Michael Schuh, DESY.
April 3 2019
ISGC, Taipei

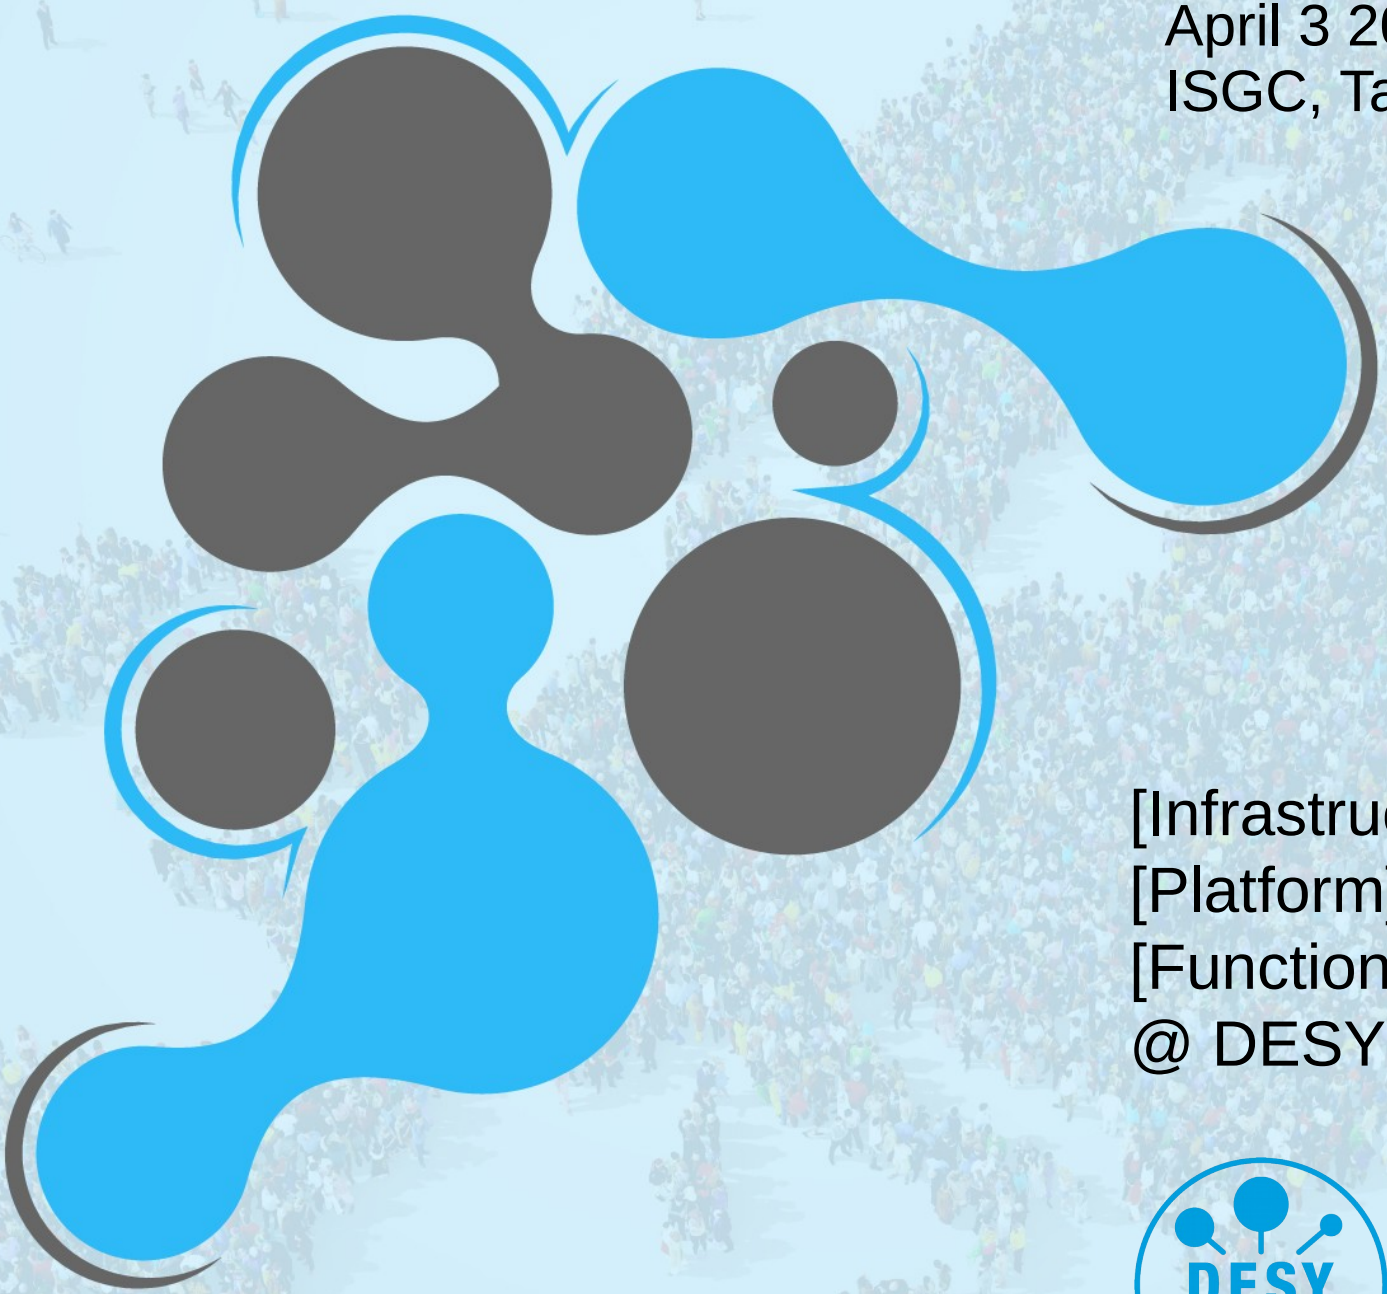

[Infrastructure]
[Platform]
[Function]-as-a-Service
@ DESY OpenStack Cloud

DESY.

EOSCpilot
The European Open Science Cloud for Research Pilot Project
www.eoscpilot.eu

Outline

- European XFEL – use case @ DESY
  - Femto-second X-ray Crystallography
- Scientific cloud computing
  - EOSC
  - Auto-scaling GitOps
- Function-as-a-Service
  - in Jupyter Notebooks
  - and event-driven automation with dCache

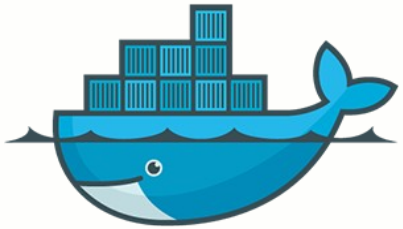

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

SECTION 1

- EU-XFEL – use case @ DESY
  - Femto-second X-ray Crystallography

DESY

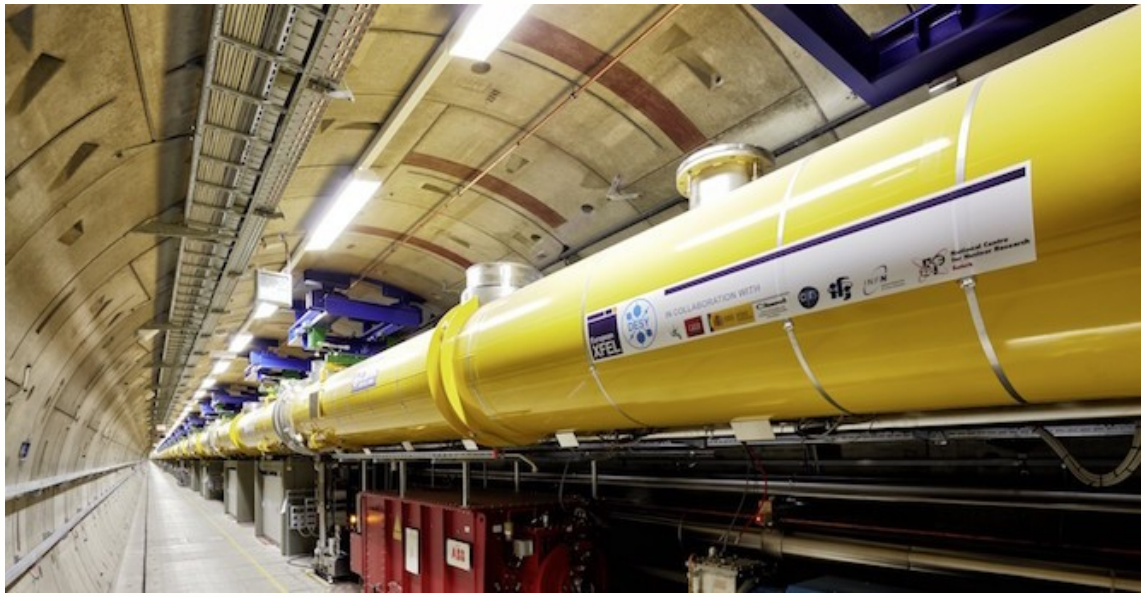

DESY Deutsches **E**lektonen-**SY**nchrotron (German Electron-Synchrotron)

