

# Xen Image Manager



escaping image storage hell

Virtualization Miniconf, linux.conf.au 2007  
Jonathan Oxer <jon@ivt.com.au>

# Storage Options



NFS

Fibrechannel

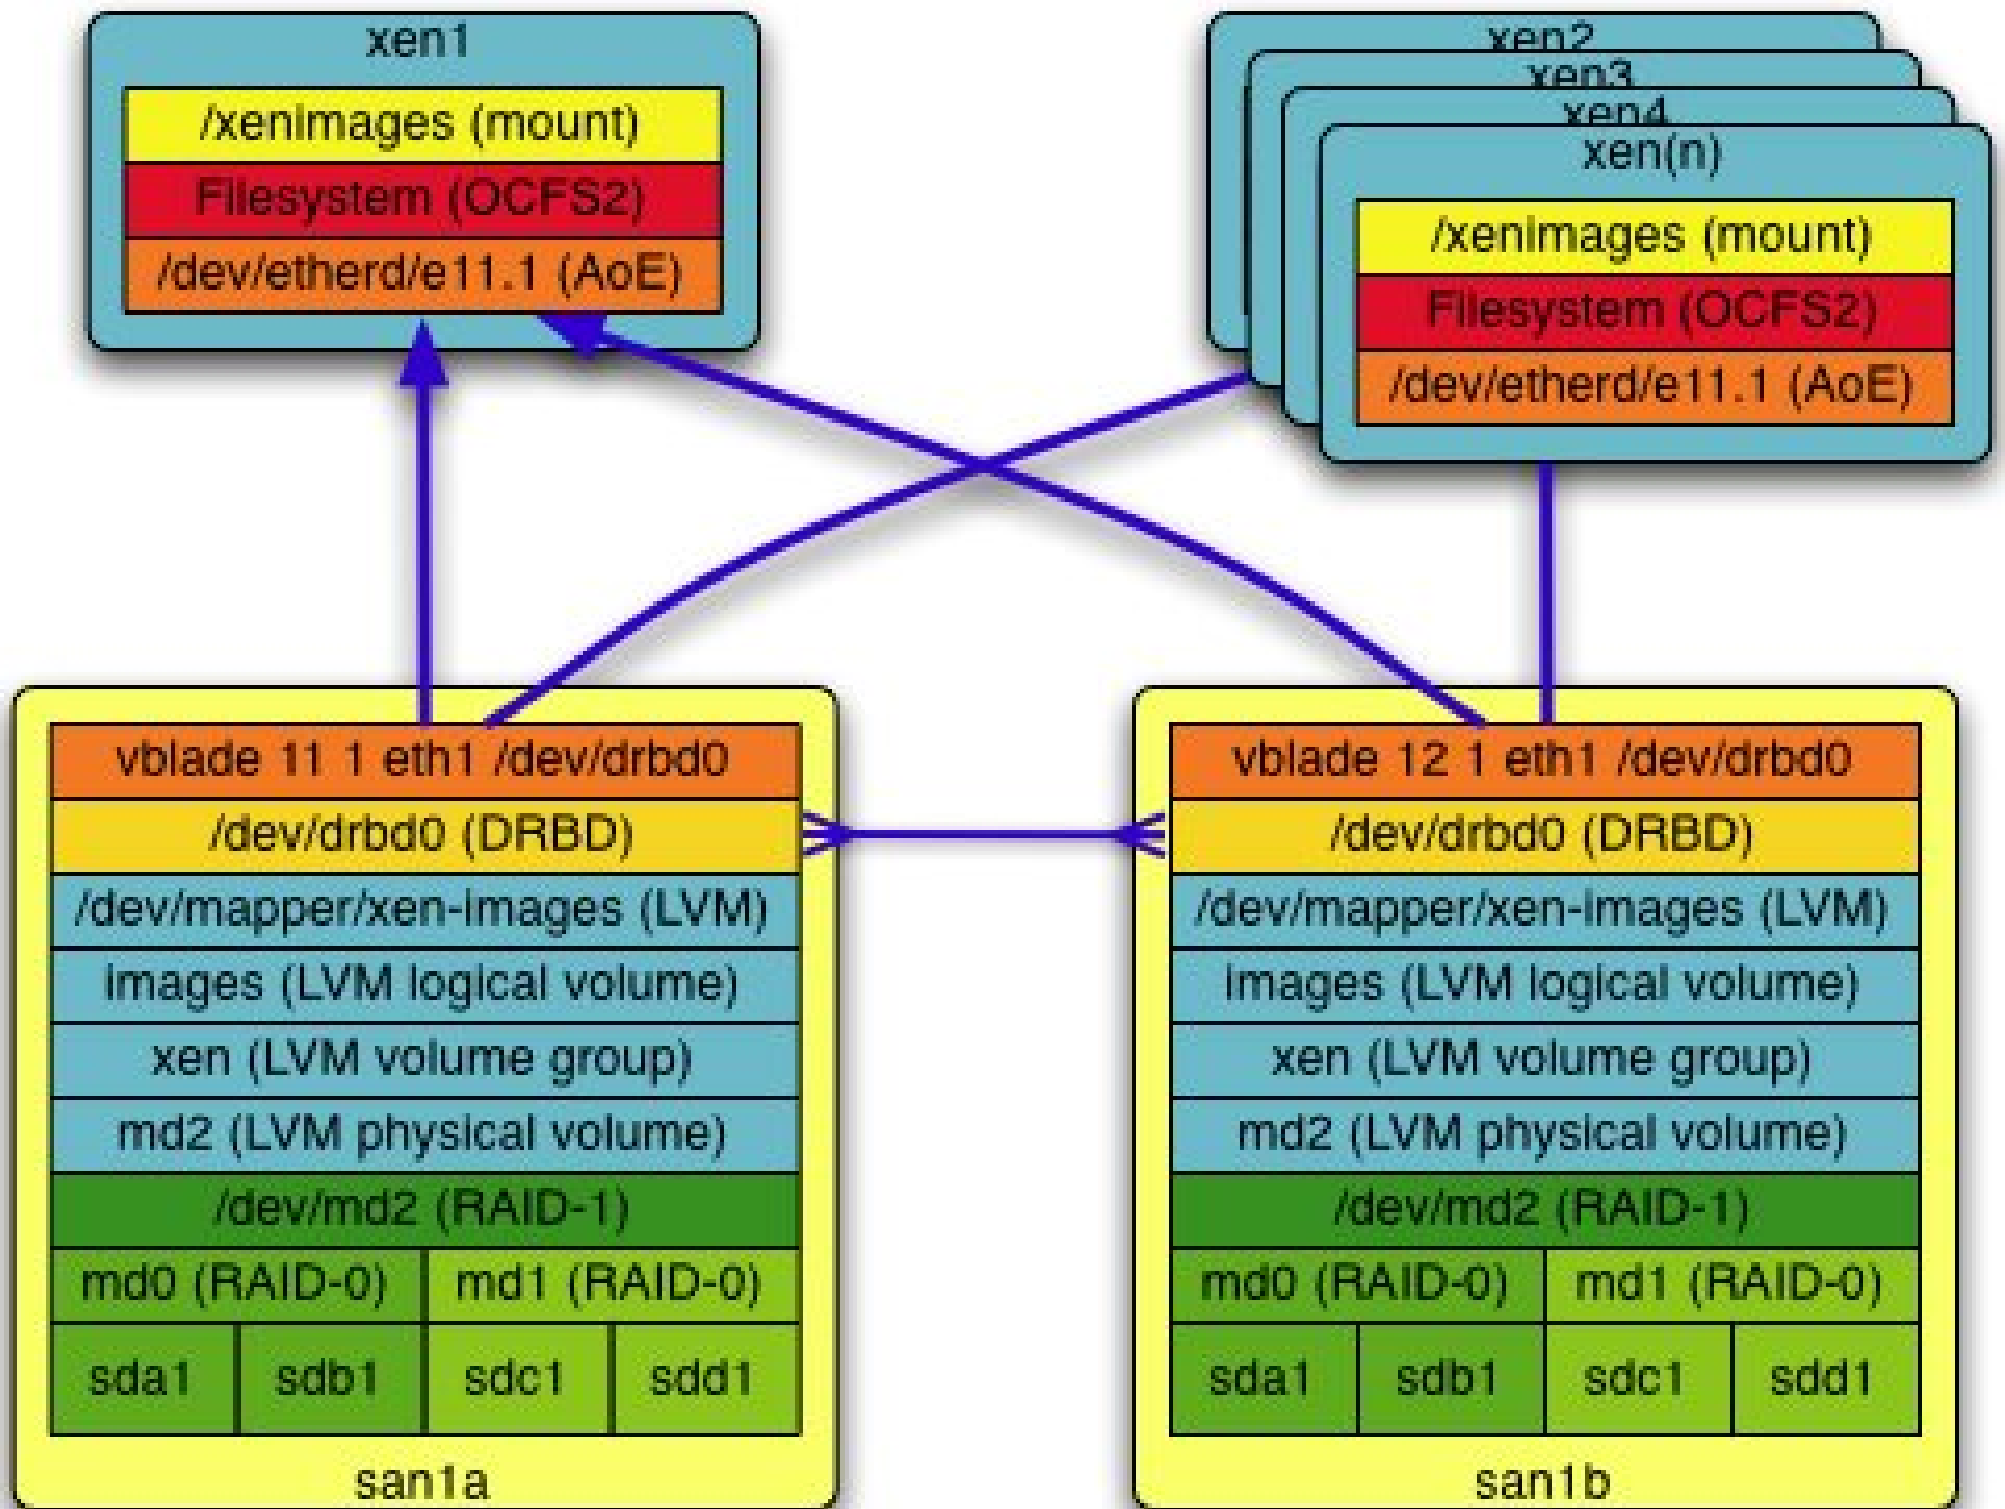
ATA-over-Ethernet

iSCSI

NBD

Loopback images

Local partitions




**Enough  
layers  
yet?**



**Breathe slowly.**

**Relax.**

**Go to your  
happy place.**



**I want my  
happy place,  
dammit!**



# Xen Image Manager



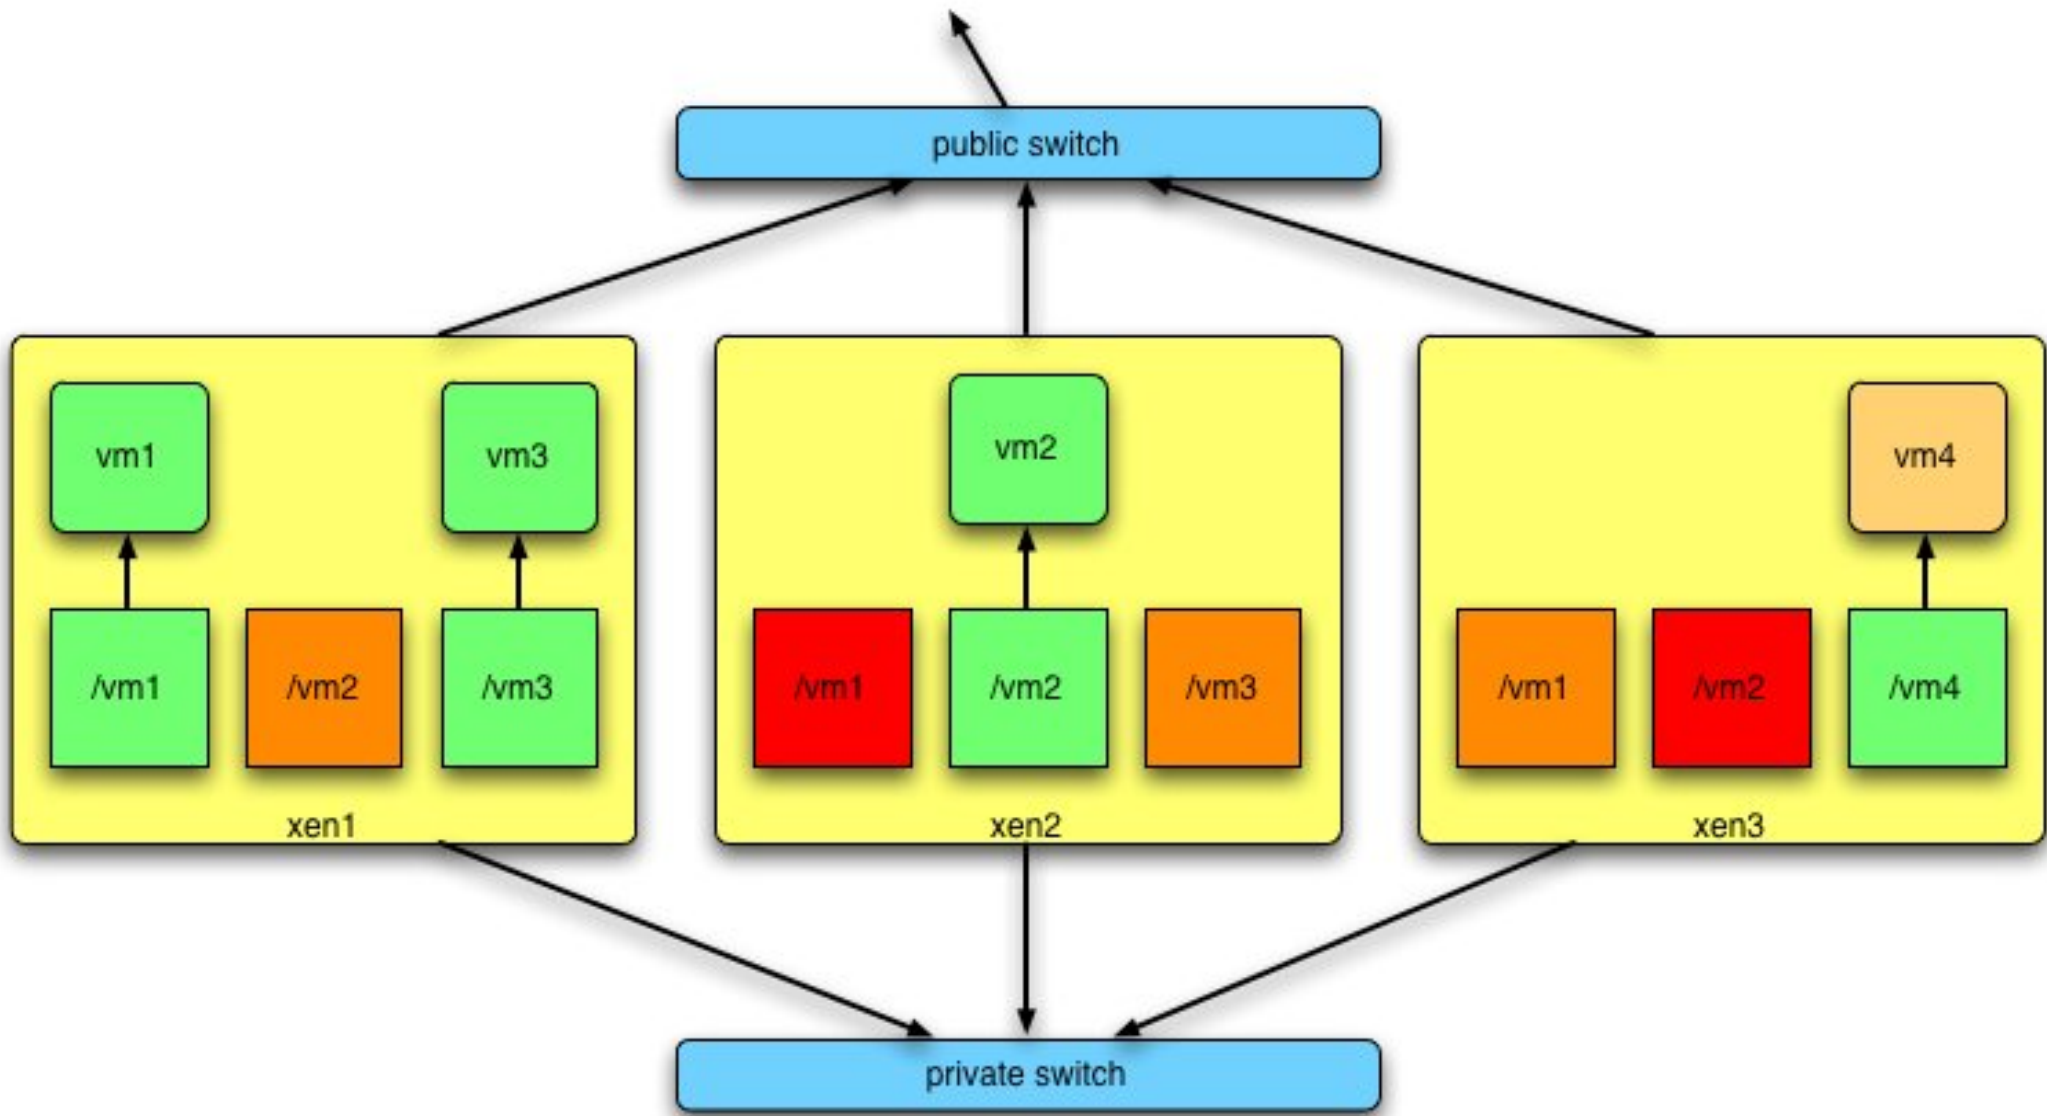
