Using Xen for Sandboxing PBS Worker Nodes

Daniel Bayer Steffen Möller

University of Lübeck, Institute for Neuro- and Bioinformatics

Virtualization Users Workshop, Hamburg
January 2007

Outline

Motivation

We want to use our institutes ordinary desktop PCs as worker nodes.

Grid Jobs

- must not have access to local network
- must not be able to see other non-grid processes
- must not slow down user processes

$\Rightarrow$ We need real virtualisation

Infrastructure Used

Currently three Desktop Computers are used:

- Debian Sid (Linux 2.6.18, Xen 3.0.3)
- Intel Pentium 4, 3 GHz (32 Bit)
- 4 GB RAM

- Debian Sid (Kernel 2.6.18, Xen 3.0.3)
- Intel Pentium 4, 3 GHz (32 Bit)
- 1 GB RAM

- Debian Etch (Linux 2.6.17, Xen 3.0.3)
- AMD Athlon 64 X2 4200+, 2.2 GHz (64 Bit, Dual Core)
- 4 GB RAM

Overview

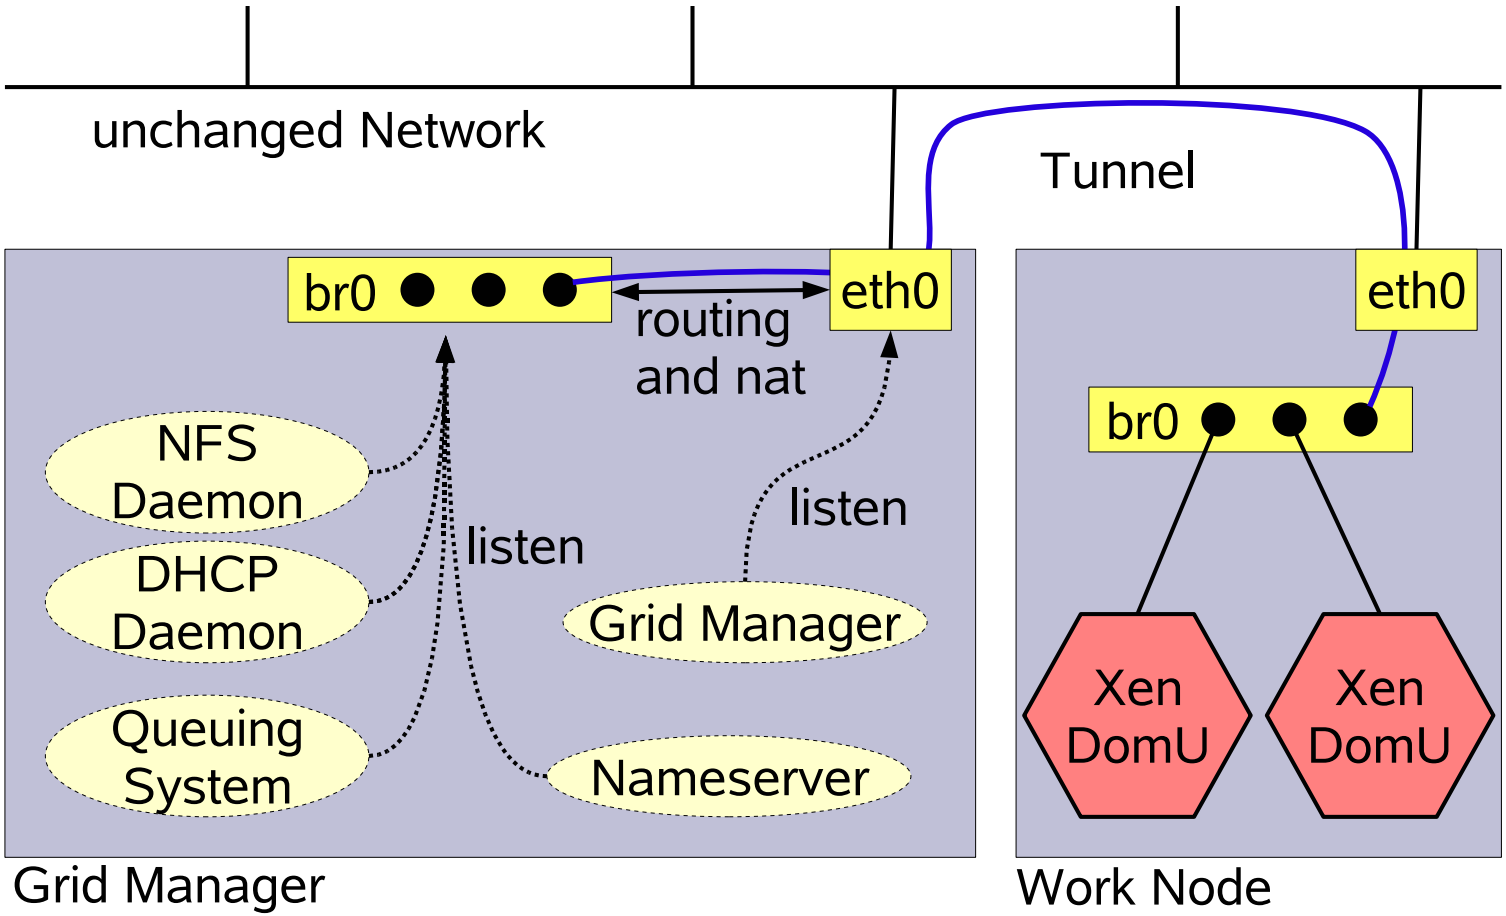
unchanged Network
Tunnel
br0
eth0
routing
and nat
NFS
Daemon
DHCP
Daemon
Queuing
System
listen
listen
Grid Manager
Nameserver
Grid Manager
eth0
br0
Xen
DomU
Xen
DomU
Work Node

