

System Administrators

The Setup and Maintenance Guide details hardware and software requirements, how to bring up a new HUB, upgrade system software, etc.

Installation

What is HUBzero?

HUBzero is a platform used to create dynamic web sites for scientific research and educational activities. With HUBzero, you can easily publish your research software and related educational materials on the web. Powerful middleware serves up interactive simulation and modeling tools via your web browser. These tools can connect you with rendering farms and powerful Grid computing resources.

Minimum System Requirements

HUBzero installations require one or more dedicated physical hosts running Debian GNU/Linux 5.0.

Other distributions might theoretically work with some modification, although they would be totally unsupported.

A typical starter HUBzero installation might consist of a single physical server with dual 64-bit quad-core CPUs, 16 Gigabytes of RAM and a terabyte of disk

It is possible to run HUBzero inside of a virtual machine such as ones created by VMware and VirtualBox. While fully functional there will be significant performance and resource limitations in such an environment.

Target Audience

This document and the installation of a HUBzero system has a target audience of experienced Linux administrators (preferably experienced with Debian GNU/Linux).

What's Included

HUBzero is composed of the following packages (subject to change):

<i>Package Name</i>	<i>Purpose</i>
hubzero-addrepo	Creates project areas for tool development
hubzero-apache2	Apache 2.2 Site Configuration Files for HUBzero
hubzero-app	HUBzero App Installer
hubzero-app-invoke	HUBzero application (tool) invocation scripts
hubzero-app-workspace	HUBzero App providing a lightweight Linux desktop, for app/tool development

SYSTEM ADMINISTRATORS

hubzero-cms	The HUBzero Content Management System (based on Joomla! framework)
hubzero-cms-joomla	Joomla! framework used by HUBzero
hubzero-cms-setup	Customizes the HUBzero CMS database
hubzero-config	Configures a HUBzero server and provides an automated installation script
hubzero-expire-sessions	Expires unused app/tool sessions
hubzero-filexfer	Transfer files between App Sessions and user's desktop
hubzero-filexfer-xlate	Support daemon for the filexfer program
hubzero-firewall	HUBzero firewall that protects app/tool sessions
hubzero-icewm	Linux ICE window manager configuration, used in workspaces
hubzero-icewm-captive	Linux ICE window manager specially crafted to support tools in a sessions
hubzero-icewm-themes	The HUBzero Linux ICE window manager theme, used in workspaces
hubzero-mw-client	HUBzero middleware - client
hubzero-mw-service	HUBzero middleware - execution host session manager
hubzero-mw-session	HUBzero middleware - per session tools
hubzero-mysql	MySQL configuration package for HUBzero
hubzero-openldap	OpenLDAP configuration package for HUBzero
hubzero-python	HUBzero python API module
hubzero-rapture	The Rapid APPLication infrastrucTURE toolkit for building scientific tools
hubzero-rapture-session	Session support packages for Rappture
hubzero-ratpoison-captive	Linux window manager, used in app/tool sessions
hubzero-submit-client	The session based part of the job submission server
hubzero-submit-distributor	Part of the job submission server
hubzero-submit-server	Part of the job submission server
hubzero-telequotad	Disk quota monitor
hubzero-texvc	Helper utility to generate math fomulas for wiki markup
hubzero-trac-mysqldatauthz	Plug-in for MySQL user auth in project development areas
hubzero-twm-captive	Linux TWM window manager, used in app/tool sessions
hubzero-use	Command for configuring the environment within a workspace
hubzero-usermap	File permission mapping FUSE filesystem used by WebDAV
hubzero-vncproxy	Routes vnc between web server and app/tool session

icewm	Modified Linux ICE window manager, used in workspaces
tightvnc-java	Modified VNC Client that receives app/tool sessions within a web browser
vnc4server	Modified VNC Server that sends app/tool sessions to the web browser

Source Code

You'll find most of the source code within the web root of a working hub. But you can get source code for the middleware and all other parts by installing source code via the package mechanism. Please refer to [Section 1.2.7](#) for instructions.

Debian Build

Install Basic Operating System

The latest version of Debian GNU/Linux 5.0 (5.0.8 as of this writing at <http://www.debian.org/releases/lenny/debian-installer/>) should be installed on each physical host used by a HUBzero installation. HUBzero has packaging support for i386 (32bit) and amd64 (64bit) Intel architectures. However the 32bit version gets more limited testing. Debian 6.0 is not supported at this time.

It is strongly recommended that you install from one of the ISO disk images listed in the above link. VM images (other than those distributed by HUBzero) and other boutique Debian installs are known to fail due to incomplete dependencies in the current HUBzero packaging.

Locale Settings

At the initial boot menu choose the command line install.

Choose your language, country, and keyboard (English, United States, and American English will be the default)

The next step is to enter the hostname and domain name for the system.

Choose the time zone for the system (the time zones will be limited to the country that was selected previously).

Disk Partitions

The recommended minimum for partitions sizes are 100 GB bootable partition for the root filesystem, 50 GB export filesystem for each HUBzero server, and an appropriately sized swap partition.

Partitioning method: (Guided - use entire disk)

Select disk to partition: (SCSI1 (0,0,0))

- Partitioning scheme: (All files in one partition)
- Select each partition and delete
- Select #1 - then Delete the partition
- Select #5 - then Delete the partition
- Select free space
- create a new partition
- set size (ex: 100 GB)

- type is primary
- location is beginning
- Done setting up partition
- Select free space
- create a new partition
- set size (ex: 1 GB)
- type is primary
- location is beginning
- Select Use As: set to swap area
- Done setting up partition
- Select free space
- create a new partition
- use all available space left
- type is primary
- Select mount point, change to do not mount it

Done setting up partition

Finish partitioning and write changes to disk

Want to return - say No

Write the changes to disks? Select Yes

Base System Accounts

The install will begin installing the base system then ask for the creation of the root password and to create a new user. It is suggested to skip the step of creating a new user.

- Skip the step of creating a new user
- Select (takes to root password screen)
- Select (takes to Debian installer main menu)
- Select Configure the package manager. (next one down)

Package Manager

Choose the package manager archive to use for downloads. ftp.us.debian.org is a good choice for the archive mirror for installs based in the U.S. and leave the HTTP proxy information blank unless your network requires it.

The install will begin retrieving a list of package options and configure base components. The

SYSTEM ADMINISTRATORS

install may prompt to participate in the package usage survey. The decision to participate or not will not affect the install and is up to the individual system administrator.

In the Software Selection option, deselect "Desktop environment" but leave "Standard system" then choose to install the GRUB boot loader to the master boot record and finish the installation (when the installation is complete your system will reboot into a minimal Debian GNU/Linux system).

Debian GNU/Linux

Install Basic Operating System

The latest version of [Debian GNU/Linux 5.0](#) (5.0.8 as of this writing) should be installed on each physical host used by a HUBzero installation.

To install Debian GNU/Linux, you can easily [obtain a copy](#), and then follow the [installation instructions](#) for your architecture.

Installing Debian GNU/Linux using a a small bootable [CD](#) is the simplest method.

When installing Debian GNU/Linux be sure to do the following:

- Ensure the disk(s) are partitioned to have at least:
 - A bootable partition at least 100.0 GB in size for the root filesystem.
 - An empty partition at least 50.0 GB in size (note the device name of this partition for later)
 - An appropriately sized swap partition.
- When prompted to select an installation package just select "Standard System", other packages will be added later

When the installation is complete your system will reboot into a minimal Debian GNU/Linux system.

Don't forget to remove your installation media and/or change your server's boot media order if you changed them prior to installation.

Set hostname

Optional. If you didn't specify the fully qualified domain name when running setup you will need to set it here.

HUBzero expects the ``hostname`` (and ``hostname -f``) command to return the fully qualified hostname for the system.

```
# hostname myhub.org
```

To make the change permanent you must also edit the file `/etc/hostname`, which can be done simply with:

```
# echo "myhub.org" > /etc/hostname
```


Fix hosts

Now edit `/etc/hosts` by making sure that a line exists that looks like

```
192.168.2.10    full-host-name optional-short-alias-of host
```

The number 192.168.2.10 must be replaced by the current IP number of your server. This is normally reported by the `'ifconfig eth0'` command.

Do not remove the line that looks like

```
127.0.0.1      localhost
```

Delete local user

If you created a local user account when installing the operating system now would be a good time to delete it before it causes you future problems.

```
# deluser username
```

Configure Networking

Optional. If you didn't configure networking during installation you will need to do so now.

For help with networking setup try this [link](#).

Setting up your IP address.

The IP addresses associated with any network cards you might have are read from the file `/etc/network/interfaces`. This file has documentation you can read with:

```
# man interfaces
```

A sample entry for a machine with a static address would look something like this:

```
# The loopback network interface
```

SYSTEM ADMINISTRATORS

```
auto lo
iface lo inet loopback

# The primary network interface
auto eth0
iface eth0 inet static
    address 192.168.1.90
    gateway 192.168.1.1
    netmask 255.255.255.0
    network 192.168.1.0
    broadcast 192.168.1.255
```

Here we've setup the IP addresss, the default gateway, and the netmask.

For a machine running DHCP the setup would look much simpler:

```
# The loopback network interface
auto lo
iface lo inet loopback

# The primary network interface - use DHCP to find our address
auto eth0
iface eth0 inet dhcp
```

(If you're using a DHCP based setup you must have a DHCP client package installed - usually one of pump or dhcp-client.)

If you make changes to this file you can cause them to take effect by running:

```
# /etc/init.d/networking restart
```

Setting up DNS

Use whatever nameserver and other options as recommended by your ISP. If you used DHCP to set up networking it is likely this has already been set.

When it comes to DNS setup Debian doesn't differ from other distributions. To cause your machine to consult with a particular server for name lookups you simply add their addresses to

/etc/resolv.conf.

For example a machine which should perform lookups from the DNS server at IP address 192.168.1.10 would have a resolv.conf file looking like this:

```
nameserver 192.168.1.10
```

Configure Advanced Package Tool

Now configure what debian distribution mirror to use and the location of the HUBzero package repository by editing /etc/apt/sources.list to look like:

```
deb http://ftp.us.debian.org/debian/ lenny main
deb-src http://ftp.us.debian.org/debian/ lenny main

deb http://security.debian.org/ lenny/updates main
deb-src http://security.debian.org/ lenny/updates main

deb http://volatile.debian.org/debian-volatile lenny/volatile main
deb-src http://volatile.debian.org/debian-volatile lenny/volatile main

deb http://packages.hubzero.org/deb buck main contrib non-free
deb-src http://packages.hubzero.org/deb buck main contrib non-free
```

You will need to get and install the hubzero archive key to be able to verify packages from the hubzero archive:

```
# wget http://packages.hubzero.org/deb/hubzero-signing-
key.asc -q -O - | apt-key add -
```

Once the public key for <http://packages.hubzero.org> has been install you can then upgrade the current packages to their latest releases.

```
# apt-get update
# apt-get upgrade
```

SSH

Next we install fail2ban and ssh

```
# apt-get install fail2ban ssh
```

At this point you can continue configuration and setup remotely if that is more convenient.

Enable OpenVZ

If you are installing this in a VirtualBox VM you must enable PAE/NX support. Go to system -> processor of your VM, select "Enable PAE/NX".

To use OpenVZ you must use an OpenVZ enabled kernel which is easily installed.

