BNL's experience with Globus-Connect

Hironori Ito

Brookhaven National Laboratory

BROOKHAVEN NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN SCIENCE ASSOCIATES

Globus Connect/Online at BNL

- BNL has some experience of using Globus Connect Service to serve as a storage site for HPC sites in US. (See Doug's Talk for any detail.)
- BNL like many other T1s use dCache as the main storage service for ATLAS. Therefore, the attempt was made to use BNL T1 dCache as the final aggregation endpoint for the outputs for jobs run at HPCs.
- Various efforts were made to make it work under dCache. Although it did work somewhat, it never achieved the level of performance and/or stability to serve as the storage for T1.
- BNL also uses Globus Connect for different user communities (eg. NSLSII light source and Nano Center) . And, BNL has paid the license fees. For this purpose, it does not use dCache as a storage.

Various attempts

- Tried regular dCache GridFTP with Globus online
  - FTP had the issue with creating directory listing.
    - The patch in dCache GridFTP was provided from the developers.
  - Go MD5 checksum has the issue.
    - Now dCache supports multiple checksum
  - Go has the issue with dCache's write once storage.
    - Retry fails because it tries to append the already closed, failed file.
- Tried Gloubs GridFTP with dCache storage mounted as NFS v4.1.
  - It worked a bit better. But, it was still not stable enough.
    - Not tested with newest dCache.
- Tried Globus GridFTP with regular storage mounted as NFS v3
  - Worked fine until recently when the load became issue with more transfers from HPC sites.
  - The performance of the NFS storage was too low.
- Tried Globus GridFTP with LustreFS/storage.
  - This works.
  - It would have worked with NFS or GPFS or any other POSIX storage.

Debugging info comparison between GO and FTS

2021-02-23 11:44 pm timeout error

View details

Globus Connect Transfer Info

Error (transfer)
Endpoint: dcdoor15 (d8981638-6af7-11e6-83d0-22000b97daec)
Server: dcdoor15.usatlas.bnl.gov:2811
Command: STOR
/pnfs/usatlas.bnl.gov/users/hiroito/globus/aglt2/h03/08/data18_13TeV.00363129.physics_Main.merge.AOD.f993_m2032._lb0342._0004.1
Message: The operation timed out

Details: Timeout waiting for response

Higher level logs seem to be available with higher level of the subscription?

srm://uct2-dc1.uchicago.edu:8443/srm/managerv2?SFN=/pnfs/uchicago.edu/atlasdatadisk/rucio/mc16_13TeV/a6/35/AOD.20093348._006960.pool.root.1

FTS Transfer Info

- Transfer host: fts307.usatlas.bnl.gov
- Staging host:
- PID: 301882
- Hash: 0CC0
- Activity: Production Input
- Selection strategy: throughput
- Attempts: 0
- Duration: 0 seconds
- Checksum: ADLER32:083623d5
- User specified size: 6807046627
- Configuration:
- Parameters: nostreams:0,timeout:13584,buffsersize:0
- Finished time: 2021-02-24T12:44:51Z
- Error reason: TRANSFER [70] TRANSFER globus_ftp_client: the server responded with an error 555 SRM upload aborted: Changing file state because request state has changed.
- Log file:

https://fts307.usatlas.bnl.gov:8449/var/log/fts3/transfers/2021-02-24/grid-srm.rzg.mpg.de__uct2-dc1.uchicago.edu/2021-02-24-1244__grid-srm.rzg.mpg.de__uct2-dc1.uchicago.edu__782154520__f636b3f0-762a-11eb-979f-b49691292f84

- Metadata:

