FTS Steering — ATLAS DDM Minutes

2019-02-19

- Protection against double submission.
    - We need a efficient way to prevent double submission on the FTS side, since we cannot ensure it 100% on the Rucio side. This is our top priority.
    - *—> Will come in 3.9 (~April) Need to agree on error. (+ transfer-duplicate for Rucio monitoring)*
- Analytics studies on FTS performance are strongly needed.
    - What we do well, what can be done better.
    - Are the tails the problems?
    - Are the median transfer time?
    - Are we exploiting fully the available bandwidth/storage capabilities?
    - We need deep study. E.g., understanding of asynch vs synch move (caches), what would happen if we have e.g. 50% of today FTS moved data managed by caches. If the caches are working so perfectly, can FTS include the movements from/to them?
    - Source selection : Which one is the best ?
    - N.b. a lot of work has been done already by Joaquin
    - How do we all together approach this topic??
    - *—> Analytics person will join FTS team (starting April)*
- Scaling of FTS: we see that in case of more than e.g. 1-10M files, FTS collapse. Plans to improve the situation?
    - *—> FTS team has a plan to assess the performance and will provide updates on the developments/improvements. Overall goal is to converge to a single instance per VO.*
- FTS3 servers Deployment and Operations: it's a pain right now.
    - The servers don't talk each others.
    - They have different configs (we discussed to push configs everywhere, not done).
    - We need each time campaign to upgrade FTSes, and we have often issues of configs. Why not moving to some containerized image , pushed centrally and just deployed at e.g. three sites?
    - Can we think about one single instance serving all sites on the medium term ?
    - *→ same as above, FTS team has discussed this already*
- Automatic session reuse, today is useless how it is done. We need more intelligence (a-la garbage collection, wait e.g. 10mins and group files also in other job ids -- job ids in theory are only logical grouping, I thought FTS work at file level)….
- Network information integration, plans for the future
    - NOTED
- Data carousel
    - Where/how to do the throttling (Rucio or FTS side)
        - 200 in bulk from FTS side (configurable)
    - Bulked FIFO mode
    - "Smart" request reshuffling

    - Writing in bulk (--> group by dataset)
- Rucio+OIDC —> 5 PM from XDC
- QoS
    - Transitions via CDMI

2018-03-14

ATLAS top priority requests

- Selection of the source.
    - Which source selection to use: we ask FTS which one to start using. The one right now is better than zero.
    - All instance should be able to exchange info between themselves.
    - Shares should be taken into account in the source selection!
    - Ask FTS: can you do the study?
- Delegation client upgrade (longer than 4 hours)
    - done
- Why the single FTS instance at CERN can't manage the whole Grid, and fill the network pipes?
    - Are we filling network pipes?
    - How difficult is to have the FTS talking each others? Maybe more difficult than just throwing more machines (or fatter ones - or whatever is needed) into an instance and make sure that one single, with hot backup, can manage anything.
    - Can we (ATLAS & FTS, FTS only, experiments -- someone?) test this?
- Modification on gfal2 for signed URL to be able to use GSC
- Support different authentications on same hostname
- Auto Load balancing for specific ens.
    - E.g.: HPC like Nersc have different DTN (gridftp). Has been noticed (thks Brian, code committed in Rucio) that if we parallelize the DTNs we can increase the throughput a lot. Should ATLAS (experiments) keeping this info, or FTS? What if FTS "poll" the IPs behind an alias (RB or LB), and automatically parallelize the transfers on the various gridftp servers? This would simplify a lot the topology need on the experiment side.
    - —> need to check;
- Bearer token support

2017-06-07

FTS priorities

- What are the FTS priority for the next years? What is the intelligence that will be in FTS and what should be on the experiments DDM? Just ideas:
    - Are you planning to work on the e.g. possibility experiment pass multi src and FTS act a-la torrent?
        - aria2 (http) + xroot, ongoing tests
        - always #1 streams on gridftp
    - Protocols for sites: do you plan to make the storage representation inside FTS site aware, or at least storage aware? E.g. EOS can transfer through xrootd, gridftp, etc. Having the storage protocols not connected is not optimal.
        -
    - Automatic session reuse . ATLAS is pushing to have more sites SRMless but with GridFTP. If we have a multifile job with gridftp, can FTS automatically decide to use the gridftp session reuse?
        - ongoing work, reduction of FTS-internal schedulers first
        - start with heuristic ? chunk/recompose jobs (job currently atomic)
        - merge on Rucio-side? PanDA workflow/flag? by activity?
    - Network: is not infinite. Is there a plan from FTS to integrate the infrastructure info, Does FTS plan to integrate SDN?
        - can interact with SDN controller
        - discuss in Manchester
    - FTS - knowledgeable of file structure (zip, root)
        - Problematic with 3rd party copy
        - archive awareness (offset/metadata?)
    - activity shares on links based on throughput rather than no of files
        - FTS will think about problem
    - ObjectStores usage
        - Can FTS manage cache storage?
        - explicitly ask for third party copy
    - Increase usage of TAPE
        - to be discussed in a pre-GDB?
- Other open question :
    - Fair-share: We need to apply it per destination and not per link, e.g., express
    - 3rd party HTTP: potential combinations (protocol, storage, failures)
        - http://fts3-docs.web.cern.ch/fts3-docs/docs/3rdpartycopy.html
    - Minimum default for #active on link? Right now it's 2 (too little?)
        - FTS will try to be a bit more aggressive

