



## FTS3 at FNAL

Lorena Lobato Pardavila  
XRootD and FTS workshop  
March 27, 2023

# Outline

- Introduction
- Configurations
  - CMS
  - Public
    - CI Logon Tokens progress
- Main differences
- Advantages & disadvantages of each configuration
- Summary

# Introduction

Fermilab (FNAL) has two configurations of FTS3 currently running.

- CMS - Server and Database running on the same host (physical server)
  - Service Manager -> Chih-Hao Huang
- Public experiments (mostly DUNE) – Server containerized deployed on OKD + database running on a VM
  - Service Manager -> Myself
- Different purposes – different configurations

What?	Used by?
<ul style="list-style-type: none"><li>+ FTS3 CERN (no F-FTS)</li><li>+ File Transfer System that efficiently schedule data transfers.</li><li>+ Composed by a Frontend server and a Database backed</li></ul>	<ul style="list-style-type: none"><li>+ CMS</li><li>+ DUNE uses for small purposes</li></ul>

# CMS

# Configuration: CMS


- A single physical server (**cmsftssrv1.fnal.gov**) is running everything including FTS services and the FTS database
- Infrastructure is under puppet control
- Taking everything out of the box with minimal local configuration
- Decision on choosing hardware: bang for the buck with more **memory (most important)**
- Another server with an identical configuration as “standby” (**cmsftssrv3.fnal.gov**). A floating IP/alias (**cmsfts3.fnal.gov**) is in front of both servers -> Helps with disaster recovery
- This switching method works very well with PhEDEx and ok with Rucio
  - PhEDEx, with agents running on the Ses, does not need FTS to complete ongoing transfers. The unfinished transfers before the switch will continue to their completion.
  - Rucio relies on FTS to check transfer states, yet unfinished transfers will fail, time out, and retry.
- Upgrade: upgrade the standby first, then switch.

# Configuration: CMS

- Usage: among 99 registered storage elements in Rucio, 28 are using CMS FTS3 FNAL instance whereas 25 are using it as the primary FTS service
- All T3\_US\_\* are using CMS FTS3 FNAL instance as the primary FTS service
- Most T\*\_US\_\* are using CMS FTS3 FNAL instance as the primary FTS service except:
  - T2\_US\_Caltech, T2\_US\_UCSD, T2\_US\_Vanderbilt – using it as secondary
  - T2\_US\_MIT and T2\_US\_MIT\_Tape – don't use cmsfts3.fnal.gov at all
- Three sites outside the US are using CMS FTS3 FNAL as the primary FTS service:
  - T2\_BR\_SPRACE, T2\_BR\_UERJ and T3\_MX\_Cinvestav

# Configuration: CMS


- Total Volume Transferred in the last Year



SOURCE: [https://monit-grafana.cern.ch/d/CljJHKdGk/fts-transfers?from=now-1y&orgId=20&to=now&var-bin=1d&var-dst\\_country=All&var-dst\\_site=All&var-fts\\_server=cmsfts3.fnal.gov&var-group\\_by=activity&var-include=&var-protocol=All&var-src\\_country=All&var-src\\_site=All&var-staging=All&var-vo=cms&viewPanel=14](https://monit-grafana.cern.ch/d/CljJHKdGk/fts-transfers?from=now-1y&orgId=20&to=now&var-bin=1d&var-dst_country=All&var-dst_site=All&var-fts_server=cmsfts3.fnal.gov&var-group_by=activity&var-include=&var-protocol=All&var-src_country=All&var-src_site=All&var-staging=All&var-vo=cms&viewPanel=14)

# Configuration: CMS

- Volumen Transferred / Number of Transfers in the last Year



SOURCE: [https://monit-grafana.cern.ch/d/CljJHKdGk/fts-transfers?from=now-1y&orgId=20&to=now&var-bin=1d&var-dst\\_country=All&var-dst\\_site=All&var-fts\\_server=cmsfts3.fnal.gov&var-group\\_by=activity&var-include=&var-protocol=All&var-src\\_country=All&var-src\\_site=All&var-staging=All&var-vo=cms&viewPanel=13](https://monit-grafana.cern.ch/d/CljJHKdGk/fts-transfers?from=now-1y&orgId=20&to=now&var-bin=1d&var-dst_country=All&var-dst_site=All&var-fts_server=cmsfts3.fnal.gov&var-group_by=activity&var-include=&var-protocol=All&var-src_country=All&var-src_site=All&var-staging=All&var-vo=cms&viewPanel=13)


