Yandex

Skygrid at Yandex

Alexander Baranov (YSDA)

Outline

- Skygrid system and its architecture
- Usage in LHCb
- Yandex Cloud resources and LHCb general overview

Skygrid basic premise

- Lets you compute **result = f(input)**
- **f** — computation function, packed in docker container
- **input, result** — datasets located on arbitrary backend
- Backends: XRootD, local fs, webdav, git, …
- Stores computations status and results locations
- Interface: RESTful HTTP API

Skygrid

- System to execute computations in docker containers
- Closest analogoues: Kubernetes, Docker Swarm, HTCondor
- Covers HEP-specific needs for data management
- Is extendable for other workloads
- Developed by YSDA and is opensource

github.com/skygrid

Skygrid architecture overview

Users

Metascheduler

Worker nodes

Job descriptors

Users

1. Generate

Job JSON descriptor

Job descriptors

Contains basic information about job:

- What to take as input? (files URI)
- What to execute? (docker image, CLI arguments)
- Where to store output? (backend URI)

URI’s can reference different backends: XRootD, local fs, git, webdav, etc.

y = f(x)
output = docker(input)

Skygrid Architecture

Users

1. Generate

Job JSON descriptor

2. Submit

Metascheduler

Metascheduler

- Queues and Jobs management service
- Queues: CRUD + push into + pull from
- Jobs: CRUD, mainly status/output updates
- Has RESTful HTTP API
- Uses MongoDB to store and index jobs

