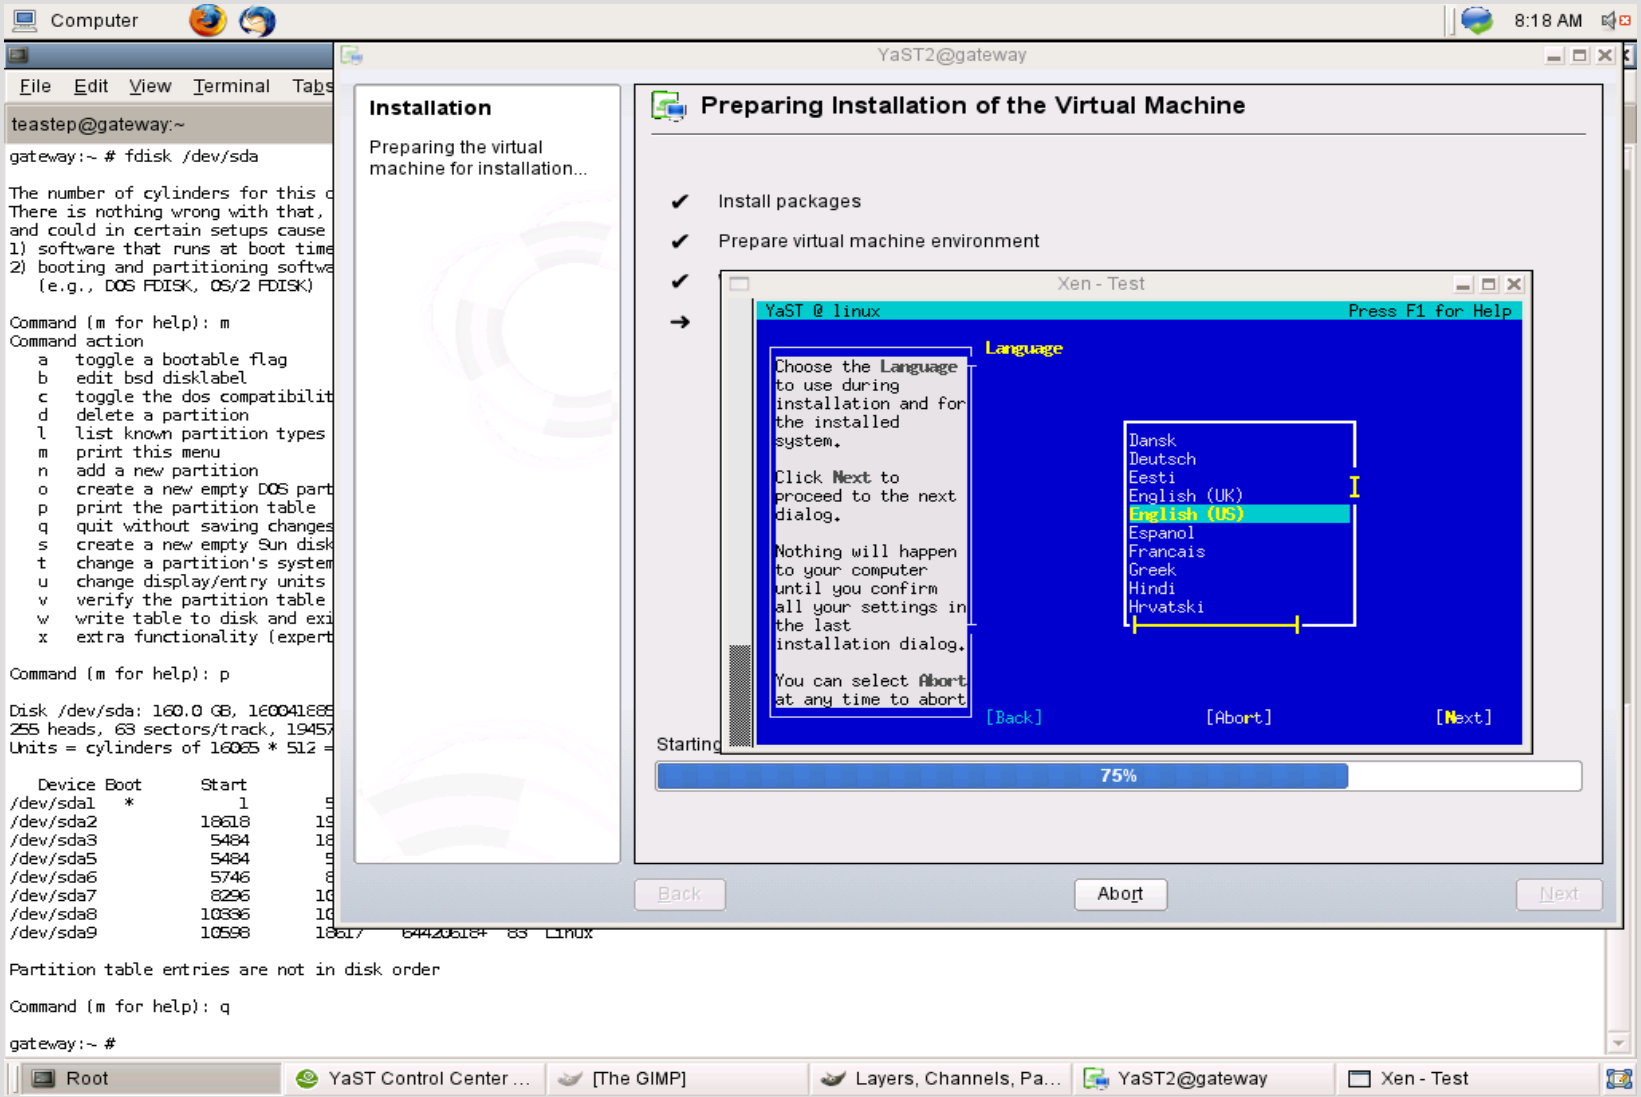

DomU Installation - Yast2

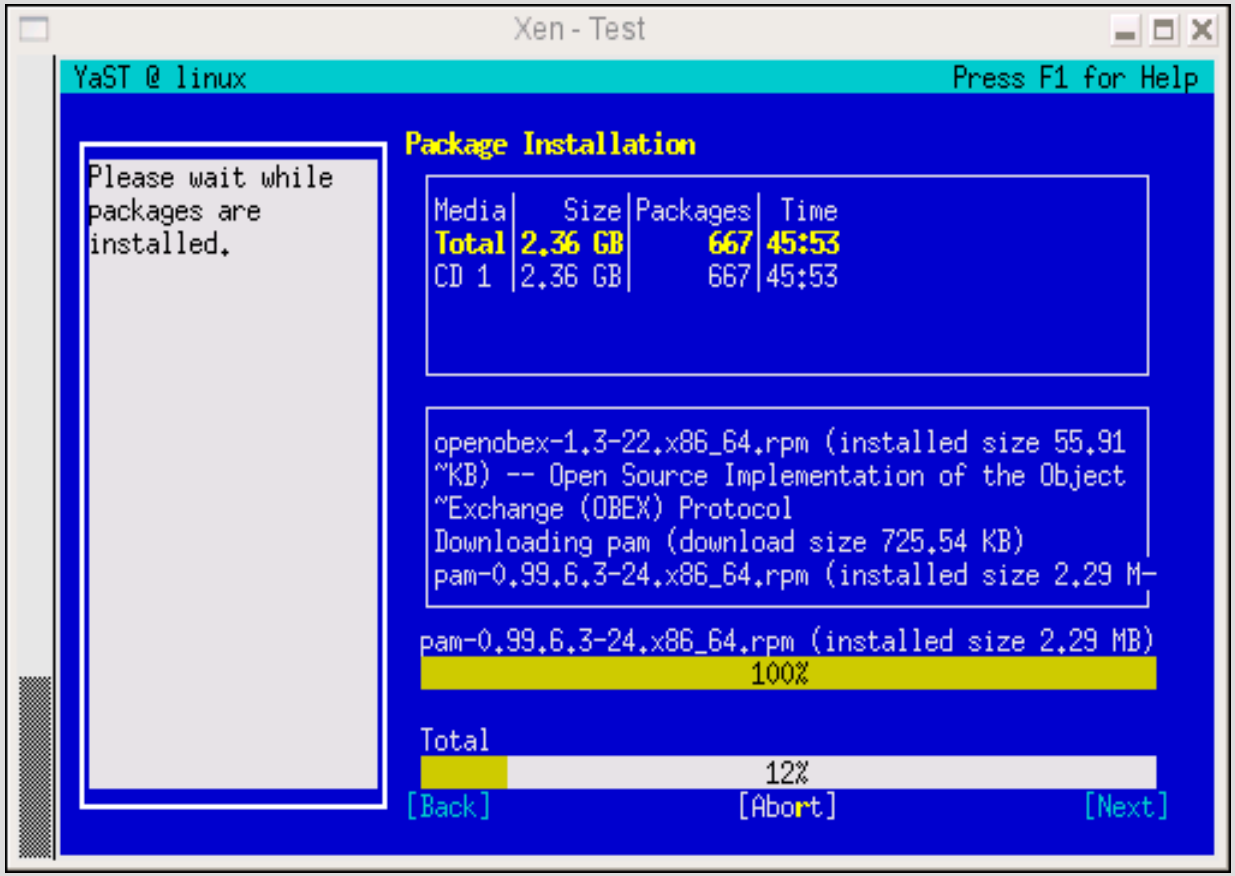

Xen Network Options

- Bridged – Default
  - See http://www.shorewall.net/Xen.html
  - Also http://www.shorewall.net/XenMyWay.html
