

Installation

Target Audience

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

Minimum System Requirements

HUBzero installations require one or more dedicated hosts running Debian GNU/Linux version 8 (jessie).

A typical starter HUBzero installation might consist of a single physical server with dual 64-bit quad-core CPUs, 24 Gigabytes of RAM and a terabyte of disk.

Production systems should try to not limit hardware resources, HUBzero is designed to run on systems with many CPU cores and lots of RAM. If you are looking for a system to run a small site with limited physical or virtual resources this is probably not the system for you. However, for demonstration or development purposes we often create VM images with less than a gigabyte of RAM and 5 gigabytes of disk. While fully functional, these virtual machines would only be suitable for a single user doing development or testing.

System Architecture

All hardware, filesystem partitions, RAID configurations, backup models, security models, etc. and base configurations of the hosts (email server, SSH server, network, etc.) are the responsibility of the system administrator managing the host.

The Hubzero software expects to be installed on a headless server from a minimal ISO with only one network interface (required by OpenVZ) with an MTU no less than '1500'. System accounts must not be created with an id of 1000 or greater - more about that in a forthcoming section.

Linux

Install Basic Operating System

Advanced Linux system administrator skills are required, please read carefully. Selecting all the default configurations during the operating system installation may not be correct.

The latest version of Debian GNU/Linux 8 64-bit should be downloaded and installed. Do NOT install any additional packages (especially Debian Desktop Environment) during installation. All packages required by HUBzero will be installed automatically as needed.

Hubzero packaging support is for the amd64 (64 bit) Intel architecture only.

System reboots are required to complete the installation. Be sure to remove the install disk or reset your server's boot media before rebooting.

The precise server configuration (such as disk partitioning, networking, etc) is dependent on how the hub is to be used and what hardware is being used, all the possible configuration options are not specifically outlined here. This installation guide outlines a very basic configuration but may not be suitable for larger sites. For larger sites, it is generally expected that the hub will be managed by an experienced Linux administrator who can help setup your site to meet your specific requirements.

All hardware, filesystem partitions, RAID configurations, backup models, security models, etc. and base configurations of the host's email server, SSH server, network, etc. are the responsibility of the system administrator managing the host.

Configure Networking and DNS

Configure you host's network as desired. A Fully Qualified Domain Name (FQDN) is required and a static IP is highly recommended (a common basic configuration for a web server).

The Hubzero middleware uses an IP addresses in the 192.168.0.0/16 subnet, do not use a conflicting host IP during network configuration.

Set hostname

Throughout this documentation you will see specific instructions for running commands, with part of the text highlighted. The highlighted text should be modified to your local configuration choices. (e.g. replace "example.com" with the fully qualified hostname of your machine).

INSTALLATION

HUBzero expects the `hostname` command to return the fully qualified hostname for the system. This step may be skipped if previously configured.

```
sudo hostname example.com
```

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

```
sudo echo "example.com" > /etc/hostname
```

Fix hosts

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

```
127.0.1.1    example.com    example
```

Any other lines with "127.0.1.1" should be removed.

Delete local users

HUBzero reserves all user ids from 1000 up for hub accounts. As part of the HUBzero middleware every account must map to a corresponding system account. Therefore when starting up a hub it is required to remove all accounts that have user ids 1000 or greater. On a new installation there is typically one such account that is created when you set up the hub, and this account can be removed as follows:

```
sudo rm -fr /home/username  
sudo deluser username
```

If you require additional system accounts, they can be numbered between 500-999 without interfering with hub operations.

INSTALLATION

Disable SELinux

Hubzero does not currently support SELinux. Since the default install of RHEL turns it on, we have to turn it off.

```
sudo sed -i 's/^SELINUX=.* /SELINUX=disabled/g' /etc/selinux/config
```

Reboot the system for this change to take effect

```
sudo reboot
```

Configure Advanced Package Tool

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

```
sudo apt-key adv --keyserver keys.gnupg.net --recv-keys 143C99EF
```

Now configure the location of the HUBzero package repository

Add the following line to `/etc/apt/sources.list`

```
deb http://packages.hubzero.org/deb julian-deb8 main
```

With the above configured, update the local package database with information about the packages now available through the new repository:

```
sudo apt-get update
```

Install Support Packages

INSTALLATION

```
sudo apt-get install -y sudo dirmngr software-properties-common
```

Update the initial OS install

```
sudo apt-get upgrade -y
```

Firewall

Install

```
sudo apt-get install -y hubzero-iptables-basic
```

```
sudo service hubzero-iptables-basic start
```

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).