- Physics with Photons, Free Electron Lasers
- Accelerator technologies
- Experimental particle physics
- Astroparticle physics

Images: http://www.desy.de/femto_eng/index_eng.html

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

DESY IT

- Scientific Computing, HPC, HTC, Grid and Cloud
- LHC Tier-2 center, large scale storage and archiving
- dCache – a peta scale storage platform
- Computing and storage provider for European XFEL

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

European XFEL

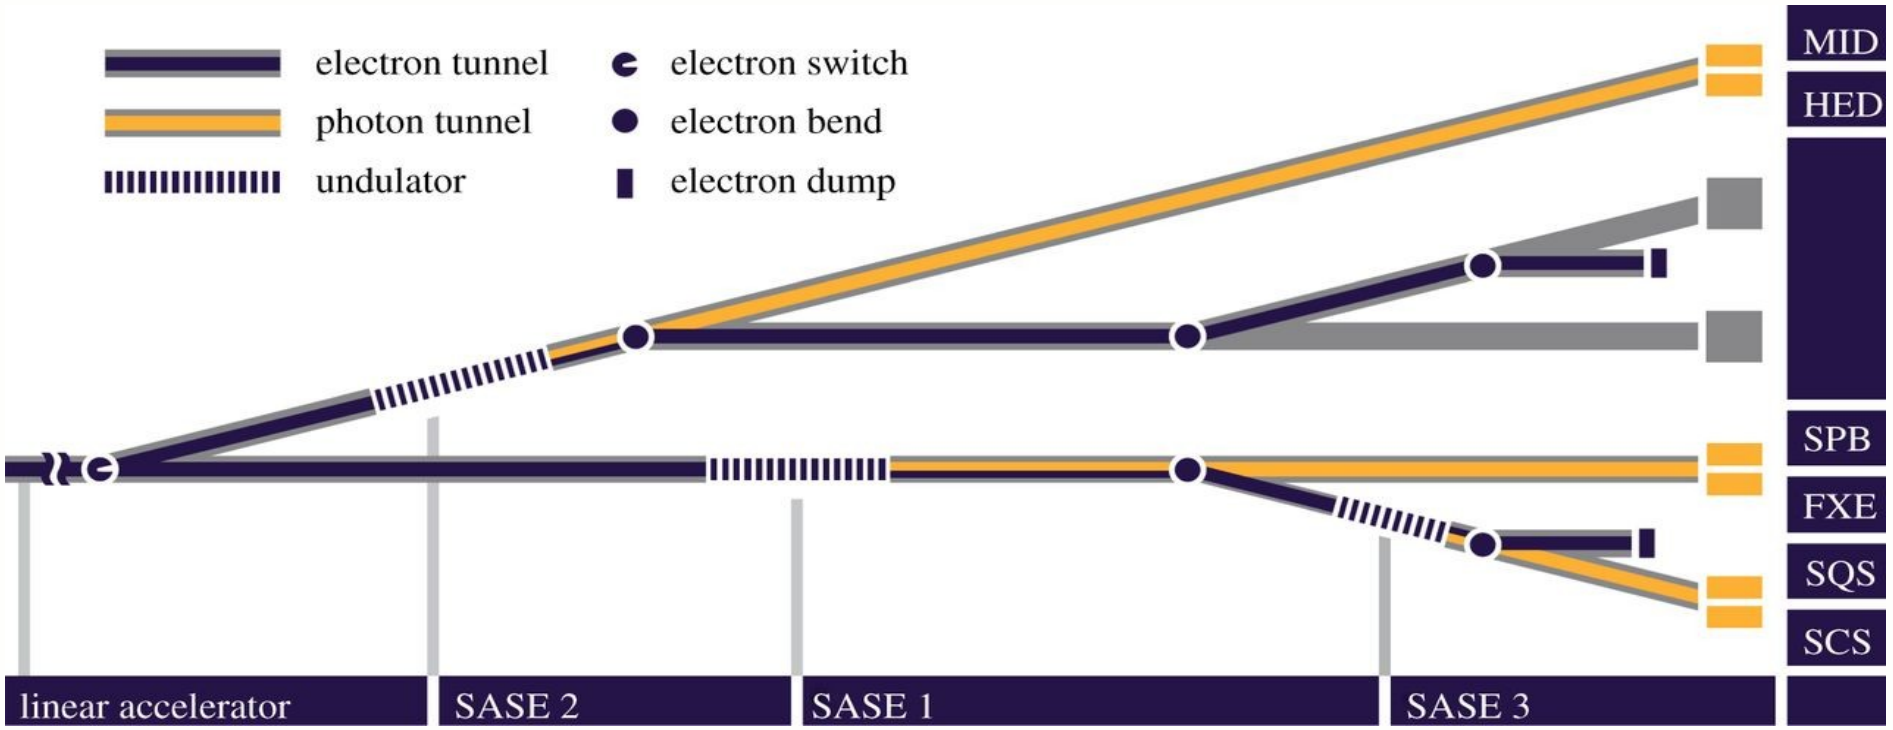

Self-
Amplified
Spontaneous
Emission

Most brilliant light source ever built
Extremely short laser pulses

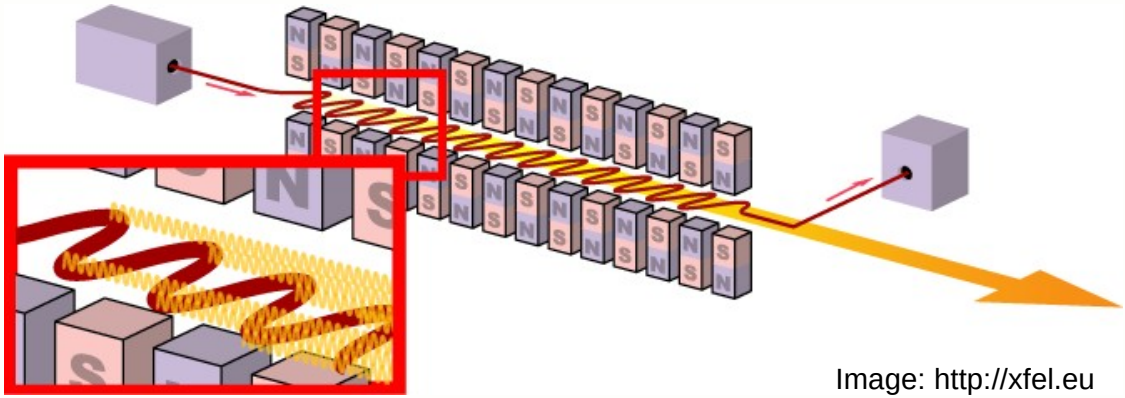

Image: http://xfel.eu

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Experiments at the European XFEL