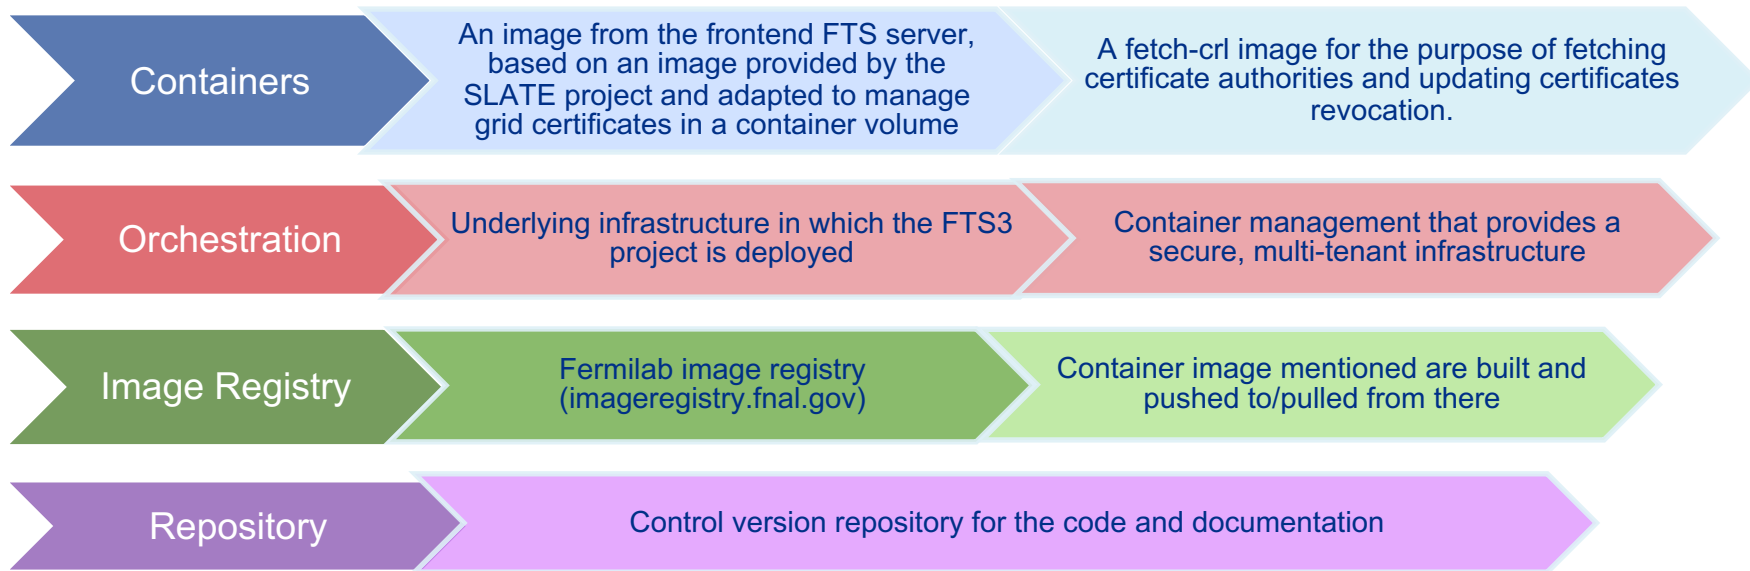


# Public

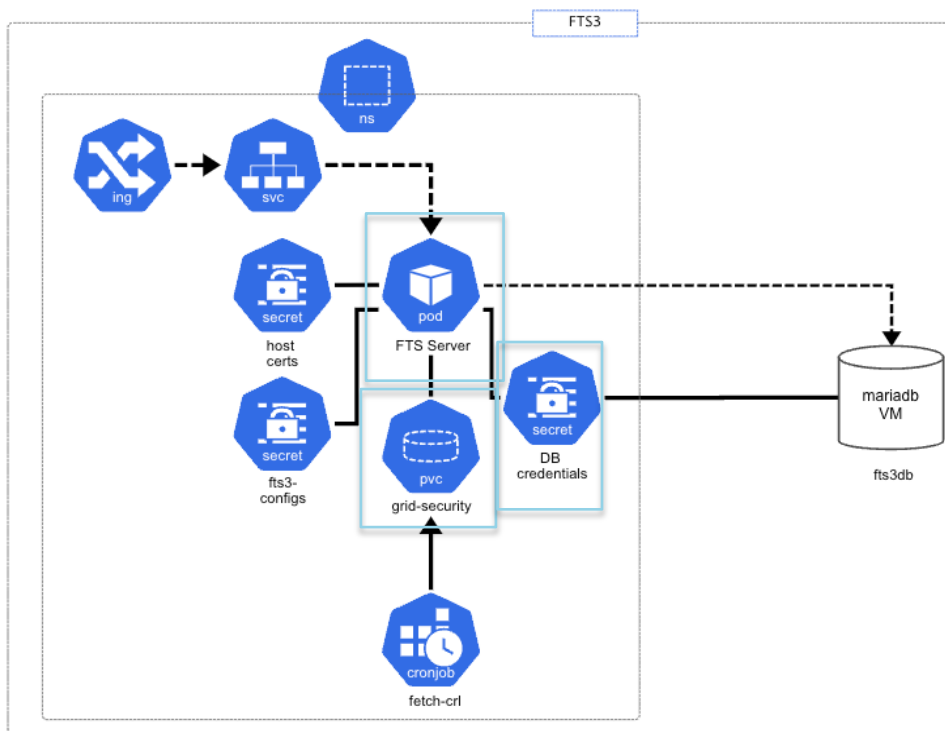
# Configuration: Public

- FTS server containers including the FTS services + K8 cronjobs running (CRLs, certs & CAs) – and deployed in OKD- and the FTS database running on a VM.
- Lower HW requirements
- Set up an instance of FTS3 at Fermilab for experiments other than CMS to use (American transfers)
- Dune uses for small purposes (mostly for NERSC and BNL) and we expect to increase the demand in the coming years.
- RUCIO uses FTS3 as transportation layer

# Configuration: Public



# Configuration: Public



# Public: CILogon Tokens progress

- CILogon was not accepted as Tokens provider
- Working in the past months with the FTS3 team, CILogon developers, and FNAL tokens experts to make this happens
- Our initial tests are working and we created/approved clients on the test machines thar work with FTS3 workflow.
- Waiting for CERN FTS3 to deploy it in production

# Main Differences

CMS(cmsftssrv1.fnal.gov)	Public(fts3-public.fnal.gov)
<p>Running on a physical server</p> <ul style="list-style-type: none"><li>• FTS 3.12.4</li><li>• 256 GB - should be able to support 15,000 fts_url_copy</li><li>• fts3db in MariaDB</li></ul>	<p>Server containerized with a mariadb running on a VM</p> <ul style="list-style-type: none"><li>• FTS 3.12.5</li><li>• 1 server running on a shared cluster (10G)</li><li>• fts3db in MariaDB (8 cores, 100G, 16 GB)</li></ul>
Demand is higher	Limited used
CERN MONIT centralized	FNAL monitoring
Recommended configuration	Challenging configuration

