Data lake dress rehearsal: Experiment tests and Logbook

Tests planned during FDR (17 Nov)

Experime nt, Point of Contact	Test name	Scope to be used	Method (cron script, interactiv e) Cron Offset	Anticipated timeline (description, time constraints) Reason for test (capabilities tested, what constitutes success?)	File size, file number	Trans fer volum e for full test	Any QoS require ments	replicas / RSEs
Expt: SKA Point of Contact: James Collinson	(1) Pulsar Observ ations injection test	SKA_SKA O_TEAM _PSS, SKA_SKA O_TEAM _PST (need to update scope names)	cron Cron offset: 05mins	 Description: For 4 hours at any point during the 24hrs, injecting new group of files in a dataset every ten minutes. Files fall into two containers, representing different SKA Projects Reason: mimics a real SKA observation use case, files grouped together to give sensible transfer sizes Success if: Rucio handles file injection rate OK, with replications to correct sites completed after (3) days. 	project 1: 50 MB, 3500 files, once per hour for 24 hours project 2: 50 MB, 400 files	projec t 1: 4.2 TB Proje ct 2: 0.48 TB	No	How many replicas? 2 Which RSEs? Or how many? 2, transfers between DESY-DCACHE and EULAKE-1

				Verification method: Rucio events Dashboard interrogation				
Expt: SKA Point of Contact: Rohini Joshi	(2) 24hr QoS lifecycle	SKA_SKA O_TEAM _QOS24	script, launched automati cally Cron offset: 10mins	 Description: 24-hr test moving data on basis of QoS class. Reason: SKA data lifecycle test, but reduced overall timescale to give results in timely fashion. Success if: QoS can be used to automatically place data in appropriate locations with rules specified at the time of data generation. Capture how long this takes Verification method: Rucio events Dashboard interrogation (ES SKA) 	100MB x 10	1GB	Yes	How many replicas? 4 Which RSEs? Or how many? QoS based - 4 in total
Expt: SKA Point of Contact: Rohini Joshi	(3) SDC1 data movem ent	SKA_SKA O_TEAM	script, launched automati cally Will run once.	 Description: QoS based movement of SDC1 data Reason: Demonstrate making SDC1 data 'available' and ready for processing by moving it to FAST storage. Success if: Replicas available one on each RSE with the label QOS=FAST Verification method: Rucio CLI 	4GB, x4 + auxiliary files	17 GB	Yes	How many replicas? Which RSEs? Or how many? QoS based - equal to the number of 'FAST' RSEs

				+ dashboard interrogation				
Expt: CTA Point of Contact: Frederic Gillardo	(4) CTA ingest	CTA_LAP P_FREDE RIC	htCondor Cron offset: 20mins	 Description: for 6 h : Ingest 500 Dataset of 10 files The ingest is performed using the CTA data workflow: First, data are uploaded on LAPP-DCACHE(simulat es site storage) Then, data are replicated in europe. (ALPAMED-DPM , PIC-DCACHE) Finally data are removed from the Canary Island 	File size : 2, 0.1 0.01, GB Number of files : 560 000	Total size : 7 TB	Yes	How many replicas? 3 Which RSEs? Or how many? 3 transfers ALPAMED-DPM, LAPP-DCACHE, PIC-DACHE
				 Reason: simulate a night data captured from telescope in Canary Island Success if: Rucio handles file injection rate OK, with replications to correct sites completed after 4 hours. 				

Expt: ATLAS Point of Contact: Stéphane Jezequel	(5) ATLAS QoS function ality test	ATLAS_ LAPP_J EZEQU EL	Still interactiv e shell script (singulari ty) with 1 thread It takes 30 seconds to upload 2 GB (~0.5 Gb/s) in french storage federatio n or Napoli and 10 seconds to register files + dataset in Rucio + creating rules Cron offset: N/A	Description: ATLAS focuses on open data workflow (500 TB DAOD_PHYSLITE 200 days each year in 2029 → 2 GB/mn). ATLAS central production workflow integrated in Rucio > 10 years Reason: Contribute to demonstrate reliability of datalake and QoS functionality Success if: ESCAPE infrastructure (Rucio and datalake) do not collapse under load Verification method: Rucio events Dashboard	9 files from H→ yy analysis uploade d alternati vely in ALPLAM ED-DPM and INFN-N A-DPM (arbitrary upload sites) File size from 30 to 700 MB (total 2 GB)	1 datas et of 2 GB per mn → 3 TB over 24 hours Per replic a	YES In additio n to ALPA MED-F R/ INFN- NA-DP M	How many replicas? 1 copy QoS=SAFE (2nd SAFE replica) Lifetime (99999 s = 1 day) 2 copies QoS =CHEAP-ANALYSI S (lifetime : 9999 s = 3 hours) Which RSEs? Or how many? Up to Rucio to decide (hope it changes at each request)
--	---	---------------------------------	--	--	--	--	--	--

Expt: MAGIC Point of CoExpt: Agustin Bruzzese(6) MAGIC testMAGIC_ PIC_BR UZZESEcron Cron offset: 30minsDescription: Through means of rucio, we want to register all the mock MAGIC data appearing at PIC_INJECT, and replicate it to PIC_DCACHEFile size 600 MbIn this test, 24 tilest, 24 tilest, 24 will be creat ed ed every hour.NoHow many replicase? PIC_INJEC (deleted or successful replicated) testBruzzese(6) MAGIC testMAGIC_ offset: 30minsCron offset: 30minsDescription: Through means of rucio, we want to register all the pIC_INJECT, and replicate it to PIC_INJECT Reason: mimics a real MAGIC observation use case, with mock data but realistic file size but real namespaceFile size test testIn this test, 24 the every the <th>y file at The ly ACHE, Es? Or ? T and HE - two</th>	y file at The ly ACHE, Es? Or ? T and HE - two
--	---