Database

Install

```
sudo DEBIAN_FRONTEND=noninteractive apt-get install -y hubzero-mysql
```

```
sudo service mysql start
```

If you leave off setting DEBIAN-FRONTED environment variable you will be prompted to enter a MySQL administrative password. This password will get reset at a later step.

Configure

Default configuration works well for starters. But for optimal performance you will need a database administrator capable of tuning your database to your hardware configuration and site usage.

Mail

Install Exim4

We need to install exim4 to enable outgoing email

```
sudo apt-get install -y exim4
```

```
sudo service exim4 start
```

Configure

```
sudo dpkg-reconfigure exim4-config
```

Configure mail as appropriate for your site and IT infrastructure. We outline a sample standalone configuration below. The requirement is for php to be able to send mail (registration confirmation and other notices need to go out) and for exim4 to receive mail (for support ticket and forum email gateway functions to work).

This is just an example of a standalone mail configuration.

General type of mail configuration

internet site; mail is sent and received directly using SMTP

Mail name

enter the fully qualified domain name (FQDN) of the host (example.com)

IP-addresses to listen on for incoming SMTP connections

leave blank (listen for connections on all available network interfaces)

Other destinations for which mail is accepted

INSTALLATION

leave blank or (equivalently) with local hostname (all local domains will be treated identically)

Domains to relay mail for

leave blank

Machines to relay mail for

leave 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?

Yes

Test

Use a real email address below so you can see if you get the email

```
sudo Mail -v someone@gmail.com
```

Web Server

Install

We need to install the Apache httpd web server

```
sudo apt-get install -y hubzero-apache2
```

```
sudo service apache2 start
```

Configure

Disable the "default" site provided by Debian

```
sudo a2dissite 000-default
```

INSTALLATION

PHP

Install

We need to install PHP

```
sudo apt-get install -y hubzero-php
```

CMS

Install

```
sudo apt-get install -y hubzero-cms-2.2
```

Configure

```
sudo hzcms install example
```

It is necessary to immediately run the updater to apply fixes that have not been incorporated into the initial installation.

```
sudo hzcms update
```

Enable the http and https hub sites

```
sudo a2ensite example example-ssl  
sudo /etc/init.d/apache2 restart
```

SSL Configuration

The default SSL certificate is meant for evaluation purposes only. For a production Hub, you will need to obtain a certificate from a provider. A certificate may contain two or three pieces: a public certificate, a private key, and sometimes an intermediate certificate.

Once you obtain the certificate, install their respective sub-directories in the **/etc/ssl** directory.

After the certificates are installed, you will need to modify the Apache configuration template located in `/etc/apache/sites-m4/{hubname}-ssl.m4`.

Replace SSLCERTFILE on line 184 with the path to your SSL certificate.

INSTALLATION

```
SSLCertificateFile /etc/ssl/path/to/cert
```

Replace SSLCERTKEYFILE on line 185 with the path to your SSL key.

```
SSLCertificateKeyFile /etc/ssl/path/to/key
```

If you have an intermediate certificate:

remove line 186: `ifdef('SSLCERTCHAINFILE',``

remove line 188: `)dnI`

Replace SSLCERTCHAINFILE with the path to your SSL key.

Once the paths have been updated, reconfigure the CMS.

```
sudo hzcms reconfigure example
sudo /etc/init.d/apache2 restart
```

If you are using the HTML5 VNC Proxy Server, [you must update your certificate settings as well.](#)

Test

The default installation of the CMS uses a self signed SSL certificate. Some browsers will not accept this certificate and not allow access to the site.

<https://support.mozilla.org/en-US/questions/1012036>

You will need to install a proper SSL certificate.

Mailgateway

Install the Hubzero Mailgateway

```
sudo apt-get install -y hubzero-mailgateway
```

Configure the Hubzero Mailgateway

```
sudo hzcms configure mailgateway --enable
```

OpenLDAP

Install HUBzero LDAP support

```
sudo apt-get install -y hubzero-openldap
```

```
sudo service slapd start
```

You will be prompted to enter a LDAP administrative password.

Some packages will ask you to configure them when you run this step

Configuring nslcd: LDAP server URI:

Enter "ldap://localhost/"

Configuring nslcd: LDAP server search base:

keep the default

Configuring libnss-ldapd

Select only "group", "passwd", "shadow"

Configure OpenLDAP Database

```
sudo hzldap init dc=example,dc=com
sudo hzcms configure ldap --enable
sudo hzldap syncusers
```

Test

```
sudo getent passwd
```

INSTALLATION

You should see an entry for user 'admin' toward the end of the list if everything is working correctly.