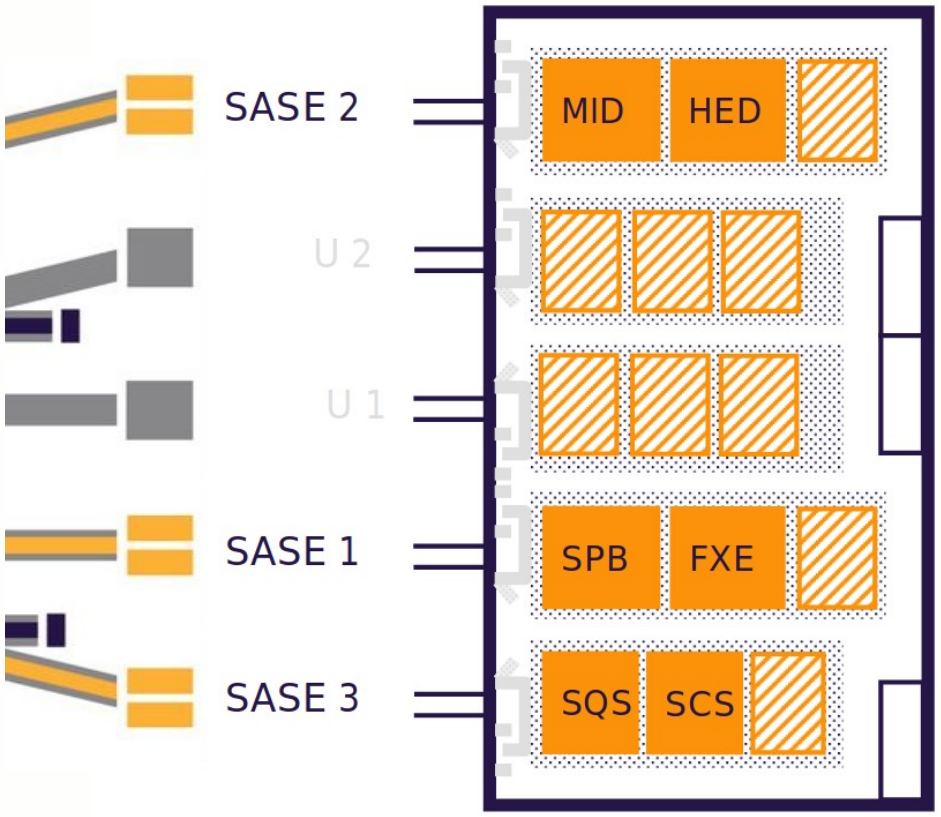

MID Materials Imaging & Dynamics

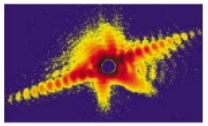

HED High Energy Density Science

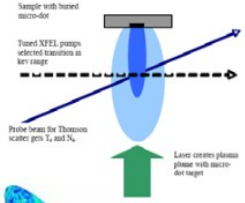

SPB Single Particle & Biomolecules

FXE Femtosecond X-ray Experiments

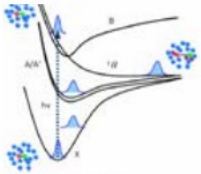

SQS Small Quantum Systems

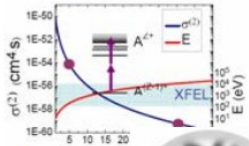

SCS Spectroscopy & Coherent Scattering

More about experiments: http://www.xfel.eu

European XFEL

Experiments in physics, chemistry, materials science, biology, nanotechnology

Source: http://xfel.eu

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Exposure time: femto-seconds (10^{-15} s)

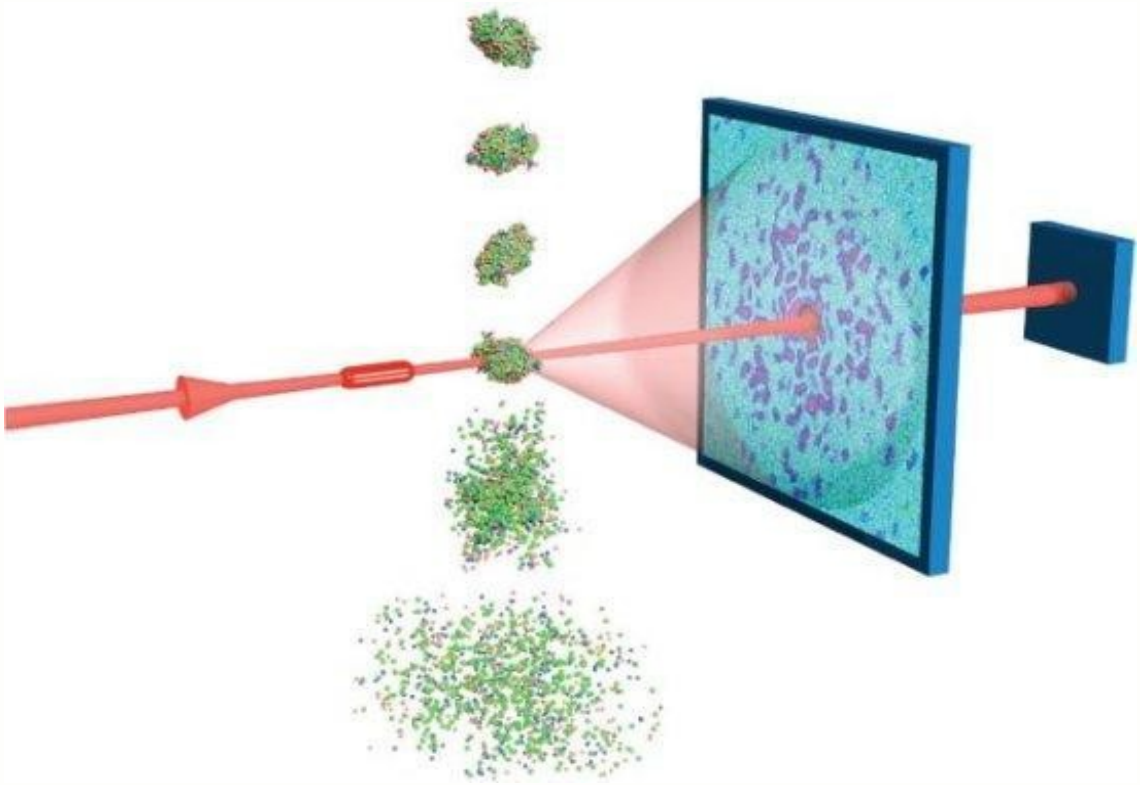

- Molecular 3D imaging at the atomic scale
- Diffraction before destruction
- Unprecedented resolution in both: time and space
- Dynamics of inner structures (not just surface) of biological objects
  - Virus
  - Cell nucleus
  - ...
- To reconstruct a 3D image, a large amount of random oriented single images have to be combined

source: https://cid.cfel.de/research/femtosecond_crystallography/

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Data rates for the FXE instrument

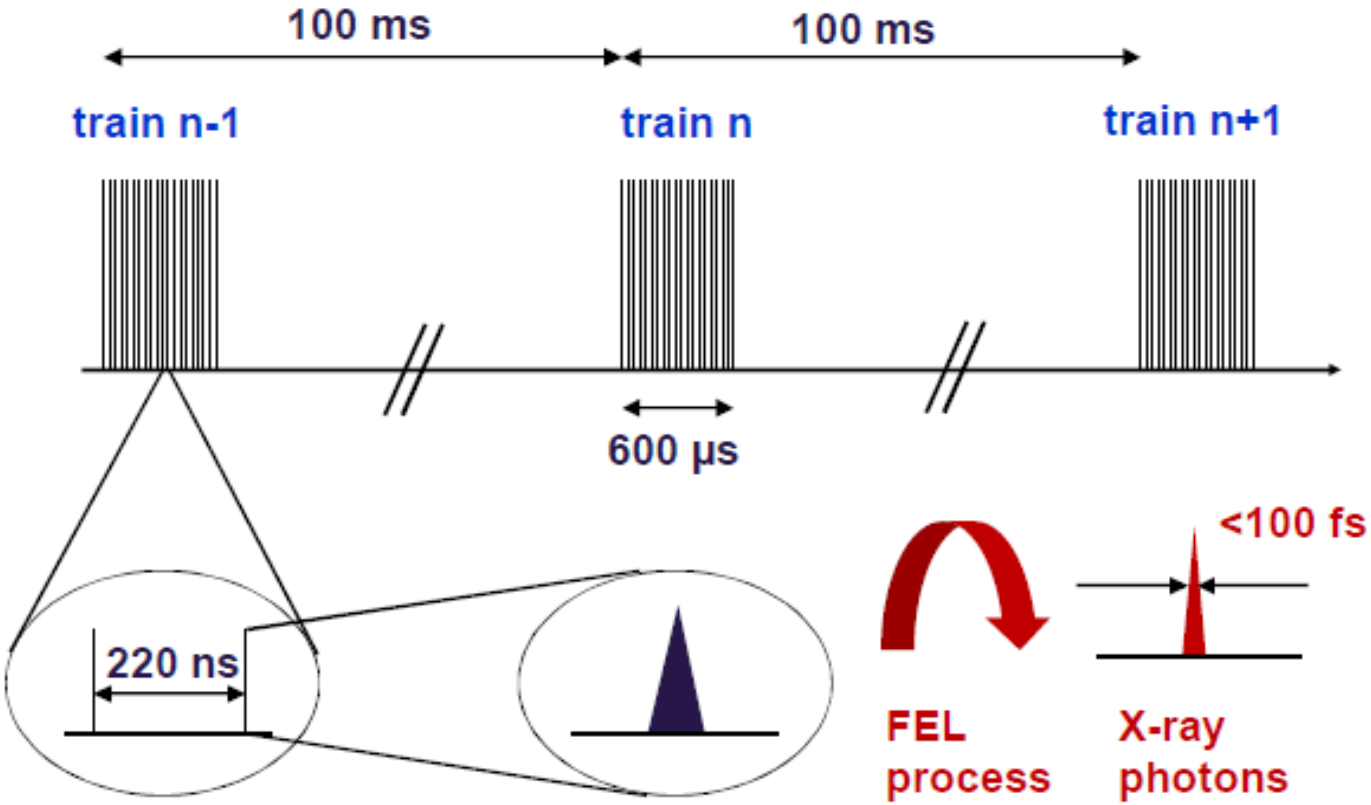

Beam time per experiment: several hours

Readout in bunch structure:

- 10 Hz train rate
- 4.5 MHz puls rate
- 2700 pulses per train

1 Mpxl detectors:
- 2MB/pulse
- 1GB/train
- 10GB/second

4 Mpxl detectors:
- 8MB/pulse
- 3GB/train
- 30GB/second

Image: http://xfel.eu

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Data processing with the CrystFEL framework