CRUD = Create, Read, Update, Delete

Skygrid Architecture

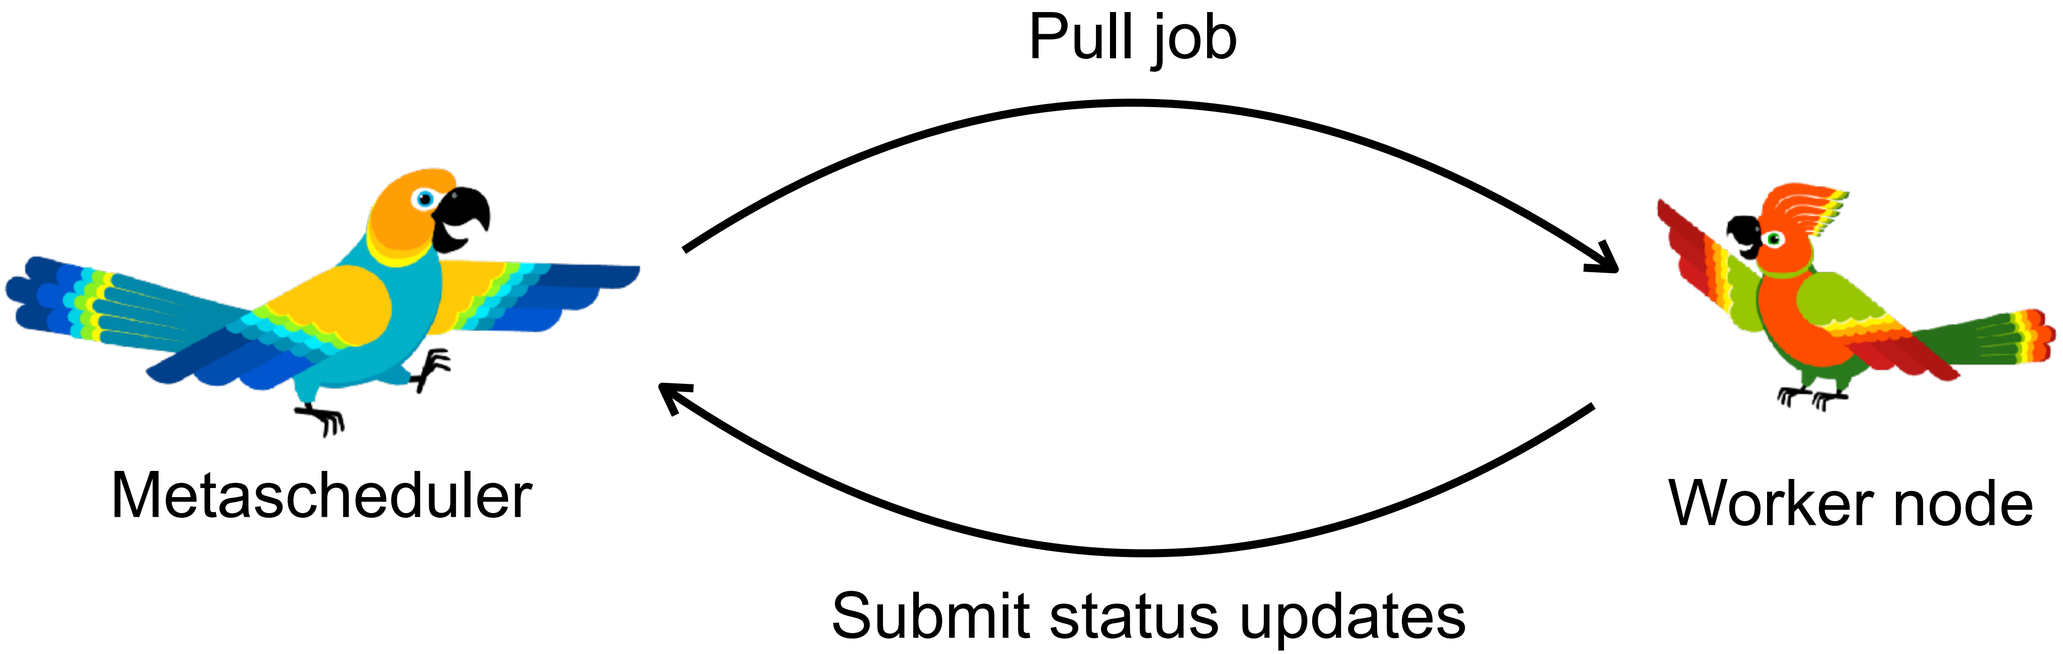

Workers

Execution algorithm:

1. Pull job descriptors from metascheduler
2. Download input data
3. Mount input into docker container, execute
4. Wait container to finish
5. Upload output data

〉(Update metascheduler info on each stage)

Workers Nodes

Only minimal software needed:

- Docker
- Skygrid worker
- Backend-dependent software(i.e. xrootd)

Skygrid architecture overview

Users

Push job

Pull updates

Metascheduler

Pull job

Push updates

Worker nodes

Skygrid for LHCb

- Yandex Skygrid cluster is opportunistically available for LHCb (1600 cores)
- Successful integration of LHCb DIRAC and Skygrid
- Paired SL6 docker image with Vcycle vm creation scripts
- Manually launched pilots on Yandex site
- Successful execution for couple month
- Temporary paused due to other activities on cluster

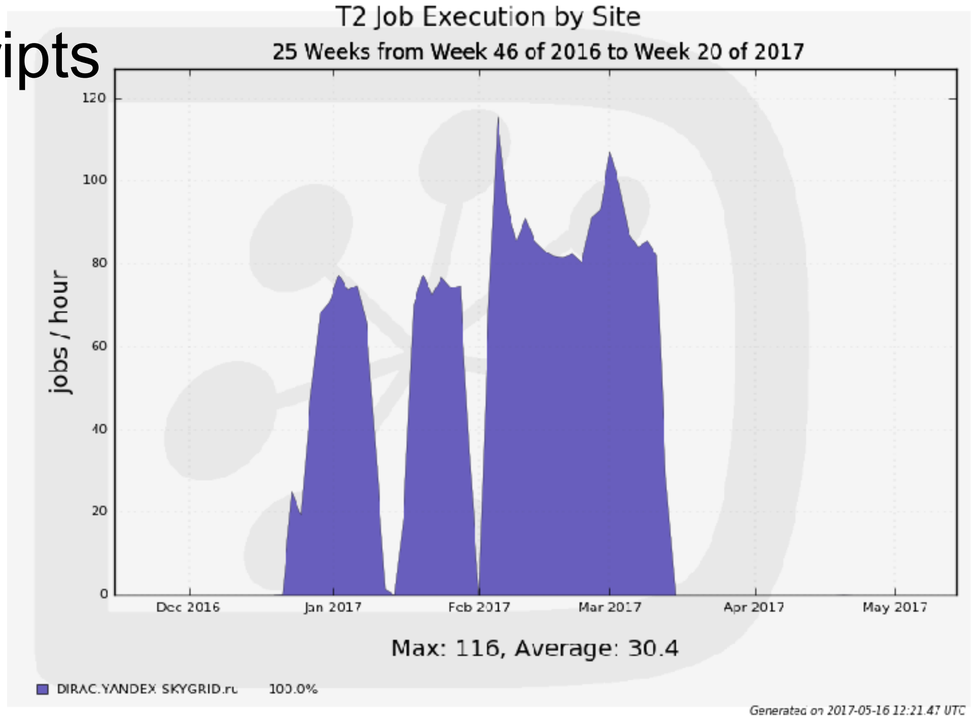

Yandex Cloud and LHCb

- Yandex Cloud — cloud.yandex.com
- Gradually becoming really public
- Already available to LHCb and CERN

- Migrating current T2 site to Yandex Cloud, usage with Vcycle
- Working PoC, now scaling resources
- Direct Yandex-CERN 10G network peering
- Containers support in Yandex Cloud in the future

Conclusion

Skygrid

- Container computations + Data backends
- Easy integration with existing technologies
- Reproducibility, automatic analysis preservation

Yandex and CERN

- Hybrid Y.Cloud+Skygrid usage — migration almost finished
- Direct 10G connectivity
- Y.Cloud considers servicing scientific community at CERN as an important use case

Questions

github.com/skygrid