- Routed
  - See http://www.shorewall.net/XenMyWay-Routed.html
- Nat

Xen Network Options - Bridged

- Default
- Pros
  - Simple
  - DomU's can use local DHCP Server
  - Works well for creating multiple servers in one physical system
  - Distro Installers assume it

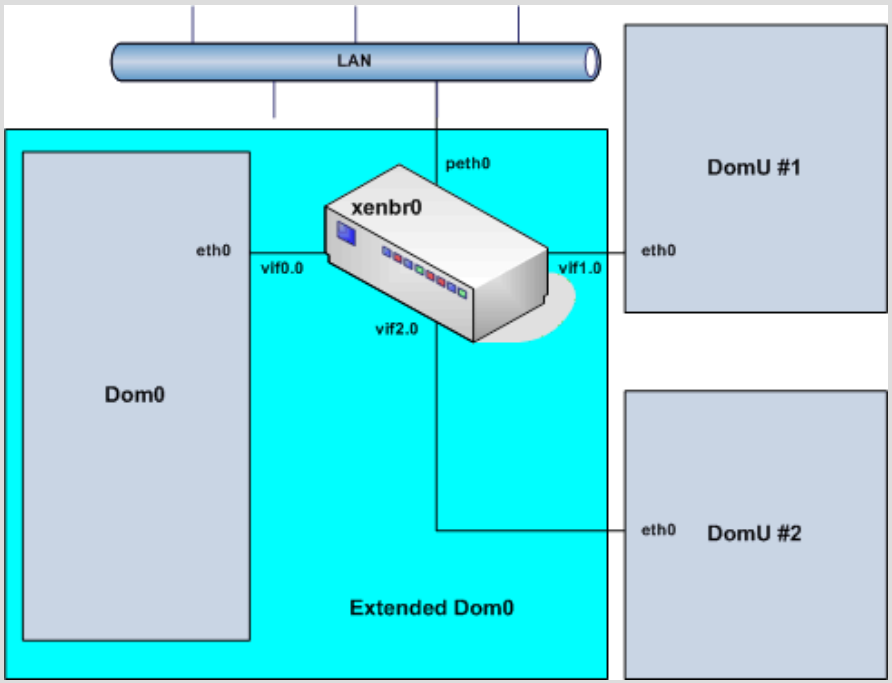

Xen Network Options - Bridged

- Cons – Awkward to firewall
  - Bridge is visible in Dom0
  - Beginning with Kernel 2.6.20, dealing with bridges gets a lot harder

