

Managing Xen with Quattor

Stephen Childs

Department of Computer Science

Trinity College Dublin

Quattor Workshop

March 12th 2007

Motivation

- Grid-Ireland has deployed service nodes on Xen since January 2005
- Configuration of Xen VMs largely manual
- Want to manage via Quattor
 - Store VM configurations centrally
 - Allow from-scratch installation of host & guests
 - Facilitate creation of replicas (in testbeds)

Components

- All integration: allow Kickstart installation of Xen VMs
- ncm-xen: Quattor component for generating guest VM configurations
- ncm-grub: support for Xen multi-boot

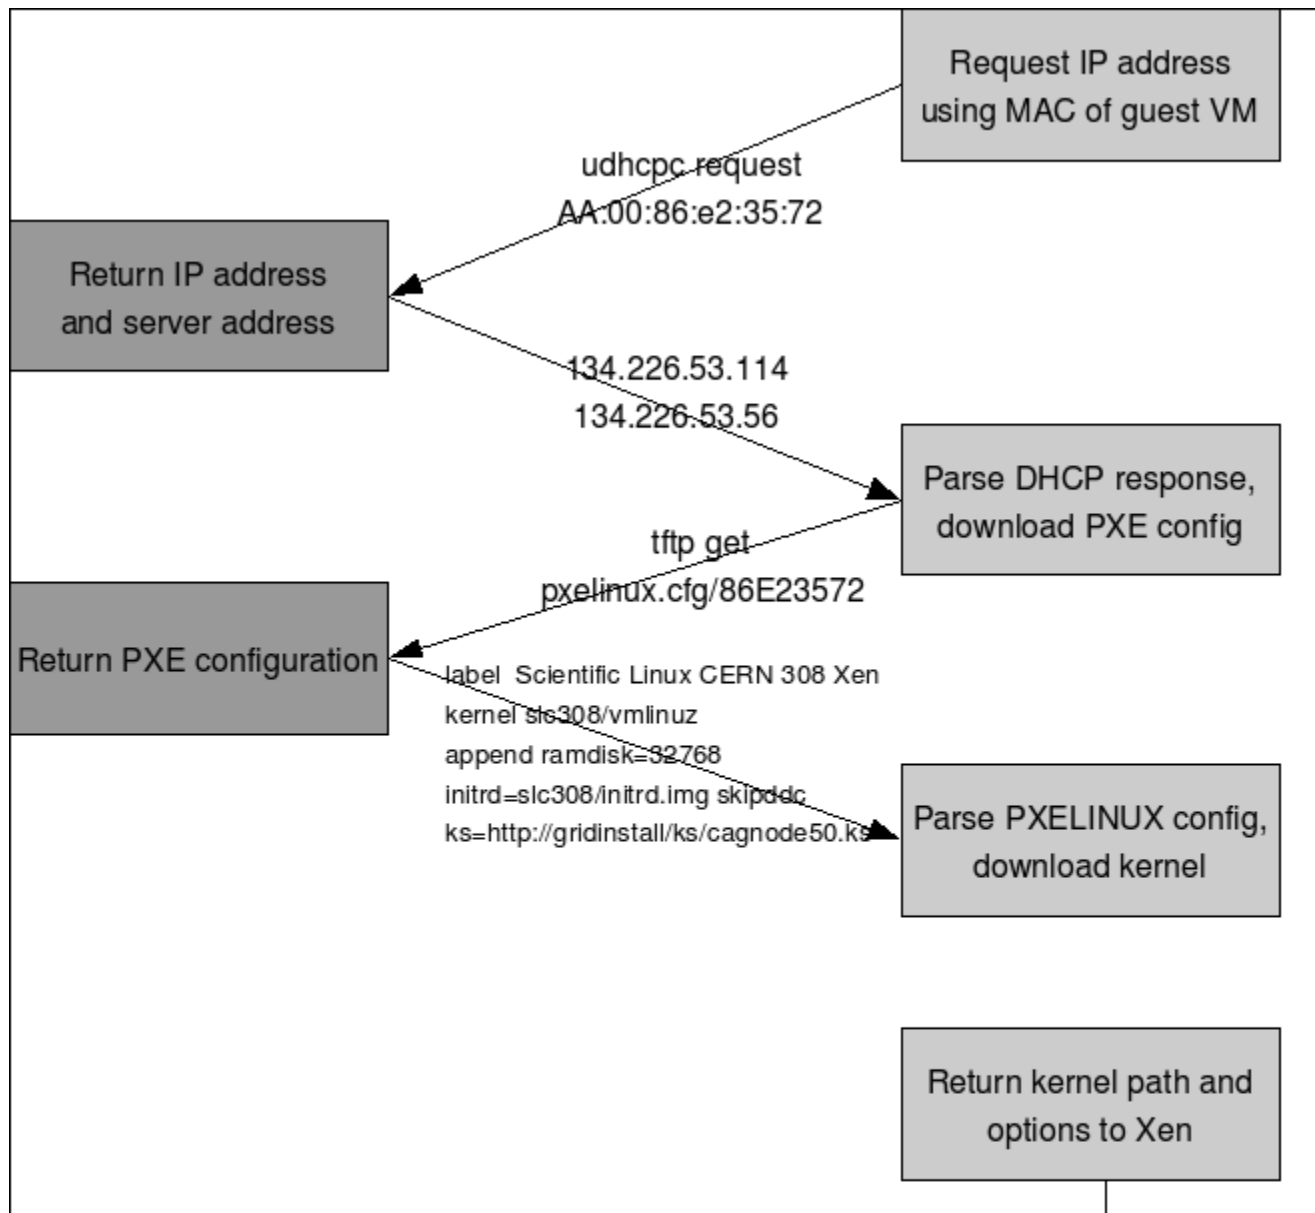
All integration

- Two types of Xen VM: para-virtualised (PVM) & hardware (fully) virtualised (HVM)
- HVMs can do native PXE with appropriate image
- PVMs can't do PXE – no network card
- Simplest option is have host do PXE on behalf of guests during startup – pypxeboot
- NB need patched Anaconda for Xen – available from Linux Support at CERN for SLC 3/4

pypxeboot

- Python script invoked by Xen during VM boot
- A Xen “bootloader” – retrieves kernel and config and outputs to Xen (see also pygrub)