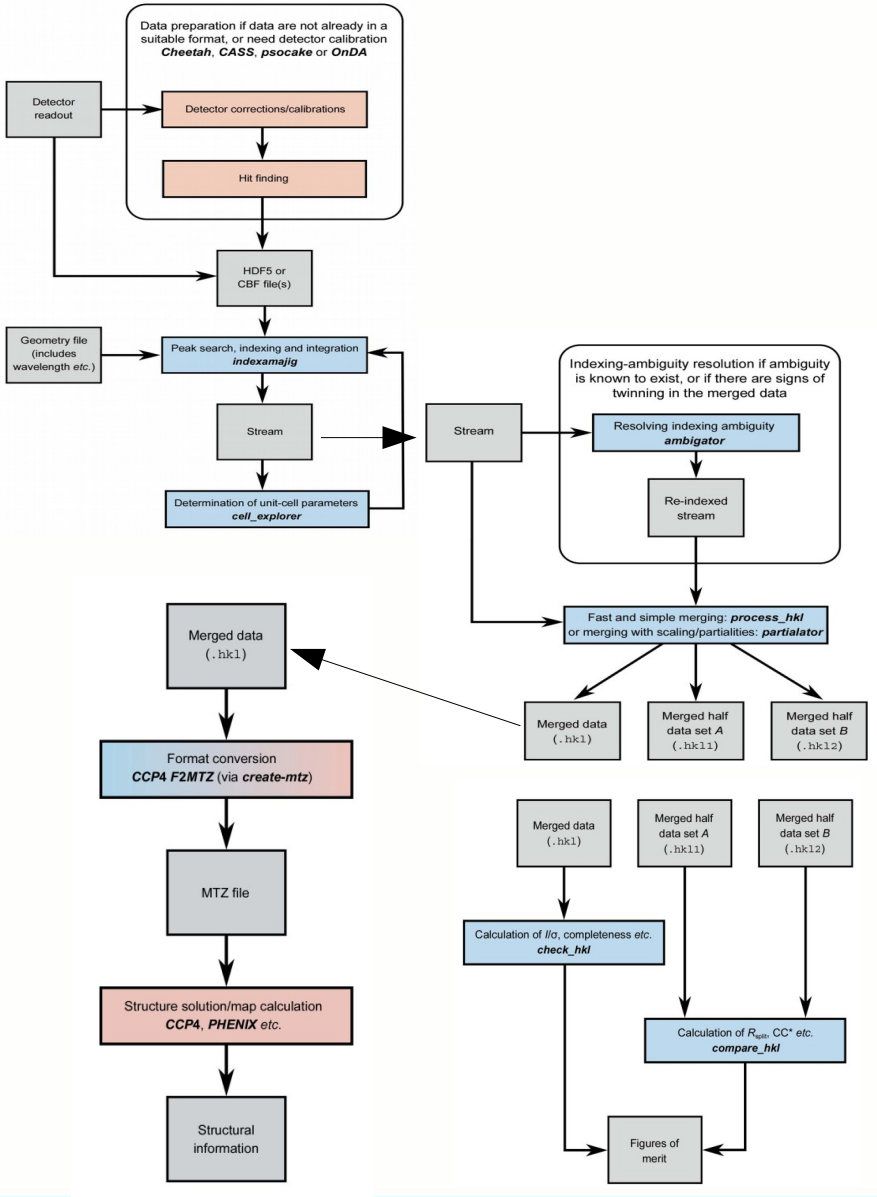

Processing serial crystallography data with **CrystFEL**: a step-by-step guide

http://journals.iucr.org/d/issues/2019/02/00/ba5291/index.html
https://www.desy.de/~twhite/crystfel

Microservices:
„CrystFEL is a suite of software comprising 15 core programs: [...] CrystFEL is primarily a command-line-driven piece of software, with some exceptions [...].”

Reproducibility challenge:
„In addition to the core programs, the CrystFEL package contains a repository of scripts which are intended to be copied to the working directory and customised to suit the individual situation.”

Image: http://journals.iucr.org/d/issues/2019/02/00/ba5291/index.html

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Compute infrastructure, Infiniband fabrics

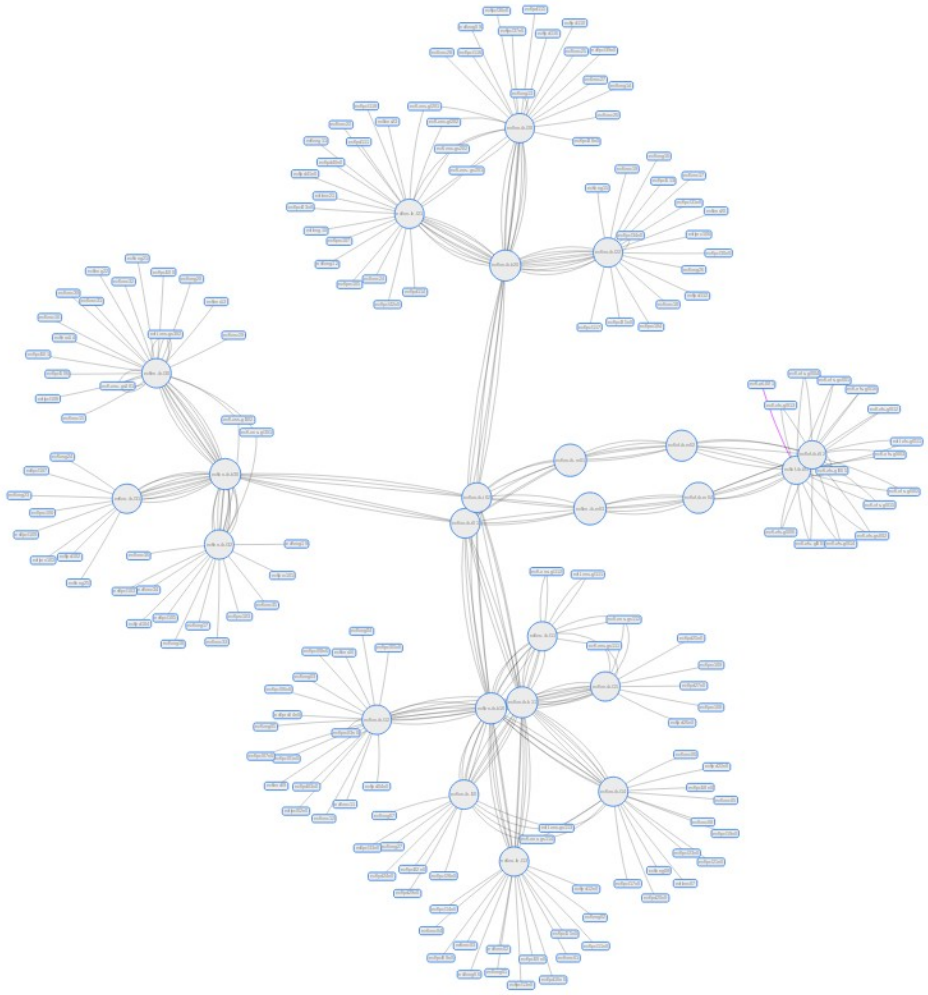

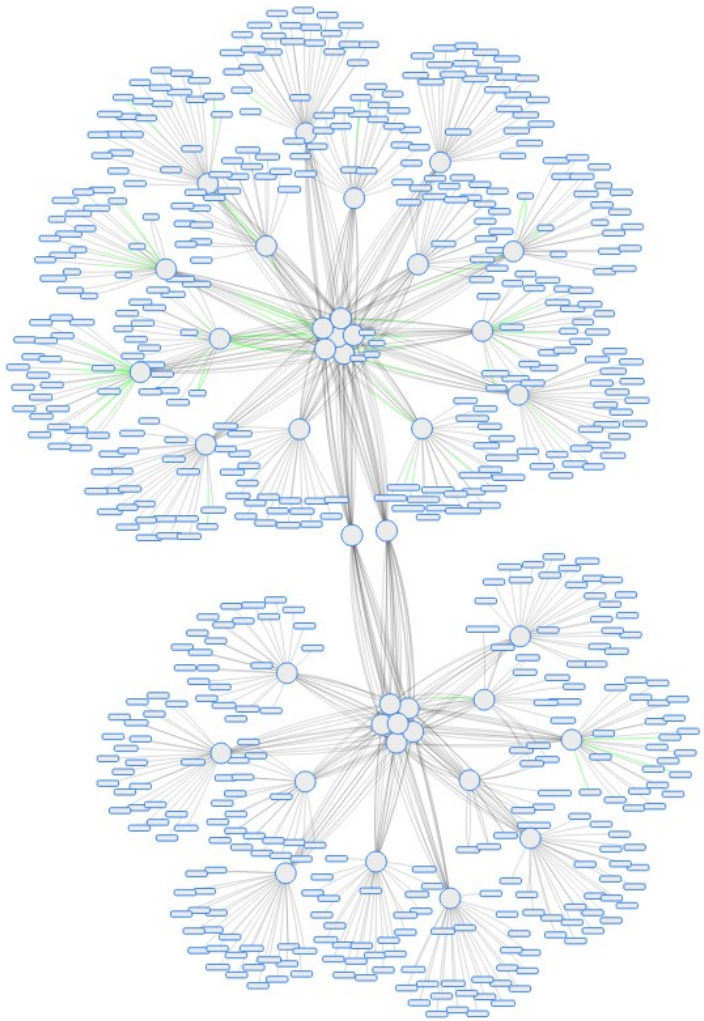

Online / Offline Compute Clusters

~28,000 CPUs, 130 GPUs, ~250 TB RAM

Infiniband fabrics

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Maxwell Jupyter Job Options

Maxwell partitions shared node on Jupyter partition

Choice of GPU none

Note: For partitions without GPUs (or choice of GPUs) the GPU selection will be set to 'none'

Job duration 8 hours

Note: on the shared Jupyter partition (jhub) the time limit is always 7 days!