# vim xim



# xemacs xemax



# Xen External Manager (Archives) Xtension



# xen image manager



Removes reliance on external storage.

Stores VM images directly on Xen nodes using LVM.

Surveys nodes to find available resources.

Propagates images between nodes using LVM snapshots.

Encapsulates xm functionality for convenience.

# what works now



Building new VM images from a master:

```
build <DomGroup> <DomName>
```

Demolishing existing Vms:

```
demolish <DomGroup> <DomName>
```

Propagating a VM snapshot to other nodes:

```
propagate <DomGroup> <DomName> <Host>
```

# what kinda works



Surveying nodes to discover stored images,  
Surveying nodes to discover running VMs:

```
survey <Host>
```

# what doesn't work



Relocation of a running VM:

```
relocate <DomGroup> <DomName> <Host>
```

Automatic migration

Resource utilisation management

Sharing status / config between nodes

# problems, problems



## Rsync

- Does device nodes (yay!)
- Can't do remote to remote (boo!)
- Honours suid bit (yay!)
- Uses lots of memory and time (boo!)

## Unison

- Can't do device nodes (boo!)
- Can do remote to remote (yay!)
- Needs patch to honour suid bit (boo!)
- Uses little memory, super fast (yay!)

# config conundrum



We want to know

## Per Node:

- RAM capacity
- Disk capacity
- CPU capacity
- Availability
- Location: rack

- Actual count
- Active instance
- Actual runstate

## Per VM:

- RAM allocation
- Disk allocation
- CPU allocation
  
- Desired count
- Runstate

## Per Image:

- Image size
  
  
  
  
  
  
  
  
  
- Location: node
- Age

# what next?



Decide on config structure

~~Finalise decentralised execution model~~

Make it work real good and stuff

# xen image manager



# Thankyou :-)

Slides: <http://jon.oxer.com.au/talks>

Project: <http://jon.oxer.com.au/xim>

Contact: Jonathan Oxeer <[jon@ivt.com.au](mailto:jon@ivt.com.au)>