- pypxeboot uses customised udhcpc to get IP address for guest's MAC address
- tftp client used to retrieve pxelinux config, kernel and initrd – downloaded and path passed to Xen
- Xen uses these to boot the VM

pypxeboot boot process



linux (kernel /var/lib/xen/kernel.9aBlaj)
(ramdisk /var/lib/xen/initrd.KDITiH)
(args 'ramdisk=32768 ksdevice=eth0
ks=http://gridinstall/ks/cagnode50.ks skipddc')



ncm-xen

- Schema describes Xen VM configuration:

```
"/software/components/xen/domains" = push(nlist(  
  "memory", 512,  
  "name", "cagnode50",  
  "disk", list(nlist("type", 'file',  
                    "path", '/var/xen-grid/cagnode50/fs/disk',  
                    "device", 'sda',  
                    "size", 10*GB,  
                    "rw", 'w'),  
              nlist("type", 'lvm',  
                    "hostdevice", 'xenvg',  
                    "hostvol", 'cagnode50-swap',  
                    "size", 6*GB,  
                    "device", 'hda3', "rw", 'w'))),  
  "vif", "[ 'mac=AA:00:86:e2:35:72' ]",  
  "bootloader", "/usr/bin/pypxeboot",  
  "bootargs", "vif[0]"  
));
```

