

Globus Procedures on LONI Systems

Alexander B. Pacheco

User Services Consultant
LSU HPC & LONI
sys-help@loni.org

LSU HPC Training Series
Louisiana State University
November 10, 2010



- LONI Systems

LINUX

- Eric
- Louie
- Oliver
- Painter
- Poseidon
- Queenbee

AIX

- Bluedawg
- Ducky
- Lacumba
- Neptune
- Zeke

- How do I manage data and jobs on the LONI systems?
- Do I need to login to all systems or
- Is there a tool to manage this



- Globus Toolkit for building computing grids.
- Provides
 - Grid Resource Allocation Manager (GRAM) for managing jobs
 - GridFTP, a more reliable and high performance file transfer for Grid computing applications

<http://www.cct.lsu.edu/~apacheco/tutorials/globus.php>
<http://www.cct.lsu.edu/~apacheco/tutorials/globus.pdf>



- Access to LONI Account
 - Apply for LONI accounts at <https://allocations.loni.org>

- Active LONI Grid Certificate
- Login to qb1.loni.org

```
[apacheco@llacinski-1 ~] ssh qb1.loni.org
Password:
Last login: Mon Nov  1 09:49:17 2010 from 130.70.53.27
```

- Add globus to your environment and resoft

```
[apacheco@qb1 ~]$ soft add +globus-4.0.8-r2
[apacheco@qb1 ~]$ resoft
```



- Request LONI Grid Certificate

```
[apacheco@qb1 ~]$ grid-cert-request
```

```
A certificate request and private key is being created.  
You will be asked to enter a PEM pass phrase.  
This pass phrase is akin to your account password,  
and is used to protect your key file.  
If you forget your pass phrase, you will need to  
obtain a new certificate.
```

```
Generating a 1024 bit RSA private key
```

```
.....++++++  
.....++++++
```

```
writing new private key to '/home/apacheco/.globus/userkey.pem'
```

```
Enter PEM pass phrase:
```

- To recreate a certificate

```
[apacheco@qb1 ~]$ grid-cert-request -force
```



- .globus directory will be created in the home directory with three files: `usercert_request.pem`, `userkey.pem` and `usercert.pm`
- email `usercert_request.pem` file to **ca@loni.org**

```
[apacheco@qb1 ~]$ cat $HOME/.globus/usercert_request.pem | mail ca@loni.org
```
- Copy the signed certificate you receive from LONI administrator to `$HOME/.globus/usercert.pem`



- Certificate Information

```
[apacheco@qb1 ~]$ grid-cert-info
```

```
Certificate:
```

```
  Data:
```

```
    Version: 3 (0x2)
```

```
    Serial Number: 340 (0x154)
```

```
    Signature Algorithm: md5WithRSAEncryption
```

```
    Issuer: C=US, O=Louisiana Optical Network Initiative, OU=loni.org, CN=LONI CA
```

```
  blah blah blah
```

- Distinguished Name

```
[apacheco@qb1 ~]$ grid-cert-info -subject
```

```
/C=US/O=Louisiana Optical Network Initiative/OU=loni.org/OU=sys.loni.org/CN=Alexander P
```

- Issuer Hash

```
[apacheco@qb1 ~]$ grid-cert-info -issuerhash
```

```
a3bf9f3c
```



- More Details/Help

```
[apacheco@qb1 ~]$ grid-cert-info -help
grid-cert-info [-help] [-file certfile] [-all] [-subject] [...]
```

```
    Displays certificate information. Unless the optional -file
    argument is given, the default location of the file containing the
    blah blah blah
```

- Generating a valid proxy.

```
[apacheco@qb1 ~]$ grid-proxy-init
```

```
Your identity: /C=US/O=Louisiana Optical Network Initiative/OU=loni.org/OU=sys.loni.org/CN=All
Enter GRID pass phrase for this identity:
Creating proxy ..... Done
Your proxy is valid until: Fri Oct 22 03:30:20 2010
```