Expt: LOFAR Point of Contact: Yan Grange/ VN.Pande y	(7) LOFAR processi ng pipeline	LOFAR_ ASTRO N_GRA NGE	Shell script Cron offset: 35mins	Description: Moving files to a faster QOS than they are, downloading subbands from an observation (few TB) and doing local processing. Re-ingesting the output into the data lake. We will repeat this a few times over the day to simulate multiple users Reason: This mimics a smaller scale user processing data. Success if: Processed data is well-ingested Verification method: Rucio dashboard(?)	Few TB, few tar files	Order 10 TB	Yes	How many replicas? 2 (of which 1 is temporary) Which RSEs? Or how many? SARA-DCACHE Some place with a cheap QOS
Expt: LOFAR Point of Contact: Yan Grange/ V.N.Pand ey	(8) LOFAR one day ingest	LOFAR_ ASTRO N_GRA NGE	Shell scripting Cron offset: 40mins	Description: We will be simulating a day of ingest in the LOFAR Long-Term archive. Data will go in two QOS levels (a fast and a slow one) and the fast one will disappear after the slow QOS is successfully copied and/or a specific time has passed. Reason: The idea is to see what would happen if the LOFAR LTA would be in the data lake. Success if: Data first exists remotely with both QOS levels, and later with only one. Verification method: Rucio events dashboard	~20TB over one day. ~1000 files	~20T B	Yes	How many replicas? 2 / 1 Which RSEs? Or how many? 2, of which SARA-DCACHE probably the primary upload point

Expt: CMS Point of Contact: Diego Ciangottin i	(9) CMS_o pendata	CMS_IN FN_DCI ANGOT	Shell/pyt hon scripts Cron offset: 45mins	Description: import in RUCIO an opendata dataset from existing grid storage outside the lake with one command line. Execute data discovery for a simple workflow asking for that dataset. Bonus: repeating the procedure on different QoS source Reason: end-to-end functionality check of all the compute and data integration Success if: the workflow succeed in reading data Verification method: exit code of the payload and copy of the file read on local disk	Few 100s of GBs O(100)fil es	Few 100s of GBs	Yes, but not essenti al	How many replicas? 1 or 2 Which RSEs? Or how many? CNAF storm + an other with an xrd protocol

Expt: LSST Point of Contact: Bastien Gounon / Fabio Hernande z	(10) LSST_d ata_ing est	LSST_C CIN2P3 _GOUN ON-test	Shell script, batch jobs	 Description: ingest the HSC RC2 dataset from CC-IN2P3 local storage to Rucio, at a realistic LSST data rate (20TB/24h), then confirm integrity and accessibility of the data via a notebook. Reason: confirm that the Rucio instance handles realistic data rates for the LSST experiment. Success if: the whole dataset (800GB) is successfully ingested in under ~1 hour, and we are able to print images in the notebook afterwards. Verification method: start and finish timestamp on the batch job in which the ingestion will be made, confirmation with Rucio CLI that the files are available, monitoring of the dCache instance to which the data is uploaded and successful execution of the FITS drawing notebook against the freshly ingested data. 	45248 x 18MB	814 GB	No	How many replicas? 1 Which RSEs? Or how many? IN2P3-CC-DCACH E

Expt: FAIR Point of Contact: Marek Szuba	(11) FAIR_in gest_m ock_cb m	FAIR_G SI_SZU BA	script, plus possibly cron / systemd timers	 Description: Upload, to GSI-ROOT, one 1-GB file every 10 minutes for the whole duration of the rehearsal. For each file, request 2 replicas in QOS=SAFE (=> 2 transfers) + 1 replica in QOS=CHEAP-ANALYSIS (=> 0 transfers - GSI-ROOT belongs to this class) Reason: The file size and the QoS tagging will roughly approximate data ingestion from CBM (i.e. the FAIR experiment expected to produce the largest volume of raw data), albeit at ~1/500 of the actual expected data rate (6 PB/month, i.e. 200 TB/day) due to space constraints. Success if: all uploads succeed and get replicated to respective QoS classes Verification method: with Rucio CLI 	144 x 1 GB	432 GB	Yes	How many replicas? 3 Which RSEs? Or how many? GSI-ROOT + 2/file from the QOS class SAFE

Expt: EGO/VIR GO Point of Contact: Pierre Chanial	(12) Virgo data real-tim e transfer	VIRGO_ EGO_C HANIAL	Python scripts, Celery	 Description: Upload 4h of Virgo public data sampled at 4kHz from an EGO server to the datalake Download them to CNAF. The data is split into 1s samples. Reason: Making available the real-time strain data to pipelines and tools assessing the data quality. Success if: All chunks have timely been transferred at CNAF. Verification method: Estimation of the latency distribution. 	4 * 3600 files of 85 kB	1.2 GB 1440 0 files	No	How many replicas? 1 Which RSEs? Or how many? EULAKE-1
---	--	---------------------------	------------------------------	---	-------------------------------	------------------------------	----	---

Expt: (spare	Description:	How many replicas?
Row)	Reason:	
Point of	Success if:	Which RSEs? Or
oomact.	Verification method:	now many :

RSE mapping for tests

Test / RSE	DESY- DCAC HE	SARA- DCAC HE	PIC-D CACH E	EULA KE-1	LAPP- DCAC HE	IN2P3- CC-D CACH E	CNAF- STOR M	ALPA MED- DPM	GSI-R OOT	INFN- NA-DP M	LAPP- WEBD AV	INFN- NA-DP M-FED	PIC-IN JECT (MAGI C ONLY)
1	5 TB (uploa d site)			5 TB									
2	1 GB (max)		1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	1 GB (max)	
3			17 GB								17 GB		
4 (CTA)			500 GB		500 GB			500 GB					
5	3 TB (max)					3 TB (max)	3 TB (safe)	3 TB (max)	3 TB (max)	3 TB (max)			
6 (MAGI C)			337.5 GB										337.5 GB (injecti on site)
7 (LOFA R Proces sing)		10 TB											