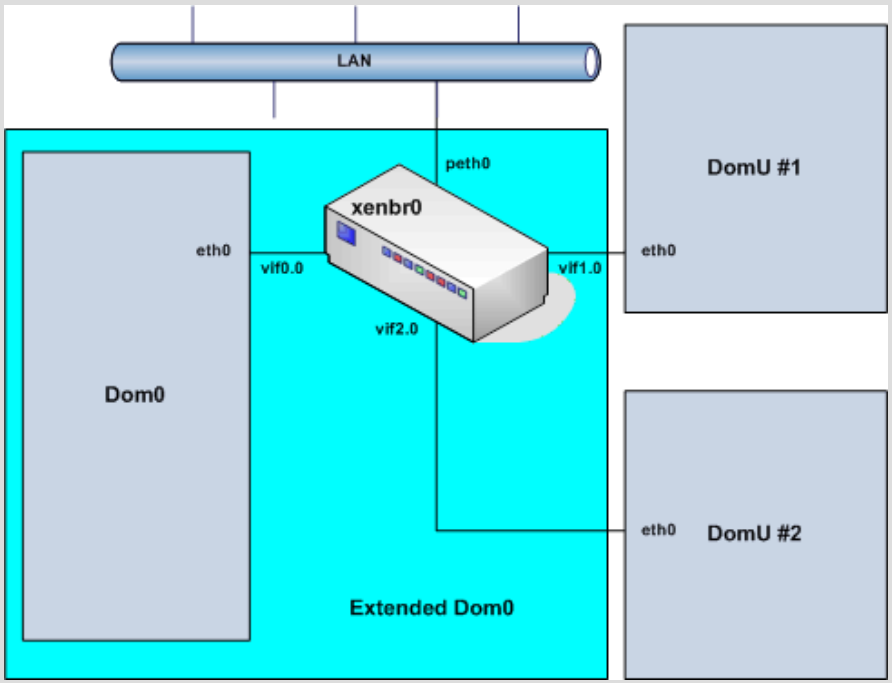

Xen Network Options - Bridged

- Firewall environment can be made more conventional if you run the firewall in a DomU
- Uses *device delegation*

