FTS3
Data movement service for WLCG

Workshop on the future of Big Data management

Michail Salichos
CERN

What is FTS?

- FTS is the service responsible for distributing the majority of LHC data across the WLCG infrastructure (transferred 25PB in 2012)
- Is a low level data movement service, responsible for moving sets of files from one site to another while allowing participating sites to control the network resource usage
- Users interact with FTS by submitting transfer jobs, that simply say "copy <source URL> to <destination URL>
  - FTS then queues, schedules and performs the transfer, retrying it if necessary

FTS3 main features

- protocols support: srm, gsiftp, http, xroot
- db back-ends: oracle, MySql
- simplified configuration using JSON formatted messages
- scales well horizontally
- job priorities
- debug mode transfer logging – gridftp control channel info
- blacklisting users (DN) and SEs
- smart transfer retry mechanism based on error classification
- transfer auto-tuning (dynamically adjusting the number of active transfers based on success/failure rate and achieved throughput)
- REST-style interface for transfer submission and status retrieval
- multiple replicas support
- session/connection reuse (gridftp, ssl, etc)
- the list goes on … https://svnweb.cern.ch/trac/fts3

FTS3 resource management

- FTS3 is designed to "fill up" and efficiently use its infrastructure
  - protection from overload of network and endpoints is implemented in two ways
    - auto-tune transfers
    - direct/manual configuration of limits

FTS3 managing resources example

Standalone endpoint configuration:

```
{
"se" : "se.cern.ch",
"active" : true,
"in" : {
      "share" : [{"cms" : 5}, {"atlas" : 5}, {"public" : 5}],
      "protocol" : [{"nostreams" : 5}, {"urlcopy_tx_to" : 3000}]
      },
"out" : {
      "share" : [{"cms" : 6}, {"atlas" : 5}, {"public" : 4}],
       "protocol":[{"nostreams" : 8}, {"urlcopy_tx_to" : 3600}]
      }
}
```

Endpoints pair:

```
{
      "symbolic_name" : "se-link",
      "source_se" : "se1.cern.ch",
      "destination_se" : "se2.cern.ch",
      "share" : [{"cms" : 1}, {"atlas" : 2}, {"public" : 3}],
      "protocol" : [{"nostreams" : 8}, {"urlcopy_tx_to" : 3600}],
      "active":true
}
```