Current Status					
Partition	**# nodes**	**# avail**	**# GPUs avail**	**# P100 avail**	**# V100 avail**
jhub	11	11	0	0	0
maxwell	61	8	0	0	0
maxgpu	19	16	16	5	9
all	327	143	0	0	0
allgpu	88	35	35	21	9
upex	183	117	14	14	0

Spawn

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Interactive HPC with Jupyter Notebooks

```
In [12]: lscpu | grep CPU | head -n 3

CPU op-mode(s):        32-bit, 64-bit
CPU(s):                80
On-line CPU(s) list:   0-79
```

```
In [13]: free -h

              total        used        free      shared  buff/cache   available
Mem:           503G         39G        389G        106M         75G        462G
Swap:          4.0G        109M        3.9G
```

```
In [14]: grep -q ^flags.*\ hypervisor /proc/cpuinfo && echo "VM: yes" || echo "VM: no"

VM: no
```

```
In [32]: scontrol show job $SLURM_JOBID  | grep EndTime
         scontrol show job $SLURM_JOBID  | grep billing

   StartTime=2019-03-31T13:48:49 EndTime=2019-03-31T14:48:49 Deadline=N/A
   TRES=cpu=80,node=1,billing=80
```

- Interactive session on HPC cluster
- Billing: 80 CPUs
- Job will terminate after 1 hour

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

SECTION 2

- Scientific cloud computing
  - European Open Science Cloud
  - Auto-scaling GitOps

European Open Science Cloud (EOSC) implementation phase 2018-2020 from Minimum Viable Research Data Ecosystem to the main Rules of Participation and the implementation of the FAIR data principles

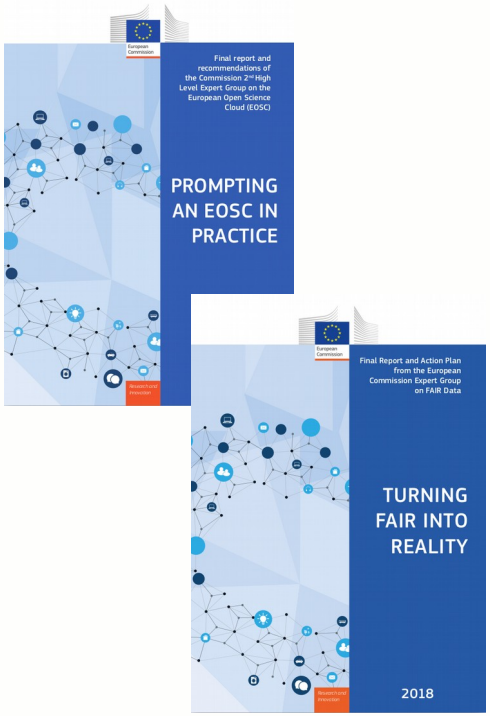

https://ec.europa.eu/info/publications/prompting-eosc-practice_en
https://ec.europa.eu/info/publications/turning-fair-reality_en
https://www.eosc-portal.eu/