Optimizer

- We suspect that decision is not taken wrt efficiency improvement but wrt nominal value of efficiency.
- We would suggest to make it more aggressive:
    - Weighted decisions according to efficiency improvement.
    - Instant probes with different parallels than the actual value.

Example of sudden change of min_parallel at BNL:

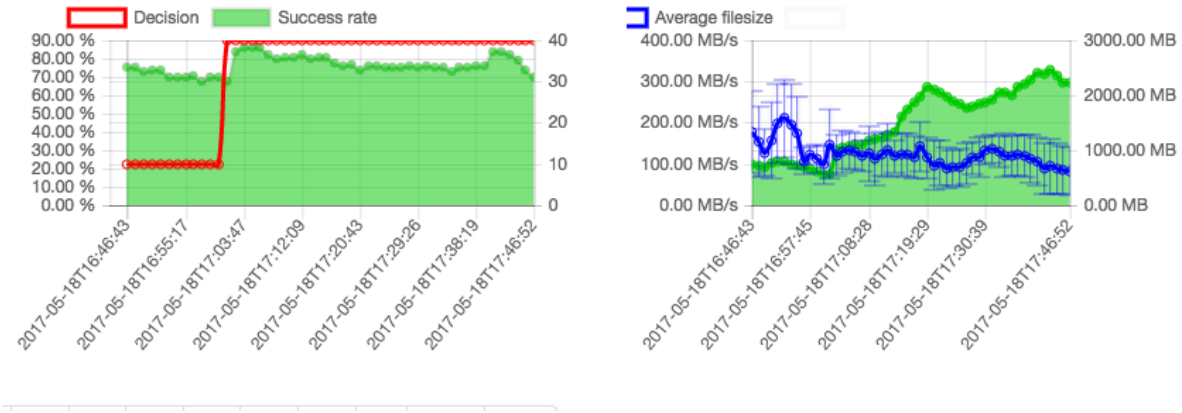

⇒ efficiency increased by 10%, transfer rate probably increased as well, but this is not evident because also file size changed.

Could optimizer sit in some local minimum?

2016-08-30

Date this part of the doc started: 26 July. Updated 30th August.

[HIGH]

- Many things will “come in next release”: can you please clarify which versions is ready to go in production, and can we know from sites running FTS what they run?
    - E.g. stale connection status: is this fixed, can we ask sites to move to that release?
    - FTS pilot: ATLAS could increase the load if needed.
    - max/min for optimiser
- Xrdcp third party transfers: is it true that other experiments are using already FTS with xrdcp? Is it mature enough for ATLAS to use it, or at least to test it for EOS to Castor RAW files transfers?

[MEDIUM]

- Fair-share: We need to apply it per destination and not per linlefik, e.g., express
- Config between various FTS servers and defaults clearly visible.
- Config between Optimizer and fixed streams
- Automatic session reuse . ATLAS is pushing to have more sites SRMless but with GridFTP. If we have a multifile job with gridftp, can FTS automatically decide to use the gridftp session reuse?

[LOW]

- Network map, at least LHCOPN (Tier-0 Tier-1s)
- Monitoring: different numbers in different places. Need to be understood.

…

2016-05-12

List of ranked priorities/points (from ATLAS/Rucio point of view):

- Fair share broken: why didn't we spot the issue for few months ? Would it be possible to show the selected share name and value in the job itself and print it also on the monitoring.
- Client side generated job id:
    - No yet in use from our side..it'll come with 1.5.11 coming
    - We still observe some stalled connections, reported by a probe
- Fair-share: We need to apply it per destination and not per link, e.g., express

Review document :
https://docs.google.com/presentation/d/1Hr2pMzXmk6Ve2gdGvLaQK1rrcgPFGTjioAP2Mrx8k38/edit?pref=2&pli=1#slide=id.p

From Rod:
If
- destination(BU) has a limit of 100 transfers
- 10 source sites have queued transfers
- only 1 source has express activity transfers queued (BNL)
- no channel is limited by the optimizer
- activity shares 50:50 express:other

Currently we get 10 transfers per channel and 5 of the BNL-BU ones are express. As a destination BU is getting 95 other and 5 express transfers. If we decide we want the shares at the destination level, then we would want 50 express. Then BNL-BU would be allowed to rise above the 10 up to the optimizer limit of 50, say, leaving 50 for all other channels.

BNL-BU 25 express:25 other
BU dest 25 express and 75 other, so close to the shares.