WebDAV

Install WebDAV

```
sudo apt-get install -y hubzero-webdav
```

Configure WebDAV

```
sudo hzcms configure webdav --enable
```

Test

```
sudo ls -l /webdav/home/admin  
total 0
```

Browse to your site's `https /webdav` address (e.g. `https://myhub/webdav`). You should get prompted for a username and password. Use the admin account. You should see an empty directory listing and no error messages.

Now test using a WebDAV client.

```
sudo apt-get install cadaver  
sudo cadaver https://localhost/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 'admin' 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

```
sudo apt-get purge cadaver
```

Troubleshooting

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

```
sudo lsmod | grep fuse  
fuse                54176  0
```

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

```
sudo modprobe fuse
```

Then try the test again

Subversion

Install

```
sudo apt-get install -y hubzero-subversion
```

Configure

```
sudo hzcms configure subversion --enable
```

INSTALLATION

Trac

Install

```
sudo apt-get install -y hubzero-trac
```

Configure

```
sudo hzcms configure trac --enable
```

INSTALLATION

Forge

Install

```
sudo apt-get install -y hubzero-forge
```

Configure

```
sudo hzcms configure forge --enable
```

OpenVZ

Install

```
sudo apt-get install -y hubzero-openvz-repo
sudo apt-get update
```

Then install

```
sudo apt-get install -y hubzero-openvz
```

Configure

```
sudo hzcms configure openvz --enable
```

If configuration is successful it should prompt you to reboot the server to activate the new kernel.

```
sudo reboot
```

Test

```
sudo vzlist
Container(s) not found
```

Or it will list the containers currently running if you check this on a running hub. The salient point being that the command doesn't issue any kind of error message.

Maxwell Client

Install

```
sudo apt-get install -y hubzero-mw2-client
sudo apt-get install -y hubzero-expire-sessions
```

Configure

```
sudo hzcms configure mw2-client --enable
sudo service expire-sessions restart
```

Maxwell File Service

Install

```
sudo apt-get install -y hubzero-mw2-file-service
```

Maxwell Service

Install

```
sudo apt-get install -y hubzero-mw2-exec-service
```

```
sudo apt-get install -y hubzero-mw2-iptables-basic
```

```
sudo service hubzero-mw2-iptables-basic start
```

Configure

```
sudo mkvztemplate amd64 wheezy ellie
```

```
sudo hzcms configure mw2-service --enable
```

```
sudo hzcms mw-  
host add localhost up openvz pubnet sessions workspace fileserver
```

Test

```
sudo 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

INSTALLATION

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-
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/mw/startxvnc: Becoming nobody.
/usr/lib/mw/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)
```

```
sudo vzlist
      VEID      NPROC STATUS  IP_ADDR      HOSTNAME
      1          6 running 10.51.0.1    -
```

INSTALLATION

```
sudo 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

```
sudo 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 =====  
=====
```

INSTALLATION

VE is unmounted

VNC Proxy Server (HTML5)

Install

```
sudo apt-get install -y hubzero-vncproxyd-ws
```

Configure

```
sudo hzvncproxyd-ws-config configure --enable
```

```
sudo service hzvncproxyd-ws start
```

Install SSL certificate files

Copy your SSL certificate files to `/etc/hzvncproxyd-ws/ssl-cert-hzvncproxyd-ws.pem` and `/etc/hzvncproxyd-ws/ssl-cert-hzvncproxyd-ws.key` and make sure they are readable by the user "hznvncproxy" and restart the hznvncproxy service. Commands are below. Be sure to replace "mycertfile" and "mycertkey" with the correct filename and path of your certificate file.

The following will accomplish this configuration:

```
sudo cp "mycertfile" /etc/hzvncproxyd-ws/ssl-cert-hzvncproxyd-ws.pem
sudo cp "mycertkey" /etc/hzvncproxyd-ws/ssl-cert-hzvncproxyd-ws.key
```

```
cd /etc/hzvncproxyd-ws
```

```
sudo chown hznvncproxy:hznvncproxy /etc/hzvncproxyd-ws/ssl-cert-hznvncproxyd-ws.pem
```

```
sudo chown hznvncproxy:hznvncproxy /etc/hzvncproxyd-ws/ssl-cert-hznvncproxyd-ws.key
```

```
sudo chmod ug+r /etc/hzvncproxyd-ws/ssl-cert-hznvncproxyd-ws.pem
```

```
sudo chmod ug+r /etc/hzvncproxyd-ws/ssl-cert-hznvncproxyd-ws.key
```

When done, the output of "ls - al" should appear like this:

INSTALLATION

```
# ls -al
total 32
drwxr-xr-x.  3 root          root          4096 Dec 18  2018 .
drwxr-xr-x. 127 root          root          12288 Sep  6 14:20 ..
-rw-r-----. 1 hzvncproxy hzvncproxy  3243 May  9  2018 ssl-cert-
hzvncproxyd-ws.key
-rw-r-----. 1 hzvncproxy hzvncproxy  4167 May  9  2018 ssl-cert-
hzvncproxyd-ws.pem
drwxr-x---.  2 hzvncproxy root           4096 Feb 13  2019 targets
```

If you are using a self-signed or otherwise invalid certificate the tool viewer will likely reject it and not work. If you are using the same certificate as your website and you allowed Chrome to use the invalid cert then the tool viewer will probably accept it. If you are using Firefox the tool viewer will always reject the invalid certificate. Always use a valid SSL certificate with hzvncproxyd-ws.

telequotad

install

```
sudo apt-get install -y hubzero-telequotad
```

```
sudo service telequotad start
```

Configure

In order for filesystems quotas to work they must be enabled when they are mounted. Determine which filesystem contains your home directories and add "quota" to the mount option of the corresponding entry in the `/etc/fstab` file. Only the filesystem with `/home` on it matters to telequotad.

If quotas weren't already in affect, the run something like the following (depending on your filesystem configuration) to start up the quota system.

```
sudo mount -oremount /
sudo /etc/init.d/quota restart
sudo hzcms configure telequotad --enable
```

Test

```
sudo repquota -a
```

Should show disk usage for all users.

Workspace

Install

```
sudo apt-get install hubzero-app
sudo apt-get install hubzero-app-workspace
sudo hubzero-
app install --publish /usr/share/hubzero/apps/workspace-1.3.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.

Filexfer

Install

```
sudo apt-get install -y hubzero-filexfer-xlate
```

Configure

```
sudo hzcms configure filexfer --enable
```

Rappture

Install

Install rappture for Debian 7 session containers

```
sudo apt-get install hubzero-rappture-deb7
```

Configure

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.

```
sudo chroot /var/lib/vz/template/debian-7.0-amd64-maxwell
sudo apt-get update
sudo apt-get upgrade
sudo apt-get install hubzero-rappture-session
sudo 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

A user must setup their runtime environment in order to use the Rappture toolkit. Run the following command before attempting to run any Rappture tests.

```
use rappture
```

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.

INSTALLATION

```
$ mkdir examples
$ cp -r /apps/share/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).

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 or executed directly on remote resources.

Installation

```
sudo apt-get install hubzero-submit-pegasus
sudo apt-get install hubzero-submit-condor
sudo apt-get install hubzero-submit-common
sudo apt-get install hubzero-submit-server
sudo apt-get install hubzero-submit-distributor
sudo apt-get install hubzero-submit-monitors
sudo hzcms configure submit-server --enable
sudo /etc/init.d/submit-server start
```

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:




```
Session Edit View Bookmarks Settings Help
$ ls -a -I "[A-Z]*.py" -I "[A-Z]*.pyc" -I "." -I ".."
.ssh          distributor.sh      monitorJob.py      monitorTunnelA.py  server.py
BatchMonitors environmentwhitelist.dat monitorJobDB        monitorTunnelD.py  sites.dat
Scripts       environmentwhitelist.dft monitorJobQ.py      monitorTunnelI.py  sites.dft
bin           etc                monitorJobR.py     monitorTunnelR.py  tools.dat
config        managers.dat       monitorJobS.py     monitorTunnelT.py  tools.dft
distributor   managers.dft       monitorJobT.py     monitors.dat        tunnels.dat
distributor.py monitorJob.dump     monitorTunnel.py   monitors.dft        tunnels.dft
```

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


INSTALLATION

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 further assumes that access to this account can be made by direct ssh login or using an ssh tunnel (port forwarding).



- Domain name
- Single user account
 - » Community user
 - » Home directory
 - » Application staging directory
 - » Scratch directory
- Batch job submission classification

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).



```
Session Edit View Bookmarks Settings Help
$ cat .ssh/submit_rsa.pub
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEA07poA77spui+ru
lXao63KsIGKwj9/ea0vVi2DgE5xUJ80bGhP/jhpXdLkxWSPLyp
4ZlACGJwv3tI0rji4UjimQW85m0XsMlLr8oIHfl1Sekg6V0V7p
gVvvp9dwL8G4upG50wXmNFJpWPGw== submit@hubzero.org
```

