

DOMA TPC

A. Forti, B. Bockelman

DOMA general

10 July 2019



TPC Phase 2

- **Phase 2** (deadline 30 June 2019): All sites providing more than 3PB of storage to WLCG experiments are required to have one non-GridFTP endpoint in production.
 - All sites have non-GridFTP endpoints used for other activities than TPC
 - xrootd
 - http(s)
 - Reason why these were chosen as alternative to gridftp
 - **Not** used for TPC
- **Phase 3** complete deployment by the end of 2019

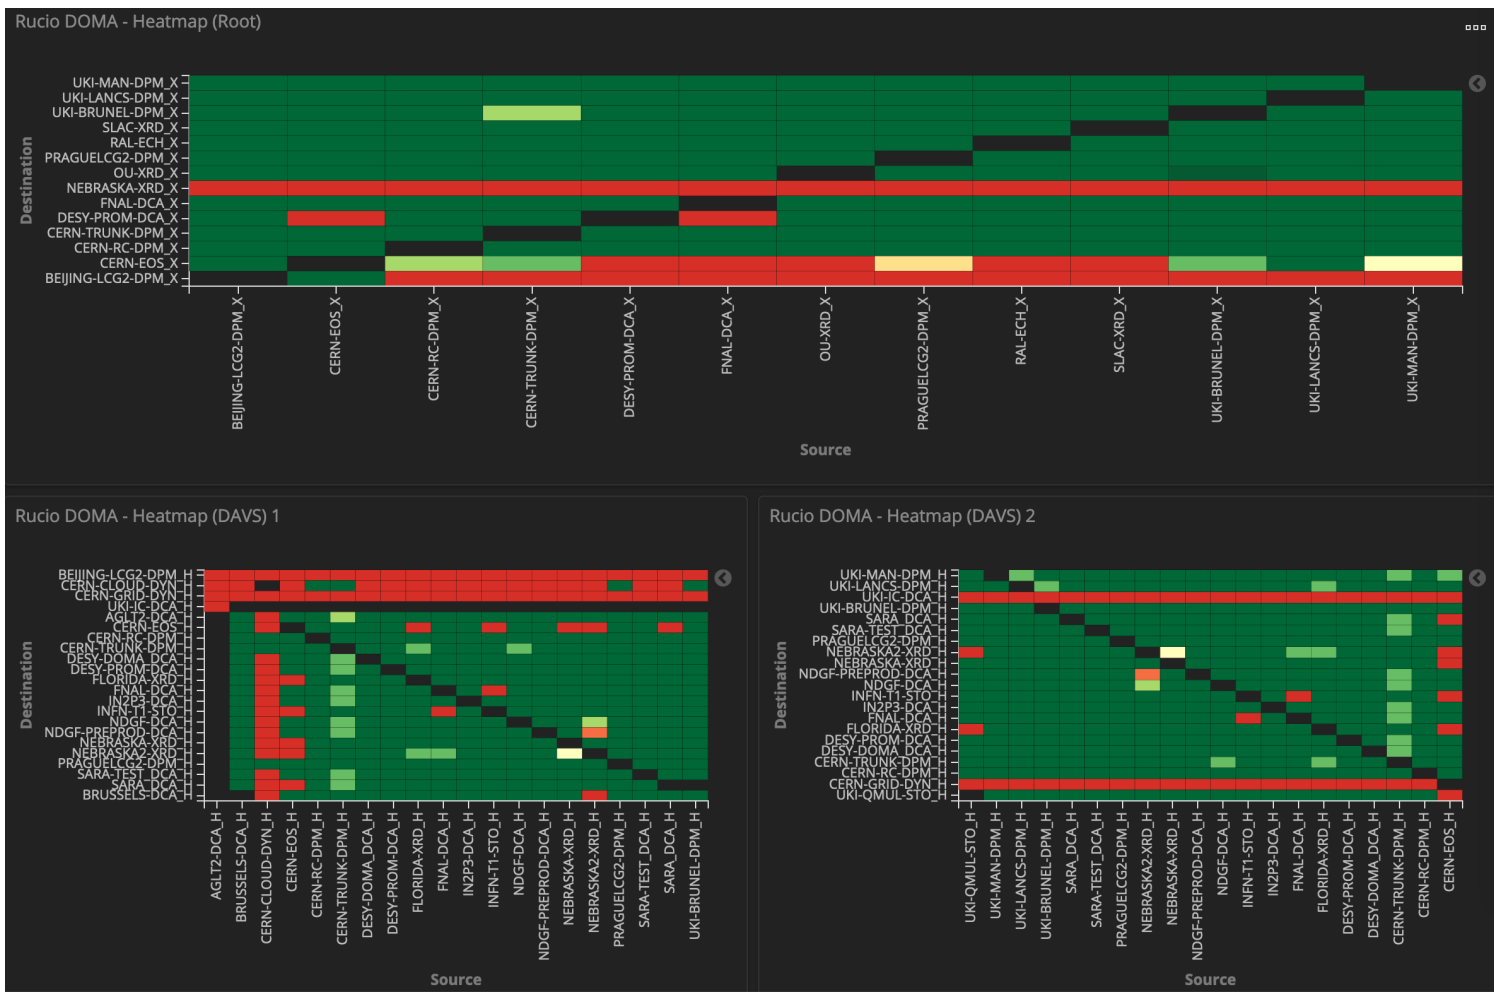


Smoke Tests

- Http smoke tests
 - Run once a day very detailed tests to pinpoint the exact problem of why the transfer might be failing
 - Push/pull
 - Capture detail diagnostic
 - Sys admins can run them manually to see the effect of changes
 - Xrootd version being worked on
 - There is now a python version that needs to be tested.




Rucio functional tests



- Error pattern indicates problems are mostly the sites rather than protocol.
 - Either standard problems or configuration ones

DOMA stress tests



- Stress tests
 - Only production sites with a baseline version of the storage
 - For http since January
 - >500TB / week across participating sites
 - Need more solid comparison with gridftp
 - Adapting now for xrootd too
 - Setting up Manchester and Prague to start with.
 - More sites when these two work



Storage baselines

- Xrootd: 4.9.1
- DPM: 1.12.0
- StoRM: 1.11.15
- Dcache: 5.2.0 is the golden release
 - **v4.1, v4.0 and v3.2 have run out of support**
- EOS: 4.4.37



Can we go SRM-less?

- SRM is used for two things
 - WAN transfers
 - Getting the space information from the storage
- ATLAS can go SRM-less using gridftp
 - Work in this direction has already started particularly with DPM DOME'd sites
- CMS doesn't require SRM it's using load balanced gridftp
- Alice is already SRM-less it's using xrootd
- LHCb is working on it and has opened tickets for sites to agree on a SRM-less future



SRM-less status

- DPM sites are an example of why it is necessary to do this
 - Legacy configuration ←