- Globus provides `globus-url-copy`, a scriptable command line tool that can do multi-protocol data movement.
- Supports: `gsiftp://` (GridFTP), `ftp://`, `http://`, `https://`, and `file://`
- Usage:

```
globus-url-copy [-help | -usage] [-version[s]] [-vb] [-dbg] [-b | -a]
                [-q] [-r] [-rst] [-f <filename>]
                [-s <subject>] [-ds <subject>] [-ss <subject>]
                [-tcp-bs <size>] [-bs <size>] [-p <parallelism>]
                [-notpt] [-nodcau] [-dcsafe | -dcpriv]
                <sourceURL> <destURL>
```

- transfer a file from queenbee to oliver

```
[apacheco@qb1 globus-tutorial]$ globus-url-copy \
  gsiftp://qb1.loni.org/home/apacheco/globus-tutorial/100mbfile \
  gsiftp://oliver1.loni.org/home/apacheco/100mbfile
```

- create a 1GB file and transfer to zeke with verbose (-vb) option

```
[apacheco@qb1 globus-tutorial]$ dd count=1000 bs=1024k if=/dev/zero \
  of=/work/apacheco/1gbfile
```



Moving Files using Globus II

```
1000+0 records in
1000+0 records out
```

```
[apacheco@qb1 globus-tutorial]$ globus-url-copy -vb \
    file:///work/apacheco/lgbfile \
    gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/lgbfile

Source: file:///work/apacheco/
Dest:   gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/
        lgbfile
        1020264448 bytes          29.87 MB/sec avg          35.00 MB/sec inst
```

- Create a directory to store a filename

```
[apacheco@qb1 ~]$ globus-url-copy -vb -cd \
gsiftp://eric1.loni.org/home/apacheco/100mbfile \
gsiftp://qb1.loni.org/home/apacheco/createdirectory/100mbfile

Source: gsiftp://eric1.loni.org/home/apacheco/
Dest:   gsiftp://qb1.loni.org/home/apacheco/createdirectory/
        100mbfile
        104857600 bytes          58.82 MB/sec avg          58.82 MB/sec inst
```

- Transfer files from a directory recursively

```
[apacheco@qb1 globus-tutorial]$ globus-url-copy -vb -r -cd \
gsiftp://qb1.loni.org/home/apacheco/globus-tutorial/ \
gsiftp://oliver1.loni.org/home/apacheco/createdirectory/
```



Moving Files using Globus III

```
Source: gsiftp://qb1.loni.org/home/apacheco/globus-tutorial/  
Dest:   gsiftp://oliver1.loni.org/home/apacheco/createdirectory/  
200mbfile  
  209715200 bytes      34.48 MB/sec avg      34.48 MB/sec inst  
100mbfile  
  104857600 bytes     27.03 MB/sec avg     27.03 MB/sec inst
```

- You can adjust the tcp buffer size and buffer size using the options `-tcp-bs` and `-bs`
- Use parallel streams for transfer with `-p` option

```
[apacheco@qb1 globus-tutorial]$ globus-url-copy -p 4 -vb \  
file:///work/apacheco/1gbfile \  
gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/1gbfile
```
- On the LONI machines, `-p 5 -tcp-bs 2097152` is recommended for a reasonable performance (not necessarily the most optimized set of option numbers).
- multiple transfers from a script file using the `-f` switch (example in job submission section)



- Job Submission using `globus-job-run`.

```
[apacheco@qb1 ~]$ globus-job-run -help
Usage: globus-job-run
        [-help|-usage]           print usage and exit
        [-version]               print version and exit
```

- Run a simple command on a LONI machine

```
[apacheco@qb1 ~]$ globus-job-run louie1.loni.org /bin/date
Thu Oct 21 14:57:19 CDT 2010
```

```
[apacheco@qb1 ~]$ globus-job-run louie1.loni.org /bin/hostname
louie1.loni.org
```