Creating the Bridge

On all computers the following was added to `/etc/network/interface`:

```
auto br0
iface br0 inet static
        pre-up brctl addbr br0
        post-down brctl delbr br0
        address 10.0.xxx.1
        network 10.0.0.0
        netmask 255.0.0.0
        broadcast 10.255.255.255
        bridge_fd 0
        bridge_hello 0
        bridge_stp off
```

Creating the Tunnel

- Tunnels are set up with `vtun`
- a script monitors and restarts the tunnel if necessary
- Other programms using tap-devices can be used.

Setting up the Grid Manager

- install ARC Middleware
  - available from `http://www.nordugrid.org`
  - rpm packages for most linux distributions are provided
  - must be accessible from outside

Setting up the Grid Manager

- install ARC Middleware
- install Queuing System
  - in Lübeck torque is used, but others are also possible
  - only listens on the virtual interface `br0`

Setting up the Grid Manager

- install ARC Middleware
- install Queuing System
- install Nameserver
  - provides DNS lookup and reverse lookup for the virtual network
  - only accessible from within the virtual network

Setting up the Grid Manager

- install ARC Middleware
- install Queuing System
- install Nameserver
- install DHCP Server
  - offers IP adresses to worker DomU
  - uses fixed IP adresses
  - only listens on virtual interface!

Setting up the Grid Manager

- install ARC Middleware
- install Queuing System
- install Nameserver
- install DHCP Server
- install NFS Daemon
  - provides work directory and software to worker DomU

Setting up the Grid Manager

- install ARC Middleware
- install Queuing System
- install Nameserver
- install DHCP Server
- install NFS Daemon

These services may run on different computers.

Configuring Xen

Changes to `/etc/xen/xend-config.sxp`:

- On computers acting as router:

```
(network-script network-nat)
(vif-script vif-bridge)
```

- On other computers:

```
(network-script network-dummy)
(vif-script vif-bridge)
```

Starting the Worker DomU

- A image was created, which is accessible by NFS.
- To start a new DomU this image is
  1. copied to the worker node
  2. customised, which involves
     - setting the hostname
     - configuring torque
     - adding grid uid to `/etc/passwd` and `/etc/groups`
     - adding some hosts to `/etc/hosts`
- a MAC adress for the DomU is calculated
- the DomU ist started

Summary

- Virtualisation with Xen allows proper separation of local environment and grid computing environment.
- The presented approach is well suited for small sites.

- For larger sites the centric architecture of the virtual network and the bridging of all DomU to a single large network segment might be a performance bottleneck.

$\Rightarrow$ use of virtual routers