 - DOME configuration →
- Many sites configured DOME but are not actually using it because experiments is still configured to use SRM.
 - Need active work with sites to remove the features and adapt the experiments frameworks
- Dcache sites can also become SRM-less using gridftp
 - Sites with dcache >v4.1 also have the SRR json
 - One more reason to upgrade
- Storm sites would have to install a vanilla gridftp to use it without SRM
 - SRR json from >1.11.13



FTS

- To enable TPC activities and monitor them correctly
 - specific versions of FTS/gfal2/xrootd and davix and new library in order to issue tokens have been installed
 - CERN FTS are already upgraded
 - Andrea has opened tickets for RAL, FNAL, BNL
 - http streaming needs to be switched off to avoid ambiguity when monitoring TPC
 - This cannot be done yet until all sites can handle TPC in production
 - Http not used for TPC yet officially but ATLAS has already had complaints.
 - If streaming switched on ATLAS has to switch it off at sites that have it listed as one of the protocols
 - Also some FTS used also for other sciences



Matching protocols

- An immediate consequence of dropping SRM
 - In ATLAS if fallback protocols are configured with a certain order rucio will try to use them
 - Rucio needs to match the protocols on the list independently from what priority they have
 - Transfers should fail only if one of the two sites misses a protocol
 - This is under development and will be ready in October
 - CMS doesn't use fall backs so they haven't observed this problem yet.
- Still there is the problem of matching protocols
 - Sites supporting only one or the other protocol are still a problem
 - ATLAS wants both protocols deployed at all sites
 - Too early to tell for CMS



Experiments tests: ATLAS

- ATLAS functional tests
 - For now tests protocol is hardcoded (both http, xrootd)
 - All T1s + DOMA TPC sites
 - Need **FTS and sites** to be upgraded
 - BNL and RAL FTS need upgrading
 - Sites a strongly encouraged to plan for an upgrade

