



FTS3: Data Movement Service in containers deployed in OKD

Lorena Lobato Pardavila
CHEP 2021 online Workshop
20 May 2021

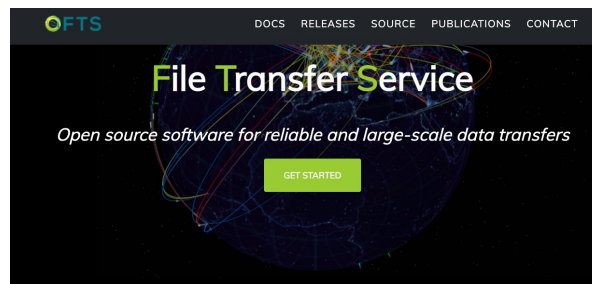
What am I going to present today?

- FTS3
- Architecture
- Certificate Management and Maintenance – Kubernetes Cronjobs
- FTS3 instances running at Fermilab
- Summary

FTS3

<https://fts.web.cern.ch/fts/>
<https://slateci.io/>

- File Transfer System that efficiently schedule data transfers
- Composed by a Frontend and a backend (database)
- Integrated with experiment frameworks: Rucio, PhEDEx and DIRAC



WHY CHOOSE US?

Developed at CERN, FTS distributes the majority of the Large Hadron Collider's data across the Worldwide LHC Computing Grid (WLCG) infrastructure.

Simplicity
Easy user interfaces for submitting transfers: Python, CLI, Python Client, WebFTS and Web Monitoring.

Reliability
Checksums and retries are provided per transfer.

Flexibility
Multiprotocol support (Webdav/https, GridFTP, xroot, SRM).

Intelligence
Parallel transfers optimization to get the most from network without burning the storages. Priorities/Activities support for transfers classification.

Multiprotocol support (WebDAV/HTTPS, GridFTP, XRootD and SRM)

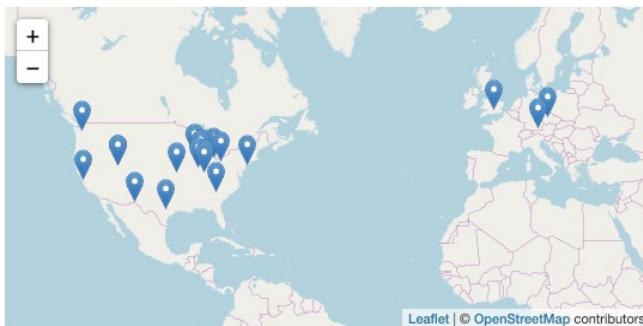
Clients can access the service (REST APIs, python bindings, CLI)

Transfers from and to different Storage Systems as EOS, DPM, Object storages such as S3, dCache, CASTOR and CTA

Support for tapes (bringonline)



Powering science collaborations with flexible and secure service deployment



Federated Operation of Science Platforms

Modern research requires collaboration across facilities, institutions and scientific domains. SLATE helps build multi-institution cyberinfrastructure using secure and declarative deployment tools.

SLATE enables a federated "NoOPs" operations model that gives cyberinfrastructure developers the flexibility to innovate at scale, expanding the reach of domain specific science gateways and multi-site research platforms.

Containerized Services



XCACHE



GLOBUS CONNECT



perfSONAR



HTCONDOR



GridFTP

Architecture




- A standard **MariaDB image for the backend database** (provided by OKD and used without modification)
- An image from the **frontend FTS3 server**, based on an image provided by the SLATE project and adapted to manage grid certificates in a container volume
- A **fetch-crl image**, for the purpose of fetching certificate authorities and updating certificate revocation




- Fermilab image registry (**imageregistry.fnal.gov**)
- Docker images mentioned are built and pushed to/pulled from there
- Backed by **Gluster storage for persistent volumes**



- Underlying infrastructure in which the FTS3 project is deployed
- Fermilab deploys **OKD**, the community distribution of OpenShift
- **Container management** that provides a secure, multi-tenant infrastructure
- OKD versus Kubernetes (DeploymentConfigs, BuildConfigs, ImageStreams)