8 (LOFA R INGES T)		20TB (uploa d site)											
9 (CMS)							1 TB (max)						
10 (LSST)						20 TB							
11 (FAIR)		0.144 TB (max)			0.144 TB (max)		0.144 TB (max)		Upload (0.144 TB)				
12 (EGO)				1.2 GB 14400 files			1.2 GB 14400 files						
Test / RSE	DESY- DCAC HE	SARA- DCAC HE	PIC-D CACH E	EULA KE-1	LAPP- DCAC HE	IN2P3- CC-D CACH E	CNAF- STOR M	ALPA MED- DPM	GSI-R OOT	INFN- NA-DP M	LAPP- WEBD AV	INFN- NA-DP M-FED	PIC-IN JECT (MAGI C ONLY)

	DESY-DC ACHE	SARA-DC ACHE	PIC-DCAC HE	EULAKE-1	LAPP-DCA CHE	IN2P3-CC- DCACHE	CNAF-ST ORM	ALPAMED -DPM	GSI-ROOT	INFN-NA- DPM	LAPP-WE BDAV	INFN-NA- DPM-FED	PIC-INJEC T (MAGIC ONLY)
QoS offered	"CHEAP-A NALYSIS"	"SAFE"	"FAST"	"OPPORT UNISTIC"	"SAFE"	"CHEAP-A NALYSIS"	"SAFE"	"CHEAP-A NALYSIS"	"CHEAP-A NALYSIS"	"CHEAP-A NALYSIS"	"FAST"	?	"SAFE"
Storage offered to WP2 (from wiki)	110 ТВ	98 TB	30 TB	350 T B	10 TB	60 TB	10 TB (100k inode limit)	100 ТВ	1 TB	100 TB	(missing from wiki table)	(missing from wiki table - asked Alessandra Doria)	(MAGIC USE ONLY)
(Approx imate usage ~early Nov 2020	8 TB	50 TB	2 TB	7 TB	4 TB	7 TB	3 ТВ	3 ТВ	0.5 TB	7 TB	0.1 TB	0 ТВ	(MAGIC USE ONLY)
Max require d space for tests	8TB	31 TB	1 TB	6 TB	0.7 TB	23 TB	5 TB	3.5 TB	3 TB (!)	3 TB	0.2 TB		
protoco ls support ed	xrootd davs gsiftp	xrootd davs gsiftp	xrootd davs	xrootd davs gsiftp	davs	xrootd gsiftp	davs	xrootd davs gsiftp	xrootd davs	xrootd davs gsiftp	davs	?	?

Rucio stats dashboard: https://monit-grafana.cern.ch/d/74yXDN2Gk/rucio-stats?orgId=51

Log Books

Test Log book page (1) Pulsar Observations injection test

SCOPES: SKA_SKA0_TEAM_PSS, SKA_SKA0_TEAM_SPS **RSEs used:** DESY-DCACHE, EULAKE-1

Notes (issues and success):

Though there was concern that on the day the server loading could increase response times and prevent the cron running hourly (test run time is ~50min so server slowdown could lead to missed crons), for the first half of the day the upload rate was sufficient to complete each test.

When the auth server was restarted (~15:15 UTC) an error was thrown by the rucio-analysis package, caused by an error when calling the listReplicationRules method. Issue reported: <u>https://github.com/ESCAPE-WP2/rucio-analysis/issues/19</u>

Dashboard views:

Full day PSS test <u>here</u> Full day SPS test <u>here</u>

Annotated graphs showing events:



Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Tests were very successful, adding an amount of data into the data lake corresponding roughly to a pulsar search use case (spread over 24 hours) and single pulse search. One of the crons failed when the auth server was restarted (around 15:15) but the next cron proceeded normally. Transfer failure rates were very low, at less than 0.1%. In dashboard view above, some of the events relating to the replication occurred after 00:00 (end of the time window), but did complete successfully. Since we were at the limit of upload speed, and wish to ideally upload this amount of data over the course of 4 hours, it would be good to run a registration-only version of this test in future. SPS test achieved 100% success rate.

Test Log book page (2) 24hr QoS lifecycle

SCOPE: SKA_SKA0_TEAM_QOS24

RSEs used: ALL RSEs except INFN-NA-DPM (bc it has no QoS label)

Notes (issues and success):

Not impacted despite issues at GSI-ROOT. ES database sync was not added till 4 pm, synced retrospectively

Dashboard views:

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?from=1605571200000&orgId=51&to=1605657599000&var-bin=\$__auto_in terval_bin&var-scopes=SKA_SKAO_TEAM_QOS24&var-filename=*&var-rses=All&var-src=All&var-dst=All&var-datasource=Monit%2 0ESCAPE%20(long%20term)

Good spread of RSEs used via QoS **Annotated graphs showing events:**

Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed: Add panel to show QoS based success/failures and something to capture data deletion once QoS based rules expire.

Summary: Brief paragraph summarising the test, the results, success/fail Success

Test Log book page (3) SDC1 data movement

SCOPE: SKA_SKA0_TEAM RSEs used: DESY-DCACHE, LAPP-WEBDAV

Notes (issues and success):

16/11

2:45 Accidentally ran test to replicate to FAST. LAPP-WEBDAV replicas made
16:00 Ran "rucio delete-rule --purge-replicas 623f433d083041d8a9a55457a08cfebc"
17/11
9:30 LAPP-WEBDAV replicas seem to be gone now as well as the rule made yesterday
10:30 Test run kicks off. Dataset level rule is in stuck state due to 'no sources' bc existing rules from data ingest are file level. File replicas created on FAST QOS despite rule being stuck
Replications kicked off on 16/11, files deleted, recreated on 17/11 can be seen here:
https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=160553520000&to=1605614399000&var-bin=10m&var-s
copes=SKA_SKAO_TEAM&var-filename=*&var-rses=All&var-src=All&var-datasource=Monit%20ESCAPE%20(long%20t
erm)
11:18 Created file level rules to replicate to FAST QOS, should not initiate any transfers(1 has been queued)

Dashboard views:

Whole day view <u>link</u>

Annotated graphs showing events:

Things to change:

Changes to the test needed:

Changes to the datalake needed: Dataset level rule is in STUCK state, even though corresponding file replicas have successfully been made.