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

CNCF Cloud Native Landscape

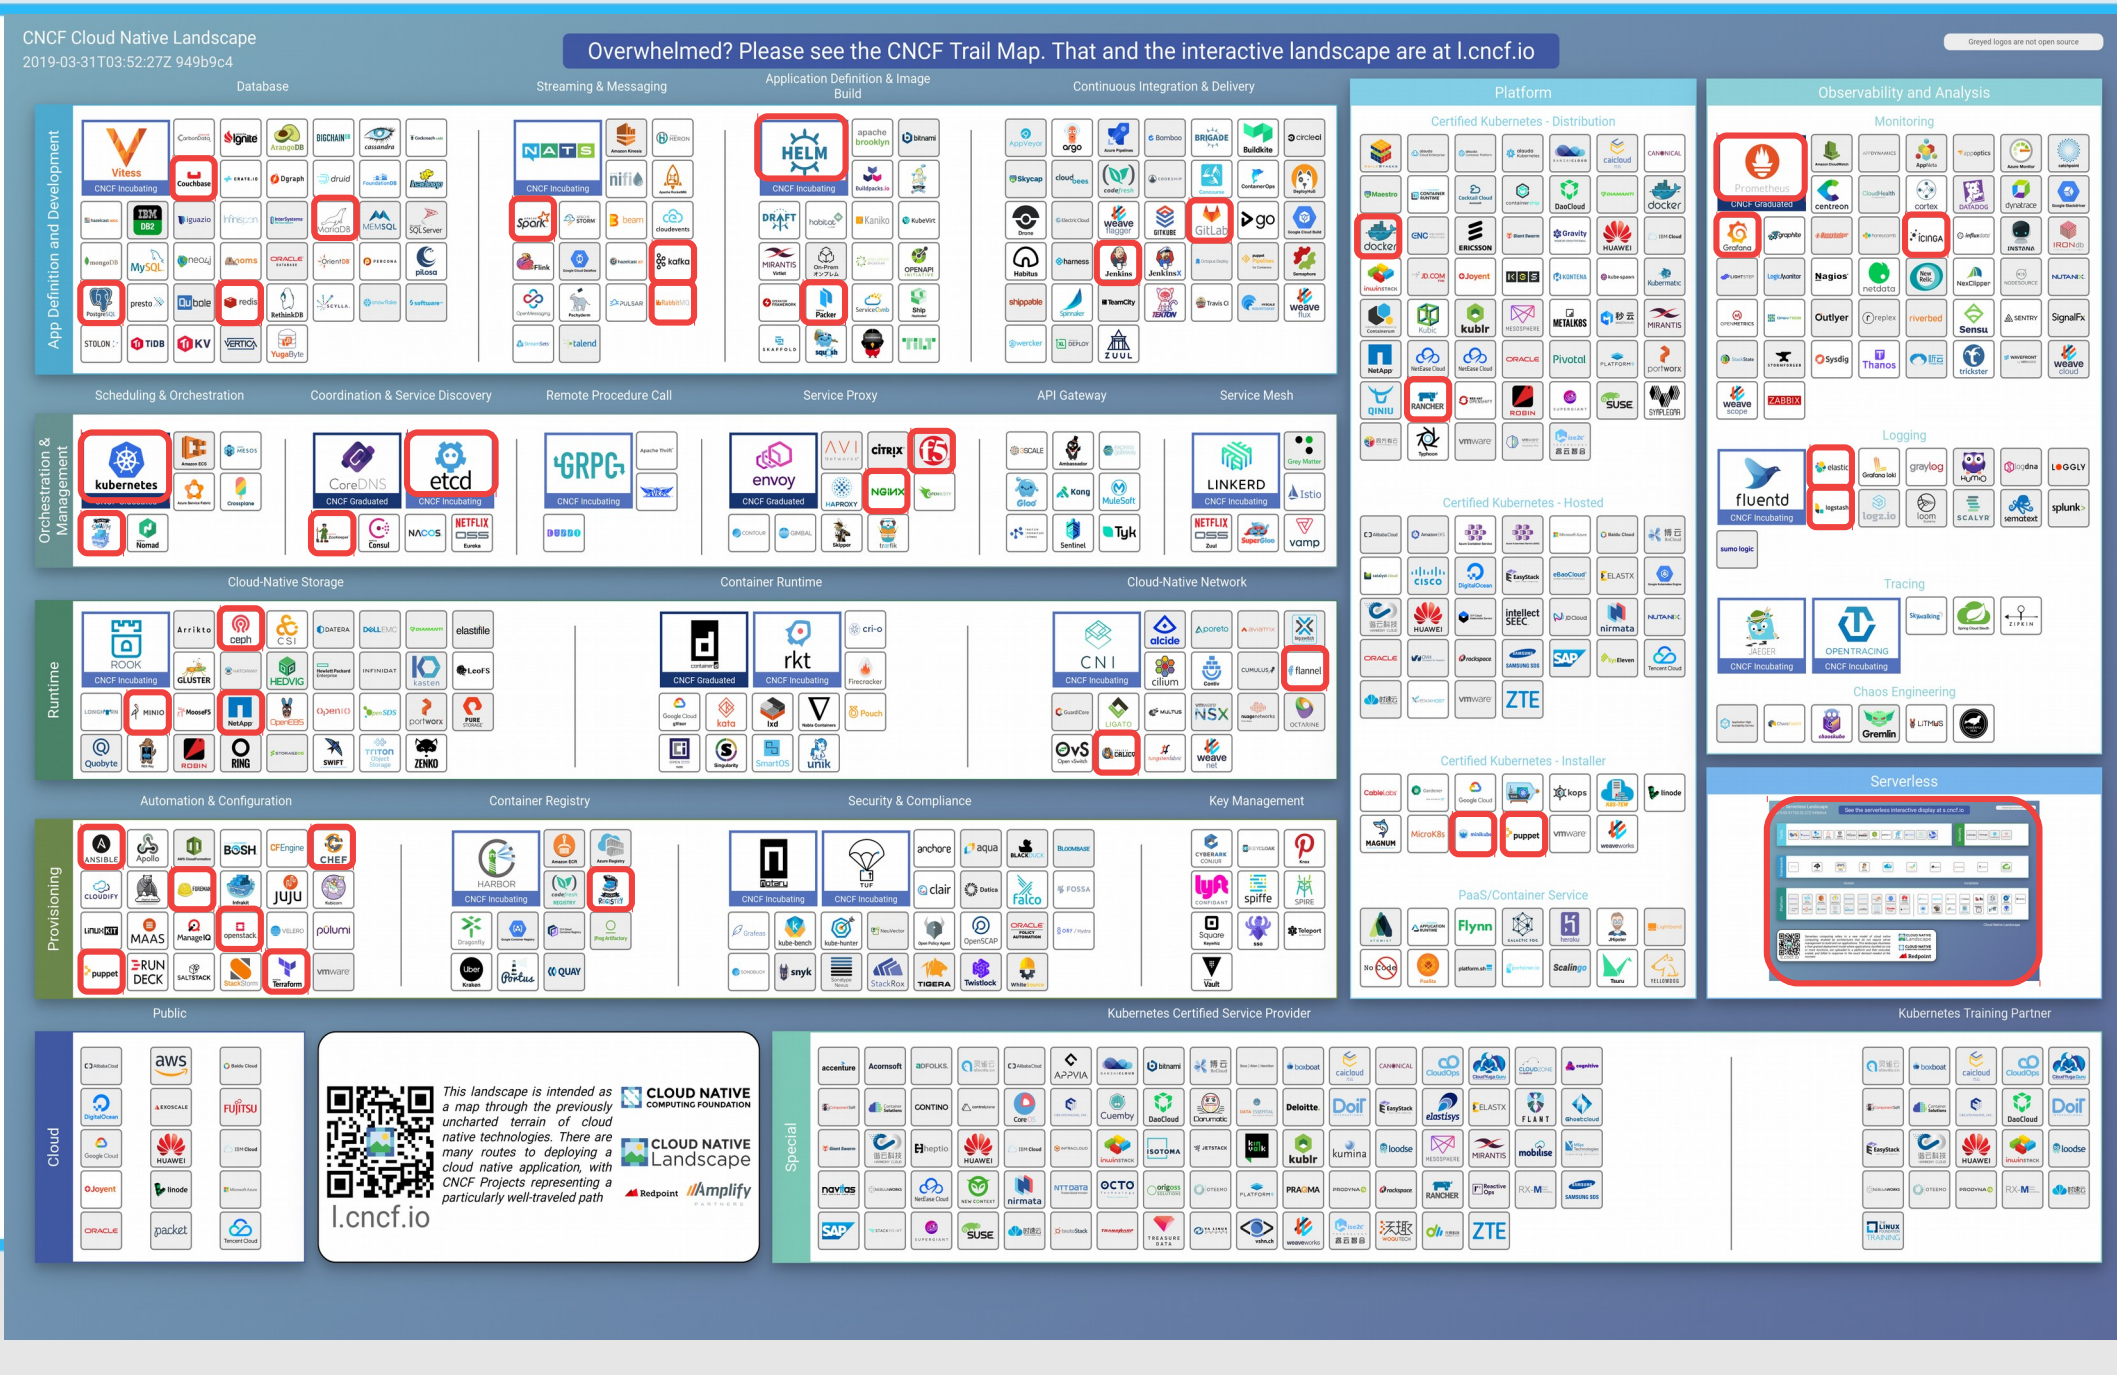

CNCF Serverless Landscape

CNCF Serverless Landscape
2019-03-31T03:52:27Z 949b9c4

See the serverless interactive display at s.cncf.io

Greyed logos are not open source

Tools

CLOUDZERO, dashbird, Epsagon, event gateway, HASURA GraphQL Engine, IO|pipe, Iron.io, Node Lambda, python-λ, SCAR, SIGMA, STACKERY, THUNDRA

Security

intrinsic, Protego, PURESEC, threat stack

Framework

APEX, Architect, aws Chalice, AWS SAM, Claudia.js, FLOGO, gun.io, serverless, SPARTA, Spring Cloud Function

Platform

Hosted

ALGORITHMIA, Alibaba Cloud Function Compute, AWS Lambda, Azure Functions, BINARIS, cloudboost, CLOUDFLARE Workers, Google Cloud Functions, HUAWEI FunctionStage, IBM Cloud Functions, kitsune, netlify, PubNub Functions, spotinst, stdlib, Tencent Cloud, twilio Functions, ZEIT

Installable

OpenWhisk, AppScale, DISPATCH, fission, fn, GALACTIC FOG, Knative, Kubeless, Kyma, LunchBadger, nuclio, OPENFAAS, PipelineAI, riff, Virtual Kubelet

Cloud Native Landscape

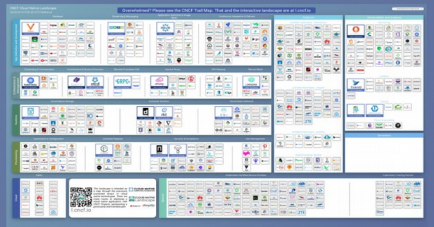

s.cncf.io

Serverless computing refers to a new model of cloud native computing, enabled by architectures that do not require server management to build and run applications. This landscape illustrates a finer-grained deployment model where applications, bundled as one or more functions, are uploaded to a platform and then executed, scaled, and billed in response to the exact demand needed at the moment

CLOUD NATIVE Landscape
CLOUD NATIVE COMPUTING FOUNDATION
Redpoint

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Virtualized compute infrastructure