FTS3 configuration model

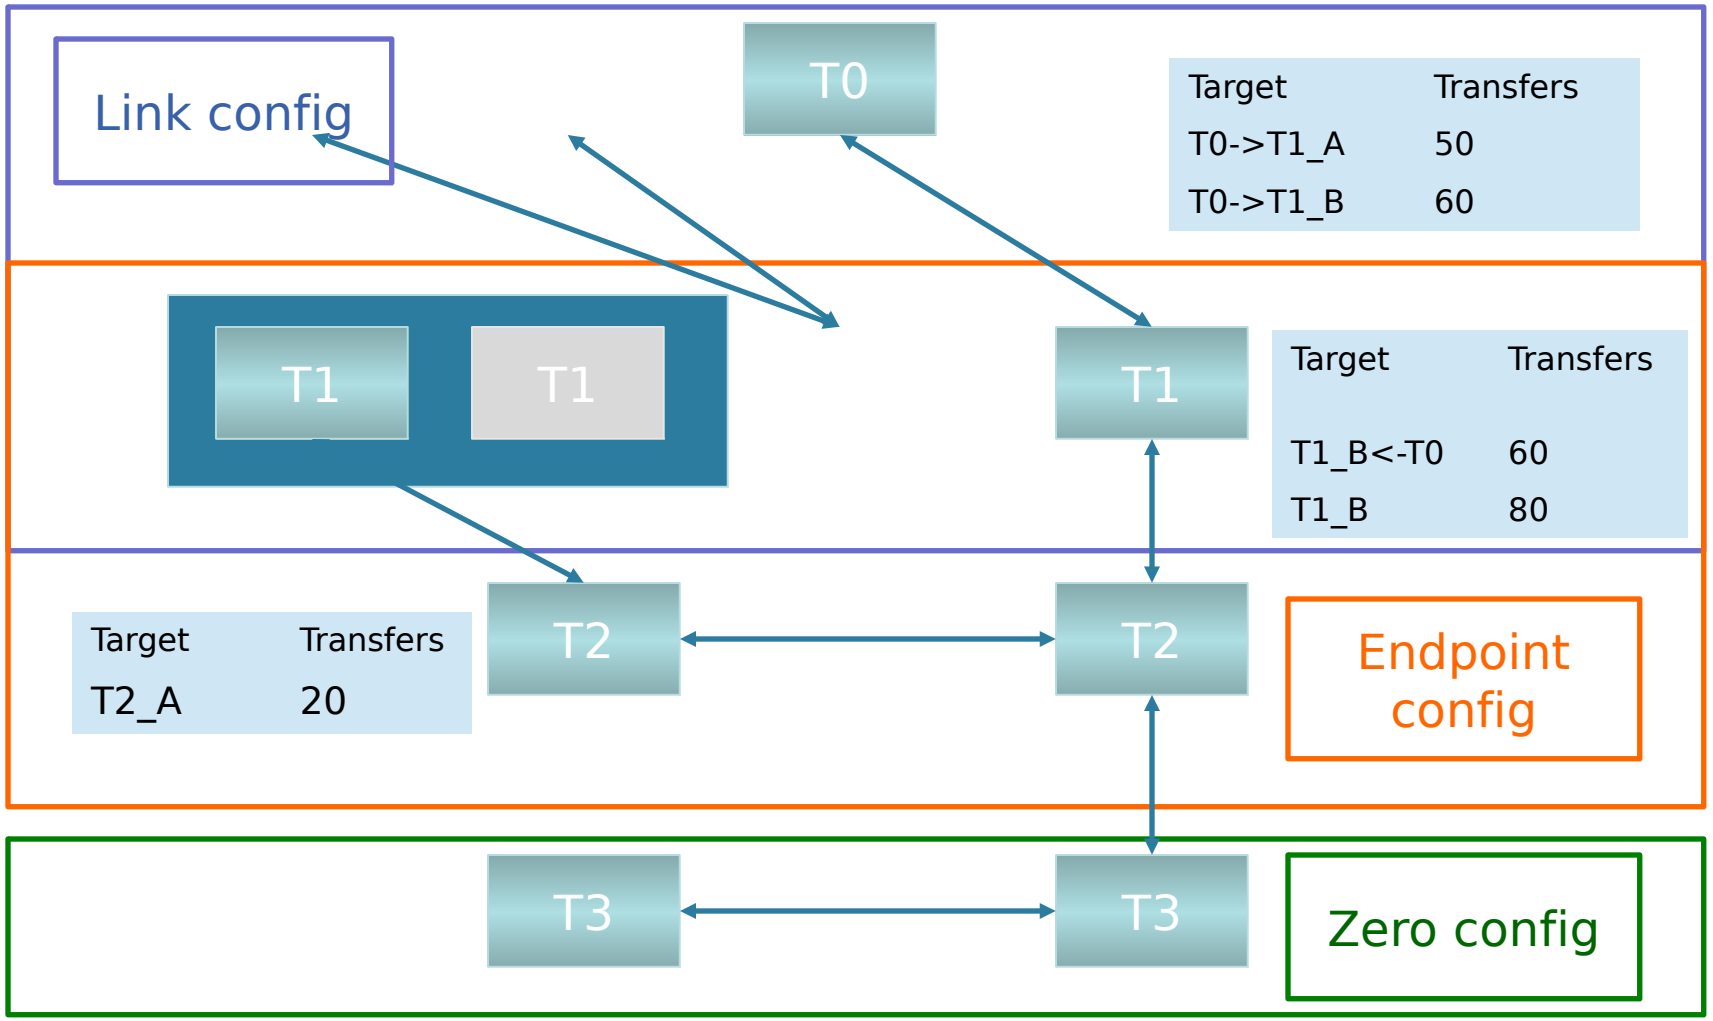

FTS3 protocols support

- FTS3 is built on top of gfal2
- gfal2
  - Provides protocol plug-ins
  - Offers a 3rd party copy extension
  - Much more

    https://svnweb.cern.ch/trac/lcgutil/wiki/gfal2

FTS3 RESTful API

- Currently, a subset of FTS3 operations can be used through the RESTful API
  - Submit jobs and get job information
- What's the added value of this?
  - Standard clients and/or libraries can be used to do these operations
    - [lib]Curl, Python's urllib2...
- Easy to install
  - yum install fts-rest
  - service httpd restart
  - You are done!
- More info https://svnweb.cern.ch/trac/fts3/wiki/FTSRest