HUBzero makes extensive use of [OpenVZ](#) containers so it is recommended to just use the OpenVZ enabled kernel on all HUBzero servers. To install a 64 bit kernel run the command:

```
# apt-get install linux-image-2.6-openvz-amd64
```

For 32 bit kernels, run the command:

```
# apt-get install linux-image-2.6-openvz-686
```

You will need to reboot the server to activate the new kernel.

```
# reboot
```

Once you have rebooted you can verify the new kernel is active

```
# uname -a
Linux myhub.hubzero.org 2.6.26-2-openvz-
amd64 #1 SMP Thu Nov 5 03:06:00 UTC 2009 x86_64 GNU/Linux
```

or for 32 bit kernels

```
# uname -a
Linux myhub.hubzero.org 2.6.26-2-openvz-686 #1 SMP Thu Nov 5 03:06:00
UTC 2009 x86_64 GNU/Linux
```

With the new kernel active you can optionally remove the old one. For 64 bit kernels run:

```
# apt-get purge linux-image-2.6.26-2-amd64
```

or for 32 bit kernels run:

```
# apt-get purge linux-image-2.6.26-2-686
```

Prepare Filesystem

The root filesystem (/) runs with quotas disabled and contains the primary operating system for the server and for each OpenVZ container hosted on the server.

Each HUBzero server may use an addition partition for use appropriate to the function of the server (web document root, project data, home directories, etc).

If you did not create an empty partition during setup, create one now using your favorite disk partitioning tool. Be sure to note the device name for the partition you create as it will be used below.

Once you have an empty partition ready we can install a filesystem. Replace "/dev/PART" with the device name for the empty partition you have created (e.g., /dev/sda2). The command "fdisk -l" will list all partitions the system knows about.

```
# mke2fs -j /dev/PART
# e2fsck -f -C 0 /dev/PART
# mkdir /export
```

Then make sure the following line appears in /etc/fstab

SYSTEM ADMINISTRATORS

```
/dev/PART    /export      ext3         defaults,quota,errors=remount-  
ro          0           2
```

Then mount the new filesystem

```
# mount /export
```

Bind mount /home

Create a 'home' directory in our new /export filesystem. move the contents of the default home directory to the new location, then bind mount new location over the old.

```
# mkdir -p /export/home/myhub  
# mv /home/* /export/home  
# mount --bind /export/home /home
```

Bind mount /opt

Currently HUBzero uses the /opt directory for storing subversion and trac data as well as some of hubzero supporting software as well. We recognize this may not be the best organization.

Create a 'opt' directory in our new /export filesystem. move the contents of the default /opt directory to the new location, then bind mount new location over the old.

```
# mkdir /export/opt  
# mv /opt/* /export/opt  
# mount --bind /export/opt /opt
```

Bind mount /apps

Currently HUBzero uses the /apps directory for storing installed tools and other software that needs to be available to each execution container.

Create a 'apps' directory in our new /export filesystem and in the root filesystem. Then bind mount /export/apps over /apps.

```
# mkdir /export/apps
# mkdir /apps
# mount --bind /export/apps /apps
```

Bind mount /www

HUBzero uses the /www directory for storing the document root and supporting directories needed by the web server.

Create a 'www' directory in our new /export filesystem and in the root filesystem. Then bind mount /export/www over /www.

```
# mkdir -p /export/www
# mkdir /www
# mount --bind /export/www /www
```

Update /etc/fstab

Now edit **/etc/fstab** with the bind mounts we created above by adding the following lines

```
/export/opt /opt          none    bind,defaults
0          0
/export/apps /apps          none    bind,defaults
0          0
/export/home /home          none    bind,defaults
0          0
/export/www  /www           none    bind,defaults
0          0
```

Hubzero Installation

Hubzero Software Installation

This section describes the automated installation of the HUBzero software. It assumes that the server has been setup with a basic Debian installation as described in the previous sections of the documentation.

Note It is very important that the server hostname be described properly in the file `/etc/hostname`, that the server's fully qualified domain name be accurately reported by the `'hostname -f'` command and that the file `/etc/hosts` contain an accurate mapping of the server's IP address to fully qualified hostname. Please refer to the previous section of the documentation for details.

To begin the process of installing HUBzero software, run the following command:

```
# apt-get install hubzero-config
```

This package will ask 12 questions - 6 pertaining to system variables that need to be set and another 6 passwords needed by the system. If the fully qualified domain name is set properly, 5 of the first 6 questions will contain the proper value by default. The other question is for an email address that will be used in several places for system notification messages.

The six passwords all require some kind of input. Note that the final password is for the super administrator *admin* user. This is the only one that will normally be used by the site maintainer.

Once the questions are answered the package will leave a script called *hz-install* in the `/root` directory. This script will use the information previously entered to fully construct a working HUBzero server. Run the command by typing:

```
# bash ./hz-install
```

The software installation process will then start and can take quite a while to finish depending on the server's network speed. Once the process is complete, you should be able to access your HUBzero server using a browser. You can log in with the *admin* account and begin customizing your site.

Mail

Install

We need to reconfigure exim4 to enable outgoing email (exim4 got installed earlier as a prerequisite for the mysql server).

```
# dpkg-reconfigure exim4-config
```

- Select "internet site; mail is sent and received directly using SMTP" then configure as appropriate for your site.
- Enter the FQDN of your site when asked
- Listen on all IP addresses (i.e., make list blank)
- Other destinations for which mail is accepted: should be made blank
- Domains to relay mail for: should be made blank
- Machines to relay mail for: should be made blank
- Keep number of DNS-queries minimal (Dial-on-Demand): No
- Delivery method for local mail: mbox format in /var/mail/
- Split configuration into small files? No

This is just an example. Mail should be configured however the site needs. The CMS just expects to be able send outgoing email.

Rappture

Install

The Rappture application is install in the apps directory along with proper links.

```
# apt-get install hubzero-rappture
```

Session Installation

Rappture is used from inside a container and needs several other packages installed to allow use of all its features. This process has been simplified by using the hubzero-rappture-session which only contains the dependencies needed to pull in these other packages.

```
# chroot /var/lib/vz/template/debian-5.0-amd64-maxwell
# apt-get update
# apt-get upgrade
# apt-get install hubzero-rappture-session
```

This is also a good time to add some default paths to your session environment so that it doesn't need to be whenever someone logs in. Modify the /etc/profile file as follows:

Change

```
if [ "`id -u`" -eq 0 ]; then
    PATH="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin"
else
    PATH="/usr/local/bin:/usr/bin:/bin:/usr/games"
fi
```

To

```
if [ "`id -u`" -eq 0 ]; then
    PATH="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin"
else
    PATH="/apps/rappture/bin:/apps/bin:
/usr/local/bin:/usr/bin:/bin:/usr/games"
```

```
fi
```

Be sure to exit the chroot environment when you are done:

```
# exit
```

A workspace may need to be opened and closed a few times before the changes to the session template appear in a workspace.

Test

Rappture comes with several demonstration scripts that can effectively test many parts of the package. These demonstrations must be copied to a user's home directory within a workspace before running.

```
$ mkdir examples
$ cp -r /apps/rappture/examples/* examples/.
$ cd examples
$ ./demo.bash
```

A window should open on the workspace showing that part of the demonstration. Close that window to see the next demonstration. Some demonstrations may need something inputted to work properly (such as the graphing calculator).

Filexfer

Install filexfer

Install the filexfer packages

```
# apt-get install hubzero-filexfer
```

```
# apt-get install hubzero-filexfer-xlate
```

Configure Apache for filexfer

Modify the hub site file at `/etc/apache2/sites-available/hub-ssl`

The apache hub site configuration files are preconfigured to support this. Uncommented the two highlighted lines as shown below. Note that the relative placement of these lines is important.

Edit `/etc/apache2/sites-available/hub-ssl`

```
<VirtualHost *:443>
    RewriteEngine    on
    RewriteMap       xlate          prg:/usr/lib/mw/bin/filexfer-
xlate
    ...
    ...
    ...
    <Directory /www/myhub>
        Options FollowSymLinks
        AllowOverride None
        Order allow,deny
        Allow from all

        RewriteEngine    On

RewriteRule    ^filexfer/(.*) ${xlate:$1|nothing} [P,QSA,L]
```

Then restart apache

SYSTEM ADMINISTRATORS

```
# /etc/init.d/apache2 restart
```

Submit

Introduction

The submit command provides a means for HUB end users to execute applications on remote resources. The end user is not required to have knowledge of remote job submission mechanics. Jobs can be submitted to traditional queued batch systems including PBS and Condor.

Installation

```
# apt-get install hubzero-app-submit
# apt-get install hubzero-submit-server
# apt-get install hubzero-submit-distributor
```

At completion of the apt-get install commands several files will be located in the directory /opt/submit. Excluding python files the directory listing should like the following:

Configuration

submit provides a mechanism to execute jobs on machines outside the HUB domain. To accomplish this feat some configuration is required on the HUB and some additional software must be installed and configured on hosts in remote domains. Before attempting to configure submit it is necessary to obtain access to the target remote domain(s). The premise is that a single account on the remote domain will serve as an execution launch point for all HUB end users. It is further assumed that access to this account can be made by direct ssh login or using an ssh tunnel (port forwarding).

Having attained account access to one or more remote domains it is possible to proceed with submit configuration. To get started the ssh public generated by the installation should be transferred to the remote domain host(s).

HUB Configuration

The behavior of submit is controlled through a set of configuration files. The configuration files contain descriptions of the various parameters required to connect to a remote domain, exchange files, and execute simulation codes. There are separate files for defining remote sites, staged tools, multiprocessor managers, permissible environment variables, remote job monitors, and ssh tunneling. Most parameters have default values and it is not required that all

parameters be explicitly defined in the configuration files. A simple example is given for each category of configuration file.

Sites

Remote sites are defined in the file sites.dat. Each remote site is defined by a stanza indicating an access mechanism and other account and venue specific information. Defined keywords are

- [name] - site name. Used as command line argument (-v/--venue) and in tools.dat (destinations)
- venues - comma separated list of hostnames. If multiple hostnames are listed one site will chosen at random.
- tunnelDesignator - name of tunnel defined in tunnels.dat.
- siteMonitorDesignator - name of site monitor defined in monitors.dat.
- venueMechanism - possible mechanisms are ssh and local.
- remoteUser - login user at remote site.
- remoteBatchAccount - some batch systems requirement that an account be provided in addition to user information.
- remoteBatchSystem - the possible batch submission systems include CONDOR, PBS, and LSF. SCRIPT may also be specified to specify that a script will be executed directly on the remote host.
- remoteBatchQueue - when remoteBatchSystem is PBS the queue name may be specified.
- remoteBatchPartition - slurm parameter to define partition for remote job
- remoteBatchPartitionSize - slurm parameter to define partition size, currently for BG machines.
- remoteBatchConstraints - slurm parameter to define constraints for remote job
- remoteBinDirectory - define directory where shell scripts related to the site should be kept.
- remoteApplicationRootDirectory - define directory where application executables are located.
- remoteScratchDirectory - define the top level directory where jobs should be executed. Each job will create a subdirectory under remoteScratchDirectory to isolated jobs from each other.
- remotePpn - set the number of processors (cores) per node. The PPN is applied to PBS and LSF job description files. The user may override the value defined here from the command line.
- remoteManager - site specific multi-processor manager. Refers to definition in managers.dat.
- remoteHostAttribute - define host attributes. Attributes are applied to PBS description files.
- stageFiles - A True/False value indicating whether or not files should be staged to remote site. If the the job submission host and remote host share a file system file staging may not be necessary. Default is True.