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Success, need to understand how collection level rules work with file level rules

Test Log book page (4) CTA ingest

SCOPE: CTA_LAPP_FREDERIC

RSEs used: LAPP-DCACHE, PIC-DCAHE, IN2P3-CC-DCACHE, ALPAMED-DPM

Dashboard views:

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605571200000&to=1605657599000&var-bin=\$__auto_in terval_bin&var-scopes=CTA_LAPP_FREDERIC&var-filename=*&var-rses=All&var-src=All&var-dst=All&var-datasource=Monit%20ES CAPE%20(long%20term)

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?from=1605571200000&orgId=51&to=1605743999000&var-bin=10m&var-s copes=CTA_LAPP_FREDERIC&var-filename=*&var-rses=All&var-src=All&var-dst=All&var-datasource=Monit%20ESCAPE%20(long %20term)&viewPanel=1438

Annotated graphs showing events:

Things to change:

Upload tests were supposed to run exclusively on lapp-dcache, but there was some connection issue so I moved the upload to IN2P3-CC-DCACHE.

Changes to the datalake needed:None

Changes to the dashboard needed:None

Summary:

The upload works OK, I had some connection issue probably because of an authentication issue on the RUCIO server. Replications rules are executed, but for now I do not have the tool to monitor that each replication rules are correctly executed (work in progress)

- 06:45: Test 1 started : 200 datasets with 20000 files with a size of 10 MB uploaded on LAPP-DCACHE
- 10:00: Test 1 is stopped : too many connection errors and client machine (lapp-esc03) is running out of CPU

- 10:48: Test 2 is started : 20 datasets with 20000 files with a size of 10 MB uploaded on IN2P3-CC-DCACHE from lappsI7e
- **13:48:** Test 3 is started : 200 datasets with 20000 files with a size of 1 MB uploaded on IN2P3-CC-DCACHE from lapp-esc03 with new configuration : 12 core 12 GB of RAM and singularity version upgraded to 3.6.4
- **15:18:** Test 4 is started : 20 datasets with 20000 files with a size of 1 MB uploaded on LAPP-DCACHE from lapps17b
- **19:37:** Test 5 is started : 5 datasets with 20 files with a size of 2 GB uploaded on LAPP-DCACHE from lappsl7g
- 21:01: Test 6 is started : 5 datasets with 200 files with a size of 2 GB uploaded on LAPP-DCACHE from lappsl7g
- 21:47: Test 7 is started : 5 datasets with 200 files with a size of 2 GB uploaded on LAPP-DCACHE from lappsI7b
- **23:00** Test 2 & 5 is finished (only replications rules remain)
- 07:53 (D+1): Test 3&4 finished (only replications rules remain)
- 09:25 (D+1): Test 6 finished (only replications rules remain)
- Replication rules, some of them stuck, not clear monitoring view, how do we monitor the final status?





Test Log book page (5) ATLAS Open Data QoS functionality test

SCOPE: ATLAS_LAPP_JEZEQUEL

RSEs used:

- Upload files from LAPP cluster (lappsl) to ALPAMED-DPM (FRANCE) and INFN-NA-DPM (ITALY)
- Request transfer to 1 RSE QoS=SAFE and 2 RSEs QoS=CHEAP-ANALYSIS

Notes (issues and success):

Dashboard views:

Rucio events (ATLAS_LAPP_JEZEQUEL scope, 17/11/2020):

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=160557120000&to=1605743999000&var-bin=\$__auto_in terval_bin&var-scopes=ATLAS_LAPP_JEZEQUEL&var-filename=*&var-rses=All&var-src=All&var-dst=All&var-datasource=Monit%20 ESCAPE%20(long%20term)

Annotated graphs showing events:

Things to change:

Changes to the test needed: No change

Changes to the datalake needed: No change

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

- Using Rucio client based on singularity instance (as normal user not admin support)
- Mostly successful although some worries about impact of network saturation on data uploading to Italy (effectively outside LAPP)



sr	DESY- DCACHE	SARA- OCACHE	PIC-OCACHE	EULAKE-1	LAPP- DCACHE	IN2P3-CD- DCACHE	CNAF-STORF	ALPANED-OPM	GSI-ROOT	INFN-NA-OPI	LAPP- WEBDAV	INFN-NA- DPM-FED
DESY-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	150s	NO DATA	100s	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
SARA-DCACHE	100%	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	100%	21s	100%	NO DATA	ND DATA
PIC-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
EULAKE-1	NO DATA	NO BATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
LAPP-DCACHE	100%	NO DATA	NO DATA	NO DATA	NO DATA	100x	NO DATA	100%	9.4s	100%	NO DATA	ND DATA
IN2P3-CC-DCACHE	NO DATA.	100s	NO DATA	NO DATA	100 x	NO DATA	100s	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
CNAF-STORM	100s	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	9.7s	100%	NO DATA	NO DATA
ALPAMED-DPM	27%	94s	NO DATA	NO DATA	100s	100%	100s	NO DATA	49s	92%	NO DATA	ND DATA
GSI-ROOT	NO DATA	88%	NO DATA	NO DATA	05	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
INFN-NA-DPM	100s	100s	NO DATA	NO DATA	99%	100%	99%	905	44s	NO DATA	NO DATA	NO DATA
LAPP-WEBDAV	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	ND DATA
INFN-NA-DPM-FED	NO DATA.	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA

~ Transfer Matrix: transfer-done/transfer-submitted

•

 During many hours on 17th November, the outgoing 10 Gb/s link from LAPP is saturated by ATLAS data rebalancing (ATLAS-LAPP storage close to be full in LAPP so some datasets are automatically moved to other ATLAS storages) on top of usual activities.