```
Session Edit View Bookmarks Settings Help
$ cat ${HOME}/.ssh/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEA07poA77spui+ru
lXao63KsIGKwj9/ea0vVi2DgE5xUJ80bGhP/jhpXdLkxWSPLyp
4ZlACGJwv3tI0rji4UjimQW85m0XsMlLr8oIHfl1Sekg6V0V7p
gVvvp9dwL8G4upG50wXmNFJpWPGw== submit@hubzero.org
```

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, file access controls, 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.



HUB

- Remote sites
- Staged tools
- Remote job monitors
- Multi-processor managers
- Permissible environment variables
- ssh tunnels

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, SGE, 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
- parallelEnvironment - sge parameter
- 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. (deprecated)
- executableClassificationsAllowed - classifications accepted by site. Classifications are set in appaccess.dat
- 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.
- minimumWallTime - minimum walltime allowed for site or queue. Time should be expressed in minutes.
- maximumWallTime - maximum walltime allowed for site or queue. Time should be expressed in minutes.
- minimumCores - minimum number of cores allowed for site or queue.
- maximumCores - maximum number of cores allowed for site or queue.
- pegasusTemplates - pertinent pegasus templates for site, rc, and transaction files.

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

INSTALLATION



```
Session Edit View Bookmarks Settings Help
$ hostname -f
cluster.campus.edu
$ whoami
yourhub
$ echo ${HOME}
/home/yourhub
$ printenv | SCRATCH
CLUSTER_SCRATCH=/scratch/yourhub
```

```
[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 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.
- 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

INSTALLATION

enabled.

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



cluster.campus.edu

```
Session Edit View Bookmarks Settings Help
$ cd ${HOME}
$ ls -R
.:
apps
./apps:
planets stars
./apps/planets:
bin
./apps/planets/bin:
earth.x jupiter.x mars.x mercury.x neptune.x saturn.x uranus.x venus.x
./apps/stars:
bin
./apps/stars/bin:
antares.x betelgeuse.x polaris.x sun.x
```

```
[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.

INSTALLATION

- remoteUser - login user at remote site.
- remoteBinDirectory - define directory where shell scripts related to the site should be kept.
- 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.



```
Session Edit View Bookmarks Settings Help
$ qstat -u yourhub
cluster.campus.edu:
Job ID      Username Queue  Jobname SessID NDS TSK Req'd Elap
-----
9823388.steele- yourhub standby earth 3508 2 16 04:00 R 02:15
9824065.steele- yourhub standby sun 9975 1 1 04:00 R 01:26
```

```
[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.

INSTALLATION

- `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.



```
Session Edit View Bookmarks Settings Help
$ module available mpich
----- /opt/modules/modulefiles -----
mpich-intel/10.1.025          mpich-intel/9.1.045
mpich-intel/11.1.038(default)  mpich2-intel/11.1.038(default)

$ module load mpich-intel/11.1.038
$ which mpirun
/apps/rhel5/mpich-1.2.7p1/p4-intel-11.1.038/bin/mpirun
```

```
[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.

File access controls

Application or file level access control is described by entries listed in the file `appaccess.dat`. The ability to transfer files from the HUB to remote sites is granted on a group basis as defined by white and black lists. Each list is given a designated priority and classification. In cases where a file appears on multiple lists, the highest priority takes precedence. Simple wildcard operators are allowed in the filename declaration allowing for easy listing of entire directories. Each site lists acceptable classification(s) in `sites.dat`. Defined keywords are

- `[group]` - group name.
- `whitelist` - comma separated list of paths. Wildcards allowed.

INSTALLATION

- blacklist - comma separated list of paths. Wildcards allowed.
- priority - higher priority wins
- classification - apps or user. user class are treated as arbitrary executables.
- state - possible values are enabled or disabled. If not explicitly set the default value is enabled.

An example file giving permissions reminiscent of those defined in earlier submit releases is presented here

```
[public]
whitelist = /apps/*.
priority = 0
classification = apps
```

```
[submit]
whitelist = ${HOME}/*.
priority = 0
classification = home
```

The group public is intended to include all users. Your system may use a different group such as users for this purpose. The definitions shown here allow all users access to files in /apps where applications are published. Additionally members of the submit group are allowed to send files from their \$HOME directory.

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.
```

INSTALLATION

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 an accepted gateway host.



```
[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-server.log, distributor.log, and monitorJob.log.

submit.log

The submit-server.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 submit-server process ID (PID) or submit ID (ID:) once one 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.