# Advantages & disadvantages

## CMS(cmsftssrv1.fnal.gov)

Having the database and the server running on the same node – Helps failure issue

## Public(fts3-public.fnal.gov)

Horizontal Scaling according to demand

Managing the FTS processes with  
supervisord

When the FTS service, for any reason, stops working, it is not very obvious since the monitoring front end (port 8449) is served by httpd

The most common reason that prevents FTS service from working correctly is the log area being filled

The main problem (currently) with the FTS service is the soci library issue

Configuration can be more challenging

# Summary

- FTS3 is a low-maintenance service if you follow the recommended configuration 😊
- For FNAL Public, a lot of work ahead. We hope to expand the idea of having a public FTS3 instance containerized at least for transfers involving Americas
- OKD instance that supports multiple tenants enabling sharing of the cluster + resilience to hardware or application faults
- FTS server is vulnerable still toward the DB – hopefully devs can work on the DB connections
- Have to learn everything the hard way, bothering FTS devs quite often 😊



# Thank you!

# Backup Slides

# Certificate Management and Maintenance - Kubernetes CronJobs

- FTS3 developers assume Certificate Authorities, certificate revocation lists (CRLs), and certificates (certs) to be in a common area
- We have installed OSG repository
  - CAs provide the trust roots for the OSG public
  - Help to have CRLs up to the date on the hosts
  - `osg-ca-manager`

# FTS3 testing instance running at Fermilab

