The XEN of LCFG

Panagiotis Kritikakos
Computing Support Officer, EPCC
pkritika@epcc..ed.ac.uk
+44 131 650 7282

Overview

- Why
- lcfg-xen component
  - What it does
  - What it will do
  - Resources
- Host configuration
- Guest configuration
- Build summary
- In future
- Questions

Why

No existing support!

- Package list exists but a bit useless without a component
- lcfg-grub component to set up the Xen kernel
- lcfg-boot component should be used to enable daemons
- lcfg-file component to manage configuration files
- A mess of different components to get Xen managed

XEN component – What it does

The current available stable release

- Installs the Xen packages
  - Xen kernel
  - Xen utilities
  - Virtual manager
- Configures the machines to boot the Xen kernel
  - Sets the right Grub resources
- Controls the *xend* and *xendomains* daemons
  - start / stop
  - runlevels

XEN component – What it will do

In addition to the existing functionality, development is undertaken for the next stable release

- Managing the guest's configuration file
  - Controlling the resources as specified in the profile
- Partly management of the guest's disk image
  - Create it if it doesn't exist
  - Keep it untouched if it exists
- Automate boot and shutdown of guests when the component starts/stops
  - A resource will specify if the machine should do so

Host configuration

– On the MDP layer, all is needed is to add a header

```
#include <mdp/options/xen-host.h>
```

– The header:

- Adds the Xen packages to the package list
- Adds the lcfg-xen package to the package list
- Configures Grub for booting the Xen kernel
- Pulls in the default resources for the Xen component

A new VM specifications

We want a new virtual machine:

- Named *testVMxen*
- With a specific *UUID*
- With memory 1024MB
- With two disk images, 12 and 13GB respectively
- Name of the second disk *testVMxen2d*
- Specific MAC address
- Enable automatic booting with the start of the component
- Will use default resources for networking and disk images storing space

A new VM resources

```
!xen.virtualmachines    mADD(testVMxen)
!xen.name_testVMxen     mSET(testVMxen)
!xen.uuid_testVMxen     mSET(0a2040b3-3d85-49d4-accc-e38bc2e51a6d)
!xen.maxmem_testVMxen   mSET(1024)
!xen.memory_testVMxen   mSET(1024)
!xen.disk_testVMxen     mSET(12)
!xen.disktype_testVMxen       mSET(image)
!xen.disksec_testVMxen        mSET(13)
!xen.disktypesec_testVMxen    mSET(image)
!xen.disknamesec_testVMxen    mSET(testVMxen2d)
!xen.mac_testVMxen            mSET(00:1E:C9:53:22:ED)
!xen.boot_testVMxen           mSET(yes)
```

The resources

- *name* – The name of the virtual machine
- *uuid* – The associated uuid for the virtual machine
- *maxmem* – The maximum memory that will be allocated
- *memory* – The initial allocated memory
- *disk* – The size of the first disk in GB (will be needed only for image)
- *disktype* – The type of the first disk (image or physical)
- *disksec* – The size of the second disk in GB
- *disktypesec* – The type of the second disk
- *disknamesec* – The name of the disk image file
- *mac* – The MAC address associated with eth0 of the virtual machine
- *boot* – Defines if a VM will boot automatically when the component starts

More resources

Default resources (default values in the parenthesis)

– *vcpus* – Virtual CPUs to be used (1)
– *diskpath* – First disk image path (/var/lib/xen/images)
– *diskpathsec* – Second disk image path (/var/lib/xen/images)
– *bridge* – Xen network interface for eth0 (xenbr0)
– *script* – Xen networking script for eth0 (vif-bridge)

Other resources

– *macsec* – The MAC address for eth1
– *bridgesec* – Xen network interface for eth1 (xenbr1)
– *scriptsec* – Xen networking script for eth1 (vif-bridge)
– *cdpath* – Sets the ISO or physical CD/DVD path

Host profile

```
#define MDP_OS sl5_64
#include <lcfg/hw/dell_poweredge_1950.h>
!hardware.modlist mADD(mptsas)
hardware.mod_mptsas alias scsi_hostadapter mptsas

#define MDP_OVERRIDE__NFSCLIENT
#include <phys/options/epcc/nis_accounts.h>

/* Xen host and virtual machines */
#include <mdp/options/xen-host.h>

/* pe2900x1 */
!xen.virtualmachines            mADD(pe2900x1)
!xen.name_pe2900x1              mSET(pe2900x1)
!xen.uuid_pe2900x1              mSET(e050b5aa-4a58-4bcc-899f-9087180f8c0c)
!xen.maxmem_pe2900x1            mSET(512)
!xen.memory_pe2900x1            mSET(256)
!xen.disk_pe2900x1              mSET(20)
!xen.disktype_pe2900x1          mSET(image)
!xen.disksec_pe2900x1           mSET(4)
!xen.disktypesec_pe2900x1       mSET(image)
!xen.mac_pe2900x1               mSET(00:15:17:34:62:01)
!xen.boot_pe2900x1              mSET(yes)

!network.interfaces     mADD(eth0)
network.device_eth0     auto
!network.hwaddr_eth0    mSET(00:15:17:34:62:9C)
!network.ipaddr_eth0                    DHCP
!network.persistentdhclient_eth0        yes
```

Guest configuration

– Needs only the appropriate hardware header

```
#include <phys/hw/xen_vm.h>
```

Guest profile

```
#define MDP_OS sl5
#include <phys/hw/xen_vm.h>
#include <mdp/options/minimal.h>

#define _MYSQL_USER     mysql
#include <lcfg/options/mysql-server.h>

#include <phys/options/epcc/packagelists.h>
!profile.pkgcppopts     mADD(-DMYSQLSERVER)

!network.interfaces     mADD(eth0)
network.device_eth0     auto
!network.hwaddr_eth0    mSET(00:15:17:34:62:01)
!network.ipaddr_eth0                    DHCP
!network.persistentdhclient_eth0        yes
```

Build summary

Host build

- EdLAN DB and DHCP entry
- Create LCFG profile including the Xen header
- Build host using PIE
- Ready to host new guests

Guest build

- EdLAN DB and DHCP entry
- Create LCFG profile including the Xen hardware header
- Create configuration file using the Xen component on the host
- Build guest using PIE
- Ready to deploy services

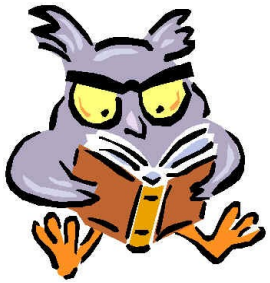

In future...

- Manage Xen network scripts
- Properly managing individual guest resources at run time
- Use of LVM for guest's storage
- Specifying, building and managing a group of guests
  - Potential for virtual clusters using current and future multi-core hosts
- PIE support for x86_64 guests
- PIE para-virtualisation support

Questions

Links

Presentation:

http://www.epcc.ed.ac.uk/~pkritika/lcfg/xen_lcfg.pdf

Wiki pages:

https://wiki.lcfg.org/bin/view/LCFG/XENonLCFG

https://www.wiki.ed.ac.uk/display/DSwiki/Xen+and+LCFG+Investigation