Getting More From Your Virtual Machine

Kyrre Begnum, Oslo Univ. College, Norway

John Sechrest, Oregon State University

Steven Jenkins, East Tennessee St. University

Introduction

Too many students, Too Few Machines

=> Toy Problems

=> Trivial learning

=> Rigid Support practices

=> Homogeneous Problems

=> High Support Overhead

Virtual Machines

- *Every physical system has multiple virtual instances running on it.*
- *Each virtual instance has network access*
- *Each instance is a full linux system*
- *Virtual services can be provided by either Xen or UML*

Benefits of Virtualization

- *One computer -> Many Virtual machines*
  (Kyrre used 1 machine for 84 virtual instances)
- *Each student can have several virtual machines*
- *Restarting a system is easy*
- *Students get full control of system*

More Benefits of Virtualization

- *Templates provide preconfigured environments*
- *You can create networks of virtual machines*
- *Automation reduces overhead*

MLN

- MLN is a wrapper around Xen or UML which automates virtual system management
  - Configure new Virtual Machines
  - Configure networks of many machines
  - Separate Spaces for each student
  - Configure many virtual machine instances over many different physical machines with one config file
  - Supports sharing networks between students

MLN projects

- *Logical Groups of machines are called Projects*
- *Individual users can create projects with full administrative rights*
- *Projects are defined by a text config file*
- *Virtual machine filesystems are created from filesystem templates*

MLN on Linux

- *MLN runs on Linux*
- *It supports several different systems:*
  - *RedHat*
  - *Debian*
  - *Ubuntu*
  - *BusyBox*

Configuration options

- *MLN configuration supports setting system variables for virtual machines*
  - *Filesystem Template*
  - *users + passwds*
  - *startup commands*
  - *network parameters*

Project in a Box

- *Provide a working environment for a team of students.*
- *Pre-configured with Web server*
- *Including Media Wiki*
- *and Blog software*
- *Student control of system config*

MLN Configuration

```
global {    project group1   }
  #include superclasses.mln
host cms-server-group1 {
   superclass projectVM
   template blimp.ext3
   root_passwd *somthinghere*1
  users {joe *apasswd* }
  network eth0 { address 192.182.19.10 }
  }
```

Student Benefits

- *Move student attention to task instead of overhead*
- *Students can choose tools they want to use*
- *Students have direct control of the working environment*
- *Students have access to services over network*

Managing Large Scale Networks

- Used in Firewalls and intrusion course
- Network + DMZ + firewall
  - 12 student groups
  - 6 virtual machines/group
  - each group reboots seperately

Large Scale Configuration

- *global { project group1*
  $owner = group1
  address = 128.193.36.17
  $mount = /home/group1/ms004a
  }
  #include common.mln

217 lines of config file produced 84 running systems

Hardware

- *Dual Xeon 3.0 Ghz*
- *6 GB of memory*
- *SCSI Raid disks*
- *direct network access*

Class Experiences

- *Run in classes at OSU and Oslo (and sweden and netherlands)*
- *Creating lab configurations is easier*
- *process scales to more groups*
- *students able to restart systems*
- *supports diverse scenarios in the same assignment*

So many machines, so little time

Time: *With physical machines, lab configuration took significant time.*

- *A lab with 15 machines used to take 2 weeks now it takes 10 minute*

Space: *The need for space for special labs is reduced*

Money: *The cost of a new configuration is greatly reduced.*

Benfits

- *Direct student control -> reduce IT staff overhead*
- *write once & execute many times -> increases reuse of activities*
- *Programatic control of systems reduces effort to support different scenarios*
- *System security can be managed easier*
- *More services for less effort*

Future Work

- *MLN is growing plug-ins to support more features*
- *Automatic configuration and self configuration tools are part of the underlying research.*

MLN Configuration

- *global { project p1 }*
  superclass host { size 2500M
  memory 128M }
  network eth0 {
  switch lan
  netmask 255.255.255. } }

host one {
superclass host
service_host backend1
roles { webserver(mysite.com) }}

Resources

- *Cfengine - http://www.cfengine.org*
- *MLN - http://www.sf.net/projects/mln*
- *Xen - http://www.xensource.com*

Sechrest@eecs.oregonstate.edu