TPC: Status, Plans, Contribution to the HL-LHC review document

MORGRIDGE
INSTITUTE FOR RESEARCH
CORE COMPUTATION

FEARLESS SCIENCE

TPC Status

- Honestly, not much big-picture has changed in the three weeks since the last DOMA general meeting presentation.
  - The EOS team was able to get HTTP-TPC working in PPS.
  - ATLAS began to move data in production with non-GridFTP.
    - Unsurprisingly, a number of bugs were discovered that were missed in integration testing.
    - **Example**: DPM always tries to delegate credentials when in "pull mode"; the testbed FTS instance would always allow fallback to "push mode" and the transfer would succeed – but this mode is disabled in ATLAS production FTS.
  - Expect a burst of tickets around this activity. Alessandra is keeping an excellent overview in an ADC ticket.

DOMA-TPC smoke test, started 2020-03-03T12:00+0100, took 67:20.

SOUND ENDPOINTS

SCORE	ENDPOINT	SOFTWARE	WORK-AROUNDS
20	AGLT2	dCache	[in 01:03]
20	BEIJING	DPM	[in 02:34]
20	BNL	dCache	[in 02:12]
20	BRUSSELS	dCache	[in 01:13]
20	CALTECH2	xrootd-R/HDFS	[in 02:17]
20	CALTECH	xrootd-D/HDFS	[in 01:44]
20	DESY-DOMA	dCache	[in 00:18]
20	DESY-PROM	dCache	[in 00:13]
20	FNAL	dCache	[in 00:48]
20	IN2P3	dCache	[in 00:21]
20	IN2P3-TEST	dCache	[in 00:26]
20	INFN-T1	StoRM	[in 00:52]
20	KIT	dCache [2]	[in 00:34]
20	LRZ-LMU	dCache	[in 00:42]
20	NDGF	dCache [2]	[in 00:39]
20	NDGF-PREPROD	dCache [2]	[in 01:15]
20	NEBRASKA2	xrootd-R/HDFS	[in 01:10]
20	NEBRASKA	xrootd-D/HDFS	[in 00:55]
20	PIC-PROD	dCache	[in 01:42]
20	PRAGUELCG2	DPM	[in 00:27]
20	PURDUE	xrootd-D/HDFS	[in 00:49]
20	SARA	dCache [2]	[in 02:50]
20	SARA-test	dCache	[in 00:18]
20	TOKYO-LCG2	DPM	[in 01:22]
20	TRIUMF-DYNAFED	DynaFed/S3	[in 01:45]
20	TRIUMF-PPS	dCache	[in 00:49]
20	TRIUMF-PROD	dCache	[in 10:02]
20	UKI-BRUNEL	DPM	[in 00:39]
20	UKI-IC	dCache	[in 00:26]
20	UKI-LANCS	DPM	[in 00:53]
20	UKI-MAN	DPM	[in 00:42]
20	UKI-MAN-PROD	DPM	[in 00:42]
20	UKI-QMUL-DEV	StoRM	[in 00:37]
20	UKI-QMUL-PROD	StoRM	[in 00:36]
20	UNI-BONN	xrootd-R/CephFS	[in 00:31]
14	PURDUE2	xrootd-R/HDFS	[in 01:30]
13	SLAC	xrootd-D/xrootd	[in 01:02]
10	SLAC2	xrootd-R/xrootd	[in 01:21]
7	UCSD	xrootd-R/HDFS	[in 01:39]
1	CERN	EOS	[in 00:31]