FTS3 RESTful example

```
$ curl --capath /etc/grid-security/certificates -E ~/proxy.pem --cacert ~/proxy.pem https://fts3-pilot.cern.ch:8446/jobs/a40b82b7-1132-459f-a641-f8b49137a713
{
  [some values omitted]
  "dest_se": "gsiftp://lxfsra10a01.cern.ch",
  "files": [
    {
      [some values omitted]
      "source_surl": "gsiftp://lxfsra04a04.cern.ch/dpm/cern.ch/home/dteam/vector-sample.txt",
      "dest_surl": "gsiftp://lxfsra10a01.cern.ch/dpm/cern.ch/home/dteam/copy.9.77",
      "file_state": "FAILED",
      "finish_time": "2013-03-22T13:37:33",
      "internal_file_params": "nostreams:1,timeout:5000,buffersize:0",
      "job_finished": "2013-03-22T13:37:33",
      "job_id": "a40b82b7-1132-459f-a641-f8b49137a713",
      "reason": "globus_ftp_client: the server responded with an error 426 Transfer failed due to unexpected exception: java.io.IOException: Connection timed out "",
      "start_time": "2013-03-22T13:37:32",
      "transferhost": "fts3src2",
    }
  ],
  "finish_time": "2013-03-22T13:37:33",
  "job_id": "a40b82b7-1132-459f-a641-f8b49137a713",
  "job_state": "FAILED",
  "overwrite_flag": null,
  "reason": "One or more files failed. Please have a look at the details for more information",
  "source_se": "gsiftp://lxfsra04a04.cern.ch"
}
```

FTS3 RESTful example (2)

```
{
 "files": [
    {
     "sources": ["root://source/file"],
     "destinations": ["root://dest/file"],
     "metadata": "User defined file metadata",
     "filesize": 1024,
     "checksum": 'adler32:1234',
    }
  ],
  "params": {
    "verify_checksum": true,
    "reuse": false,
    "job_metadata": "User defined job metadata",
    "overwrite": false,
  }
}
```

FTS3 monitoring

- Using messaging for getting a global picture
  - Each FTS3 server publishes messages to a message bus to report transfer status and state transitions
- Web interface for individual FTS3 server monitoring