	AGLT2_SCRATCHDISK	AUSTRALIA-ATLAS_SCRATCHDISK	BNL-OSG2_SCRATCHDISK	CA-VICTORIA-WESTGRID-T2_SCRATCHDISK	CERN-PROD_SCRATCHDISK	DESY-HH_SCRATCHDISK	UNI-BONN_LOCALGROUPDISK
AGLT2_SCRATCHDISK	-	100%	-	100%	100%	100%	100%
AUSTRALIA-ATLAS_SCRATCHDISK	0%	-	0%	0%	0%	0%	0%
BNL-OSG2_SCRATCHDISK	-	100%	-	100%	100%	75%	100%
CA-VICTORIA-WESTGRID-T2_SCRATCHDISK	100%	100%	100%	-	100%	0%	100%
CERN-PROD_SCRATCHDISK	0%	0%	0%	0%	-	0%	0%
DESY-HH_SCRATCHDISK	0%	0%	0%	0%	0%	-	0%
PRAGUELCG2_SCRATCHDISK	0%	0%	0%	0%	0%	0%	-
UKI-NORTHGRID-MAN-HEP_SCRATCHDISK	100%	100%	100%	100%	83%	83%	100%
UNI-BONN_LOCALGROUPDISK	100%	67%	0%	0%	0%	50%	100%



Petr's tests

- Petr has also doing tests checking the AGIS configuration for each site
- Underneath it uses gfal-copy and it tests all protocols with direct transfers and TPC in two directions + it tries to guess the version

ATLAS Rucio configuration & TPC (Wed Jun 26 11:42:16 2019)

endpoint src -> dst	tier	size [TB]	protocol[version]	copy	copy	delete	tpc	tpc	tpc
				local dst	dst local	dst	src dst	dst src	dst dst
prague1cg2 AGLT2	2	6863	srm[dCache/5.0.5]	ok/1	ok/1	ok/3	ok/2	ok/1	ok/1
			gsiftp[FTP door ready]		ok/4	ok/1			
			root[xrootd/0x89020000]		ok/2		ok/3	error/3	error/3
			davs[dCache/5.0.5]	ok/2	ok/3	ok/2	ok/4	ok/2	ok/2
prague1cg2 ANLASC	3	13	gsiftp[gt6.0/12.9]	test failed					
			root[xrootd]	test failed					
prague1cg2 Australia-ATLAS	2	1432	srm[DPM/1.9.0-1]	ok/1	ok/3	ok/2	ok/2	ok/1	ok/1
			root[xrootd/0x10030000]		ok/1		ok/3	error/3	error/3
			davs[unknown]	error/2	ok/2	ok/1	ok/4	error/2	error/2
prague1cg2 Australia-ATLAS	2	1432	srm[unknown]	test failed					
prague1cg2 AYDIN	3	0	gsiftp[unknown]	test failed					
prague1cg2 BEIJING-LCG2	2	388	srm[DPM/1.10.0-1]	ok/1	ok/1		ok/2	ok/1	ok/1
			root[xrootd/0x40000]		ok/2	ok/2	ok/3	error/3	error/3
			davs[Apache/2.2.15]		error/3	ok/1	error/4	error/2	error/2
prague1cg2 BNL-ATLAS	1	21172	srm[dCache/4.2.23]	ok/1	ok/1	ok/2	ok/2	ok/1	ok/1
			gsiftp[FTP door ready]		ok/3	ok/3			
			root[xrootd/0x40000]	error/3	error/2		error/3	error/3	error/3
			root[xrootd/0x40000]	error/3	error/2		error/3	error/3	error/3
			root						
davs[unknown]	ok/2	ok/4	ok/1	ok/4	ok/2	ok/2			
prague1cg2 CA-SFU-T2	2	3220	srm[dCache/3.2.53]	ok/1	ok/1		ok/2	ok/1	ok/1
			root[xrootd/0x89020000]		error/3		error/3	error/3	error/3
			davs[unknown]	ok/2	ok/2	ok/1	error/4	ok/2	ok/2



More sites

- We since March we have added a number of extra sites particularly to the functional tests
 - We need more production sites willing to upgrade to the baseline and enable the required features
 - Particularly DPM sites, for which there is already a WLCG organised TF sites can refer to move to DOME.
 -but not only
 - Dcache sites can participate to http tests with almost every major release because most changes are backported.
 - For xrootd we need v5.2
 - Xrootd, EOS and StoRM are needed too



Conclusions

- Getting there
 - Need more tests for xrootd
- Sites are strongly encouraged to upgrade to baseline releases
 - It will help to plug them in a more straightforward manner into the tests
- In ATLAS becoming SRM-less is a related activity
- Goal is still to have an alternative TPC non-gridftp protocol at each site by the end of the year
 - Possibly both to avoid mismatches