- `passUseEnvironment` - A True/False value indicating whether or not the HUB 'use' environment should be passed to the remote site. Default is False. True only makes sense if the remote site is within the HUB domain.
- `arbitraryExecutableAllowed` - A True/False value indicating whether or not execution of arbitrary scripts or binaries are allowed on the remote site. Default is True. If set to False the executable must be staged or emanate from `/apps`.
- `members` - a list of site names. Providing a member list gives a layer of abstraction between the user facing name and a remote destination. If multiple members are listed one will be randomly selected for each job.
- `state` - possible values are enabled or disabled. If not explicitly set the default value is enabled.
- `failoverSite` - specify a backup site if site is not available. Site availability is determined by site probes.
- `checkProbeResult` - A True/False value indicating whether or not probe results should determine site availability. Default is True.
- `restrictedToUsers` - comma separated list of user names. If the list is empty all users may garner site access. User restrictions are applied before group restrictions.
- `restrictedToGroups` - comma separated list of group names. If the list is empty all groups may garner site access.
- `logUserRemotely` - maintain log on remote site mapping HUB id, user to remote batch job id. If not explicitly set the default value is False.
- `undeclaredSiteSelectionWeight` - used when no site is specified to choose between sites where selection weight > 0 .

An example stanza is presented for a site that is accessed through ssh.

```
[cluster]
venues = cluster.campus.edu
remotePpn = 8
remoteBatchSystem = PBS
remoteBatchQueue = standby
remoteUser = yourhub
remoteManager = mpich-intel64
venueMechanism = ssh
remoteScratchDirectory = /scratch/yourhub
siteMonitorDesignator = clusterPBS
```

Tools

Staged tools are defined in the file tools.dat. Each staged tool is defined by a stanza indicating an where a tool is staged and any access restrictions. The existence of a staged tool at multiple sites can be expressed with multiple stanzas or multiple destinations within a single stanza. If the tool requires multiprocessors a manager can also be indicated. Defined keywords are

- [name] - tool name. Used as command line argument to execute staged tools. Repeats are permitted to indicate staging at multiple sites.
- destinations - comma separated list of destinations. Destination may exist in sites.dat or be a grid site defined by a ClassAd file.
- executablePath - path to executable at remote site. The path may be given as an absolute path on the remote site or a path relative to remoteApplicationRootDirectory defined in sites.dat.
- restrictedToUsers - comma separated list of user names. If the list is empty all users may garner tool access. User restrictions are applied before group restrictions.
- restrictedToGroups - comma separated list of group names. If the list is empty all groups may garner tool access.
- environment - comma separated list of environment variables in the form e=v.
- softenvExtensions - per site softenv environment declaration for TeraGrid sites running GRAM4.
- remoteManager - tool specific multi-processor manager. Refers to definition in managers.dat. Overrides value set by site definition.
- state - possible values are enabled or disabled. If not explicitly set the default value is enabled.

An example stanza is presented for a staged tool maintained in the yourhub account on a remote site.

```
[earth]
destinations = cluster
executablePath = ${HOME}/apps/planets/bin/earth.x
remoteManager = mpich-intel
```

```
[sun]
destinations = cluster
executablePath = ${HOME}/apps/stars/bin/sun.x
remoteManager = mpich-intel
```

Monitors

Remote job monitors are defined in the file `monitors.dat`. Each remote monitor is defined by a stanza indicating where the monitor is located and to be executed. Defined keywords are

- `[name]` - monitor name. Used in `sites.dat` (`siteMonitorDesignator`)
- `venue` - hostname upon which to launch monitor daemon. Typically this is a cluster headnode.
- `venueMechanism` - monitoring job launch process. The default is `ssh`.
- `tunnelDesignator` - name of tunnel defined in `tunnels.dat`.
- `remoteUser` - login user at remote site.
- `remoteMonitorCommand` - command to launch monitor daemon process.
- `state` - possible values are `enabled` or `disabled`. If not explicitly set the default value is `enabled`.

An example stanza is presented for a remote monitor tool used to report status of PBS jobs.

```
[clusterPBS]
venue = cluster.campus.edu
remoteUser = yourhub
remoteMonitorCommand = ${HOME}/SubmitMonitor/monitorPBS.py
```

Multi-processor managers

Multiprocessor managers are defined in the file `managers.dat`. Each manager is defined by a stanza indicating the set of commands used to execute a multiprocessor simulation run. Defined keywords are

- `[name]` - manager name. Used in `sites.dat` and `tools.dat`.
- `computationMode` - indicate how to use multiple processors for a single job. Recognized values are `mpi`, `parallel`, and `matlabmpi`. Parallel application request multiprocess have there own mechanism for inter process communication. `Matlabmpi` is used to enable the an Matlab implementation of MPI.
- `preManagerCommands` - comma separated list of commands to be executed before the manager command. Typical use of pre manager commands would be to define the environment to include a particular version of MPI amd/or compiler, or setup MPD.
- `managerCommand` - manager command commonly `mpirun`. It is possible to include strings that will be sustituted with values defined from the command line.
- `postManagerCommands` - comma separated list of commands to be executed when the manager command completes. A typical use would be to terminate an MPD setup.
- `mpiRankVariable` - define environment variable set by manager command to define

process rank. Recognized values are: MPIRUN_RANK, GMPI_ID, RMS_RANK, MXMPI_ID, MSTI_RANK, PMI_RANK, and OMPI_MCA_ns_nds_vpid. If no variable is given an attempt is made to determine process rank from command line arguments.

- environment - comma separated list of environment variables in the form e=v.
- moduleInitialize - initialize module script for sh
- modulesUnload - modules to be unloaded clearing way for replacement modules
- modulesLoad - modules to load to define mpi and other libraries
- state - possible values are enabled or disabled. If not explicitly set the default value is enabled.

An example stanza is presented for a typical MPI instance. The given command should be

suitable for /bin/sh execution.

```
[mpich-intel]
preManagerCommands = . ${MODULESHOME}/init/sh, module load mpich-
intel/11.1.038
managerCommand = mpirun -machinefile ${PBS_NODEFILE} -np NPROCESSORS
```

The token NPROCESSORS is replaced by an actual value at runtime.

Environment variables

Legal environment variables are listed in the file environmentwhitelist.dat. The objective is to

prevent end users from setting security sensitive environment variables while allowing application specific variables to be passed to the remote site. Environment variables required to define multiprocessor execution should also be included. The permissible environment variables should be entered as a simple list - one entry per line. An example file allowing use of a variables used by openmp and mpich is presenter here.

```
# environment variables listed here can be specified from the command
line with -e/--env option. Attempts to specify other environment varia
bles will be ignored and the values will not be passed to the remote s
ite.
```

```
OMP_NUM_THREADS
MPICH_HOME
```

Tunnels

In some circumstances access to clusters is restricted such that only a select list of machines is allowed to communicate with the cluster job submission node. The machines that are granted such access are sometimes referred to as gateways. In such circumstances ssh tunneling or port forwarding can be used to submit HUB jobs through the gateway machine. Tunnel definition is specified in the file tunnels.dat. Each tunnel is defined by a stanza indicating gateway host and port information. Defined keywords are

- [name] - tunnel name.
- venue - tunnel target host.
- venuePort - tunnel target port.
- gatewayHost - name of the intermediate host.
- gatewayUser - login user on gatewayHost.
- localPortOffset - local port offset used for forwarding. Actual port is localPortMinimum + localPortOffset

An example stanza is presented for a tunnel between the HUB and a remote venue by way of

```
[cluster]
venue = cluster.campus.edu
venuePort = 22
gatewayHost = gateway.campus.edu
gatewayUser = yourhub
localPortOffset = 1
```

Initialization Scripts and Log Files

The submit server and job monitoring server must be started as daemon processes running on the the submit host. If ssh tunneling is going to be used an addition server must be started as a daemon process. Each daemon process writes to a centralized log file facilitating error recording and debugging.

Initialize daemon scripts

Scripts for starting the server daemons are provided and installed in /etc/init.d. The default settings for when to start and terminate the scripts are adequate.

Log files

Submit processes log information to files located in the /var/log/submit directory tree. The exact location varies depending on the vintage of the installation. Each process has its own log file. The three most important log files are submit.log, distributor.log, and monitorJob.log.

submit.log

The submit.log file tracks when the submit server is started and stopped. Each connection from the submit client is logged with the command line and client ip address reported. All log entries are timestamped and reported by client ip address or submit ID once an ID has been assigned. Entries from all jobs are simultaneously reported and intermingled. The submit ID serves as a good search key when tracing problems. Examples of startup, job execution, and termination are given here. The job exit status and time metrics are also recorded in the MySQL database JobLog table.

```
[Sat Jan 21 14:32:39 2012] Startup: Using configdir /opt/submit
[Sat Jan 21 14:32:39 2012] Startup: Backgrounding process.
[Sat Jan 21 14:32:39 2012] Startup: Listening: protocol='tcp', host=''
```

SYSTEM ADMINISTRATORS

```
, port=830
[Sat Jan 21 14:32:39 2012] Startup: Listening: protocol='tls', host='',
, port=831

[Thu Feb  2 11:55:57 2012] 128.46.19.176: Connection to tls://:831 from
('128.46.19.176', 49737)
[Thu Feb  2 11:55:57 2012] 128.46.19.176: Server will time out in 60 s
econds.
[Thu Feb  2 11:55:57 2012] 128.46.19.176: Server will time out in 60 s
econds.
[Thu Feb  2 11:55:58 2012] 128.46.19.176: Args are:['/apps/bin/submit'
, '-n', '8', '-v', 'coates', '-w', '4.000000:00:00', '-i', 'BSLAB_512_
RUN', 'bandstrlab-r2091', '/home/nanohub/clarksm/data/sessions/460378L
/BSLAB_512_RUN/OMEN_input_1_512.cmd']
[Thu Feb  2 11:55:58 2012] 2190962: The filesystem is shared.
[Thu Feb  2 12:07:58 2012] 2190962: Job Status: venue=1:sshPBS:3934644
:coates.rcac.purdue.edu status=0 cpu=200.610000 real=36.000000
wait=619.000000
[Thu Feb  2 12:07:58 2012] 2190962: Server exiting.

[Thu Feb  2 09:38:30 2012] Startup: Server was terminated by a signal.
[Thu Feb  2 09:38:30 2012] Startup: Job Status: venue=any status=65534
cpu=0.000000 real=0.000000 wait=0.000000
[Thu Feb  2 09:38:30 2012] Startup: EXCEPTION IN MAINLOOP: int argumen
t required
[Thu Feb  2 09:38:30 2012] Startup: Server fell out of mainloop().
[Thu Feb  2 09:38:30 2012] Startup: Server exiting.
```

distributor.log

The distributor.log file tracks each job as it progresses from start to finish. Details of remote site assignment, queue status, exit status, and command execution are all reported. All entries are timestamped and reported by submit ID. The submit ID serves as the key to join data reported in submit.log. An example for submit ID 2190962 is listed here. Again the data for all jobs are intermingled.

```
[Thu Feb  2 11:55:59 2012] 2190962: command = tar vchf 02190962_01_inp
```