FTS3 global monitoring

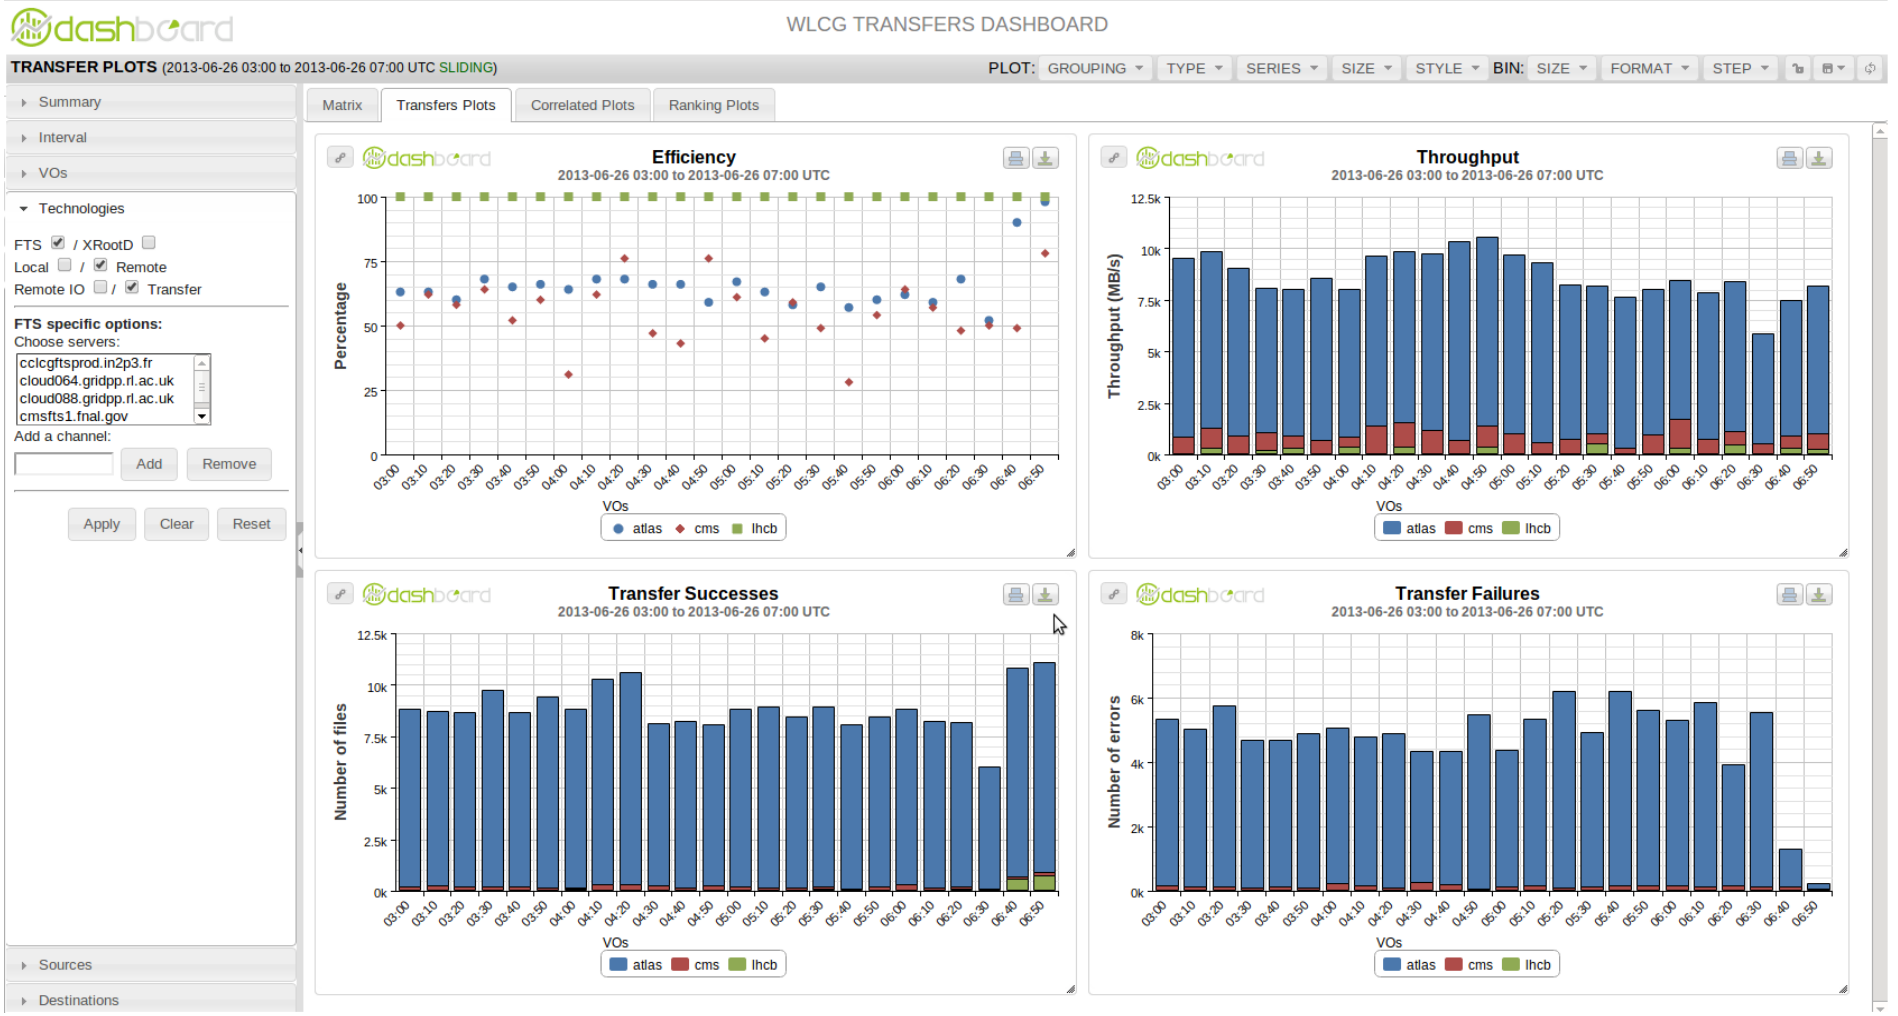

FTS3 global monitoring (2)