- According to ATLAS monitoring, only LAPP has such large transfer activity today. And FTS just tries to fill the pipe as much as possible until some transfers fails (indication that pipe is full).
 - \circ Noticed because upload to Napoli was slower than usual \rightarrow Upload AT RISK
 - Only slowing down FTS transfers from ALPAMED as expected
- Few errors on data distribution through FTS ALPAMED-DPM → DESY-DCACHE (only small fraction of specific input files) : To be understood
 - Since input was replicated in other RSEs, why not switch to use other RSEs as input ?
 - No issue if source INFN-NA-DPM

sr	DESY- DCACHE	SARA- DCACHE	PIC-OCACHE	EULAKE-1	LAPP- DCACHE	IN2P3-CC- DCACHE	CNAF-STORF	ALPAMED-OPM	GSI-ROOT	INFN-NA-OPI	LAPP- WEBDAV	INFN-NA- DPM-FED
DESY-DCACHE	NO DATA	NO DATA	NO DATA	ND DATA	150%	NO DATA	100s	NO DATA	NO DATA	ND DATA	NO DATA	NO DATA
SARA-DCACHE	100%	NO DATA	NO DATA.	NO DATA	NO DATA	NO DATA	NO DATA	100%	21%	100s	NO DATA	ND DATA
PIC-DCACHE	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	ND DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
EULAKE-1	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
LAPP-DCACHE	100s	NO DATA	NO DATA	NO DATA	NO DATA	100x	NO DATA	100s	9.4s	1005	NO DATA	NO DATA
IN2P3-CC-DCACHE	NO DATA	100s	NO DATA	NO DATA	100%	NO DATA	100%	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
CNAF-STORM	100s	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	ND DATA	NO DATA	9.7s	100s	NO DATA	NO DATA
ALPAMED-DPM	27x	94s.	NO DATA.	NO DATA	100s	100%	100%	NO DATA	49%	92%	NO DATA	ND DATA
GSI-ROOT	NO DATA	88%	NO DATA	NO DATA	0%	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
INFN-NA-DPM	100s	100s	NO DATA	NO DATA	99s	100%	99x	90%	44s	NO DATA	NO DATA	NO DATA
LAPP-WEBDAV	NO DATA	NO DATA	NO DATA	ND DATA	NO DATA	NO DATA	NO DATA	ND DATA	NO DATA	ND DATA	NO DATA	ND DATA
INFN-NA-DPM-FED	ND DATA.	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA

Transfer Matrix: transfer-done/transfer-submitted

Test Log book page (6) MAGIC Injection test

SCOPE: MAGIC_PIC_BRUZZESE

RSEs used:

- 1 temporary file at PIC-INJECT (deleted once successfully replicated)
- 1 permanent copy at PIC-DCACHE

Notes (issues and success):

Dashboard views:

1- <u>https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605484800000&to=1605743999000</u>

2- https://monit-grafana.cern.ch/d/oDm2GWoMz/magic-dashboard?orgId=51

Annotated graphs showing events:

Things to change:

Changes to the test needed: No changes needed on this regard

Changes to the datalake needed: No changes needed on this regard

Changes to the dashboard needed:



Summary: Brief paragraph summarising the test, the results, success/fail

Success

- Our tests started at 9:30 in the morning.
- Briefly, 24 files of 600mb each are replicated every hour from PIC-INJECT to PIC-DCACHE (Fig 1).
- This replication workflow was maintained throughout the dress rehearsal
- Files were replicated through datasets



Fails

- We had some connection problems between 3:00 pm 4:00 pm and 5:00 pm 6:00 pm as you can see in the image above (Fig 1).
- It was observed that some rules were not removed when the transfers finished successfully. It's probably our fault and has nothing to do with Rucio's instance. I have to investigate about it
- Files not deleted from storage, need to be tracked. Rizart mentioned some weird number of files reported by the PIC RSEs (negative values, etc.). To be investigated.

Test Log book page (7) LOFAR processing pipeline

SCOPE: LOFAR_ASTRON_GRANGE **RSEs used:**

Notes (issues and success):

Dashboard views:

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605484800000&to=1605743999000

Annotated graphs showing events:

(Open Nebula keeps track of the speeds. I am sure the API can give me a more complete plot, but this is already nice)



NET DOWNLOAD SPEED

Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

There were some unexpected changes of the RSE to use. During the day SARA-DCACHE was available, but since we prepared to use EULAKE-1 we tried using it. Downloads from EULAKE-1 had several issues. We need to investigate what went wrong and why, and repeat the experiment when we know.

The Download went on quite fast (1Gbps) somewhere in the afternoon and then later during the night we just could not download the data from EULAKE-1.

- 1. Download one data set (6hour observation) from EULAKE-1 (6GB)
- 2. UNTAR IT
- 3. CUT OUT 2 hours of data, FLAG and AVERAGE Data GET A NEW MS (Measurement SET).
- 4. Calibrate NewMS.
- 5. Apply Calibration.
- 6. Image the NEWMS .. (FITS) One Sample Image is shown below of a Radio Source 3C196.
- 7. UPLOAD the NEWMS and FITSIMAGE Back to EULAKE-1. (LIFE TIME 2-3 DAYS)
- 8. REPEAT FROM 1 to 6 for 23 DATA SETS... (That did not work as we could only get one data set)
- 9. IMAGE all the 23 NEWMS together...

Well we could only do from 1-6 for one data set. After that the experiment was not further possible due to inability to download the data. We need to look into it.



An radio image of the astronomical radio source 3C196 made using LOFAR data. The raw visibility data was downloaded via rucio from the EULAKE-1 (at CERN and is a part of DATALAKE comprising of infrastructure at several RSEs). It was processed on open nebula cloudbox at surfsara using the container based LOFAR software.

The kind of error message are like below..