SYSTEM ADMINISTRATORS

```
ut.tar --exclude='*.svn*' -C /home/nanohub/clarksm/data/sessions/46037
8L .__local_jobid.02190962_01 BSLAB_512_RUN -C /home/nanohub/clarksm/d
ata/sessions/460378L/BSLAB_512_RUN OMEN_input_1_512.cmd
[Thu Feb  2 11:55:59 2012] 2190962: remoteCommand bandstrlab-
r2091 ./BSLAB_512_RUN/OMEN_input_1_512.cmd
[Thu Feb  2 11:55:59 2012] 2190962: command = genuserid
[Thu Feb  2 11:55:59 2012] 2190962: IDENTITY = /tmp/id.uBYdxy4FUw
[Thu Feb  2 11:55:59 2012] 2190962: command = update-known-
hosts coates.rcac.purdue.edu
[Thu Feb  2 11:55:59 2012] 2190962: workingDirectory /scratch/lustreA/
n/nano0/nanoHUBjobs/1328219759_02190962_01
[Thu Feb  2 11:55:59 2012] 2190962: command = tar vrhf 02190962_01_inp
ut.tar --exclude='*.svn*' -C /home/nanohub/clarksm/data/sessions/46037
8L/02190962_01 02190962_01.pbs 02190962_01.sh
[Thu Feb  2 11:55:59 2012] 2190962: command = nice -n 19 gzip 02190962
_01_input.tar
[Thu Feb  2 11:55:59 2012] 2190962: command = cat /home/nanohub/clarks
m/data/sessions/460378L/02190962_01/02190962_01_input.tar.gz | ssh -T
-x -a -i /tmp/id.uBYdxy4FUw nano0@coates.rcac.purdue.edu "${HOME}/bin/
receiveinput.sh /scratch/lustreA/n/nano0/nanoHUBjobs/1328219759_021909
62_01 .__timestamp_transferred.02190962_01"
[Thu Feb  2 11:56:01 2012] 2190962: .__local_jobid.02190962_01
[Thu Feb  2 11:56:01 2012] 2190962: command = ssh -T -x -a -i /tmp/id.
uBYdxy4FUw nano0@coates.rcac.purdue.edu "${HOME}/bin/submitbatchjob.sh
/scratch/lustreA/n/nano0/nanoHUBjobs/1328219759_02190962_01 ./0219096
2_01.pbs"
[Thu Feb  2 11:56:01 2012] 2190962: remoteJobId = 3934644.coates-
adm.rcac.purdue.edu
[Thu Feb  2 11:56:01 2012] 2190962: confirmation: S(1):N Job
[Thu Feb  2 11:56:01 2012] 2190962: status:Job N coates
[Thu Feb  2 11:56:06 2012] 2190962: status:Simulation Q coates
[Thu Feb  2 12:07:42 2012] 2190962: status:Simulation D coates
[Thu Feb  2 12:07:42 2012] 2190962: waitForBatchJobs: nCompleteRemoteJ
obIndexes = 1, nIncompleteJobs = 0, abortGlobal = False
[Thu Feb  2 12:07:42 2012] 2190962: command = ssh -T -x -a -i /tmp/id.
uBYdxy4FUw nano0@coates.rcac.purdue.edu "${HOME}/bin/transmitresults.s
h /scratch/lustreA/n/nano0/nanoHUBjobs/1328219759_02190962_01" | tar x
zmf --ignore-case --exclude '*hub-
proxy.*' -C /home/nanohub/clarksm/data/sessions/460378L/02190962_01
[Thu Feb  2 12:07:57 2012] 2190962: command = ssh -T -x -a -i /tmp/id.
uBYdxy4FUw nano0@coates.rcac.purdue.edu "${HOME}/bin/cleanupjob.sh /sc
cratch/lustreA/n/nano0/nanoHUBjobs/1328219759_02190962_01"
[Thu Feb  2 12:07:58 2012] 2190962: venue=1:sshPBS:3934644:coates.rcac
.purdue.edu status=0 cputime=200.610000 realtime=36.000000 waittime=61
9.000000 ncpus=8
```


SYSTEM ADMINISTRATORS

monitorJob.log

The monitorJob.log file tracks the invocation and termination of each remotely executed job monitor. The remote job monitors are started on demand when job are submitted to remote sites. The remote job monitors terminate when all jobs complete at a remote site and no new activity has been initiated for a specified amount of time - typically thirty minutes. A typical report should look like:

```
[Thu Feb 2 11:05:53 2012] (22140) *****
[Thu Feb 2 11:05:53 2012] (22140) * distributor job monitor started *
[Thu Feb 2 11:05:53 2012] (22140) *****
[Thu Feb 2 11:05:53 2012] (22140) loading active jobs
[Thu Feb 2 11:05:53 2012] (22140) 73 jobs loaded from DB file
[Thu Feb 2 11:05:53 2012] (22140) 73 jobs loaded from dump file
[Thu Feb 2 11:05:53 2012] (22140) 2 jobs purged
[Thu Feb 2 11:05:53 2012] (22140) 71 monitored jobs
[Thu Feb 2 11:10:33 2012] (22311) Launching coates
[Thu Feb 2 11:10:33 2012] (22140) 72 monitored jobs
[Thu Feb 2 11:10:44 2012] (22140) Update message received from coates
[Thu Feb 2 11:12:14 2012] (22140) Update message received from coates
[Thu Feb 2 11:18:22 2012] (22629) Launching steele-fe01
[Thu Feb 2 11:18:22 2012] (22140) 73 monitored jobs
[Thu Feb 2 11:19:53 2012] (22140) Update message received from steele-
fe01
[Thu Feb 2 11:21:28 2012] (22140) Update message received from steele-
fe01
[Thu Feb 2 11:50:02 2012] (22629) Closing steele-fe01
[Thu Feb 2 11:51:28 2012] (1420) *****
[Thu Feb 2 11:51:28 2012] (1420) * distributor job monitor stopped *
[Thu Feb 2 11:51:28 2012] (1420) *****
```

It is imperative that the job monitor be running in order for notification of job progress to occur. If users report that their job appears to hang check to make sure the job monitor is running. If necessary take corrective action and restart the daemon.

monitorTunnel.log

The monitorTunnel.log file tracks invocation and termination of each ssh tunnel connection. If users report problems with job submission to sites accessed via an ssh tunnel this log file should be checked for indication of any possible problems.

Remote Domain Configuration

For job submission to remote sites via ssh it is necessary to configure a remote job monitor and a set of scripts to perform file transfer and batch job related functions. A set of scripts can be used for each different batch submission system or in some cases they may be combined with appropriate switching based on command line arguments. A separate job monitor is need for each batch submission system. Communication between the HUB and remote resource via ssh

requires inclusion of a public key in the `authorized_keys` file.

Job monitor daemon

A remote job monitor runs a daemon process and reports batch job status to a central job monitor located on the HUB. The daemon process is started by the central job monitor on demand. The daemon terminates after a configurable amount of inactivity time. The daemon code needs to be installed in the location declared in the `monitors.dat` file. The daemon requires some initial configuration to declare where it will store log and history files. The daemon does not require any special privileges any runs as a standard user. Typical configuration for the daemon looks like this:

The directory defined by `MONITORLOGLOCATION` needs to be created before the daemon is started. Sample daemon scripts used for PBS, LSF, Condor, Load Leveler, and Slurm batch systems are included in directory `BatchMonitors`.

File transfer and batch job scripts

The simple scripts are used to manage file transfer and batch job launching and termination. The location of the scripts is entered in `sites.dat`.

Examples scripts suitable for use with PBS, LSF, Condor, Load Leveler, and Slurm are included

SYSTEM ADMINISTRATORS

in directory Scripts. After modifications are made to monitors.dat the central job monitor must be notified. This can be accomplished by stopping and starting the submon daemon or a HUP signal can be sent to the monitorJob.py process.

File transfer - input files

Receive compressed tar file containing input files required for the job on stdin. The file transferredTimestampFile is used to determine what newly created or modified files should be returned to the HUB.

```
receiveinput.sh jobWorkingDirectory transferredTimestampFile
```

Batch job script - submission

Submit batch job using supplied description file. If arguments beyond job working directory and batch description file are supplied an entry is added to the remote site log file. The log file provides a record relating the HUB end user to the remote batch job identifier. The log file should be placed at a location agreed upon by the remote site and HUB.

```
submitbatchjob.sh jobWorkingDirectory jobDescriptionFile
```

The jobId is returned on stdout if job submission is successful. For an unsuccessful job submission the returned jobId should be -1.

File transfer - output files

Return compressed tar file containing job output files on stdout.

```
transmitresults.sh jobWorkingDirectory
```

File transfer - cleanup

Remove job specific directory and any other dangling files

```
cleanupjob.sh jobWorkingDirectory
```

Batch job script - termination

Terminate given remote batch job. Command line arguments specify job identifier and batch system type.

```
killbatchjob.sh jobId jobClass
```

Access Control Mechanisms

By default tools and sites are configured so that access is granted to all HUB members. In some cases it is desired to restrict access to either a tool or site to a subset of the HUB membership. The keywords `restrictedToUsers` and `restrictedToGroups` provide a mechanism to apply restrictions accordingly. Each keyword should be followed by a list of comma separated values of `userids` (logins) or `groupids` (as declared when creating a new HUB group). If user or group restrictions have been declared upon invocation of `submit` a comparison is made between the restrictions and `userid` and `group` memberships. If both user and group restrictions are declared the user restriction will be applied first, followed by the group restriction.

In addition to applying user and group restrictions another mechanism is provided by the boolean keyword `arbitraryExecutableAllowed` in the sites configuration file. In cases where the executable program is not pre-staged at the remote sites the executable needs to be transferred along with the user supplied inputs to the remote site. Published tools will have their executable program located in the `/apps/tools/revision/bin` directory. For this reason submitted programs that reside in `/apps` are assumed to be validated and approved for execution. The same cannot be said for programs in other directories. The common case where such a situation arises is when a tool developer is building and testing within the HUB workspace environment. To grant a tool developer the permission to submit such arbitrary applications the site configuration must allow arbitrary executables and the tool developer must belong the system group `submit`.

