FTS

File Transfer Service

Following the K8s trail

pre-GDB on Kubernetes - 7th of June 2022

Mihai Patrascoiu
on behalf of the FTS team

Preface

- Current FTS@CERN deployment is done via OpenStack VMs
- Every FTS node deploys all services (Transfer, Submission and Web Monitoring)
- Operation procedures are being run via Puppet and `wassh`
- Service monitoring is done:
  - At the host level (collectd and GNI alarms – Puppet provided)
  - FTS node level (*profiling logs* + Grafana dashboards)
  - FTS instance level (e.g.: FTS3-Atlas) (instance probes + transfers dashboards)
- No intention to change deployment setup in the mid-term

The K8s trail

1. Image building via CI
2. Master FTS image
3. CI testing based on master image
4. Reduce coupling of components
5. Test de-coupled deployment
6. K8s deployment

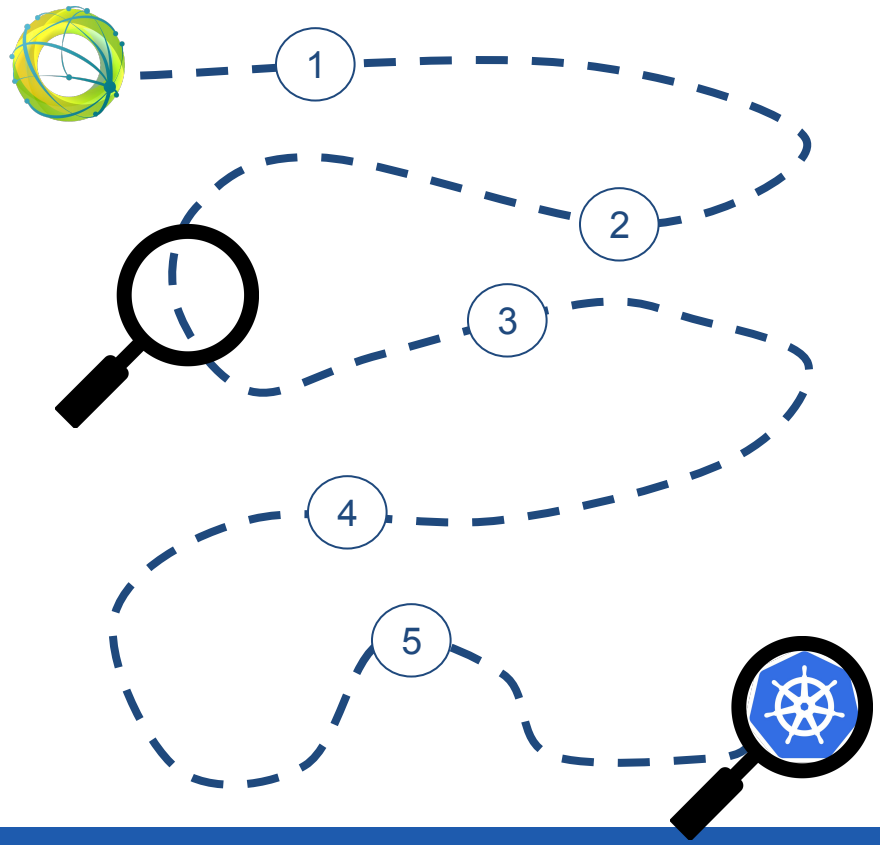

The K8s trail (detailed)

1. Fully-automatized building of Docker images via the CI for each FTS component (Transfer, Submission, Web Monitoring)
2. Master FTS image containing all services
3. Move CI testing to use FTS master image
4. Reduce *hidden* coupling between the different FTS components
5. Test a granular system where the FTS components don't reside on the same node
6. K8s deployment

Medium term

Long term

1. Image building via CI

- Build Docker images for individualized components
  - On commit: latest (only from `develop` branch)
  - On tags (e.g.: 3.11.2)
- Already available for Transfer and Web Monitoring components
- Work needed for the upcoming Python3 Submission server

```
gitlab-registry.cern.ch/fts/fts3:3.11.2
gitlab-registry.cern.ch/fts/fts-monitoring:3.11.0
gitlab-registry.cern.ch/fts/fts-rest-flask:3.12.0 (upcoming)
```

2. Master FTS image

- Build Docker image containing all FTS components
  - A stand-alone deployment of this image → deployment of FTS (sans database)
- One image available for minor release series
  - Example: v3.12 image will contain latest v3.12.x Transfer, Submission and Web Monitoring component
- Build image via CI, rebuild when each sub-component pushes a new tag
- Similar image created by SLATE

3. CI Testing based on master image

- Main motivation: deploy and test an FTS instance in the CI
  - Stand-alone instance (running the FTS probes and functional tests)
  - Together with another system (main use-case is with CTA)
- Similar work is already done by the Rucio team (Radu Carpa)

4. Reduce coupling of components

- Long-term, FTS wants to reduce the *hidden* coupling of the various FTS components
- Main motivation:
  - The coupling has proved painful across service operations,
    showing up in unexpected places→ strong reason to remove it
  - No-coupling allows horizontal scaling of individual components
  - Allows split of functionality. E.g.: fts3-atlas and fts3-atlas-monitoring
- This is a long-term direction rather than a fixed goal. The FTS service
  will evolve organically in this direction (timeline of 1y+)
  - The Python3 port of the Submission server is a big step in this direction

6. Kubernetes deployment

- Currently, not an objective for the FTS team
- Once the software removes the coupling of components, interesting tests can be done with horizontally scaling individual components based on load
- Large architectural changes need to be done to allow FTS to fully leverage K8s deployment (e.g.: the queueing system relies heavily on the DB, which K8s deployment cannot help)
- Good motivation for containerized CI testing to better mimic production
    - Docker images might turn out to be enough for this

Recap

- FTS@CERN operations will remain VM-based in the medium term
- The FTS team wants to provide a master image in the medium term
  - This may be used for both deployment or CI testing
  - Will immediately benefit OSG, hoping other communities as well
- FTS Development is heading towards better containerized CI testing
- FTS makes a good candidate for K8s deployment
  - …once more work is done on the scheduler part to truly benefit from K8s auto-scaling
- Due to effort constraints, the FTS team won't be able to explore FTS w/ K8s in the mid-term

FTS Kubernetes pioneers

- Lorena Lobato Pardavila (Fermi National Accelerator Laboratory) *et al.*
  - FTS3: Data Movement Service in containers deployed in OKD, CHEP 2021 (link)
  - EPJ Web Conf., 251 (2021) 02058 (link)
- Radu Carpa, Rucio team
  - Rucio testing (including FTS deployment) via Kubernetes
- The SLATE Project (link)
  - Provides an FTS *master* image