______ rucio -vvv download LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:23:41,795 INFO Processing 1 item(s) for input 2020-11-18 09:23:41,796 DEBUG downloadclient.py download dids num unmerged items=1; num dids=1; num merged items=1 2020-11-18 09:23:41,796 INFO Getting sources of DIDs 2020-11-18 09:23:41.796 DEBUG downloadclient.pv _get_sources schemes: None 2020-11-18 09:23:41,796 DEBUG downloadclient.py rse expression: *\istape=true get sources 2020-11-18 09:23:41,796 DEBUG downloadclient.pv num DIDs for list replicas call: 1 get sources get sources num resolved files: 1 2020-11-18 09:23:43,057 DEBUG downloadclient.py 2020-11-18 09:23:43,073 DEBUG downloadclient.py is useable "unzip -v" returned with exitcode 127 2020-11-18 09:23:43,086 DEBUG downloadclient.py is useable "tar --version" returned with exitcode 0 2020-11-18 09:23:43,086 DEBUG downloadclient.py prepare items for download Queueing file: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:23:43,086 DEBUG downloadclient.py prepare items for download real parents: set() 2020-11-18 09:23:43,086 DEBUG downloadclient.py prepare items for download options: {'LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar': {'destinations': {('.', False)}, 'ignore checksum': False, 'transfer timeout': 360}} 2020-11-18 09:23:43,086 DEBUG downloadclient.py prepare items for download Prepared sources: num sources=3/3; num non cea sources=3; num cea ids=0 2020-11-18 09:23:43,086 INFO Using main thread to download 1 file(s) 2020-11-18 09:23:43,086 DEBUG downloadclient.py download worker Start processing queued downloads Preparing download of LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:23:43,087 INFO 2020-11-18 09:23:43.228 INFO Trying to download with gsiftp from EULAKE-1: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:23:43,306 DEBUG gfal.py connect connecting 2020-11-18 09:23:43,463 DEBUG downloadclient.py download item Deleting existing temporary file: /data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:23:43,465 DEBUG gfal.py get path gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio_test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.d ppp.tar 2020-11-18 09:23:43,466 DEBUG gfal.py gfal2 copy src: gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.d ppp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:29:44.203 DEBUG downloadclient.pv download item The requested service is not available at the moment. Details: An unknown exception occurred.

Details: Transfer canceled because the timeout expired 2020-11-18 09:29:44.203 WARNING Download attempt failed. Trv 1/2 2020-11-18 09:29:44,317 DEBUG downloadclient.py download item Deleting existing temporary file: /data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:29:44,375 DEBUG gfal.py get path gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio_test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.d ppp.tar 2020-11-18 09:29:44,375 DEBUG gfal.py __gfal2_copy src: gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio_test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.d ppp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:35:45,138 DEBUG downloadclient.py download item The requested service is not available at the moment. Details: An unknown exception occurred. Details: Transfer canceled because the timeout expired 2020-11-18 09:35:45,139 WARNING Download attempt failed. Try 2/2 2020-11-18 09:35:45,265 DEBUG gfal.py close closing protocol connection 2020-11-18 09:35:45,269 INFO Trying to download with root from EULAKE-1: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:35:45,269 DEBUG gfal.py connect connecting 2020-11-18 09:35:45.272 DEBUG downloadclient.pv download item Deleting existing temporary file: /data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:35:45,375 DEBUG gfal.py get path root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio_test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.d ppp.tar 2020-11-18 09:35:45,375 DEBUG gfal.py gfal2 copy src: root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.d ppp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:42:45,376 DEBUG gfal.py gfal2 cancel cancelling gfal proc. download item The requested service is not available at the moment. 2020-11-18 09:42:47,612 DEBUG downloadclient.py Details: An unknown exception occurred. Details: Error on XrdCI::CopyProcess::Run(): [ERROR] Operation interrupted: operation canceled: The copy-job has been cancelled! 2020-11-18 09:42:47,612 WARNING Download attempt failed. Try 1/2 2020-11-18 09:42:47.745 DEBUG downloadclient.pv download item Deleting existing temporary file: /data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:42:47.891 DEBUG afal.pv aet path root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.d ppp.tar 2020-11-18 09:42:47,892 DEBUG gfal.py __gfal2_copy src: root://eoseulake.cern.ch:1094//eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.d ppp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part

2020-11-18 09:49:47,893 DEBUG gfal.py gfal2 cancel cancelling gfal proc. download item The requested service is not available at the moment. 2020-11-18 09:49:55.734 DEBUG downloadclient.pv Details: An unknown exception occurred. Details: Error on XrdCI::CopyProcess::Run(): [ERROR] Operation interrupted: operation canceled: The copy-job has been cancelled! 2020-11-18 09:49:55.734 WARNING Download attempt failed. Try 2/2 2020-11-18 09:49:55,840 DEBUG gfal.py close closing protocol connection 2020-11-18 09:49:55,840 INFO Trying to download with days from EULAKE-1: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 09:49:55,840 DEBUG gfal.py connect connecting 2020-11-18 09:49:55.843 DEBUG downloadclient.pv download item Deleting existing temporary file: /data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:49:55,909 DEBUG gfal.py get path davs://eoseulake.cern.ch:443/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dp pp.tar 2020-11-18 09:49:55,909 DEBUG gfal.py gfal2 copy src: davs://eoseulake.cern.ch:443/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dp pp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 09:55:56,552 DEBUG downloadclient.py download item The requested service is not available at the moment. Details: An unknown exception occurred. Details: Transfer canceled because the timeout expired 2020-11-18 09:55:56,552 WARNING Download attempt failed. Try 1/2 2020-11-18 09:55:56,659 DEBUG downloadclient.py download item Deleting existing temporary file: /data/pandev/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandev.5ch4s.dppp.tar.part 2020-11-18 09:55:56,722 DEBUG gfal.py get path davs://eoseulake.cern.ch:443/eos/eulake/tests/rucio_test/eulake_1/LOFAR_ASTRON_GRANGE/43/a5/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.dp pp.tar 2020-11-18 09:55:56,722 DEBUG gfal.py gfal2 copy src: davs://eoseulake.cern.ch:443/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dp pp.tar dest: file:///data/pandey/datalake lofar rehersal/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 10:01:56,102 DEBUG downloadclient.py download item The requested service is not available at the moment. Details: An unknown exception occurred. Details: Transfer canceled because the timeout expired 2020-11-18 10:01:56,103 WARNING Download attempt failed. Try 2/2 2020-11-18 10:01:56.231 DEBUG afal.pv close closing protocol connection 2020-11-18 10:01:56,233 ERROR Failed to download file LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 10:01:56,233 ERROR None of the requested files have been downloaded.