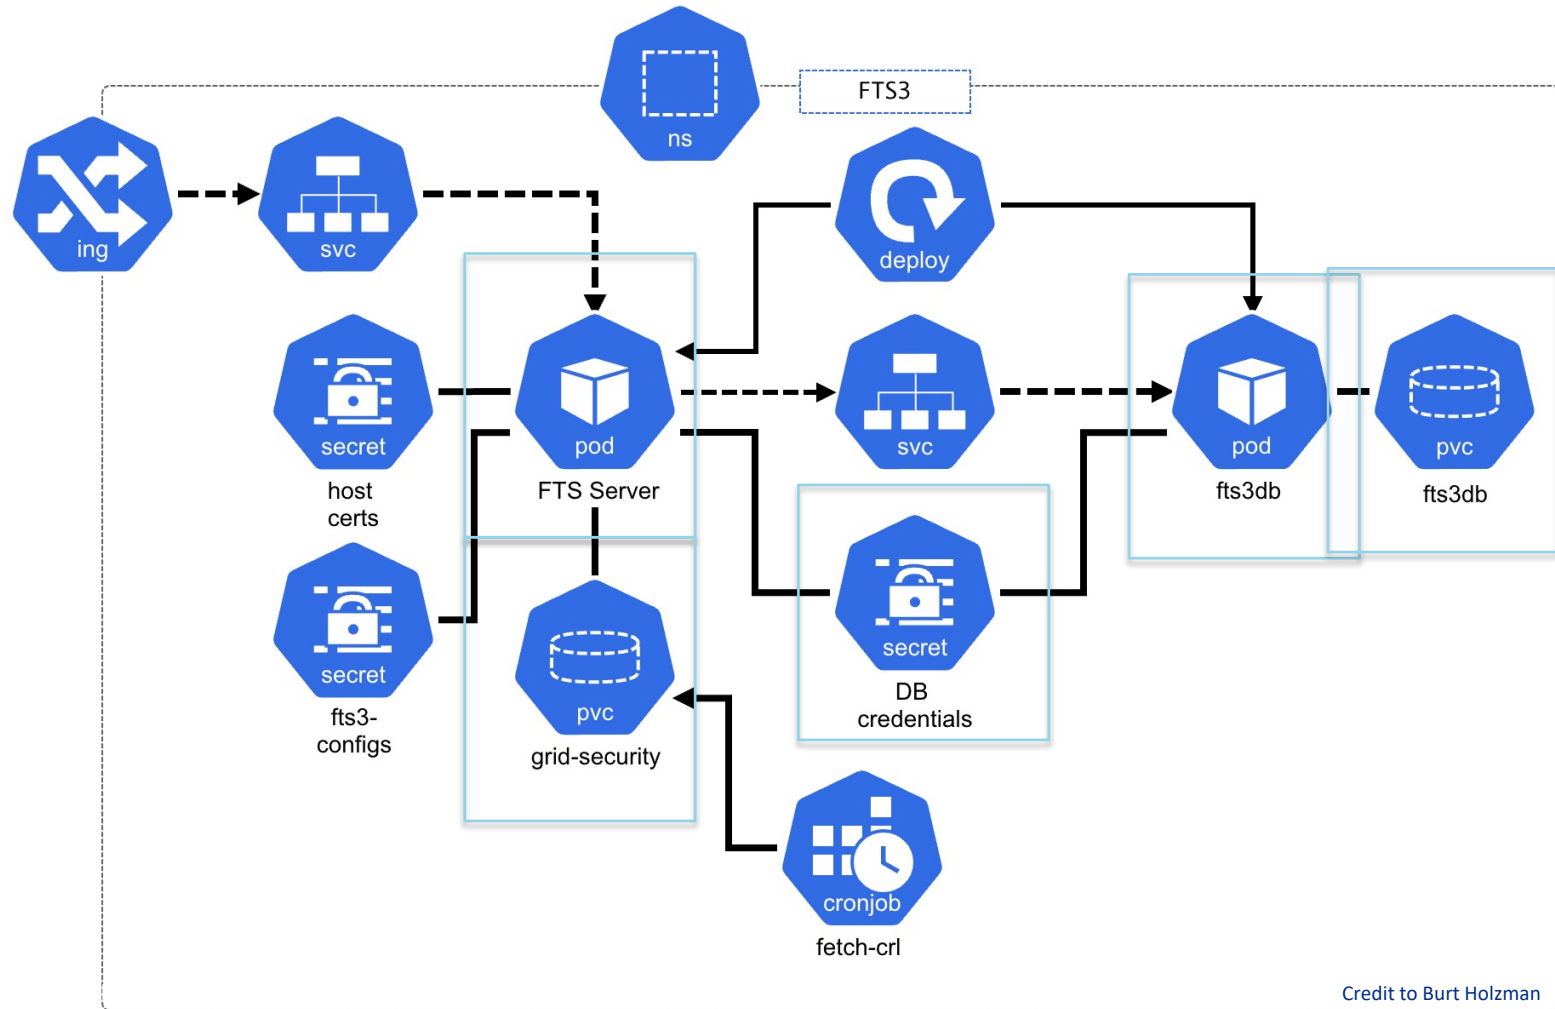
Control version repository for the code and documentation



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 public instances running at Fermilab



Summary

- ✓ Have a public FTS3 instance at least for transfers involving Americas
 - ❑ Containerized deployment -> Horizontal Scaling according to demand
- ✓ OKD instance that supports multiple tenants enabling sharing of the cluster
- ✓ Resilience to hardware or application faults
- ✓ We keep working..
 - ❑ DUNE: RUCIO tests (petabytes of data in the next few months) + monitoring
 - ❑ Production coming soon + JWT tokens provided by 2 different Entity Providers

Questions?
llobato@fnal.gov