

User Guide for Brecca *[Revised: 20080603]*

Introduction

Brecca is a cluster of Intel 32-bit machines with CentOS5, installed at Monash Clayton campus, and managed by VPAC. It is accessible only using Globus or equivalent client tools. Command-line versions of such tools are installed on the Monash Sun Grid head-node `hn3.its.monash.edu.au`, or you may install them on your own workstation; further details about these options can be found on the MSG User Guide at: <http://www.monash.edu.au/eresearch/services/mcg/msg.html>

This document provides a brief outline of the procedures which are applicable for Monash people using Brecca.

Obtaining a Personal Certificate

A personal certificate is used to identify you when you use Globus client tools to submit a job to a host which may be at another site. Such a host will need to have a trust-relationship with the issuer of your personal certificate. It is therefore recommended that you obtain your personal certificate from a Certificate Authority (CA) accredited by the International Grid Trust Federation (IGTF).

One such IGTF-accredited CA is that operated by the ARCS Certificate Authority. Some instructions requesting an APACGrid certificate can be found at: <http://www.arcs.org.au/GridGrix>

Arranging Virtual Organisation Membership

When you identify yourself to Brecca with your Personal Certificate, it needs to ascertain the role you are playing during your current session so that it can determine the Linux userid to which you should be mapped, and how you should be billed.

It does this by making callouts to a GUMS server, which in turn gets information about you from the ARCS VOMS server.

You can become a member of an appropriate Virtual Organisation as shown at: <http://gridaus.org.au/node/28>

Alternatively, you can ensure that your Personal Certificate has been installed in your browser, then visit: <https://vomrs.arcs.org.au/vomrs/ARCS/vomrs>

Expand “Member Info”, select “Select Groups and Roles”, tick “Monash” (or other appropriate Group), and hit “Submit”. Then await an email from the Group Manager concerning your next step.

Submitting Simple Globus Jobs

You should ensure that you have a personal key and its associated certificate installed on your workstation or submit-host in your `~/.globus` directory (as `userkey.pem` and `usercert.pem`), and that the directory `~/.globus/scratch` exists. If you are logged in to `hn3.its.monash.edu.au` you can proceed as follows:

- `./opt/vdt/setup.sh ..` establishes the environment so that Globus client tools can be found. On some machines (e.g. `hn3.its.monash.edu.au`) `vdt` has been defined as a system-wide alias for this, so you need only enter: `vdt`
- `voms-proxy-init -voms APACGrid -order /ARCS/Monash ..` creates a proxy valid (by

default) for 12 hours for you to use the 'Monash' role; use a different role here according to the tasks you may be undertaking.

- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft Fork -c /bin/mkdir -pv .globus/scratch ..` submits a WS-GRAM Fork job and awaits return of output.
- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft PBS -c /bin/hostname ..` submits a WS-GRAM PBS job and awaits return of output.
- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft Fork -c /bin/pwd ..` prints your home directory on the cluster

More Complex Globus Jobs

You can copy to your scratch directory a Fortran program which calculates π using MPI, then compile and execute it as follows:

- `globus-url-copy gsiftp://gridftp.its.monash.edu.au/opt/MSG/Demo/pi3a.f \ gsiftp://ng2.monash.edu.au/HomeDirectoryAsAbove/.globus/scratch/pi3a.f`
- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft PBS -f gt4-f77.rsl ..` where `gt4-f77.rsl` is a file containing:

```
<job>
  <executable>mpif77</executable>
  <directory>${GLOBUS_SCRATCH_DIR}</directory>
  <argument>pi3a.f</argument>
  <jobType>single</jobType>
  <fileCleanUp>
    <deletion><file>file:///${GLOBUS_SCRATCH_DIR}/pi3a.f</file></deletion>
  </fileCleanUp>
  <extensions>
    <module>openmpi</module>
  </extensions>
</job>
```

- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft PBS -f gt4-pi3a.rsl ..` where `gt4-pi3a.rsl` is a file containing:

```
<job>
  <executable>${GLOBUS_SCRATCH_DIR}/a.out</executable>
  <argument>120</argument>
  <count>2</count>
  <queue>run_1_hour</queue>
  <jobType>mpi</jobType>
  <fileCleanUp>
    <deletion><file>file:///${GLOBUS_SCRATCH_DIR}/a.out</file></deletion>
  </fileCleanUp>
  <extensions>
    <module>openmpi</module>
  </extensions>
</job>
```

The line containing the queue-name is optional, and in many instances your needs may better be served by declaring a maximum allowed walltime, as shown in the next example.

When there are several users whose Certificates are mapped to a common `userid`, there is a danger that different users will be using the same file names at the same time. You can get a program which

demonstrates some mechanisms to circumvent this as follows:

```
globus-url-copy gsiftp://gridftp.its.monash.edu.au/opt/MSG/Demo/Ggzip \  
file:///tmp/
```

- **Then:** `chmod a+rx /tmp/Ggzip`
- **And:** `/tmp/Ggzip -g ng2.vpac.monash.edu.au -j PBS \
gsiftp://hn3.its.monash.edu.au/etc/termcap \
gsiftp://hn3.its.monash.edu.au/tmp/tc.gz3`

A Java Example

Here's an example of a Java program which downloads an executable and uses it to (wait for it!) compute the value of 'pi' using a Brecca execute node:

- `globusrun-ws -submit -s -S -F ng2.vpac.monash.edu.au -Ft PBS -f gt4-pi.rsl ..`
where `gt4-pi.rsl` is a file containing:

```
<job>  
<executable>java</executable>  
<directory>${GLOBUS_SCRATCH_DIR}</directory>  
<argument>-jar</argument>  
<argument>bigpi.jar</argument>  
<argument>1024</argument>  
<maxWallTime>15</maxWallTime>  
<maxMemory>2000</maxMemory>  
<jobType>single</jobType>  
<fileStageIn>  
  <transfer>  
    <sourceUrl>gsiftp://gridftp.its.monash.edu.au/opt/MSG/Demo/bigpi.jar</sourceUrl>  
    <destinationUrl>file:///${GLOBUS_SCRATCH_DIR}/bigpi.jar</destinationUrl>  
  </transfer>  
</fileStageIn>  
<fileCleanUp>  
  <deletion><file>file:///${GLOBUS_SCRATCH_DIR}/bigpi.jar</file></deletion>  
</fileCleanUp>  
<extensions>  
  <module>java</module>  
</extensions>  
</job>
```

This example also shows how you can (if you wish) specify a maximum allowed walltime of 15 minutes, and a required memory of 2000M (since some nodes have less than that).

If the executable is located in your bin subdirectory, you can use it by replacing the source URL with:
`gsiftp://hn3.its.monash.edu.au/HomeDirectory/bin/bigpi.jar`

Compiling on Brecca

A compile node is available for those who need to compile their own Fortran and C programs. You can use it as shown in the following example:

- `gjenkins@hn3 src> gsiscp pi3a.f brecca.vpac.monash.edu.au:/tmp`
- `gjenkins@hn3 src> gsissh brecca.vpac.monash.edu.au`
- `[Graham Jenkins@brecca ~]$ module load openmpi`
- `[Graham Jenkins@brecca ~]$ mpif77 -o pi3a /tmp/pi3a.f`
- `[Graham Jenkins@brecca ~]$ pwd`
- `/home/grid-monash/C_AU_O_APACGrid_OU_Monash_University_CN_Graham_Jenkins`

You can then access your compiled program by submitting a Globus job through `ng2.vpac.monash.edu.au` as shown in the earlier paragraph.

Using NAMD on Brecca

The `NAMD` parallel molecular dynamics package is available on `Brecca`. The script hereunder illustrates how the `alanin` example found under the `NAMD` installation directory can be submitted to it; notice how it handles copying of files from your machine (e.g. `hn3.its.monash.edu.au`), and back to it.

```
#!/bin/sh
# Gnamd          Submits a sample NAMD program to a designated GT4 server.
#               <Graham.Jenkinss@its.monash.edu.au> Nov. 2007. Rev'd: 20071113.

#
# Default values for Gateway, Jobmanager
Gateway=hn3; Jobman=SGE

#
# Adjust options, check usage, build the job
while getopts g:j:w: Option; do
  case $Option in
    g) Gateway=$OPTARG;;
    j) Jobman=$OPTARG;;
    w) Walltime=$OPTARG;;
    \?) Bad=Y;;
  esac
done
shift `expr $OPTIND - 1`
if [ \( -z "$1" \) -o \( -n "$Bad" \) ]; then
  (echo "Usage: `basename $0` [-g Gateway] [-j Jobman] [-w Walltime] Config"
  echo "  e.g.: `basename $0` -g ng2.vpac.monash.edu.au -j PBS \\"
  echo "          gsiftp://hn3.its.monash.edu.au/tmp/Fred/alanin"
  echo " Note: Input files must reside in same directory as Config file.") >&2
  exit 2
fi
SrcDir=`dirname $1`
SrcFil=`basename $1`
Id=`basename $0`.$HOSTNAME.$$
trap 'rm -f /tmp/$Id exit 0' 1 2 3 14 15
echo "Gateway: $Gateway Walltime: $Walltime"
echo "Output-Dir: ${SrcDir}.$Id"
[ -n "$Walltime" ] && wString="<maxWallTime>$Walltime</maxWallTime>" ||
wString="<!-- -->"

cat >/tmp/$Id <<EOF
<job>
  <executable>namd2</executable>
  <directory>/\${GLOBUS_SCRATCH_DIR}</directory>
  <argument>dir.$Id/$SrcFil</argument>
  <count>4</count>
  $wString
  <jobType>mpi</jobType>
  <fileStageIn>
    <transfer>
      <sourceUrl>$SrcDir</sourceUrl>
      <destinationUrl>file:///\/\${GLOBUS_SCRATCH_DIR}/dir.$Id/</destinationUrl>
    </transfer>
  </fileStageIn>
  <fileStageOut>
    <transfer>
      <sourceUrl>file:///\/\${GLOBUS_SCRATCH_DIR}/dir.$Id/</sourceUrl>
      <destinationUrl>${SrcDir}.$Id</destinationUrl>
    </transfer>
</EOF>
```

```
</fileStageOut>
<fileCleanUp>
  <deletion><file>file:///\\${GLOBUS_SCRATCH_DIR}/dir.$Id/</file></deletion>
</fileCleanUp>
<extensions>
  <module>namd</module>
</extensions>
</job>
EOF

#
# Submit the job, then cleanup and exit.
globusrun-ws -submit -s -S -F $Gateway -Ft $Jobman -f /tmp/$Id
rm -f /tmp/$Id
exit 0
```

Other Major Applications

VPAC will be offering courses in using applications like Underworld, and in some instances these will be accessible through with GUI front-ends. Please advise them of your needs in this regard.