After sometime... The same thing worked.. So it even makes it more important to see what is going wrong (including possibility that we are doing something wrong at our end)

rucio -vvv download LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2>&1 | tee out rucio test.txt 2020-11-18 11:52:06,875 INFO Processing 1 item(s) for input 2020-11-18 11:52:06,876 DEBUG downloadclient.py download dids num unmerged items=1; num dids=1; num merged items=1 2020-11-18 11:52:06,876 INFO Getting sources of DIDs 2020-11-18 11:52:06,876 DEBUG downloadclient.py schemes: None get sources 2020-11-18 11:52:06,876 DEBUG downloadclient.py get sources rse expression: *\istape=true 2020-11-18 11:52:06,876 DEBUG downloadclient.py num DIDs for list replicas call: 1 get sources 2020-11-18 11:52:07,803 DEBUG downloadclient.py get sources num resolved files: 1 2020-11-18 11:52:07,821 DEBUG downloadclient.py is useable "unzip -v" returned with exitcode 127 is useable "tar --version" returned with exitcode 0 2020-11-18 11:52:07,831 DEBUG downloadclient.py 2020-11-18 11:52:07,832 DEBUG downloadclient.py prepare items for download Queueing file: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 11:52:07.832 DEBUG downloadclient.pv prepare items for download real parents: set() 2020-11-18 11:52:07,832 DEBUG downloadclient.py prepare items for download options: {'LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar': {'destinations': {('.', False)}, 'ignore checksum': False, 'transfer timeout': 360}} 2020-11-18 11:52:07,833 DEBUG downloadclient.py Prepared sources: num sources=3/3; prepare items for download num_non_cea_sources=3; num_cea_ids=0 2020-11-18 11:52:07,834 INFO Using main thread to download 1 file(s) 2020-11-18 11:52:07.834 DEBUG downloadclient.pv Start processing queued downloads download worker 2020-11-18 11:52:07,834 INFO Preparing download of LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 11:52:07,968 INFO Trying to download with gsiftp from EULAKE-1: LOFAR ASTRON GRANGE:L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar 2020-11-18 11:52:08,057 DEBUG gfal.py connect connecting 2020-11-18 11:52:08,188 DEBUG gfal.py get path gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.d ppp.tar 2020-11-18 11:52:08.189 DEBUG gfal.pv gfal2 copy src: gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio test/eulake 1/LOFAR ASTRON GRANGE/43/a5/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.d ppp.tar dest: file:///data/pandey/TEST_DOWNLOAD/LOFAR_ASTRON_GRANGE/L431602_SAP000_SB120_uv.MS.pandey.5ch4s.dppp.tar.part 2020-11-18 11:57:03,295 DEBUG gfal.py close closing protocol connection 2020-11-18 11:57:03,299 DEBUG downloadclient.py download item renaming '/data/pandey/TEST DOWNLOAD/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar.part' to '/data/pandey/TEST DOWNLOAD/LOFAR ASTRON GRANGE/L431602 SAP000 SB120 uv.MS.pandey.5ch4s.dppp.tar'

2020-11-18 11:57:03,432 INFO File LOFAR_ASTRON_GRANGE:L431602_SAP000_SB120_uv.MS.pandey.5ch4s.dppp.tar successfully downloaded. 6.503 GB in 279.48 seconds = 23.27 MBps

Download summary

DID LOFAR_ASTRON_GRANGE:L431602_SAP000_SB120_uv.MS.pandey.5ch4s.dppp.tar Total files: 1 Downloaded files: 1 Files already found locally: 0 Files that cannot be downloaded: 0 Completed in 296.6013 sec. [Wed Nov 18-11:57:03-pandey@145: /data/pandey/TEST_DOWNLOAD]>

Test Log book page (8) LOFAR one day ingest

SCOPE: LOFAR_ASTRON_GRANGE **RSEs used:**

Notes (issues and success):

Dashboard views:

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605484800000&to=1605743999000

Annotated graphs showing events:

(Open Nebula keeps track of the speeds. I am sure the API can give me a more complete plot, but this is already nice)





Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Uploads to EULAKE-1 went smooth, but for larger files (60GB), the checksum computation probably got in the way because there seems to be a timeout of 420 seconds after the data arrived. On the local machine, the time it takes rucio between the upload command and the actual data streaming is about 10 minutes. If we assume this is checksum computation, this is obviously more than

the 420 seconds timeout. 30 GB files go through without problems. The bandwidth gets up to 3-4Gbps to this location, without much network tuning, which is quite good.

Uploads to SARA-DCACHE went through smoothly without timeouts (until the quota limit was reached, after which all transfers failed). The bandwidth to SARA-DCACHE is around 9Gbps which is more or less max line speed (but since this is within SURFsara it is probably to be expected).

Transferring a whole observation with a single Rucio command is probably not a good idea, since rucio spends a lot of time before transferring any data then (we killed the job after a few hours). The assumption is that it first computes checksums for all the files before starting transfers. We should investigate if a way can be found to do this somehow in parallel. For now, the solution we found is to send every file independently in a script, and run four threads of rucio in parallel.

Test Log book page (9) CMS_opendata

SCOPE: CMS_INFN_DCIANGOT RSEs used: CNAF EULAKE-1

Notes (issues and success):

- 2 minor issues
 - No access to dashboard below (fix in progress)
 - Error on Checksum during upload to cnaf-storm (leading zeros issue on adler 32 checksums)
 - to be investigated yet
 - So I used eulake-1 as entrypoint for importing data into the datalake

O(300GB) open data correctly imported and analyzed using rucio for **data discovery** of the location.

- 2 imports at different time, to test the append data to dataset

2 replicas one with qos opportunistic

Dashboard views:

https://monit-grafana.cern.ch/d/74yXDN2Gk/rucio-stats?orgId=51&from=1605571200000&to=1605657599999&var-datasourc e=Monit%20ESCAPE%20(short%20term)&var-rses=All&var-scopes=All&var-experiments=CMS&var-bin=15m&viewPanel=24

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605484800000&to=1605743999000

Annotated graphs showing events:

Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Test Log book page (10) LSST_data_ingest

SCOPE: LSST_CCIN2P3_GOUNON-test **RSEs used:** IN2P3-CC-DCACHE

Notes (issues and success):

- Problem with the batch system at IN2P3-CC, no file uploads. Interactive run to be able to run the jobs.
- Singularity container image ready to run elsewhere, no dependencies. All uploads done by this singularity container w/davs.