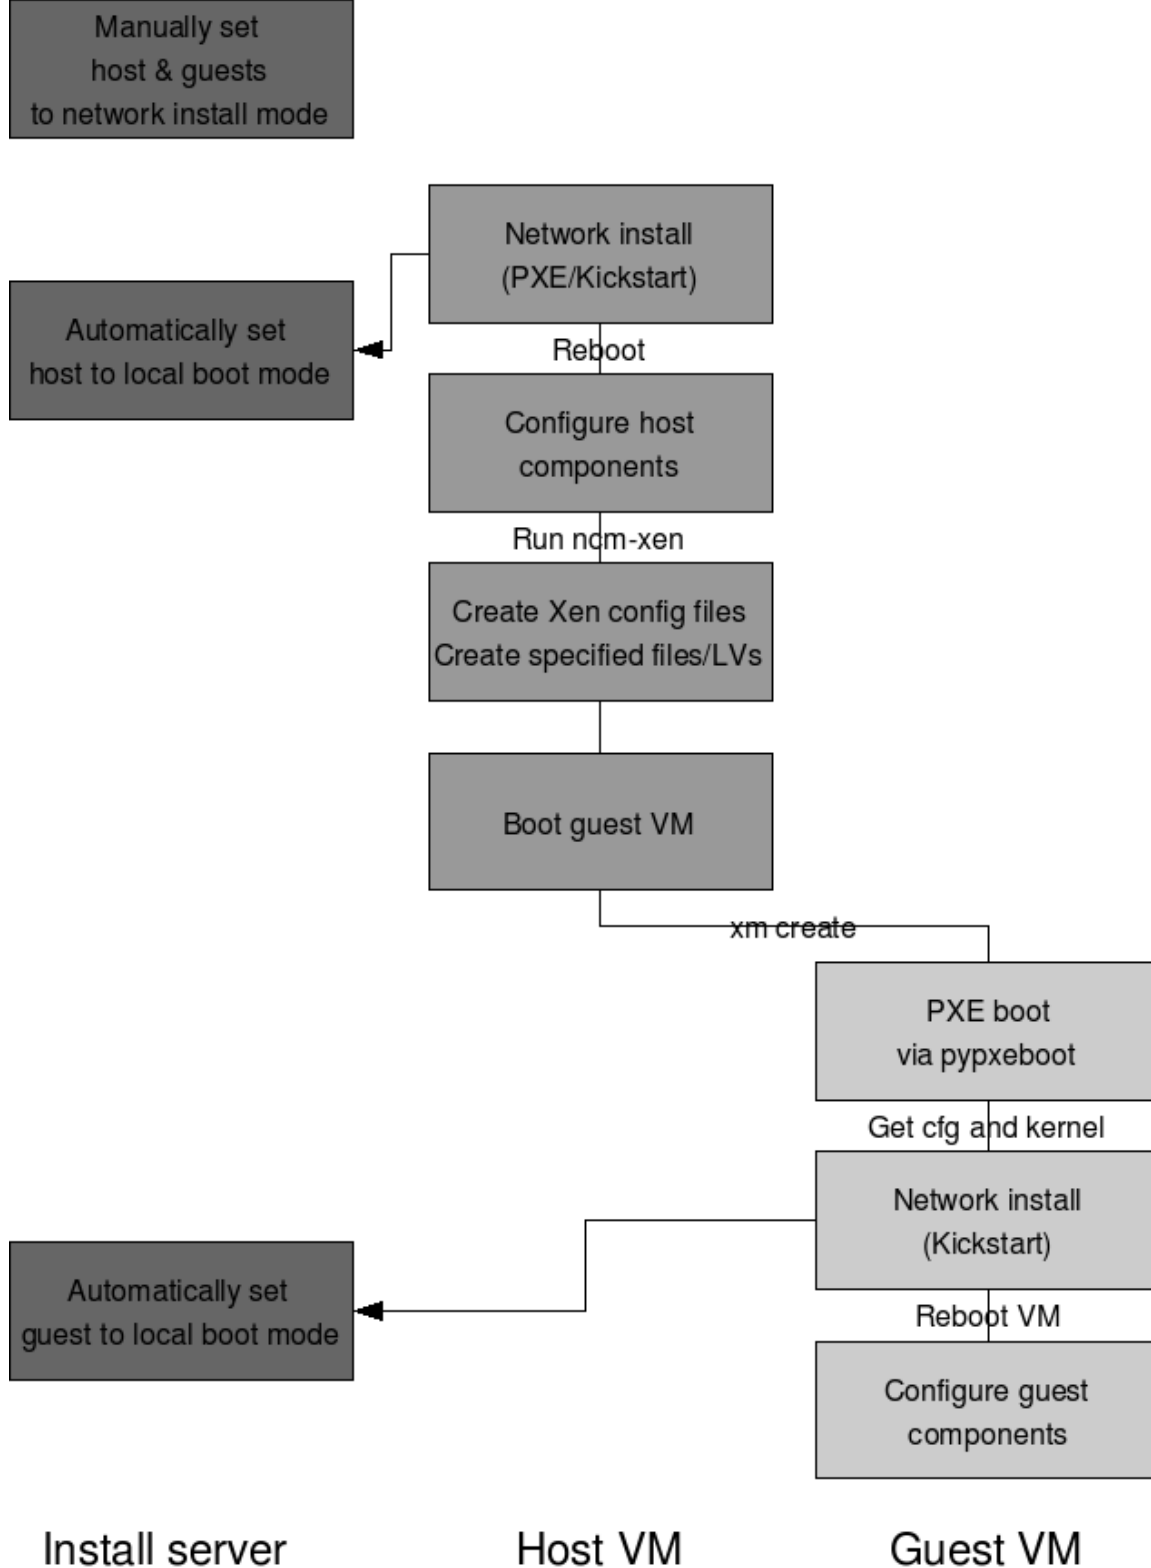
ncm-xen

- Most entries translated directly to Xen configuration options
- Disk entries processed differently
 - Xen configuration entries generated
 - Files or LVs specified but not existing are created

Automatic installation

- Combination of pypxeboot and ncm-xen allows full unattended install of host and guests
- Host boots first, installs via AII/Kickstart
- ncm-xen runs on host
 - Sets up specified files and LVs for VM FS
 - Creates Xen VM config files
 - Boots each VM – these install themselves via AII

Complete boot cycle



Summary

- pypxeboot submitted for inclusion in Xen tree
- ncm-xen 0.1.0 released for testing
- First serious use in Grid-Ireland soon (re-installation of complete gateway node)
- Paper submitted to VHPC '07
- Feedback welcome