{"dst_type": "DISK", "name": "AOD.20093348._006960.pool.root.1", "adler32": "083623d5", "source_globus_endpoint_id": null, "src_rse": "MPPMU_DATADISK", "activity": "Production Input", "src_rse_id": "71fc14be0b0b457497fec857a1b7ae88", "request_id": "fdedf20a99b84f8bbe3e4f7a1454cbf5", "request_type": "TRANSFER", "filesize": 6807046627, "dest_rse_id": "f4da89c7cef54133a45c22712dda45d5", "scope": "mc16_13TeV", "dst_rse": "MWT2_DATADISK", "src_type": "DISK", "dest_globus_endpoint_id": null, "verify_checksum": "both", "md5": null}

```
INFO    Wed, 24 Feb 2021 07:33:54 -0500; bytes: 3887866618, avg KB/sec:1205, inst KB/sec:898, elapsed:3151
INFO    Wed, 24 Feb 2021 07:35:04 -0500; bytes: 3953052921, avg KB/sec:1198, inst KB/sec:909, elapsed:3221
INFO    Wed, 24 Feb 2021 07:36:14 -0500; bytes: 4016303503, avg KB/sec:1192, inst KB/sec:882, elapsed:3291
INFO    Wed, 24 Feb 2021 07:37:24 -0500; bytes: 4082491686, avg KB/sec:1186, inst KB/sec:923, elapsed:3361
INFO    Wed, 24 Feb 2021 07:38:34 -0500; bytes: 4125278871, avg KB/sec:1174, inst KB/sec:596, elapsed:3431
INFO    Wed, 24 Feb 2021 07:39:44 -0500; bytes: 4224018222, avg KB/sec:1178, inst KB/sec:1377, elapsed:3501
INFO    Wed, 24 Feb 2021 07:40:54 -0500; bytes: 4381185479, avg KB/sec:1198, inst KB/sec:2192, elapsed:3571
INFO    Wed, 24 Feb 2021 07:42:04 -0500; bytes: 4495721901, avg KB/sec:1206, inst KB/sec:1597, elapsed:3641
INFO    Wed, 24 Feb 2021 07:43:14 -0500; bytes: 4592063670, avg KB/sec:1208, inst KB/sec:1345, elapsed:3711
INFO    Wed, 24 Feb 2021 07:44:24 -0500; bytes: 4662702427, avg KB/sec:1204, inst KB/sec:984, elapsed:3781
WARNING Wed, 24 Feb 2021 07:44:50 -0500; Timeout stopped
ERR     Wed, 24 Feb 2021 07:44:50 -0500; Recoverable error: [70] TRANSFER  globus_ftp_client: the server res
because request state has changed.
```

FTS Transfer Log

Conclusion

- Globus Connect service expects the storage to be posix-ish.
  - Tried to append a file for the failed transfer.
  - dCache needs to handle better (e.g. delete and restart the transfer from the beginning…)
- Globus Connect lacks the adlere32 checksum.
  - dCache can support both md5 and adler32. But, it won't be able to do the on-the-fly checksum for md5.
- Globus Connect lacks the detail transfer logs.
- Globus Connect works with regular posix storage.
  - Works with AWS S3 and any other S3(?) (I have not tested.)
  - It would be much nicer and simpler if it supports **dCache**.
- It would be even nicer if **FTS supports Globus an endpoint**.

Levels of GO Subscription

Application Integration Support: This is intended for organizations that integrate Globus data management functionality into their own applications. It entitles you to support when building web applications, such as data portals or science gateways, that use Globus file transfer and sharing services via our REST API or the Globus user-interface widgets.

Additional authentication assurance: Access to protected data endpoints may require a user to authenticate within a session using an identity from a specific identity provider. Institutions define the identity required and the length of time before the user is asked to re-authenticate. Globus enforces institutional policy to ensure that correct and valid credentials are presented.

Comprehensive audit logging: The Globus Connect software generates a detailed audit trail on your local storage, allowing you to reconstruct data access and user activities. Audit logs record details of all data access events as well as activities such as login and resource management.

Subscriptions for Non-Profit Research and Education

If you wish to use Globus in a commercial setting, you must have a commercial subscription.

Features (click ℹ for description)	**Basic**	**Starter**	**Standard**	**High Assurance** ℹ	**HIPAA BAA** ℹ
File transfer ℹ	Unlimited	Unlimited	Unlimited	Unlimited	Unlimited
Managed endpoints ℹ	None	1	Unlimited	Unlimited	Unlimited
Management console ℹ	–	✓	✓	✓	✓
Usage reports ℹ	–	✓	✓	✓	✓
Support for Globus Connect, Web, CLI ℹ	–	✓	✓	✓	✓
File sharing ℹ	–	✓	✓	✓	✓
Globus Plus users ℹ	–	–	✓	✓	✓
Application integration support ℹ	–	–	✓	✓	✓
HTTPS support ℹ	–	–	✓	✓	✓
Session/Device Isolation ℹ	–	–	–	✓	✓
Additional authentication assurance ℹ	–	–	–	✓	✓
Comprehensive audit logging ℹ	–	–	–	✓	✓
Business Associate Agreement ℹ	–	–	–	–	✓
Support service level ℹ	–	Monday-Friday, 9am-5pm Central; 1-business day response			
Named support contacts ℹ	–	1	5		
Pricing	Free	Contact us for subscription pricing details			