Dashboard views:

https://monit-grafana.cern.ch/d/DZsMzeFMk/rucio-events?orgId=51&from=1605484800000&to=1605743999000

(This scope only) Rucio events dash

Rucio stats for this experiment

Annotated graphs showing events:

Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Test Log book page (11) FAIR_ingest_mock_cbm SCOPE: FAIR_GSI_SZUBA RSEs used:

Notes (times are approximate and in UTC):

- 23:10: test starts running
- 08:00: first check shows all uploads to GSI-ROOT stuck as REPLICATING and 0% success rate for inbound replications. Server found to be massively overloaded (5-minute system load consistently not dropping below 40, with spikes of over 120), likely due to I/O issues because CPU load seemed to be under control (300-400% on an 8-core system). Furthermore, at some point overnight the server ran out of storage space (with ESCAPE_CERN_TEAM-noise as the main consumer - around 750 GB).
- **09:00:** GSI-ROOT has now got a high-water storage mark set to 99% of total capacity. Having temporarily removed the QOS attribute has resulted in the average system load dropping to normal levels, however there are still periodic spikes of over 120.
- **10:15:** disabling FTS tests for GSI-ROOT eliminated load spikes, however we continue to see HTTP Error 500 for inbound WebDAV transfers (which started, for unknown reasons, on Monday afternoon) and uploads. Server logs show no errors.
- **12:30:** set 'skip_upload_stat: True' (nb. according to comments in rucio-clients source code running stat on uploaded files is now obsolete) and 'verify_checksum: False' (no-op for uploads when the above is in effect but might still be relevant for replication) on GSI-ROOT. Local WebDAV uploads begin to succeed.
- **12:45:** confirmed inbound transfers have begun to succeed as well, re-set QOS=CHEAP-ANALYSIS on GSI-ROOT
- 13:30: reduced the interval between upload attempts from 10 to 5 minutes
- **16:00:** the last 7 upload attempts failed due to TLS errors connecting to the Rucio auth server but now everything works again, with no intervention from our side
- **23:40:** noticed that replication from GSI-ROOT to QOS=SAFE had stopped working around 22:50. The dashboard shows no transfer requests having even been submitted since then, anywhere in the data lake.
- 01:00: test terminates
- 10:00: observed that all previously hanging replication requests have eventually succeeded overnight

Dashboard views:

Rucio events dash for this scope (FAIR_GSI_SZUBA)

Rucio stats: Rucio stats for FAIR experiment



Annotated graphs showing events:

(09:50: reaper kicks in. 12:30: uploads begin to succeed. 13:30: we double upload frequency. 15:30-16:00: TLS errors connecting to the Rucio servers)

Things to change:

Changes to the test needed: none

Changes to the datalake needed:

set Rucio free-space threshold on GSI-ROOT to 10 GB

- disable periodic FTS-transfer tests for GSI-ROOT
- disable checksum verification for files uploaded/replicated to GSI-ROOT

Changes to the dashboard needed: none

Summary:

Test overall successful, not in the least because it shows the data lake's capability for quickly reacting to attempts of repairing/mitigating failures; on that note care should be taken to make as many of such changes possible without human intervention on the CERN side (many thanks to Riccardo and Rizart for quickly making changes in CRIC and to FTS configuration - but unless we have a permanent on-call team, this is obviously not sustainable).

The massive overload of GSI-ROOT was unexpected and we are now in contact with xrootd developers to figure out the best way of limiting concurrent client access to our RSE, as without it our server would be vulnerable to Denial of Service attacks (accidental or otherwise) even after the upcoming upgrade of the storage back-end.

It is still unknown why GSI-ROOT began having problems with WebDAV file checksumming on Monday, we will investigate it after the dress rehearsal.

Test Log Book Page (12) Virgo data real-time transfer

SCOPE: VIRGO_EGO_CHANIAL **RSEs used:**

Notes (issues and success):

Dashboard views:

Rucio events dash for this scope does not show much data...

Rucio stats: Rucio stats for VIRGO experiment



Annotated graphs showing events:

Things to change:

Changes to the test needed:

Changes to the datalake needed:

Changes to the dashboard needed:

Summary: Brief paragraph summarising the test, the results, success/fail

Network issues :

- Since 2:00 am, the outgoing 10 Gb/s link from LAPP is saturated by ATLAS data rebalancing (ATLAS-LAPP storage close to be full in LAPP so some datasets are automatically moved to other ATLAS storages) on top of usual activities. According to ATLAS monitoring, only LAPP has such large transfer activity today. And FTS just tries to fill the pipe as much as possible until some transfers fails (indication that pipe is full). This might affect FTS efficiency of ESCAPE FTS transfers.



General Comments :

- 1. Storage space issue at GSI-ROOT shows the need to have thresholds set to activate reaper demon (@riccardo) all sites should set their free space limit
- we should probably update the FTS configuration so that it only schedules (max) 2 (or so) concurrent transfers. apropos (a.p.millar@gmail.com)
- 3. Would it be an idea to use the GOCdb to announce downtimes?
- 4. Should JIRA be evolved into GGUS (WLCG ticketing system)
- 5. Faster upload (and download) access

- 6. Auth server had to be restarted in order to set the limits higher next time ensure these are set high enough before testing period. (Riccardo / Rizart, more details here on how things were looking and what projections to make in future)

 a. Outcome/solution: Run multiple auth server pods, load sharing
- 7. There is a need for a Kubernetes monitoring dashboard https://jira.skatelescope.org/browse/EDLK-100
- 8. Dark data check?
- 9. Manual vs. Automated fixes to discuss in the workshop

Summary - All scopes

Rucio Transfers done during the 24hr FDR



Transfers failed