```
[Sun Aug 26 17:28:24 2012] 0: #####  
#####
```

```
[Sun Aug 26 17:28:24 2012] 0: Backgrounding process.
```

```
[Sun Aug 26 17:28:24 2012] 0: Listening: protocol='tcp', host='', port  
=830
```

```
[Sun Sep 23 12:33:28 2012] (1154) =====  
=====
```

INSTALLATION

```
[Sun Sep 23 12:33:28 2012] (1154) Connection to tcp://:830 from ('192.168.224.14', 38770)
[Sun Sep 23 12:33:28 2012] 0: Server will time out in 60 seconds.
[Sun Sep 23 12:33:28 2012] 0: Cumulative job load is 0.84. (Max: 510.00)
[Sun Sep 23 12:33:28 2012] 1670: Args are:['/usr/bin/submit', '--local', '-p', '@iv=-3:1.5:3', '/home/hubzero/user/hillclimb/bin/hillclimb1.py', '--seed', '10', '--initialvalue', '@iv', '--lowerbound', '-3', '--upperbound', '3', '--function', 'func2', '--solutionslog', 'solutions.dat', '--bestresultlog', 'best.dat']
[Sun Sep 23 12:33:28 2012] 1670: Server stopping.
[Sun Sep 23 12:33:28 2012] 1670: Server(JobExecutor) exiting(2).
[Sun Sep 23 12:33:38 2012] (1154) =====
[Sun Sep 23 12:33:38 2012] (1154) Connection to tcp://:830 from ('192.168.224.14', 38774)
[Sun Sep 23 12:33:38 2012] 0: Server will time out in 60 seconds.
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue=1:local status=0 cpu=0.030000 real=0.000000 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue=2:local status=0 cpu=0.040000 real=0.000000 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue=3:local status=0 cpu=7.050000 real=7.000000 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue=4:local status=0 cpu=0.080000 real=0.000000 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue=5:local status=0 cpu=0.020000 real=1.000000 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Job Status: venue= status=0 cpu=10.428651 real=9.561828 wait=0.000000
[Sun Sep 23 12:33:38 2012] 1670: Server(JobExecutor) exiting(0).
[Sun Sep 23 12:48:44 2012] (1154) =====
[Sun Aug 26 17:28:17 2012] 0: Server(10836) was terminated by a signal 2.
[Sun Aug 26 17:28:17 2012] 0: Server(Listener) exiting(130).
```

distributor.log

The distributor.log file tracks each job as it progresses from start to finish. Details of remote site

INSTALLATION

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-server.log. An example for submit ID 1659 is listed here. Again the data for all jobs are intermingled.

```
[Sun Sep 23 00:04:21 2012] 0: quotaCommand = quota -w | tail -n 1
[Sun Sep 23 00:04:21 2012] 1659: command = tar vchf 00001659_01_input.
tar --exclude='*.svn*' -C /home/hubzero/user/data/sessions/3984L .__lo
cal_jobid.00001659_01 sayhiinquire.dax
[Sun Sep 23 00:04:21 2012] 1659: remoteCommand pegasus-
plan --dax ./sayhiinquire.dax
[Sun Sep 23 00:04:21 2012] 1659: workingDirectory /home/hubzero/user/d
ata/sessions/3984L
[Sun Sep 23 00:04:21 2012] 1659: command = tar vrhf 00001659_01_input.
tar --exclude='*.svn*' -C /home/hubzero/user/data/sessions/3984L/00001
659/01 00001659_01.sh
[Sun Sep 23 00:04:21 2012] 1659: command = nice -n 19 gzip 00001659_01
_input.tar
[Sun Sep 23 00:04:21 2012] 1659: command = /opt/submit/bin/receiveinpu
t.sh /home/hubzero/user/data/sessions/3984L/00001659/01 /home/hubzero/
user/data/sessions/3984L/00001659/01/.__timestamp_transferred.00001659
_01
[Sun Sep 23 00:04:21 2012] 1659: command = /opt/submit/bin/submitbatch
job.sh /home/hubzero/user/data/sessions/3984L/00001659/01 ./00001659_0
1.pegasus
[Sun Sep 23 00:04:23 2012] 1659: remoteJobId = 2012.09.23 00:04:22.996
EDT: Submitting job(s).
2012.09.23 00:04:23.002 EDT: 1 job(s) submitted to cluster 946.
2012.09.23 00:04:23.007 EDT:
2012.09.23 00:04:23.012 EDT: -----
-----
2012.09.23 00:04:23.017 EDT: File for submitting this DAG to Condor
: sayhi_inquire-0.dag.condor.sub
2012.09.23 00:04:23.023 EDT: Log of DAGMan debugging messages
: sayhi_inquire-0.dag.dagman.out
2012.09.23 00:04:23.028 EDT: Log of Condor library output
: sayhi_inquire-0.dag.lib.out
2012.09.23 00:04:23.033 EDT: Log of Condor library error messages
: sayhi_inquire-0.dag.lib.err
2012.09.23 00:04:23.038 EDT: Log of the life of condor_dagman itself
: sayhi_inquire-0.dag.dagman.log
2012.09.23 00:04:23.044 EDT:
2012.09.23 00:04:23.049 EDT: -----
-----
2012.09.23 00:04:23.054 EDT:
2012.09.23 00:04:23.059 EDT: Your Workflow has been started and runs
```

INSTALLATION