OpenLDAP

OpenLDAP Installation

The base OpenLDAP package is obtained from debian.org and installed using the command

```
# apt-get install slapd
```

This adds the following files to the file system:

```
/etc/ldap/sasl2/  
/etc/ldap/schema/nis.ldif  
/etc/ldap/schema/cosine.ldif  
/etc/ldap/schema/duaconf.schema  
/etc/ldap/schema/core.ldif  
/etc/ldap/schema/nadf.schema  
/etc/ldap/schema/openldap.ldif  
/etc/ldap/schema/inetorgperson.ldif  
/etc/ldap/schema/cosine.schema  
/etc/ldap/schema/nis.schema  
/etc/ldap/schema/core.schema  
/etc/ldap/schema/openldap.schema  
/etc/ldap/schema/collective.schema  
/etc/ldap/schema/dyngroup.schema  
/etc/ldap/schema/ppolicy.schema  
/etc/ldap/schema/java.schema  
/etc/ldap/schema/misc.schema  
/etc/ldap/schema/README  
/etc/ldap/schema/inetorgperson.schema  
/etc/ldap/schema/corba.schema  
/etc/init.d/slapd  
/etc/default/slapd  
/usr/lib/libslapi-2.4.so.2.1.0  
/usr/lib/ldap/auditlog-2.4.so.2.1.0  
/usr/lib/ldap/back_bdb.la  
/usr/lib/ldap/refint.la  
/usr/lib/ldap/auditlog.la  
/usr/lib/ldap/memberof.la  
/usr/lib/ldap/back_null.la  
/usr/lib/ldap/dynlist.la  
/usr/lib/ldap/valsort-2.4.so.2.1.0  
/usr/lib/ldap/dynlist-2.4.so.2.1.0
```

SYSTEM ADMINISTRATORS

/usr/lib/ldap/back_passwd-2.4.so.2.1.0
/usr/lib/ldap/back_relay-2.4.so.2.1.0
/usr/lib/ldap/seqmod.la
/usr/lib/ldap/back_monitor.la
/usr/lib/ldap/pcache.la
/usr/lib/ldap/back_sock-2.4.so.2.1.0
/usr/lib/ldap/back_perl.la
/usr/lib/ldap/syncprov-2.4.so.2.1.0
/usr/lib/ldap/memberof-2.4.so.2.1.0
/usr/lib/ldap/back_perl-2.4.so.2.1.0
/usr/lib/ldap/rwm-2.4.so.2.1.0
/usr/lib/ldap/ppolicy-2.4.so.2.1.0
/usr/lib/ldap/accesslog.la
/usr/lib/ldap/accesslog-2.4.so.2.1.0
/usr/lib/ldap/back_dnssrv.la
/usr/lib/ldap/back_dnssrv-2.4.so.2.1.0
/usr/lib/ldap/dyngroup-2.4.so.2.1.0
/usr/lib/ldap/syncprov.la
/usr/lib/ldap/back_passwd.la
/usr/lib/ldap/back_relay.la
/usr/lib/ldap/back_monitor-2.4.so.2.1.0
/usr/lib/ldap/retcode.la
/usr/lib/ldap/back_sock.la
/usr/lib/ldap/pcache-2.4.so.2.1.0
/usr/lib/ldap/constraint-2.4.so.2.1.0
/usr/lib/ldap/back_ldap.la
/usr/lib/ldap/back_sql-2.4.so.2.1.0
/usr/lib/ldap/back_hdb.la
/usr/lib/ldap/dds.la
/usr/lib/ldap/back_ldap-2.4.so.2.1.0
/usr/lib/ldap/back_hdb-2.4.so.2.1.0
/usr/lib/ldap/unique.la
/usr/lib/ldap/back_sql.la
/usr/lib/ldap/back_meta.la
/usr/lib/ldap/ppolicy.la
/usr/lib/ldap/back_bdb-2.4.so.2.1.0
/usr/lib/ldap/seqmod-2.4.so.2.1.0
/usr/lib/ldap/translucent.la
/usr/lib/ldap/back_shell.la
/usr/lib/ldap/back_meta-2.4.so.2.1.0
/usr/lib/ldap/rwm.la
/usr/lib/ldap/translucent-2.4.so.2.1.0
/usr/lib/ldap/unique-2.4.so.2.1.0
/usr/lib/ldap/refint-2.4.so.2.1.0
/usr/lib/ldap/dyngroup.la
/usr/lib/ldap/valsort.la

SYSTEM ADMINISTRATORS

```
/usr/lib/ldap/back_shell-2.4.so.2.1.0
/usr/lib/ldap/retcode-2.4.so.2.1.0
/usr/lib/ldap/back_null-2.4.so.2.1.0
/usr/lib/ldap/dds-2.4.so.2.1.0
/usr/lib/ldap/constraint.la
/usr/sbin/slappcat
/usr/sbin/slappasswd
/usr/sbin/slapdn
/usr/sbin/slapauth
/usr/sbin/slapd
/usr/sbin/slapadd
/usr/sbin/slapindex
/usr/sbin/slapacl
/usr/sbin/slaptest
/usr/share/lintian/
/usr/share/lintian/overrides/
/usr/share/man/man8/slappauth.8.gz
/usr/share/man/man8/slappcat.8.gz
/usr/share/man/man8/slapacl.8.gz
/usr/share/man/man8/slapadd.8.gz
/usr/share/man/man8/slaptest.8.gz
/usr/share/man/man8/slapindex.8.gz
/usr/share/man/man8/slappasswd.8.gz
/usr/share/man/man8/slapd.8.gz
/usr/share/man/man8/slapdn.8.gz
/usr/share/man/man5/slapd-monitor.5.gz
/usr/share/man/man5/slapd-sock.5.gz
/usr/share/man/man5/slapo-accesslog.5.gz
/usr/share/man/man5/slapd-dnssrv.5.gz
/usr/share/man/man5/slapo-dds.5.gz
/usr/share/man/man5/slapo-constraint.5.gz
/usr/share/man/man5/slapd-ldbm.5.gz
/usr/share/man/man5/slapd-config.5.gz
/usr/share/man/man5/slapo-dyngroup.5.gz
/usr/share/man/man5/slapo-memberof.5.gz
/usr/share/man/man5/slapo-auditlog.5.gz
/usr/share/man/man5/slapo-translucent.5.gz
/usr/share/man/man5/slapd-ldif.5.gz
/usr/share/man/man5/slapd.backends.5.gz
/usr/share/man/man5/slapd.access.5.gz
/usr/share/man/man5/slapd-meta.5.gz
/usr/share/man/man5/slapd.conf.5.gz
/usr/share/man/man5/slapd-shell.5.gz
/usr/share/man/man5/slapd-sql.5.gz
/usr/share/man/man5/slapd-passwd.5.gz
```

SYSTEM ADMINISTRATORS

```
/usr/share/man/man5/slapd-bdb.5.gz
/usr/share/man/man5/slapd-hdb.5.gz
/usr/share/man/man5/slapo-pcache.5.gz
/usr/share/man/man5/slapo-valsort.5.gz
/usr/share/man/man5/slapo-chain.5.gz
/usr/share/man/man5/slapo-unique.5.gz
/usr/share/man/man5/slapo-retcode.5.gz
/usr/share/man/man5/slapd-perl.5.gz
/usr/share/man/man5/slapd-null.5.gz
/usr/share/man/man5/slapo-ppolicy.5.gz
/usr/share/man/man5/slapd.plugin.5.gz
/usr/share/man/man5/slapd-ldap.5.gz
/usr/share/man/man5/slapo-rwm.5.gz
/usr/share/man/man5/slapo-syncprov.5.gz
/usr/share/man/man5/slapd.overlays.5.gz
/usr/share/man/man5/slapo-dynlist.5.gz
/usr/share/man/man5/slapd-relay.5.gz
/usr/share/man/man5/slapo-refint.5.gz
/usr/share/doc/slapd/copyright
/usr/share/doc/slapd/changelog.Debian.gz
/usr/share/doc/slapd/NEWS.Debian.gz
/usr/share/doc/slapd/README.Debian.gz
/usr/share/doc/slapd/examples/
/usr/share/doc/slapd/examples/slapd.backup
/usr/share/doc/slapd/TODO.Debian
/usr/share/doc/slapd/README.DB_CONFIG.gz
/usr/share/slapd/ldiftopasswd
/usr/share/slapd/slapd.conf
/usr/share/slapd/DB_CONFIG
/var/lib/slapd/
/usr/lib/ldap/back_meta.so
/usr/lib/ldap/dynlist.so
/usr/lib/ldap/memberof.so
/usr/lib/ldap/accesslog-2.4.so.2
/usr/lib/ldap/seqmod.so
/usr/lib/ldap/seqmod-2.4.so.2
/usr/lib/ldap/ppolicy.so
/usr/lib/ldap/back_dnssrv.so
/usr/lib/ldap/back_bdb.so
/usr/lib/ldap/back_ldap.so
/usr/lib/ldap/back_ldap-2.4.so.2
/usr/lib/ldap/back_sql.so
/usr/lib/ldap/back_null.so
/usr/lib/ldap/refint.so
/usr/lib/ldap/refint-2.4.so.2
/usr/lib/ldap/back_relay-2.4.so.2
```


SYSTEM ADMINISTRATORS

```
/usr/lib/ldap/back_hdb-2.4.so.2
/usr/lib/ldap/back_meta-2.4.so.2
/usr/lib/ldap/back_shell.so
/usr/lib/ldap/ppolicy-2.4.so.2
/usr/lib/ldap/back_sock-2.4.so.2
/usr/lib/ldap/auditlog-2.4.so.2
/usr/lib/ldap/dds-2.4.so.2
/usr/lib/ldap/dyngroup-2.4.so.2
/usr/lib/ldap/back_sql-2.4.so.2
/usr/lib/ldap/constraint-2.4.so.2
/usr/lib/ldap/rwm.so
/usr/lib/ldap/back_dnssrv-2.4.so.2
/usr/lib/ldap/auditlog.so
/usr/lib/ldap/syncprov.so
/usr/lib/ldap/syncprov-2.4.so.2
/usr/lib/ldap/rwm-2.4.so.2
/usr/lib/ldap/memberof-2.4.so.2
/usr/lib/ldap/valsort.so
/usr/lib/ldap/back_perl-2.4.so.2
/usr/lib/ldap/dyngroup.so
/usr/lib/ldap/back_passwd-2.4.so.2
/usr/lib/ldap/back_monitor.so
/usr/lib/ldap/valsort-2.4.so.2
/usr/lib/ldap/back_relay.so
/usr/lib/ldap/back_perl.so
/usr/lib/ldap/unique-2.4.so.2
/usr/lib/ldap/translucent.so
/usr/lib/ldap/back_null-2.4.so.2
/usr/lib/ldap/accesslog.so
/usr/lib/ldap/back_passwd.so
/usr/lib/ldap/back_shell-2.4.so.2
/usr/lib/ldap/constraint.so
/usr/lib/ldap/back_sock.so
/usr/lib/ldap/pcache.so
/usr/lib/ldap/pcache-2.4.so.2
/usr/lib/ldap/translucent-2.4.so.2
/usr/lib/ldap/dynlist-2.4.so.2
/usr/lib/ldap/back_monitor-2.4.so.2
/usr/lib/ldap/unique.so
/usr/lib/ldap/dds.so
/usr/lib/ldap/back_hdb.so
/usr/lib/ldap/retcode-2.4.so.2
/usr/lib/ldap/back_bdb-2.4.so.2
/usr/lib/ldap/retcode.so
/usr/lib/libslapi-2.4.so.2
/usr/share/doc/slaped/examples/slaped.conf
```

`/usr/share/doc/slaped/examples/DB_CONFIG`

OpenLDAP Configuration

The `/etc/ldap/slaped.conf` is edited and the base DN changed to `myhub.org` (we are assuming that `myhub.org` was the DN specified initially when asked during the initial six questions step).