It gets complicated if one worries about how the BNL-BU optimizer limit of 50 affects the source. It reduced the number of transfers to be shared out across channels with BNL as source. Some of those can have express transfers too.
Are we going to create dead channels this way, that never get a transfer starting?

- Automatic cancellation of requests which stay for too long in the queue
    - It was a bug for test fts @ ral but do we need this feature ?
- Rest interface to get the load of all channels or some channels

    - Rucio internal queue is based on rucio db. It's not the realtime status of FTS. If the poller is delayed, some channels may have no queue requests and some others may queue a lot of requests.
    - It's good to include the limitation of active transfers. For example, one channel can be limited to only few active transfers because of too many failures. We should not queue many requests to that channel.
- BNL monitor still broken, takes many reloads to work
- Duplicate transfer issue - status? (we still don't consume messages)

2015-12-02

List of ranked priorities (from ATLAS/Rucio point of view):

[HIGH] Transactional submission behavior when submitting transfers

- If not possible, we need a way to check that a job has been submitted for a destination url(s) as a recovery procedure
- Because of this issue, we have disabled the usage of time-outs when submitting transfers but we got the 'stalled' connection issue (Cf. https://its.cern.ch/jira/browse/FTS-343)
- → There will be a possibility to set a uuid when creating a transfer (subset of uuid space assigned to the project) This UUID can be queried to look for existing transfer;

[HIGH] Issue with sequences of messages.

- Wrong FAILED ↔ SUCCESS transition and race conditions
- We don't consume messages anymore and only poll which is less performant and put more load.
- Is it fixed and validated ?
    - we need to have Lionel reenabling the MSG queue

>>>----??????---->> can we restart the consumers???

[**MEDIUM**] expire the FTS jobs, because we have some t0_export jobs staying in the fts for more than one week.

- When submitting jobs, client set the lifetime
- If in the lifetime, fts fails to finish this job, fts should expire this job.

>>>----??????---->> Now Rucio has a queue limit per src - dst, we really need to be sure jobs are respecting their timeouts. An example is being searched.

[MEDIUM] Monitoring links grouped by src-dst-activity

- from webUI we want to be able to see the queue for link per activity
- With statistics: done, throughput per min/hour ?
    - would be nice to have a filter by activity in this table https://fts3.cern.ch:8449/fts3/ftsmon/#/?

>>>----??????---->> is it available?? when?

[MEDIUM] FTS set priority to boost transfers

- clients need to be fixed
- -> but do we want to use it ? I though the express share should be enough. I'm not sure it's properly implemented now on the server.

>>>----??????---->> when??

[LOW] Multiple sources in one bulk job

- Today we submit one multi source transfer per job

>>>----??????---->> when?

[LOW] get error information from FTS with messages and by polling:

- Recoverable error or not
- Possibility to provide regexps to classify the error.

>>>----??????---->> when?

2015-10-28

List of ranked priorities (from ATLAS/Rucio point of view):

[HIGH] Transactional submission behavior when submitting transfers

- If not possible, we need a way to check that a job has been submitted for a destination url(s) as a recovery procedure
- Because of this issue, we have disabled the usage of time-outs when submitting transfers but we got the 'stalled' connection issue (Cf. https://its.cern.ch/jira/browse/FTS-343)

[HIGH] Issue with sequences of messages.

- Wrong FAILED ↔ SUCCESS transition and race conditions
- We don't consume messages anymore and only poll which is less performant and put more load.
- Is it fixed and validated ?
    - we need to have Lionel reenabling the MSG queue

TBC

[**MEDIUM**] expire the FTS jobs, because we have some t0_export jobs staying in the fts for more than one week.

- When submitting jobs, client set the lifetime
- If in the lifetime, fts fails to finish this job, fts should expire this job.

[MEDIUM] Performance issue when getting all requests done for the last 2 hours

- 1 hour to get more than 10k requests for the last 2 hours. All FTS servers are concerned.

Currently the poll time for last 2 hours is about 20 minutes. The biggest number I saw is 27 minutes. But I didn't saw a lot of finished requests. I will check what will happen if there are a lot of finished requests in last 2 hours.

- Still valid ?
  → we need to check if they fixed, because also if the messages will be ok we want to be able to poll in case of troubles.

[MEDIUM] Monitoring links grouped by src-dst-activity

- from webUI we want to be able to see the queue for link per activity
- With statistics: done, throughput per min/hour ?
    - would be nice to have a filter by activity in this table https://fts3.cern.ch:8449/fts3/ftsmon/#/?

[MEDIUM] FTS set priority to boost transfers

- clients need to be fixed
- -> but do we want to use it ? I though the express share should be enough. I'm not sure it's properly implemented now on the server.

[LOW] Multiple sources in one bulk job

- Today we submit one multi source transfer per job

[LOW] get error information from FTS with messages and by polling:

- Recoverable error or not
- Possibility to provide regexps to classify the error.