



FTS3: Data Movement Service in containers deployed in OKD

Lorena Lobato Pardavila

HEPiX Spring 2021 online Workshop

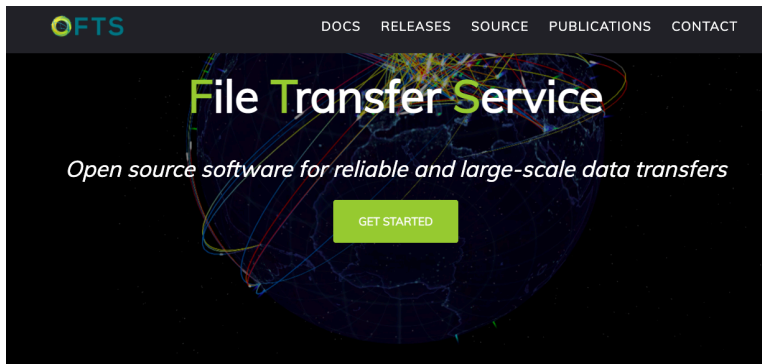
18 March 2021

What am I going to present today?

- FTS3: Data movement Service
- Architecture
- Certificate Management and Maintenance – Kubernetes Cronjobs
- FTS3 instances running at Fermilab
- Future plans
- Summary

FTS3

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



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)

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.

What is SLATE?



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

Technologies used



Containers

Provides a standard way to package application code, configurations and dependencies into a single object



Orchestration

Containers orchestration that defines tools that automate operations of container-based applications



Image Registry

Container image registry that secures images with role-based access control



Repository

DevOps platform that stores the codebase of the project and the documentation related to operations



Three Docker images, of which two are customized for the Fermilab installation:

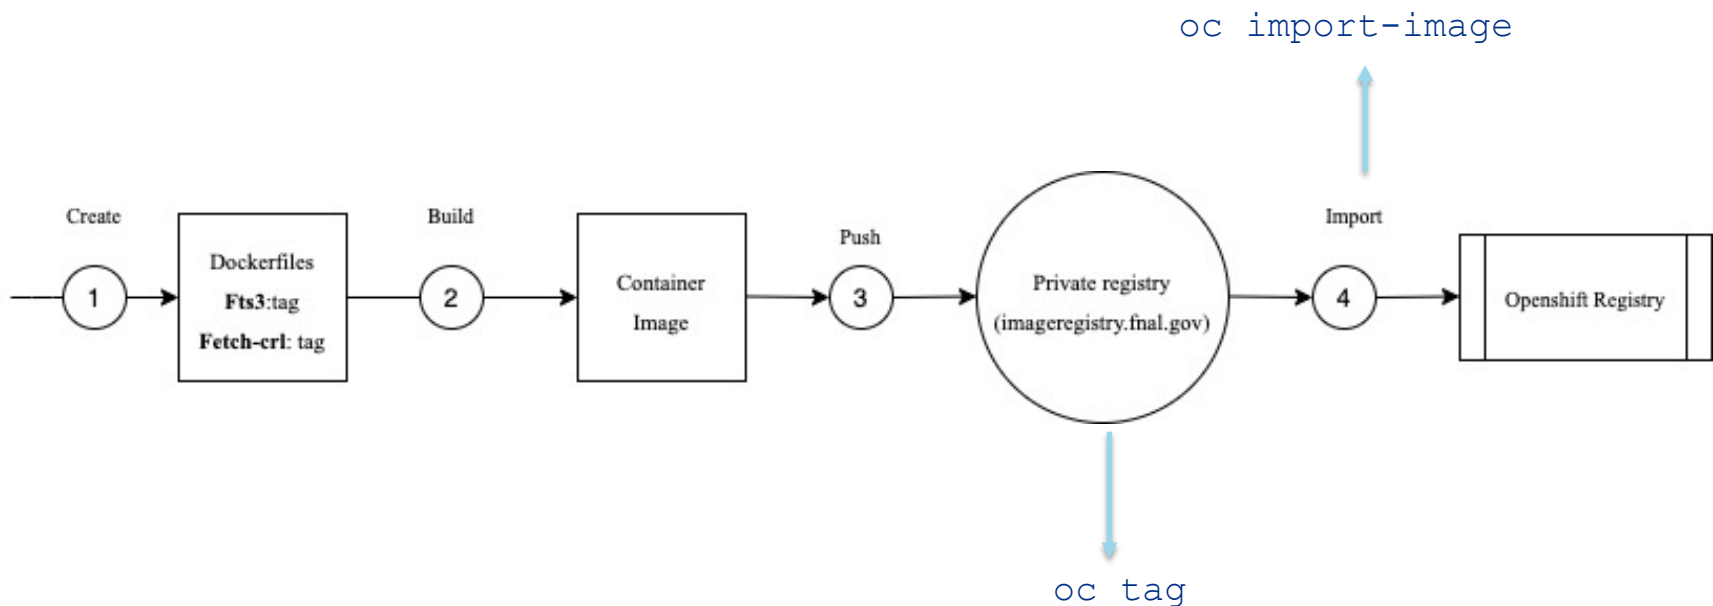
- A standard MariaDB image for the backend database (provided by OKD and used without modification)
- An image from the frontend FTS3 server, based on an image provided by the SLATE project and adapted to manage grid certificates in a container volume
- A fetch-crl image, for the purpose of fetching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

- Underlying infrastructure in which the FTS3 project is deployed.
- Fermilab deploys OKD, the community distribution of K8
- Container management that provides a secure, multi-tenant infrastructure
- OKD versus Kubernetes
 - DeploymentConfigs
 - BuildConfigs
 - ImageStreams





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



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`

Future Plans

- JWT tokens support (supported by the latest FTS3 version (3.10). Provided by 2 different Entity Providers (CILogon and WLCG IAM)
- RUCIO (petabytes of data in the next few months)
- DUNE tests

Take away

- ✓ Have a public FTS3 instance at least for transfers involving Americas. DUNE + expected more experiments
- ✓ Containerized deployment -> Horizontal Scaling according to demand
- ✓ Runs on an OKD instance that supports multiple tenants enabling sharing of the cluster
- ✓ Resilience to hardware or application faults
- ✓ RUCIO tests
- ✓ Production coming soon

Questions?