```
in base directory given below
2012.09.23 00:04:23.064 EDT:
2012.09.23 00:04:23.070 EDT:   cd /home/hubzero/user/data/sessions/398
4L/00001659/01/work/pegasus
2012.09.23 00:04:23.075 EDT:
2012.09.23 00:04:23.080 EDT:   *** To monitor the workflow you can run
***
2012.09.23 00:04:23.085 EDT:
2012.09.23 00:04:23.090 EDT:   pegasus-status -l /home/hubzero/user/da
ta/sessions/3984L/00001659/01/work/pegasus
2012.09.23 00:04:23.096 EDT:
2012.09.23 00:04:23.101 EDT:   *** To remove your workflow run ***
2012.09.23 00:04:23.106 EDT:   pegasus-remove /home/hubzero/user/data/
sessions/3984L/00001659/01/work/pegasus
2012.09.23 00:04:23.111 EDT:
2012.09.23 00:04:23.117 EDT:   Time taken to execute is 0.993 seconds
[Sun Sep 23 00:04:23 2012] 1659: confirmation: S(1):N Job
[Sun Sep 23 00:04:23 2012] 1659: status:Job N WF-DiaGrid
[Sun Sep 23 00:04:38 2012] 1659: status:DAG R WF-DiaGrid
[Sun Sep 23 00:10:42 2012] 0: quotaCommand = quota -w | tail -n 1
[Sun Sep 23 00:10:42 2012] 1660: command = tar vchf 00001660_01_input.
tar --exclude='*.svn*' -C /home/hubzero/clarksm .__local_jobid.0000166
0_01 noerror.sh
[Sun Sep 23 00:10:42 2012] 1660: remoteCommand ./noerror.sh
[Sun Sep 23 00:10:42 2012] 1660: workingDirectory /home/hubzero/clarks
m
[Sun Sep 23 00:10:42 2012] 1660: command = tar vrhf 00001660_01_input.
tar --exclude='*.svn*' -C /home/hubzero/clarksm/00001660/01 00001660_0
1.sh
[Sun Sep 23 00:10:42 2012] 1660: command = nice -n 19 gzip 00001660_01
_input.tar
[Sun Sep 23 00:10:42 2012] 1660: command = /opt/submit/bin/receiveinpu
t.sh /home/hubzero/clarksm/00001660/01 /home/hubzero/clarksm/00001660/
01/.__timestamp_transferred.00001660_01
[Sun Sep 23 00:10:42 2012] 1660: command = /opt/submit/bin/submitbatch
job.sh /home/hubzero/clarksm/00001660/01 ./00001660_01.condor
[Sun Sep 23 00:10:42 2012] 1660: remoteJobId = Submitting job(s).
1 job(s) submitted to cluster 953.
[Sun Sep 23 00:10:42 2012] 1660: confirmation: S(1):N Job
[Sun Sep 23 00:10:42 2012] 1660: status:Job N DiaGrid
[Sun Sep 23 00:11:47 2012] 1660: status:Simulation I DiaGrid
[Sun Sep 23 00:12:07 2012] 1660: Received SIGINT!
[Sun Sep 23 00:12:07 2012] 1660: waitForBatchJobs: nCompleteRemoteJobI
ndexes = 0, nIncompleteJobs = 1, abortGlobal = True
[Sun Sep 23 00:12:07 2012] 1660: command = /opt/submit/bin/killbatchjo
b.sh 953.0 CONDOR
```

INSTALLATION

```
[Sun Sep 23 00:12:07 2012] 1660: Job 953.0 marked for removal

[Sun Sep 23 00:12:07 2012] 1660: status:Simulation I DiaGrid
[Sun Sep 23 00:12:52 2012] 1660: status:Simulation D DiaGrid
[Sun Sep 23 00:12:52 2012] 1660: venue=1:localCONDOR:953.0:DiaGrid sta
tus=258 cputime=0.000000 realtime=0.000000 waittime=0.000000 ncpus=1
[Sun Sep 23 00:28:14 2012] 1659: status:DAG D WF-DiaGrid
[Sun Sep 23 00:28:14 2012] 1659: waitForBatchJobs: nCompleteRemoteJobI
ndexes = 1, nIncompleteJobs = 0, abortGlobal = False
[Sun Sep 23 00:28:14 2012] 1659: command = /opt/submit/bin/cleanupjob.
sh /home/hubzero/user/data/sessions/3984L/00001659/01
[Sun Sep 23 00:28:15 2012] 1659:
*****SUMMARY*****
*****