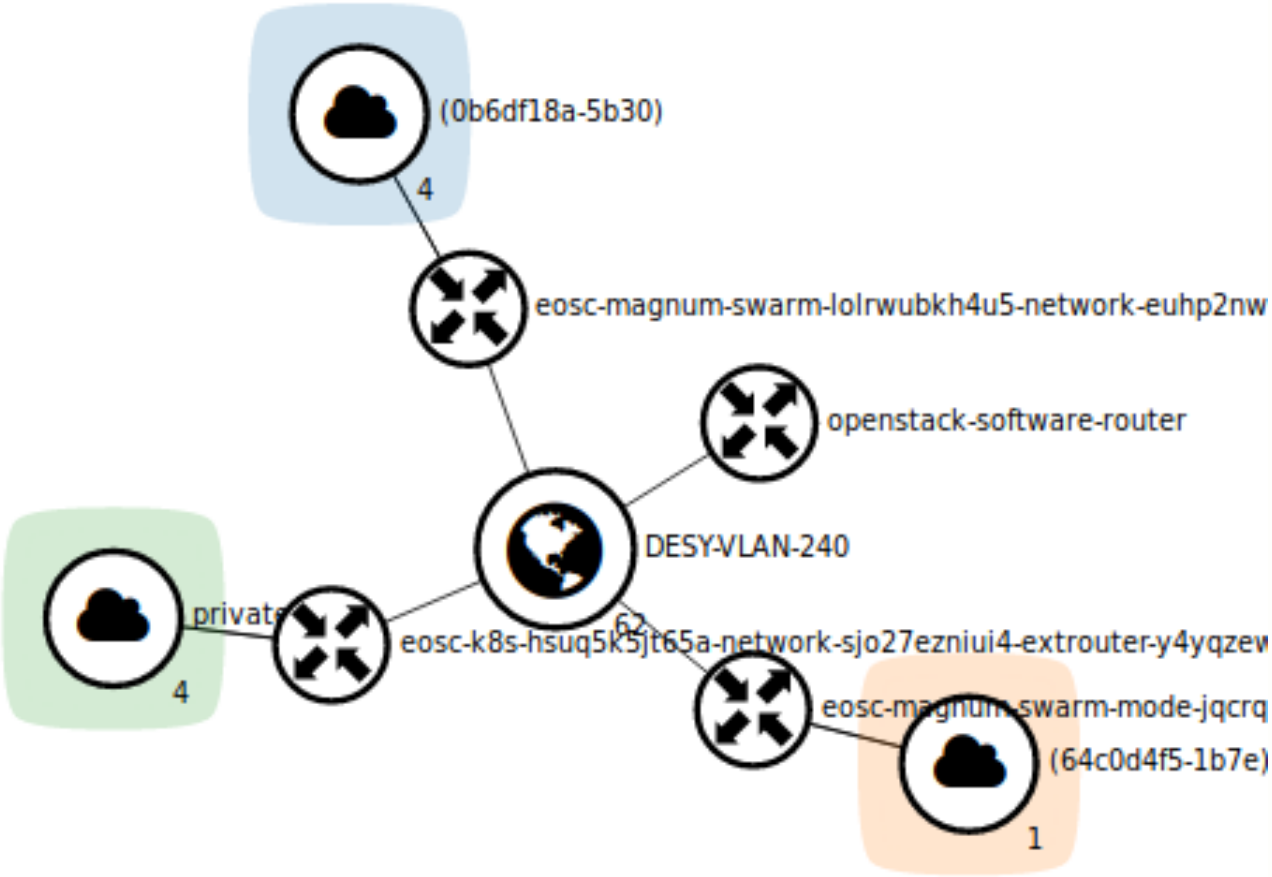

- Vms on OpenStack
- Software defined networking
- Kubernetes Clusters in private networks
- High level of isolation
- Multi-tenancy
- Scalability
- Rapid service deployments
- Here: OpenStack Magnum
- Also: rke, kubespray, ...

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

CI/CD for user defined software stacks

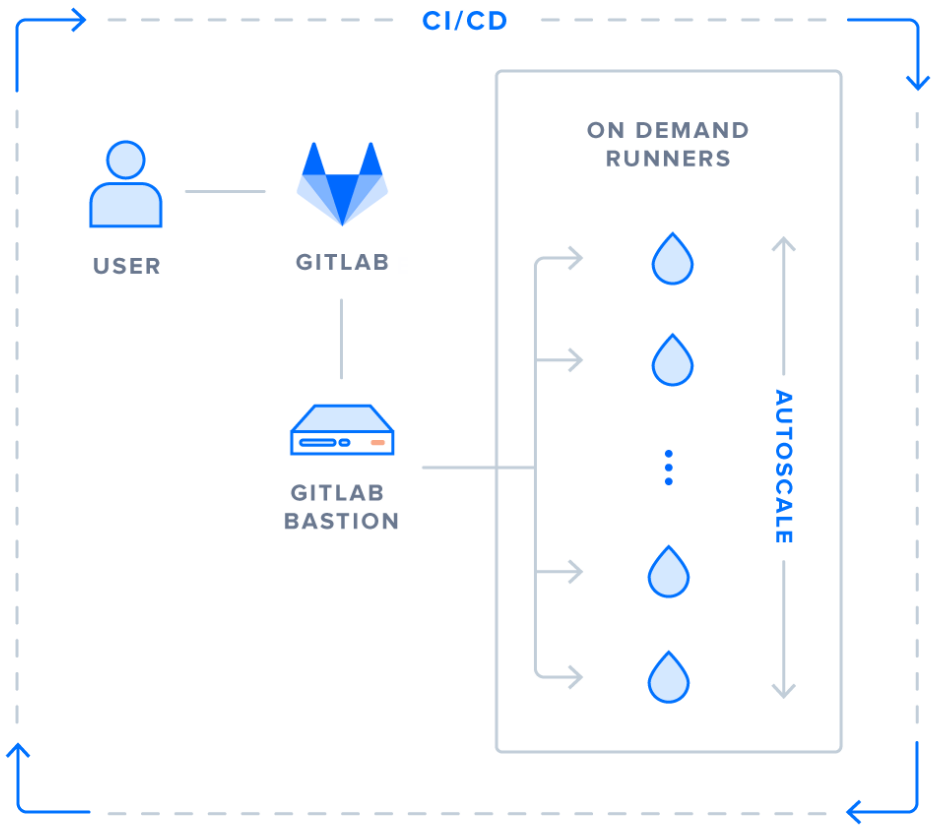

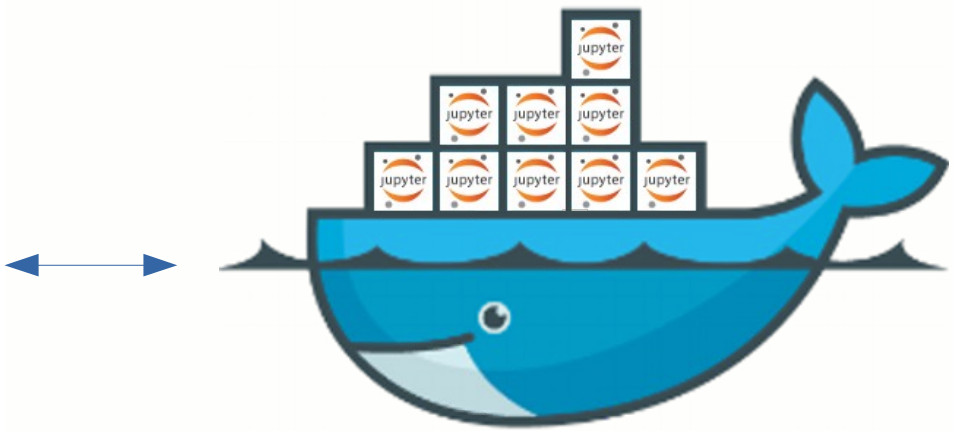

- HEAT templates for OpenStack deployment
- Creates and destroys VMs with Docker-Machine using OpenStack driver
- Users define their environments for build and tests as docker containers

Image CI/CD: https://about.gitlab.com/2018/06/19/autoscale-continuous-deployment-gitlab-runner-digital-ocean/

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

User environments in Jupyter Hub

Add requirements.txt

4 jobs from `master` in 6 minutes and 59 seconds (queued for 3 seconds)

latest

3d73cf79

Pipeline | Jobs 4

Build	Test	Release
build	test	release

pan › eosc-pan-wsk-on-jhub › **Container Registry**

Container Registry

With the Docker Container Registry integrated into GitLab, every project can have its own space to store its Docker images.

Learn more about Container Registry.

pan/eosc-pan-wsk-on-jhub

Tag	Tag ID	Size	Created
latest	a1557b47c	1.55 GiB	1 week ago
master	a1557b47c	1.55 GiB	1 week ago

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

User environments in Jupyter Hub

Add requirements.txt

4 jobs from `master` in 6 minutes and 59 seconds (queued for 3 seconds)

latest

3d73cf79

Pipeline Jobs 4

Build	Test	Release
build	test	release

pan › eosc-pan-wsk-on-jhub › **Container Registry**

Container Registry

With the Docker Container Registry integrated into GitLab, every project can have its own space to store its Docker images.

Learn more about Container Registry.

pan/eosc-pan-wsk-on-jhub

Tag	Tag ID	Size	Created
latest	a1557b47c	1.55 GiB	1 week ago
master	a1557b47c	1.55 GiB	1 week ago

Spawner Options

- **EOSC PaN FaaS image**
  Image with OpenWhisk client for function-as-a-service, based on the scientific python notebook: eosc-pan-git.desy.de:5555/pan/eosc-pan-wsk-on-jhub
- **Scientific Python environment**
  Default image with scipy and other utils
- **Extended Datascience environment**
  Additional packages for Python, R, and Julia.
- **Spark environment**
  The Jupyter Stacks spark image!
- **nbgitpuller-dev**
  We are working on importing your projects from shared links -dev.

Spawn

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Jupyter Notebooks in EOSC

https://eosc-pan-jhub.desy.de

```
In [1]: 1 lscpu | grep CPU | head -n 3

CPU op-mode(s):      32-bit, 64-bit
CPU(s):              8
On-line CPU(s) list: 0-7

In [2]: 1 free -h

              total        used        free      shared  buff/cache   available
Mem:            15G        4.2G        457M        409M         10G         10G
Swap:           0B          0B          0B

In [3]: 1 grep -q ^flags.*\ hypervisor /proc/cpuinfo && echo "VM: yes" || echo "VM: no"

VM: yes
```

https://training.notebooks.egi.eu

```
[3]: ! lscpu | grep CPU | head -n 3

CPU op-mode(s):      32-bit, 64-bit
CPU(s):              8
On-line CPU(s) list: 0-7

[4]: ! free -h

              total        used        free      shared  buff/cache   available
Mem:            31G        1.7G         17G         17M         11G         29G
Swap:           0B          0B          0B

[6]: ! grep -q ^flags.*\ hypervisor /proc/cpuinfo && echo "VM: yes" || echo "VM: no"

VM: yes
```


The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

SECTION 3

- Function-as-a-Service
  - … in Jupyter Notebooks
  - and event-driven automation with dCache