Next, the LDAP admin password is set to that specified in the six questions.

nscd Installation

The Name Service Cache Daemon `nscd` is downloaded from `debian.org` and installed using the following command:

```
# apt-get install nscd
```

This adds the following files to the file system:

```
/etc/init.d/nscd
/etc/nscd.conf
/usr/share/doc/nscd/
/usr/share/doc/nscd/changelog.Debian.gz
/usr/share/doc/nscd/NEWS.Debian.gz
/usr/share/doc/nscd/copyright
/usr/share/man/man8/nscd.8.gz
/usr/share/man/man5/nscd.conf.5.gz
/usr/sbin/nscd
/var/run/nscd/
/var/cache/nscd/
```

HUBzero LDAP Schema Installation

The HUBzero OpenLDAP package is downloaded from `packages.hubzero.org` and installed using the following command:

```
# apt-get install hubzero-openldap
```

This adds the following files to the file system:

```
/etc/ldap/schema/hub.schema
/usr/share/doc/hubzero-openldap/changelog.Debian.gz
/usr/share/doc/hubzero-openldap/copyright
/usr/share/hubzero-openldap/do-hz-openldap-install.tpl
/usr/share/hubzero-openldap/HUB-INIT-SLAPD.tpl
```

Next, the **/etc/ldap/slapd.conf** is edited and the following line added as the last "include" statement.

```
include          /etc/ldap/schema/hub.schema
```

Then OpenLDAP is restarted with

```
# /etc/init.d/slapd restart
```

PAM Configuration to Use LDAP

The PAM module for LDAP is downloaded from debian.org and installed using the following command:

```
# apt-get install libpam-ldap
```

This adds the following files to the file system:

```
/usr/share/doc/libpam-ldap/LDAP-Permissions.txt
/usr/share/doc/libpam-ldap/README.Debian
/usr/share/doc/libpam-ldap/changelog.gz
/usr/share/doc/libpam-ldap/copyright
/usr/share/doc/libpam-ldap/ldapns.schema
/usr/share/doc/libpam-ldap/AUTHORS
/usr/share/doc/libpam-ldap/examples/
/usr/share/doc/libpam-ldap/examples/pam.d/
/usr/share/doc/libpam-ldap/examples/pam.d/halt
/usr/share/doc/libpam-ldap/examples/pam.d/linuxconf-pair
/usr/share/doc/libpam-ldap/examples/pam.d/rsh
/usr/share/doc/libpam-ldap/examples/pam.d/gdm
/usr/share/doc/libpam-ldap/examples/pam.d/samba
```

SYSTEM ADMINISTRATORS

```
/usr/share/doc/libpam-ldap/examples/pam.d/ftp
/usr/share/doc/libpam-ldap/examples/pam.d/other
/usr/share/doc/libpam-ldap/examples/pam.d/mcserv
/usr/share/doc/libpam-ldap/examples/pam.d/reboot
/usr/share/doc/libpam-ldap/examples/pam.d/pop
/usr/share/doc/libpam-ldap/examples/pam.d/ppp
/usr/share/doc/libpam-ldap/examples/pam.d/vlock
/usr/share/doc/libpam-ldap/examples/pam.d/poweroff
/usr/share/doc/libpam-ldap/examples/pam.d/xdm
/usr/share/doc/libpam-ldap/examples/pam.d/xlock
/usr/share/doc/libpam-ldap/examples/pam.d/imap
/usr/share/doc/libpam-ldap/examples/pam.d/linuxconf
/usr/share/doc/libpam-ldap/examples/pam.d/xscreensaver
/usr/share/doc/libpam-ldap/examples/pam.d/passwd
/usr/share/doc/libpam-ldap/examples/pam.d/chsh
/usr/share/doc/libpam-ldap/examples/pam.d/login
/usr/share/doc/libpam-ldap/examples/pam.d/su
/usr/share/doc/libpam-ldap/examples/pam.d/ssh
/usr/share/doc/libpam-ldap/examples/pam.d/chfn
/usr/share/doc/libpam-ldap/examples/pam.d/shutdown
/usr/share/doc/libpam-ldap/examples/pam.d/rexec
/usr/share/doc/libpam-ldap/examples/pam.d/rlogin
/usr/share/doc/libpam-ldap/examples/pam.d/kde
/usr/share/doc/libpam-ldap/examples/pam.d/xserver
/usr/share/doc/libpam-ldap/examples/chsh
/usr/share/doc/libpam-ldap/examples/chfn
/usr/share/doc/libpam-ldap/examples/pam.conf
/usr/share/doc/libpam-ldap/changelog.Debian.gz
/usr/share/doc/libpam-ldap/buildinfo.gz
/usr/share/doc/libpam-ldap/README.gz
/usr/share/libpam-ldap/ldap.conf
/usr/share/man/man5/pam_ldap.conf.5.gz
/lib/security/pam_ldap.so
```

Next, **/etc/pam.d/common-auth** is modified to allow authentication against LDAP so it contains

```
#auth required pam_unix.so nullok_secure

auth sufficient pam_unix.so nullok_secure
auth sufficient pam_ldap.so try_first_pass
auth required pam_deny.so
```

Also, `/etc/pam_ldap.conf` is modified by adding the following section.

```
# HUBzero Mappings
nss_base_passwd ou=users,?one
nss_base_shadow ou=users,?one?host=web
pam_filter host=web
pam_password crypt
nss_map_attribute uniqueMember member
nss_base_group ou=groups,dc=myhub,dc=org?sub
```

NSS Configuration to Use OpenLDAP

The NSS module for using LDAP as a naming service is downloaded from debian.org and installed using the following command:

```
# apt-get install libnss-ldapd
```

This adds the following files to the file system:

```
/usr/sbin/nslcd
/usr/share/doc/libnss-ldapd/copyright
/usr/share/doc/libnss-ldapd/NEWS.gz
/usr/share/doc/libnss-ldapd/AUTHORS
/usr/share/doc/libnss-ldapd/README.gz
/usr/share/doc/libnss-ldapd/examples/
/usr/share/doc/libnss-ldapd/examples/nss-ldapd.conf.gz
/usr/share/doc/libnss-ldapd/changelog.gz
/usr/share/man/man8/nslcd.8.gz
/usr/share/man/man5/nss-ldapd.conf.5.gz
/lib/libnss_ldap.so.2
```

Next, the file `/etc/nss-ldapd.conf` is created with `root:nslcd` ownership. The only lines that are enabled are

```
uid nslcd
gid nslcd
uri ldap://127.0.0.1
base dc=myhub,dc=hubzero,dc=org
```

Finally, NSS configuration file **/etc/nsswitch.conf**

is modified by adding ldap to the first three lines:

```
passwd:          compat ldap
group:           compat ldap
shadow:         compat ldap
```

```
hosts:          files dns
networks:       files
```

```
protocols:     db files
services:      db files
ethers:        db files
rpc:           db files
```

```
netgroups:     nis
```

LDAP Services Restart

All LDAP services are restarted:

```
# /etc/init.d/slapd restart
# /etc/init.d/nslcd restart
# /etc/init.d/nscd restart
```

Test

```
# getent passwd
```

To test configuration. You should see entries for users 'hubrepo' and 'apps' toward the end of the list if everything is working correctly.

Troubleshooting

If you have a problem with the system apparently not recognizing up to date account or group information (eg., in the next section some people report receiving an error about unknown username 'hubadmin') you can nscd to flush its data cache and restart using the following commands:

SYSTEM ADMINISTRATORS

```
# nscd -i passwd
# nscd -i group
# /etc/init.d/nscd restart
# getent passwd
```

If you still don't see the hubadmin account listed then re-read the instructions and check your work very carefully. These instructions assume a fresh install, if you are working with an existing LDAP/PAM/NSS installation you will have to do more advanced troubleshooting outside the scope of this documentation.

Home Directories Creation

Home directories for the apps and admin users are created:

```
# mkdir /home/myhub/apps
# chown apps.public /home/myhub/apps
# chmod 0700 /home/myhub/apps
# mkdir -p /home/myhub/admin
# chown admin.public /home/myhub/admin
# chmod 0700 /home/myhub/admin
```

MySQL

Installation

The default MySQL server on Debian is used but the HUBzero MySQL package is downloaded from packages.hubzero.org and installed using the following command:

```
# apt-get install hubzero-mysql
```

This adds the following files to the file system:

```
/usr/share/doc/hubzero-mysql/changelog.Debian.gz  
/usr/share/doc/hubzero-mysql/copyright  
/usr/share/hubzero-mysql/hubzero-fix-debian-maint.sql
```

Next, the MySQL root password is set to the one specified in the six questions.

Bug Fix

A bug in the Debian installation that disallows `debian-sys-maint` account to GRANT privileges is fixed via the following command:

```
# /usr/bin/mysql --defaults-file=/etc/mysql/debian.cnf -D mysql
```


Apache

Installation & Configuration

The default Apache on Debian is used but the `hubzero-apache2` HUBzero Apache package is downloaded from packages.hubzero.org and installed using the following command:

```
# apt-get install hubzero-apache2
```

This adds the following files to the file system:

```
/var/log/apache2/daily/imports  
/usr/share/doc/hubzero-apache2/changelog.Debian.gz  
/usr/share/doc/hubzero-apache2/README.Debian  
/usr/share/doc/hubzero-apache2/copyright  
/etc/apache2/sites-available/hub  
/etc/apache2/sites-available/hub-ssl
```

Next, the hub Apache configuration files are enabled:

```
# /usr/sbin/a2ensite hub  
# /usr/sbin/a2ensite hub-ssl
```

Finally, Apache is restarted:

```
apache2ctl restart
```

PHP

Configuration

The following lines in the PHP configuration file `/etc/php5/apache2/php.ini` is modified from

```
display_errors = On  
log_errors = Off
```

to

```
display_errors = Off  
log_errors = On
```

and permission/ownership changed as follows.

```
# chmod 0644 /etc/php5/apache2/php.ini  
# chown root.root /etc/php5/apache2/php.ini
```

Test

```
# echo "<?php phpinfo();?>" > /var/www/index.php
```

Go to `/index.php` on your site and you should see a php status page.

Delete the test page when you are done.

```
# rm /var/www/index.php
```

CMS

Global HUBzero Configuration

Create the file `/etc/hubzero.conf` using values from the initial six questions (assuming that the user chose the default "myhub" as the name for the hub):

```
[default]
site=myhub

[myhub]
DocumentRoot=/www/myhub
HubName=myhub
HubHost=myhub.org
BaseDN=dc=myhub,dc=org
Org=dc=myhub.org
EmailAddress=
ContainerArch=i386
ContainerDebianRepo=lenny
ContainerHubzeroRepo=buck
```

CMS Installation

The Joomla! content management system is downloaded from packages.hubzero.org and installed using the following command:

```
# apt-get install hubzero-cms
```

and the web publishing area created:

```
# mkdir /www/myhub
# cp -rp /usr/lib/hubzero/cms/* /www/myhub
# chown -R www-data:www-data /www/myhub
```

Database Creation

MySQL databases for the CMS and for HUBzero metrics are created:

```
CREATE DATABASE `myhub`;
CREATE DATABASE `myhub_metrics`;
```

```
GRANT ALL PRIVILEGES ON `myhub`.* TO \'myhub
\'@\'%\' IDENTIFIED BY \\'<password>\';
GRANT ALL PRIVILEGES ON `myhub_metrics`.* TO \'myhub
\'@\'%\' IDENTIFIED BY \\'<password>\';
FLUSH PRIVILEGES;
```

where `password` is the CMS "admin" user password specified in the six questions step.