Job instance statistics           : /home/hubzero/user/data/sessions/3
984L/00001659/01/work/pegasus/statistics/jobs.txt

*****
*****

[Sun Sep 23 00:28:15 2012] 1659: venue=1:localPEGASUS:946.0:WF-DiaGrid
status=0 cputime=1.430000 realtime=2.000000 waittime=0.000000 ncpus=1
[Sun Sep 23 00:28:15 2012] 1659: venue=2:PEGASUS:952.0:DiaGrid status=
0 cputime=0.003000 realtime=0.000000 waittime=681.000000 ncpus=1 event
=/sayhi_inquire-sayhi-1.0
[Sun Sep 23 00:28:15 2012] 1659: venue=3:PEGASUS:954.0:DiaGrid status=
0 cputime=0.003000 realtime=0.000000 waittime=631.000000 ncpus=1 event
=/sayhi_inquire-inquire-1.0
```

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:

```
[Sun Aug 26 17:29:16 2012] (1485) *****
[Sun Aug 26 17:29:16 2012] (1485) * distributor job monitor started *
[Sun Aug 26 17:29:16 2012] (1485) *****
[Sun Aug 26 17:29:16 2012] (1485) loading active jobs
[Sun Aug 26 17:29:16 2012] (1485) 15 jobs loaded from DB file
```

INSTALLATION

```
[Sun Aug 26 17:29:16 2012] (1485) 15 jobs loaded from dump file
[Sun Aug 26 17:29:16 2012] (1485) 4 jobs purged
[Sun Aug 26 17:29:16 2012] (1485) 11 monitored jobs
[Sun Aug 26 18:02:04 2012] (24250) Launching wf-diagrid
[Sun Aug 26 18:02:04 2012] (1485) 12 monitored jobs
[Sun Aug 26 18:02:15 2012] (1485) Update message received from wf-
diagrid
[Sun Aug 26 18:03:15 2012] (1485) Update message received from wf-
diagrid
[Sun Aug 26 18:06:43 2012] (1485) 13 monitored jobs
...
[Thu Sep 17 17:32:51 2011] (21095) Received SIGTERM!
[Thu Sep 17 17:32:51 2011] (21095) Send TERM to child ssh process
[Thu Sep 17 17:32:51 2011] (21095) distributor site monitor stopped
[Thu Sep 17 17:32:51 2011] (17348) Send TERM to child site steele proc
ess
[Thu Sep 17 17:32:51 2011] (17348) *****
[Thu Sep 17 17:32:51 2011] (17348) * distributor job monitor stopped *
[Thu Sep 17 17:32:51 2011] (17348) *****
```


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.




- ssh public key
- Batch job monitoring daemon
- File transfer scripts
- Batch job scripts

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:



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



```
Session Edit View Bookmarks Settings Help
$ cd ${HOME}/Submit
$ ls -R
.:
LogMessage.py LogMessage.pyc logs monitorPBS.history monitorPBS.py
./logs:
monitorPBS.log
```

```
$ cat ${HOME}/Submit/monitorPBS.py
...
SITEDESIGNATOR = "clusterPBS"
MONITORROOT = os.path.join(os.sep, 'home', 'yourhub', 'Submit')
QSTATCOMMAND = "/usr/pb/bin/qstat -u yourhub"
MONITORLOGLOCATION = os.path.join(MONITORROOT, 'logs')
...
```

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

File transfer and batch job scripts

INSTALLATION

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.



```
$ cat sites.dat
[clusterPBS]
venue = cluster.campus.edu
remoteUser = yourhub
remoteBinDirectory = ${HOME}/bin
```



```
Session Edit View Bookmarks Settings Help
$ cd ${HOME}/bin
$ ls
cleanupjob.sh  receiveinput.sh  reportbatchjobstatus.sh  transmitresults.sh
killbatchjob.sh  remotematlab.sh  submitbatchjob.sh  waitforbatchjob.sh
```

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

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

INSTALLATION

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 jobScratchDirectory jobClass
```

Batch job script - termination

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

```
killbatchjob.sh jobId jobClass
```

Batch job script - post process

For some jobClasses it is appropriate to preform standard post processing actions. An example of such a jobClass is Pegasus.

```
postprocessjob.sh jobWorkingDirectory jobScratchDirectory 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.

INSTALLATION

In addition to applying user and group restrictions another mechanism is provided by the `executableClassificationsAllowed` keyword 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 be granted permission to send files from their `$HOME` directory. Discrete permission can be granted on a file by file basis in `appaccess.dat`.