Contributions to Review Document

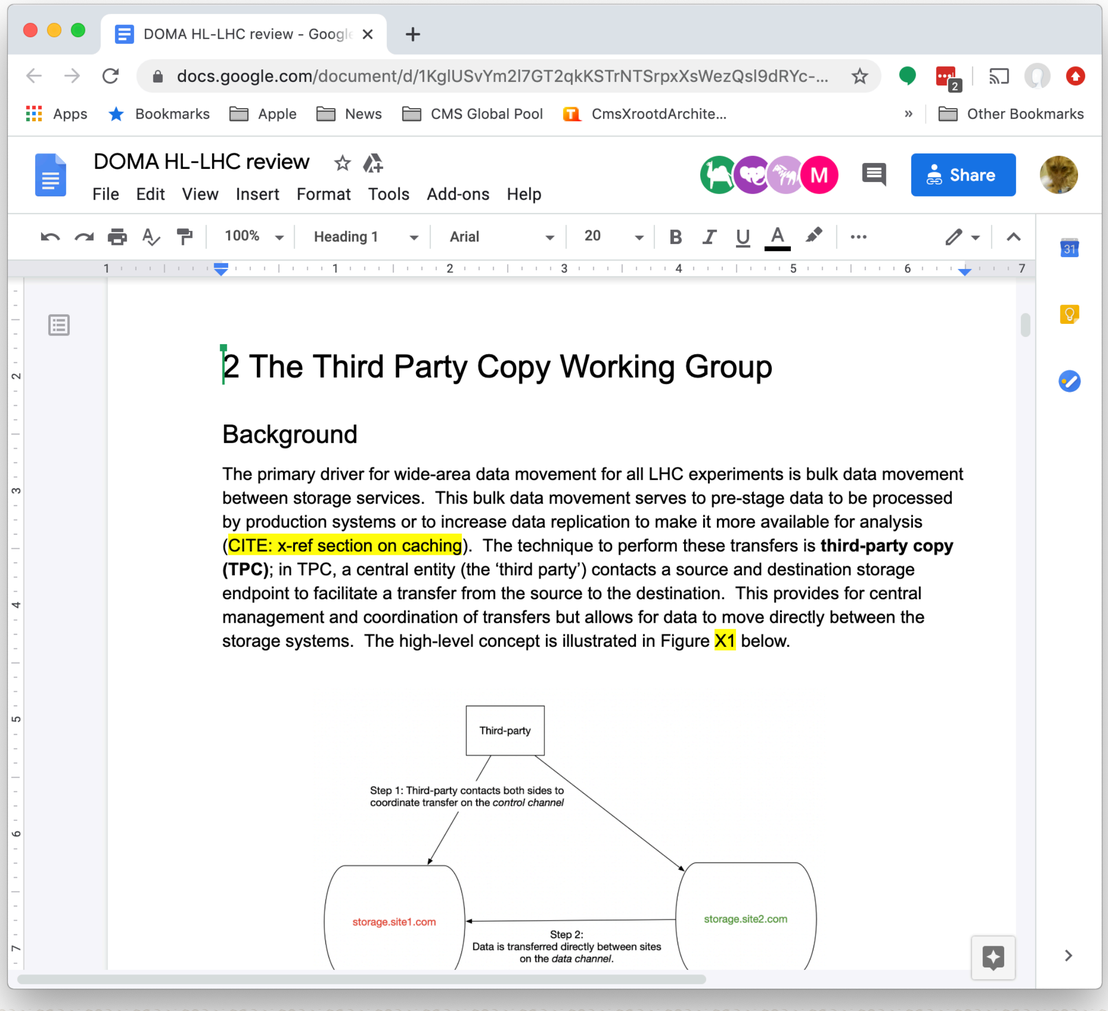

- A first draft of the TPC working group section is available! 🎉 This covers:
  - The background of why TPC is important and why it needs to evolve.
  - Current status of both HTTP-TPC and XRootD-TPC.
  - Recommendations for Run 3.
  - Notable R&D for HL-LHC.

Preliminary Items

(Take these with a grain of salt / illustrative examples – we've not been able to discuss them as a full group)

Run 3 Recommendations:

1. All storage services should have a **robust plan to retire the GridFTP** protocol and support for GSI, including an end-of-life date for GridFTP. These plans should be well-communicated to the WLCG sites and community.
2. The WLCG should promote non-GSI mechanisms for transfer authorization. We recommend **deploying production IAM instances** for the WLCG VOs and the appropriate versions of storage services, FTS3, and Rucio to do non-GSI transfers in production.
3. Experiments should **migrate at least 10% of their production transfer traffic** from GridFTP.

Preliminary Items

(Take these with a grain of salt / illustrative examples – we've not been able to discuss them as a full group)

R&D Needed Prior to HL-LHC Era:

Perform a data challenge in LS3 demonstrating bulk data movement at the scale needed for HL-LHC startup.

1. The WLCG community should **completely transition from the GridFTP protocol**.
2. For data transfers, the WLCG community should completely **transition from GSI** to the new WLCG Common JWT profile.
3. Improve HTTP-TPC implementations to ensure they utilize multiple, pipelined HTTP requests per TCP stream and manage the queue of requests at the endpoint.
4. **Investigate the use of non-TCP-based transfer protocols** between the active and passive endpoints. Some of these protocols promise significant improvements to transfer rates for data channels (although it's not obvious whether this is a worthwhile improvement over many concurrent transfers); we note this R&D can be done while keeping either HTTP-TPC or Xrootd-TPC as the TPC protocol.