- Run a shell script with arguments using `globus`

```
[apacheco@qb1 ~]$ globus-job-run oliver1.loni.org -s random-generator.sh 4 7
2377875
2910657
8458138
1522082
```

```
[apacheco@qb1 ~]$ globus-job-run bluehawk.loni.org -s random-generator.sh 4 7
5138706
1757255
3086336
5345316
```



Job Submission II

- -np N : number of processing elements
- -stdin [-l|-s] file : standard input
- -stdout [-l|-s] file : standard output
- -l[ocal] : file is relative to working directory of job (DEFAULT)
- -s[tag] : file relative to job request is staged to job host
- -x rsl-clause : RSL extension capability



- Execute commands on remote machine using input from staging machine

```
[apacheco@qb1 ~]$ globus-job-run louie1.loni.org -stdin -s /etc/hosts -l /bin/cat
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1      localhost.localdomain localhost
204.90.40.60  eta.hpc.lsu.edu eta
208.100.92.71 1611u01.sys.loni.org
.....
10.192.92.247  qbib20-1
10.192.92.248  qbib20-2
```

- Execute commands on remote machine using input from local machine

```
[apacheco@qb1 ~]$ globus-job-run louie1.loni.org -stdin -l /etc/hosts -l /bin/cat
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1      localhost.localdomain localhost
204.90.40.60  eta.hpc.lsu.edu eta
208.100.92.71 1611u01.sys.loni.org
208.100.92.72 1611u02.sys.loni.org
.....
172.17.73.227 louie127h
172.17.73.228 louie128h
```



- Execute commands and write output on staging machine

```
[apacheco@qb1 ~]$ globus-job-run louie1.loni.org -stdout -s hellohosts \  
-stdin -s /etc/hosts -l /bin/cat  
[apacheco@qb1 ~]$ head -4 hellohosts ; tail -2 hellohosts  
  
# Do not remove the following line, or various programs  
# that require network functionality will fail.  
127.0.0.1      localhost.localdomain localhost  
204.90.40.60  eta.hpc.lsu.edu eta  
10.192.92.247 qbib20-1  
10.192.92.248 qbib20-2
```

- Job submission using `globus-job-submit`
- submit a job to the scheduling jobmanager

```
[apacheco@qb1 ~]$ globus-job-submit oliver1.loni.org/jobmanager /bin/date
```

```
https://oliver1.loni.org:50886/26507/1288375416/
```

- `globus-job-status` checks status of jobmanager

```
[apacheco@qb1 ~]$ globus-job-status https://oliver1.loni.org:50886/26507/1288375416/
```

```
DONE
```

- `globus-job-get-output` gets output of the job

```
[apacheco@qb1 ~]$ globus-job-get-output https://oliver1.loni.org:50886/26507/1288375416/
```

```
Fri Oct 29 13:03:36 CDT 2010
```



- Submit an openmp parallelized code on

- eric (PBS)

```
[apacheco@qbl ~]$ globus-job-run eric1.loni.org/jobmanager-pbs \  
-np 1 -m 10 -p loni_loniadmin1 -q single -stdin \  
-s /home/apacheco/dft/flux-full.dat -stdout \  
-s flux-eric1-parallel.dat -s /home/apacheco/dft/dft.parallel
```