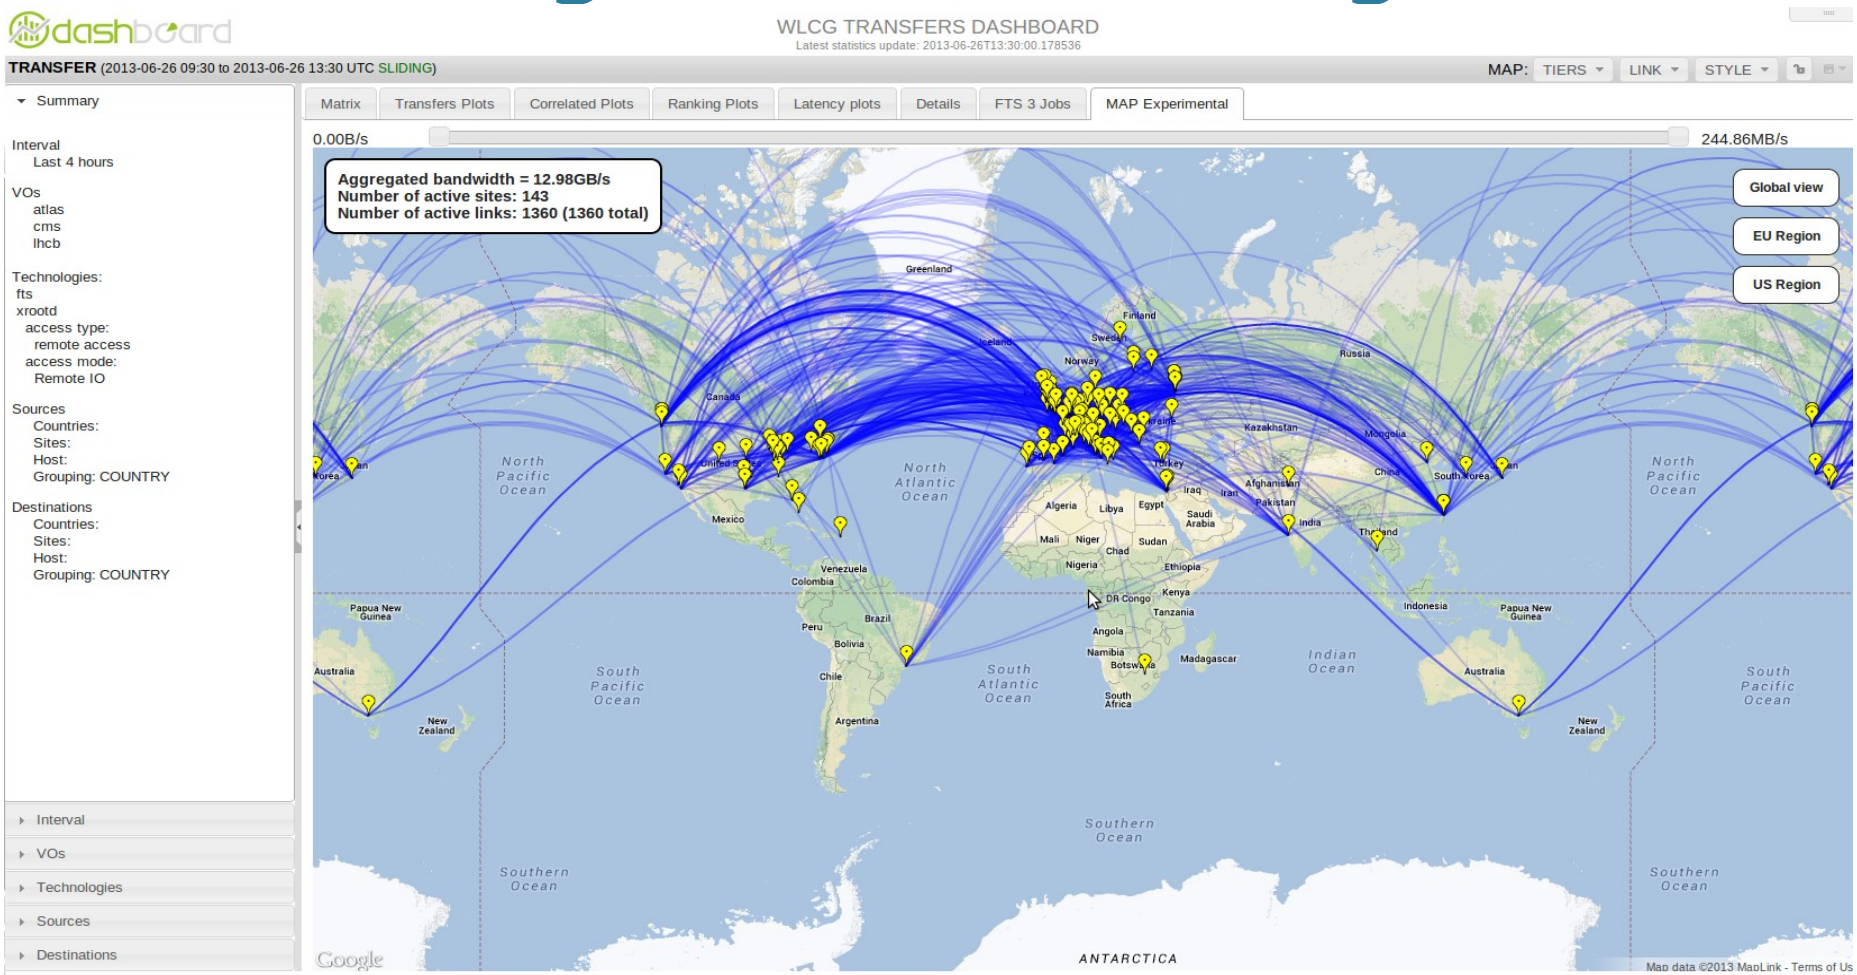

FTS3 standalone monitoring

Pilot FTS3 Monitoring

Search job

Jobs Queue Staging Optimizer Error reasons Statistics Configuration audit Archive

Statistics

Overview | FTS3 Servers | Per pair | Per vo

- **Total transfers queued:** 844 Submitted
- **Total transfers active:** 0 Ready, 30 Active, 0 Staging
- **Total transfers succeeded:** 25561 Finished
- **Total transfers failed:** 3085 Failed, 48 Canceled
- **Success rate in the last hour:** 46.17%

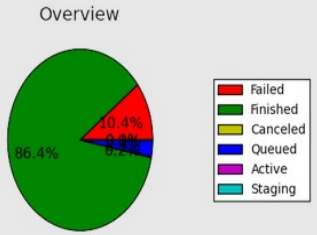

FTS3 standalone monitoring (2)

Pilot FTS3 Monitoring

Search job

Jobs Queue Staging Optimizer Error reasons Statistics Configuration audit Archive

Transfer '1759f299-9a6c-4d80-8122-0b903ec2d308' FINISHED

Submitted by '/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/CN=531497/CN=Robot: ATLAS Data Management'
VO: atlas
Received by fts3-pilot.cern.ch
Overwrite flag: Y
Reuse sessions:

The job contains 5 transfers

FILE ID	TRANSFER HOST	SOURCE URL	
119863	fts4dae3173	srm://srm.grid.sara.nl:8443/srm/managerv2?SFN=/pnfs/grid.sara.nl/data/atlas/atlasdatadisk/rucio/step09/72/b1/step09.20201326010086.physics_A.recon.ESD.closed._lb0004._0001_1372245124	SFN=/t1.grid.kiae.ru/data/atlas/atlas
119864	fts4dae3173	srm://srm.grid.sara.nl:8443/srm/managerv2?SFN=/pnfs/grid.sara.nl/data/atlas/atlasdatadisk/rucio/step09/54/1e/step09.20201326010086.physics_A.recon.ESD.closed._lb0001._0001_1372245124	SFN=/t1.grid.kiae.ru/data/atlas/atlas
119865	fts4dae3173	srm://srm.grid.sara.nl:8443/srm/managerv2?SFN=/pnfs/grid.sara.nl/data/atlas/atlasdatadisk/rucio/step09/81/f5/step09.20201326010086.physics_A.recon.ESD.closed._lb0005._0001_1372245124	SFN=/t1.grid.kiae.ru/data/atlas/atlas
119866	fts4dae316f	srm://srm.grid.sara.nl:8443/srm/managerv2?SFN=/pnfs/grid.sara.nl/data/atlas/atlasdatadisk/rucio/step09/f5/6d/step09.20201326010086.physics_A.recon.ESD.closed._lb0002._0001_1372245124	SFN=/t1.grid.kiae.ru/data/atlas/atlas
119867	fts4dae316f	srm://srm.grid.sara.nl:8443/srm/managerv2?SFN=/pnfs/grid.sara.nl/data/atlas/atlasdatadisk/rucio/step09/94/a9/step09.20201326010086.physics_A.recon.ESD.closed._lb0003._0001_1372245124	SFN=/t1.grid.kiae.ru/data/atlas/atlas

- **Attempts:** 0
- **Duration:** 2.701 seconds
- **Checksum:** ADLER32:5749e419
- **User specified size:** 0.0
- **Configuration:**
- **Parameters:** nostreams:1,timeout:903,buffersize:0
- **Log files:**
  - /var/log/fts3//2013-06-26/srm.grid.sara.nl__sdrm.t1.grid.kiae.ru/2013-06-26-1336__srm.grid.sara.nl__sdrm.t1.grid.kiae.ru__119863__1759f299-9a6c-4d80-8122-0b903ec2d308

FTS3 status

- On its way to EPEL testing, order of a couple of weeks
- Will be in production soon
- WLCG FTS3 task-force actively involved
- Installed at CERN, RAL, PIC, ASGC & BNL as a pilot service
  - Transfer volume last 3 months: ~766T
- Roadmap https://svnweb.cern.ch/trac/fts3/roadmap

FTS3 evaluation

- It is being evaluated by EGI/EUDAT
- Tested against globus GridFTP, dCache GridFTP and GridFTP interface for iRODS (Griffin)
- Many VOs already tested it successfully
  - atlas, lhcb, cms, snoplus.snolab.ca, ams02.cern.ch, vo.paus.pic.es, magic, T2K, NA62, etc
- Evaluate it yourself
  - https://svnweb.cern.ch/trac/fts3/wiki/AdminGuide

FTS3 outstanding features

- Global scheduling and shared VO configuration across distributed FTS3 servers
- Multi-hop transfers
- Bulk file deletions
- Web interface for transfer submission and status retrieval

FTS3 summary

- works out for the box, no configuration needed
  - If needed, endpoint-centric config is supported
- is a light-weight service for heavy-duty job
- **aims to become the new data movement service for the WLCG infrastructure!**

FTS3 resources

- Monitoring
  - http://www-ftsmon.gridpp.rl.ac.uk/fts3/ftsmon/jobs
  - http://vt-092.grid.sinica.edu.tw/fts3/ftsmon/
  - http://fts3.pic.es/fts3/ftsmon/jobs
  - http://dashb-wlcg-transfers.cern.ch/ui/
- Wiki
  - https://svnweb.cern.ch/trac/fts3/wiki
  - fts3-steering@cern.ch
  - Subscribe - http://cern.ch/go/99Gg

Thank you!
Questions?