Container as a function

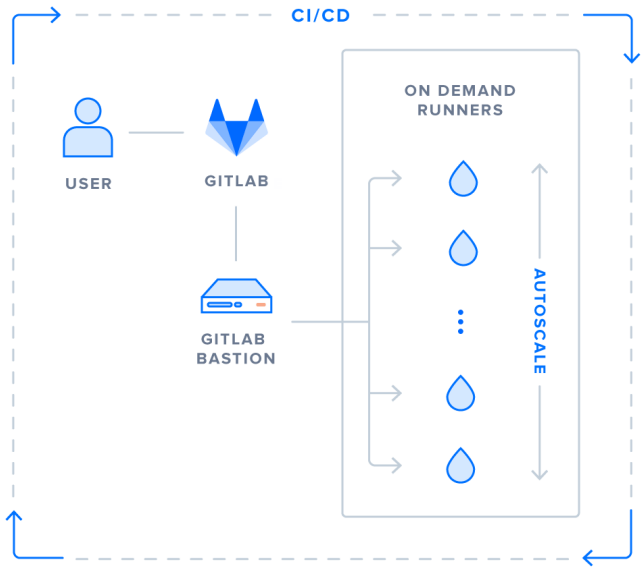

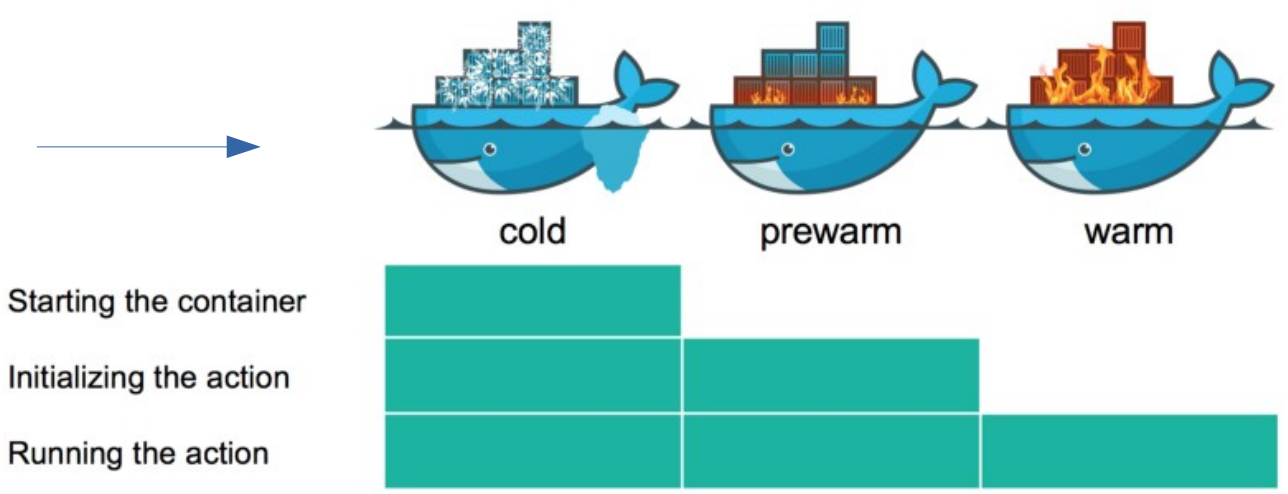

Cloud functions: No infrastructure management by the user
Efficient scaling per-function, rapid provisioning

Number of services >> 1
Number of requests of a service = arbitrary function of time

https://medium.com/openwhisk/squeezing-the-milliseconds-how-to-make-serverless-platforms-blazing-fast-aea0e9951bd0

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Event driven computation

- Execute code in response to events
  - incoming data in dCache
  - Data restaged from tape
  - new messages in dedicated queues
  - Produce derived data
  - Extract metadata
  - Manage data locality

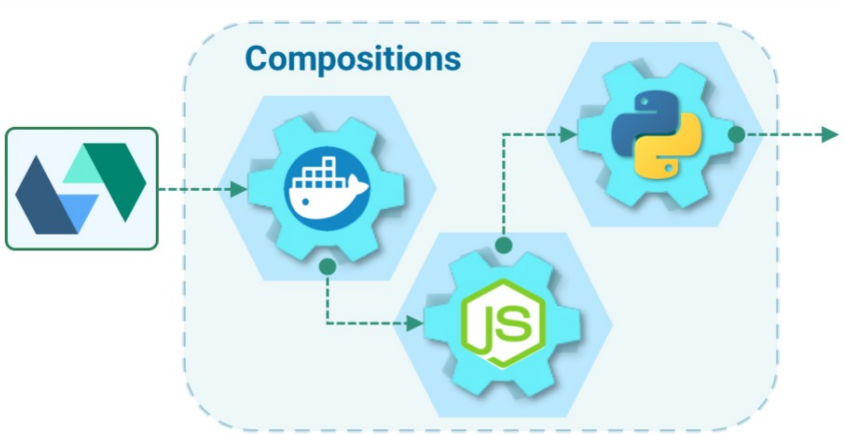

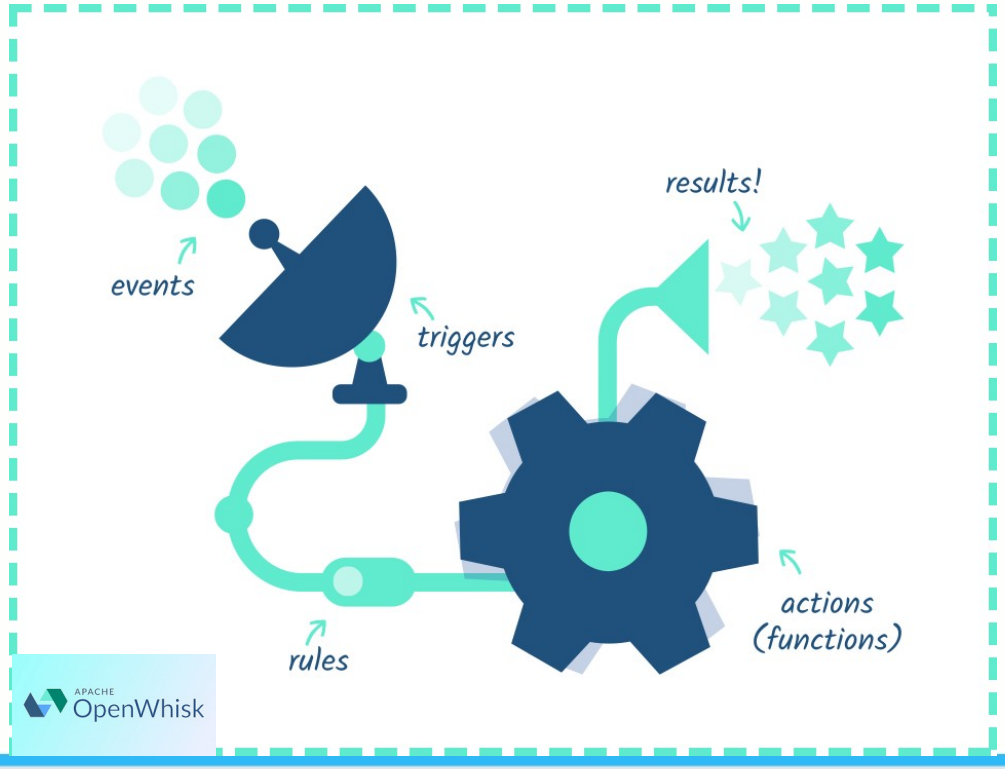

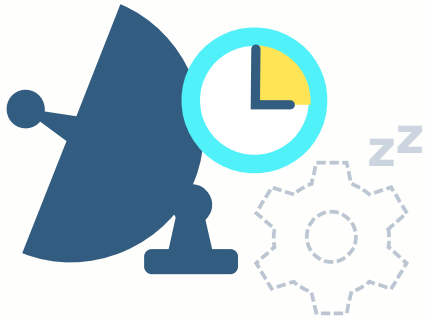

Images: https://openwhisk.apache.org/

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Conclusion

Portability	Cloud Orchestration templates, dockerized processes
Accessibility	Federated AAI, OIDC
Interoperability	Standard interfaces REST, JSON
User-friendlyness	Server-hidden Industry leading projects: GitLab, Jupyter Notebooks
Reproducibility	Everything in version control and CI/CD hashsums/pids for data, infrastructure-as-code .. and also for deployed functions and publications
Scalability	Auto-scaling Vms and Containers on OpenStack + Kubernetes

The European Open Science Cloud for Research pilot project is funded by the European Commission, DG Research & Innovation under contract no. 739563

Thank you for your attention!

Thanks to dCache people:
Tigran Mkrtchyan, DESY
Paul Millar, DESY

Open Source Software:
OpenStack
Kubernetes
Docker
dCache
Kafka
Project Jupyter
GitLab
OpenWhisk
Python
Linux
…

The European Open Science Cloud for Research pilot project
is funded by the European Commission, DG Research & Innovation
under contract no. 739563