- zeke (loadleveler)

```
[apacheco@qbl ~]$ globus-job-run zeke.loni.org/jobmanager-loadleveler \  
-env "GBLL_NETWORK_MPI=sn_all,not_shared,US" -stdin \  
-s /home/apacheco/dft/flux-full.dat -x "&(jobType=mpi) (count=2) \  
(maxWallTime=10) (queue=checkpt) (stdout=/mnt/lpfs.nfs302/apacheco/dft.out) \  
(stderr=/mnt/lpfs.nfs302/apacheco/dft.err) (hostCount=1)" \  
-s /home/apacheco/dft/dft.xlf  
[apacheco@qbl ~]$ cat dft.output  
gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/dft.out \  
file:///home/apacheco/dft.out  
gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/dft.err \  
file:///home/apacheco/dft.err  
  
[apacheco@qbl ~]$ globus-url-copy -vb -f dft.output  
Source: gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/  
Dest:   file:///home/apacheco/  
        dft.out  
  
Source: gsiftp://zeke.loni.org/mnt/lpfs.nfs302/apacheco/  
Dest:   file:///home/apacheco/  
        dft.err
```



Submit Parallel Jobs II

- copy `helloworld.c` from `queenbee.loni.org:/home/apacheco/helloworld.c`
- compile on `bluedawg/zeke/(any other P5 machine)` using

```
mpcc helloworld.c -o helloworld
```

- Submit an MPI job
 - on `bluedawg/zeke/(any other P5 machine)`

```
[apacheco@qb1]$ globus-job-run bluedawg.loni.org/jobmanager-loadleveler \  
-env "GBLL_NETWORK_MPI=sn_all,not_shared,US" -x "&(jobType=mpi) (count=8) \  
(maxWallTime=10) (queue=checkpt) (hostCount=1)" -l /home/apacheco/helloworld  
Hello world from 0 out of 8  
Hello world from 1 out of 8  
Hello world from 2 out of 8  
Hello world from 3 out of 8  
Hello world from 4 out of 8  
Hello world from 5 out of 8  
Hello world from 6 out of 8  
Hello world from 7 out of 8  
ATTENTION: 0031-408 8 tasks allocated by LoadLeveler, continuing...
```

- on `eric/oliver/(any other intel machine)`

```
[apacheco@qb1 ~]$ globus-job-run eric1.loni.org/jobmanager-pbs \  
-x "(jobType=mpi) (hostCount=1) (maxWallTime=10) (queue=checkpt) \  
(stdout=/home/apacheco/gljobrun.out)" -s /home/apacheco/helloworld
```



Submit Parallel Jobs III

```
[apacheco@eric2 ~]$ cat gljobrun.out
-----
Running PBS prologue script
-----
User and Job Data:
-----
Job ID:      262113.eric2
Username:    apacheco
Group:       loniadmin
Date:        27-Oct-2010 15:44
Node:        eric052 (7032)
-----
PBS has allocated the following nodes:

eric052

A total of 4 processors on 1 nodes allocated
-----
Check nodes and clean them of stray processes
-----
Checking node eric052 15:44:34
Done clearing all the allocated nodes
-----
Concluding PBS prologue script - 27-Oct-2010 15:44:34
-----
Hello world from 0 out of 4
Hello world from 2 out of 4
Hello world from 1 out of 4
Hello world from 3 out of 4
```



- Job submission using `globusrun` with Resource Specification Language (RSL) commands in a file

```
[apacheco@qb1 ~]$ cat eric.rsl
```

```
& (jobType=mpi)
(hostCount=1)
(maxWallTime=10)
(queue=checkpt)
(executable=/home/apacheco/helloworld)
(stdout=/home/apacheco/gljobrun.out)
```

```
[apacheco@qb1 ~]$ globusrun -r eric1.loni.org/jobmanager-pbs -f eric.rsl -b
globus_gram_client_callback_allow successful
GRAM Job submission successful
https://eric1.loni.org:50886/779/1288371858/
GLOBUS_GRAM_PROTOCOL_JOB_STATE_PENDING
```



- **Globus Toolkit**
<http://www.globus.org/toolkit/>
- **Loni Docs: Globus Tutorial using LONI resources**
https://docs.loni.org/wiki/Globus_Tutorial_using_LONI_resources
- **Moodle: HPC105 - Getting Started with Globus Toolkit**
<https://docs.loni.org/moodle/>