Big Picture Items worth Repeating

It's been ~3 years since Globus announced the retirement of Globus Toolkit

It's now time to start moving folks!

GridFTP

- Globus toolkit not supported since 2017
  - Grid Community Toolkit (EGI)
    - no plans from WLCG/OSG for major development / new features
    - support for limited time (e.g. OSG plans)
      - GridFTP/GSI can start to disappear in 2021
- WLCG DOMA – data handling evolution
  - DOMA TPC – new TPC protocols
    - GridFTP old / no active development
    - better support for industry standards
    - requirements & desirables
      - security, multi-VO, multi-impl, documentation, non-X.509, …
    - XRootD, WebDAV (COPY)

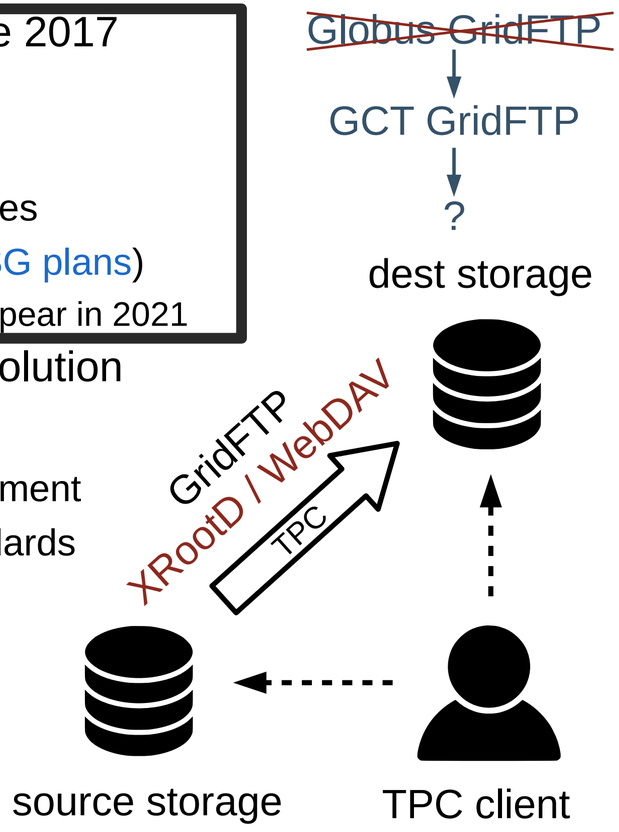

Big Picture Items worth Repeating

WLCG DOMA TPC Tests

- WLCG DOMA TPC testbed - interoparability, functional and stress
  - 250 transfers 4GB files scheduled every hour between each site
  - 1.25PB transferred every week (more than 300k transfers)
  - ~ 5% of average transfer volume within single LHC experiment
  - reaching up to 50Gb/s hourly transfer rate
  - performance for XRootD WebDAV (better then GridFTP and iperf)

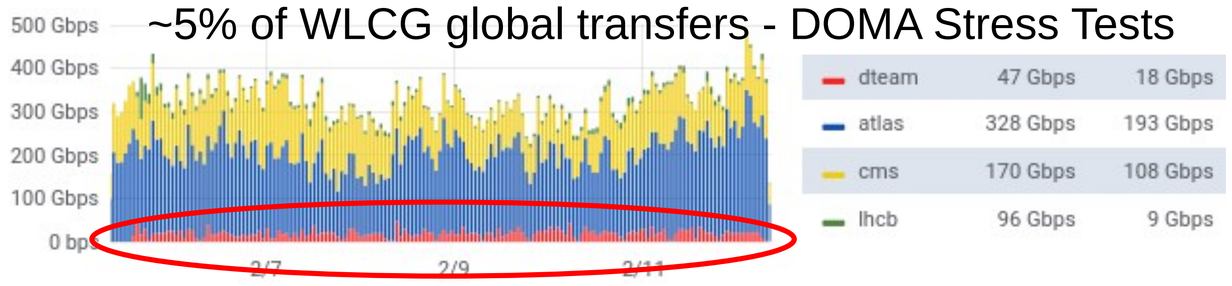

We have shown these new protocols at scale – only way up from here is to start moving data in production!

MORGRIDGE

INSTITUTE FOR RESEARCH

CORE COMPUTATION

morgridge.org

FEARLESS SCIENCE