CMS Configuration

1. The HUBZero package `hubzero-cms-setup` is downloaded from packages.hubzero.org and installed using the command

```
# apt-get install hubzero-cms-setup
```

This adds the following files to the file system:

```
/usr/share/hubzero-cms-setup/old-default-hub-db.sql
/usr/share/hubzero-cms-setup/admin_user.sql
/usr/share/hubrepo_user.sql
/usr/share/groups.sql
/usr/share/configuration.php.tpl
/usr/share/make_joomla_password.php.tpl
/usr/share/get_tools_params.sql
/usr/share/get_contrib_params.sql
/usr/share/hubzero-cms-setup/clean_params.sh
/usr/share/doc/hubzero-cms-setup/copyright
/usr/share/doc/hubzero-cms-setup/changelog.Debian.gz
/usr/share/doc/hubzero-cms-setup/README.Debian
```

2. A number of SQL files are used to populate MySQL tables. These include

```
/usr/share/hubzero-cms-setup/default-hub-db.sql
/usr/share/hubzero-cms-setup/admin_user.sql
/usr/share/hubzero-cms-setup/hubrepo_user.sql
/usr/share/hubzero-cms-setup/groups.sql
/usr/share/hubzero-cms-setup/get_tools_params.sql
/usr/share/hubzero-cms-setup/get_contrib_params.sql
```

```
/usr/share/hubzero-cms-joomla-installation/cms/installation/sql/m  
ysql/hz_sample_data.sql (where does /h-c-j-i dir come from?)  
/usr/share/hubzer-cms-joomla-  
installation/cms/installation/sql/mysql/hubzero.sql
```

3. The CMS installation directory **/var/www/myhub/installation** is removed.
4. The hubname and password is set in MySQL for the tools and contribtools components using the values specified in the six questions step initially.
5. `/var/www/myhub/hubconfiguration.php`, `configuration.php` is generated with values from the initial, six questions step.

Subversion

Configuration

This step executes the following:

```
# install --owner www-data --group www-  
data --mode 0770 -d /opt/svn/tools  
# touch /etc/apache2/svn.conf /etc/apache2/svn.bak  
# chown www-data /etc/apache2/svn.conf /etc/apache2/svn.bak
```

Test

```
# svnadmin create /opt/svn/tools/test --fs-type fsfs  
# chown -R www-data.www-data /opt/svn/tools/test  
# echo "<Location /tools/test/svn>  
    DAV svn  
    SVNPath /opt/svn/tools/test  
    AuthType Basic  
    AuthBasicProvider ldap  
    AuthName "Test"  
    AuthzLDAPAuthoritative on  
    AuthLDAPGroupAttributeIsDN on  
    AuthLDAPGroupAttribute owner  
    AuthLDAPGroupAttribute member  
    AuthLDAPURL ldap://localhost/ou=users,dc=myhub,dc=org  
    Require ldap-group gid=apps,ou=groups,dc=myhub,dc=org  
</Location>" > /etc/apache2/svn.conf  
# /etc/init.d/apache2 restart
```

Be sure to the BASEDN in the above to match that used by your configuration.

Now browse to "/tools/test/svn" using an https connection and you should get prompted for a username and password, use the apps account you created earlier when you installed LDAP. You should see "svn - Revision 0: /".

Delete test file.

SYSTEM ADMINISTRATORS

```
# echo "" > /etc/apache2/svn.conf  
# rm -fr /opt/svn/tools/test  
# /etc/init.d/apache2 restart
```

WebDAV

WebDAV Modules Deployment

```
# a2enmod dav_fs
# a2enmod dav_lock
# /etc/init.d/apache2 restart
```

WebDAV Configuration

The `/etc/apache2/sites-available/hub-ssl`

file is edited to specify the hub DN. The relevant line in the files looks like:

```
<Directory /webdav>
...
```

```
AuthLDAPURL ldap://localhost/ou=users,dc=myhub,dc=org?uid
```

```
...
</Directory>
```

Apache is restarted to enable changes.

```
# /etc/init.d/apache2 restart
```

Test

```
# install --owner www-data --group www-
data --mode 0770 -d /webdav/home/apps
# touch /webdav/home/apps/test
```

Browse to your site's `https /webdav` address (e.g. `https://myhub/webdav`). You should get prompted for a username and password. Use the `apps` account. You should see a directory listing including the file `"test"`.

Now test using a WebDAV client.

```
# apt-get install cadaver  
# cadaver https://myhub.org/webdav
```

You will be prompted to accept self signed certificate (if it is still installed) and then to enter your username and password. Use the 'apps' account again to test. When you get the "dav:/webdav/>" prompt just enter "ls" and it should show the test file.

Finally clean up test case

```
# apt-get purge cadaver  
# rm /webdav/home/apps/test  
# rmdir /webdav/home/apps /webdav/home /webdav
```

Usermap

Installation

The `hubzero-usermap` package is downloaded from `hubzero-usermap` `packages.hubzero.org` and installed using the following command:

```
# apt-get install hubzero-usermap
```

which adds the following files to the file system:

```
/sbin/mount.usermap  
/usr/share/doc/hubzero-usermap/  
/usr/share/doc/hubzero-usermap/changelog.Debian.gz  
/usr/share/doc/hubzero-usermap/copyright
```

Configuration

1. The installation scripts add the following lines to the `/etc/auto.master` file:

```
/webdav/home /etc/auto.webdav --timeout=60
```

2. `/etc/auto.webdav` is configured with the following line:

```
* -fstype=usermap,user=www-data,allow_other :&
```

3. `autofs` is started:

```
# /etc/init.d/autofs restart
```

4. **fuse** is added to the `/etc/modules` file so that it is loaded upon startup.

This automounts a `usermap` filesystem of a users home directory inside of `/webdav/home` on demand. This version of the users home directory is owned and accessible to the user `www-data` which allows WebDAV to serve its contents.

Test

```
# touch /home/myhub/apps/test
# ls -l /webdav/home/apps
```

You should see a list of files in apps's home directory ("test") which will appear to be owned by www-data.www-data

```
# mount -l
```

You should see something like:

```
mount.usermap on /webdav/home/apps type fuse.mount.usermap
(rw,nosuid,nodev,allow_other)
```

Finally clean up.

```
# umount -f /webdav/home/apps
# rm /webdav/home/apps/test
```

Troubleshooting

If the test doesn't work, check if the fuse kernel module is loaded

```
# lsmod | grep fuse
fuse                54176  0
```

If there is no output then try starting the kernel module manually

```
# modprobe fuse
```

SYSTEM ADMINISTRATORS

Then try the test again

Trac

Authentication Plugin Installation

The MySQL authentication plugin for Trac is installed from packages.hubzero.org using:

```
# apt-get install hubzero-trac-mysqlauthz
```

and Apache restarted:

```
# /etc/init.d/apache2 restart
```

Test

```
# svnadmin create /opt/svn/tools/test --fs-type fsfs
# chown -R www-data.www-data /opt/svn/tools/test
# trac-admin /opt/trac/tools/test initenv "Test" "sqlite:db/trac.db" "
svn" "/opt/svn/tools/test"
# chown -R www-data.www-data /opt/trac/tools/test
```

Now browse to `/tools/test/wiki` using an https connection; you should see a default Trac project page.

Delete test data.

```
# rm -fr /opt/svn/tools/test
# rm -fr /opt/trac/tools/test
# /etc/init.d/apache2 restart
```

addrepo

Installation

The hubzero-addrepo package is installed from packages.hubzero.org:

```
# apt-get install hubzero-addrepo
```

which adds the following files:

```
/opt/svn/tools/  
/opt/trac/tools/  
/etc/apache2/svn.conf  
/etc/apache2/svn.bak  
/usr/share/doc/hubzero-addrepo/  
/usr/share/doc/hubzero-addrepo/changelog.Debian.gz  
/usr/share/doc/hubzero-addrepo/copyright  
/usr/share/hubzero-addrepo/images.txt.in  
/usr/share/hubzero-addrepo/image.py.in  
/usr/share/hubzero-addrepo/link.py.in  
/usr/share/hubzero-addrepo/templates/  
/usr/share/hubzero-addrepo/templates/site_footer.cs  
/usr/share/hubzero-addrepo/templates/site_header.cs  
/usr/share/hubzero-addrepo/templates/site_css.cs  
/usr/share/hubzero-addrepo/templates/site.html  
/usr/share/hubzero-addrepo/svn/  
/usr/share/hubzero-addrepo/svn/branches/  
/usr/share/hubzero-addrepo/svn/trunk/  
/usr/share/hubzero-addrepo/svn/trunk/doc/  
/usr/share/hubzero-addrepo/svn/trunk/examples/  
/usr/share/hubzero-addrepo/svn/trunk/rappture/  
/usr/share/hubzero-addrepo/svn/trunk/src/  
/usr/share/hubzero-addrepo/svn/trunk/bin/  
/usr/share/hubzero-addrepo/svn/trunk/data/  
/usr/share/hubzero-addrepo/svn/trunk/middleware/  
/usr/share/hubzero-addrepo/svn/tags/  
/usr/share/hubzero-addrepo/cover.txt.in  
/usr/share/hubzero-addrepo/getstart.txt.in  
/usr/share/hubzero-addrepo/ldapconf.in  
/usr/bin/gensvnapache  
/usr/bin/addrepo
```

Configuration

sudo privileges are configured by adding the following lines to `/etc/sudoers`:

```
www-  
data          ALL=(apps)NOPASSW  
D:/bin/bash  /www/myhub/components/com_contribtool/scripts/*tool.php *  
www-  
data          ALL=(apps)N  
OPASSWD:/usr/bin/expect /www/myhub  
/components/com_contribtool/scripts/*tool.php *  
www-data      ALL=NOPASSWD:/etc/init.d/apache2
```

This is so that the web process can run a number of scripts as the "apps" user.

iptables

Installation

The HUBzero firewall package `hubzero-firewall` is installed from packages.hubzero.org:

```
# apt-get install hubzero-firewall
```

which installs the following files:

```
/usr/share/doc/hubzero-firewall/copyright  
/usr/share/doc/hubzero-firewall/changelog.Debian.gz  
/usr/share/doc/hubzero-firewall/README.Debian  
/etc/hubzero/firewall_on  
/etc/hubzero/firewall_off
```

HUBzero requires the use of `iptables` to route network connections between application sessions and the external network. The scripts controlling this can also be used to manage basic firewall operations for the site. The basic scripts installed here block all access to the host except for those ports required by HUBzero (`http`, `https`, `http alt`, `ldap`, `ssh`, `smtp`, `mysql`, `submit`, etc).