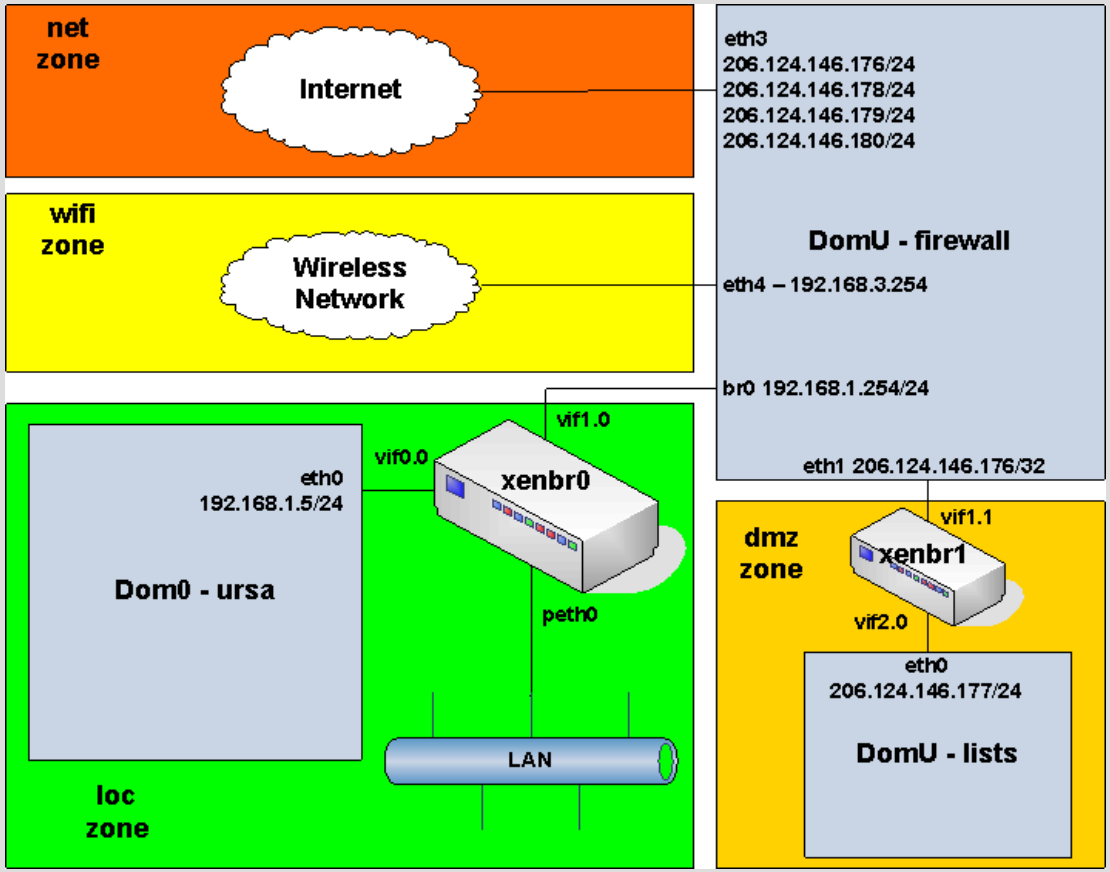

Xen Network Options - Routed

- Dom0 acts as a router for DomUs
- Pros
  - No Bridges
- Cons
  - Little Documentation
  - DomU Installation awkward

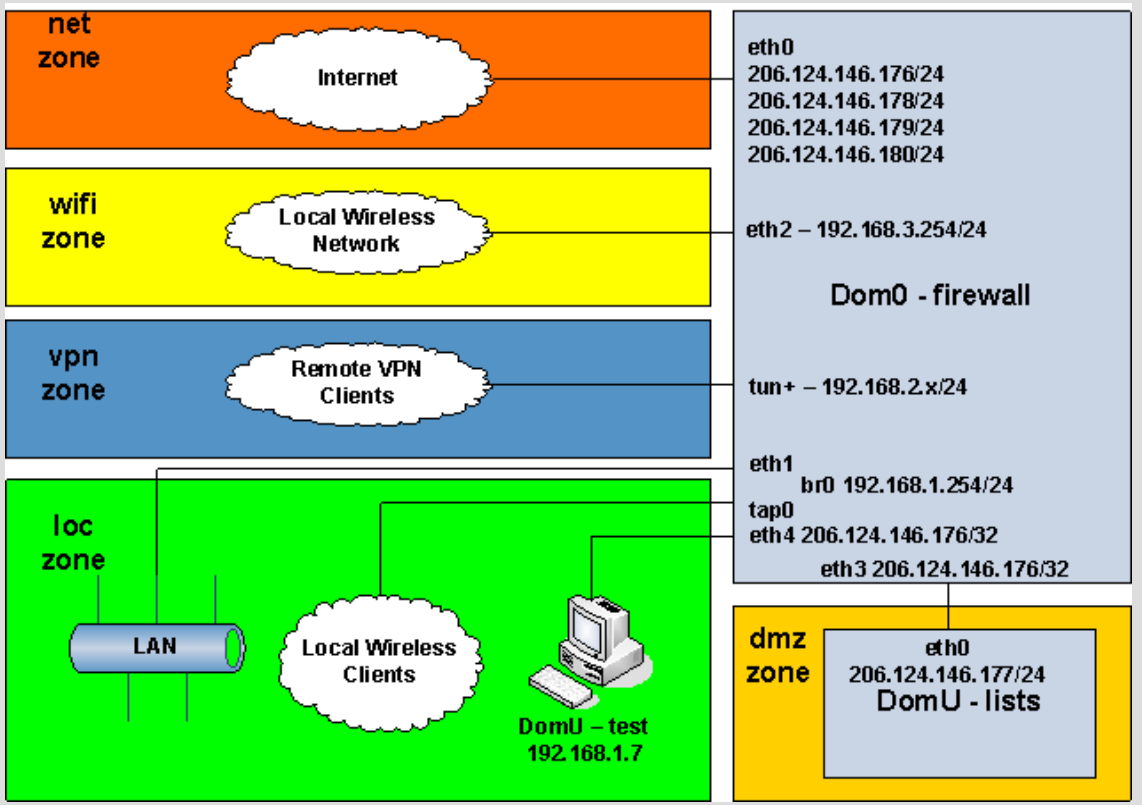

Xen Netowrk Options - Nat

- Like the routed configuration except that NAT is used rather than routing.
- I haven't used this option since Shorewall configures NAT

Q & A