Maxwell Service

Installation

The HUBzero Maxwell middleware services package `hubzero-mw-service` is installed from packages.hubzero.org:

```
# apt-get install hubzero-mw-service
```

which adds the following files:

```
/usr/lib/hubzero/bin/maxwell_service
/usr/lib/hubzero/bin/set_quotas
/etc/vz/conf/hub-session-4.0-i386.conf
/etc/vz/conf/hub-session-4.0-i386.mount
/etc/vz/conf/hub-session-4.0-i386.umount
/etc/vz/conf/hub-session-5.0-i386.conf
/etc/vz/conf/hub-session-5.0-i386.mount
/etc/vz/conf/hub-session-5.0-i386.umount
/etc/vz/conf/hub-session-4.0-amd64.conf
/etc/vz/conf/hub-session-4.0-amd64.mount
/etc/vz/conf/hub-session-4.0-amd64.umount
/etc/vz/conf/hub-session-5.0-amd64.conf
/etc/vz/conf/hub-session-5.0-amd64.mount
/etc/vz/conf/hub-session-5.0-amd64.umount
/var/log/hubzero/sessions
/usr/lib/hubzero/bin/mkvztemplate
/etc/hubzero/quota.conf
/etc/rc.boot
```

Configuration

1. A self-signed SSL key for use by `xvnc` is created:

```
# cd /etc/hubzero
# openssl req -new -x509 -days 365 -nodes -out xvnc.pem.cert -key
out xvnc.pem.key
```

2. The Maxwell service configuration file is created and installed in `/etc/hubzero/maxwell.conf`:

```
mysql_host = "localhost"
mysql_user = "myhub"
mysql_password = "<myhub user's pw>"
mysql_db = "myhub"
filexfer_decoration=""filexfer_sitelogo { <h1><a href="http://hubzero.org/" title="HUBzero home page"><span>HUBzero.org: online simulations and more</span></a></h1> }
filexfer_stylesheet http://$huburl/templates/filexfer/upload.css

hub_name="myhub"
hub_url="http://myhub.org"
hub_homedir="/home/myhub"
hub_template="hubbasic"
```

3. The host table is initialized in the database:

```
MySQL> INSERT INTO 'host' VALUES ('localhost',27,'up',1,1);
```

4. The OpenVZ template for creating containers dynamically is created:

```
# /usr/lib/mw/bin/mkvztemplate amd64 lenny
```

Test

```
# /usr/lib/hubzero/bin/maxwell_service startvnc 1 800x600 24
```

Enter an 8 character password when prompted (e.g., "testtest")

This should result in a newly create OpenVZ session with an instance of a VNC server running inside of it. The output of the above command should look something like:

```
Reading passphrase:
testtest
===== begin /etc/vz/conf/hub-
session-5.0-amd64.umount =====
```

```
Removing /var/lib/vz/root/1 :root etc var tmp dev/shm dev
===== end /etc/vz/conf/hub-
```

SYSTEM ADMINISTRATORS

```
session-5.0-amd64.umount =====
stunnel already running
Starting VE ...
===== begin /etc/vz/conf/1.mount =====
=====
Removing and repopulating: root etc var tmp dev
Mounting: /var/lib/vz/template/debian-5.0-amd64-maxwell home apps
===== end /etc/vz/conf/1.mount =====
=====
VE is mounted
Setting CPU units: 1000
Configure meminfo: 2000000
VE start in progress...
TIME: 0 seconds.
Waiting for container to finish booting.
/usr/lib/hubzero/startxvnc: Becoming nobody.
/usr/lib/hubzero/startxvnc: Waiting for 8-byte vncpasswd and EOF.
1+0 records in
1+0 records out
8 bytes (8 B) copied, 3.5333e-05 s, 226 kB/s
Got the vncpasswd
Adding auth for 10.51.0.1:0 and 10.51.0.1/unix:0
xauth: creating new authority file Xauthority-10.51.0.1:0
Adding IP address(es): 10.51.0.1
if-up.d/mountnfs[venet0]: waiting for interface venet0:0 before doing
NFS mounts (warning).
WARNING: Settings were not saved and will be resetted to original valu
es on next start (use --save flag)

# vzlist
      VEID      NPROC STATUS  IP_ADDR      HOSTNAME
      1          6 running 10.51.0.1    -

# openssl s_client -connect localhost:4001
```

This should report an SSL connection with a self signed certificate and output text should end

with:

RFB 003.008

If you see this then you successfully connected to the VNC server running inside the newly created OpenVZ session.

Clean up

```
# /usr/lib/hubzero/bin/maxwell_service stopvnc 1
```

Which should give output similar to:

```
Killing 6 processes in veid 1 with signal 1
Killing 7 processes in veid 1 with signal 2
Killing 5 processes in veid 1 with signal 15
Got signal 9
Stopping VE ...
VE was stopped
===== begin /etc/vz/conf/1.umount =====
=====
Unmounting /var/lib/vz/root/1/usr
Unmounting /var/lib/vz/root/1/home
Unmounting /var/lib/vz/root/1/apps
Unmounting /var/lib/vz/root/1/.root

Removing /var/lib/vz/root/1 :root etc var tmp dev/shm dev
Removing /var/lib/vz/private/1: apps bin emul home lib lib32 lib64 mnt
  opt proc sbin sys usr .root
===== end /etc/vz/conf/1.umount =====
=====
VE is unmounted
```

Maxwell Client

Installation

The Maxwell client is installed through the hubzero-mw-client package from packages.hubzero.org:

```
# apt-get install hubzero-mw-client
```

This adds the following files:

```
/usr/share/hubzero-mw-client  
/usr/lib/hubzero/bin/maxwell  
/usr/share/hubzero-mw-client/maxwell.conf-dist
```

Configuration

1. SSH keys for the Maxwell client are created and installed in `/etc/hubzero`:

```
# ssh-keygen -t rsa -f /etc/hubzero/maxwell.key -N '' -C www-  
data@`hostname`  
# ssh-keygen -t rsa -f /etc/hubzero/notify.key -N '' -C root@`hos  
tname`  
  
# chown www-data:www-data /etc/hubzero/maxwell.key  
# chown www-data:www-data /etc/hubzero/maxwell.key.pub  
# chown www-data:www-data /etc/hubzero/notify.key  
# chown www-data:www-data /etc/hubzero/notify.key.pub  
  
# chmod 0400 /etc/hubzero/maxwell.key  
# chmod 0400 /etc/hubzero/maxwell.key.pub  
# chmod 0400 /etc/hubzero/notify.key  
# chmod 0400 /etc/hubzero/notify.key.pub
```

2. SSH is configured to allow notify and Maxwell keyed clients to connect as `www-data` and `root`:

```
# mkdir -p /root/.ssh  
# cat /etc/mw/maxwell.key.pub >> /root/.ssh/authorized_keys  
  
# mkdir -p ~www-data/.ssh
```

```
# grep -q -f /etc/hubzero/notify.key.pub ~www-data/.ssh/authorized_
d_keys || (echo -n "COMMAND=\"/usr/lib/hubzero/bin/maxwell notify
\" \" >> ~www-data/.ssh/authorized_keys
# cat /etc/hubzero/notify.key.pub >> ~www-
data/.ssh/authorized_keys
# chown www-data.www-data ~www-data/.ssh
# chmod 0700 ~www-data/.ssh
# chown www-data.www-data ~www-data/.ssh/authorized_keys
# chmod 0400 ~www-data/.ssh/authorized_keys
```

3. The sample maxwell.conf file is deployed:

```
# cp /usr/lib/mw/maxwell.conf-dist /etc/hubzero/maxwell.conf
```

4. /etc/mw/maxwell.conf is modified to configure the following variables (using values specified in the six questions step initially):

```
mysql_host = "localhost"
mysql_user="myhub"
mysql_password="<MySQL password>"
mysql_db="myhub"

default_vnc_timeout=86400
session_suffix="L"

fi
lexf
er_decor
ation="" "filexfe
r_sitelogo { <h1><a href="http:/
/myhub.org/" title="myHUB home page"><span>myhub.org
: online simulations and more</span></a></h1> }
filexfer_stylesheet http://myhub.org
/templates/filexfer/upload.css"
" " "
hub_name="myhub"
hub_url="http://myhub.org/"
hub_homedir="/home/myhub"
hub_template="hubbasic"
```

- Insert values for localhost in the MySQL 'host' table. ???

Test

```
# su www-data
$ ssh -i /etc/mw/maxwell.key root@localhost ls
The authenticity of host 'localhost (127.0.0.1)' can't be established.
RSA key fingerprint is e5:3c:7d:41:71:0b:0f:2a:0c:0e:bb:15:4d:e7:2f:08.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'localhost' (RSA) to the list of known hosts.
list of files
$ exit
#
```

VNC Client

Installation

The VNC client is downloaded from packages.hubzero.org and installed:

```
# apt-get install tightvnc-java
```

This adds the following files:

```
/usr/share/man/man1/jtightvncviewer.1.gz  
/usr/share/doc/tightvnc-java/copyright  
/usr/share/doc/tightvnc-java/changelog.Debian.gz  
/usr/share/doc/tightvnc-java/README.gz  
/usr/share/tightvnc-java/index.vnc  
/usr/share/tightvnc-java/*.class (32 files)  
/usr/share/java/tightvncviewer-1.3.8.jar  
/usr/share/java/signed-tightvncviewer-1.3.8.jar  
/usr/bin/jtightvncviewer  
/usr/share/tightvnc-java/SignedVncViewer.jar  
/usr/share/tightvnc-java/VncViewer.jar  
/usr/share/java/tightvncviewer.jar./usr/share/java/signed-  
tightvncviewer.jar
```


vncproxy

Installation

HUBzero vncproxy packages are installed from packages.hubzero.org:

```
# apt-get install hubzero-vncproxy
# apt-get install hubzero-vncproxy2
```

This adds the following files:

```
/usr/share/doc/hubzero-vncproxy2/copyright
/usr/share/doc/hubzero-vncproxy2/changelog.Debian.gz
/usr/share/doc/hubzero-vncproxy2/README.Debian
/usr/lib/hubzero/bin/vncproxy2-helper
/usr/lib/hubzero/bin/vncproxy2
/etc/default/vncproxy2
/etc/init.d/vncproxy2
/usr/share/doc/hubzero-vncproxy/changelog.Debian.gz
/usr/share/doc/hubzero-vncproxy/copyright
/usr/bin/
/usr/bin/vncproxy
/usr/lib/hubzero/bin/vncproxy.py
/etc/init.d/vncproxy
```

expire

Installation

The hubzero-expire package is installed from packages.hubzero.org:

```
# apt-get install hubzero-mw-expire
```

This adds the following files:

```
/usr/share/doc/hubzero-mw-expire/changelog.Debian.gz  
/usr/share/doc/hubzero-mw-expire/copyright  
/usr/lib/hubzero/bin/expire-sessions.py  
/etc/init.d/expire-sessions
```

telequotad

Installation

The `hubzero-telequotad` package is installed from packages.hubzero.org:

```
# apt-get install hubzero-mw-telequotad
```

This compiles `telequotad` and installs the following files:

```
/usr/share/doc/hubzero-telequotad/changelog.Debian.gz  
/usr/share/doc/hubzero-telequotad/copyright  
/usr/sbin/telequotad  
/usr/lib/hubzero/bin/telequotad  
/etc/hubzero/telequotad.conf  
/usr/lib/hubzero/bin/trim  
/etc/init.d/telequotad
```

App Container

Configuration

Users in the "apps" group are granted permission to manage HUBzero Apps by adding the following lines to `/etc/sudoers`:

```
%apps  ALL=NOPASSWD:/bin/su - apps
```

Workspace

Installation

The workspace app (tool) is installed using packages from packages.hubzero.org:

```
# apt-get install hubzero-app
# apt-get install hubzero-app-workspace
```

This adds the following files:

```
/usr/share/doc/hubzero-app/copyright
/usr/share/doc/hubzero-app/changelog.Debian.gz
/usr/bin/hubzero-app
```

The Workspace tool is then published:

```
# hubzero-app setup
# hubzero-
app install --publish /usr/lib/hubzero/apps/workspace-1.0.hza
```

Test

You should then be able to log in to the site and see the "Workspace" tool in the tool list and launch it in your browser.