
CERN


CHEP 2018
23RD INTERNATIONAL CONFERENCE ON
COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS
9-13 July 2018
National Palace of Culture
Sofia, Bulgaria

HNSciCloud

a Hybrid Cloud for Science

CHEP 2018 Conference, Sofia, Bulgaria

9th July 2018

João Fernandes

CERN

IT department

Helix Nebula – The Science Cloud with Grant Agreement 687614 is a Pre-Commercial Procurement Action funded by H2020 Framework Programme

The Helix Nebula Science Cloud

- To **provide a common cloud platform for the European research community**

- Via a collective effort of 10 procurer Research Organisations forming the **Buyers Group**

Expressing the need to increase the analysis capability and capacity offered to their users to keep pace with the growth in scientific data

Helix Nebula Hybrid Cloud Model

Bringing together:

- Research Organisations
- Data Providers
- Publicly funded e-infrastructures
- Commercial cloud providers

with:

Procurement and Governance suitable for the dynamic cloud market

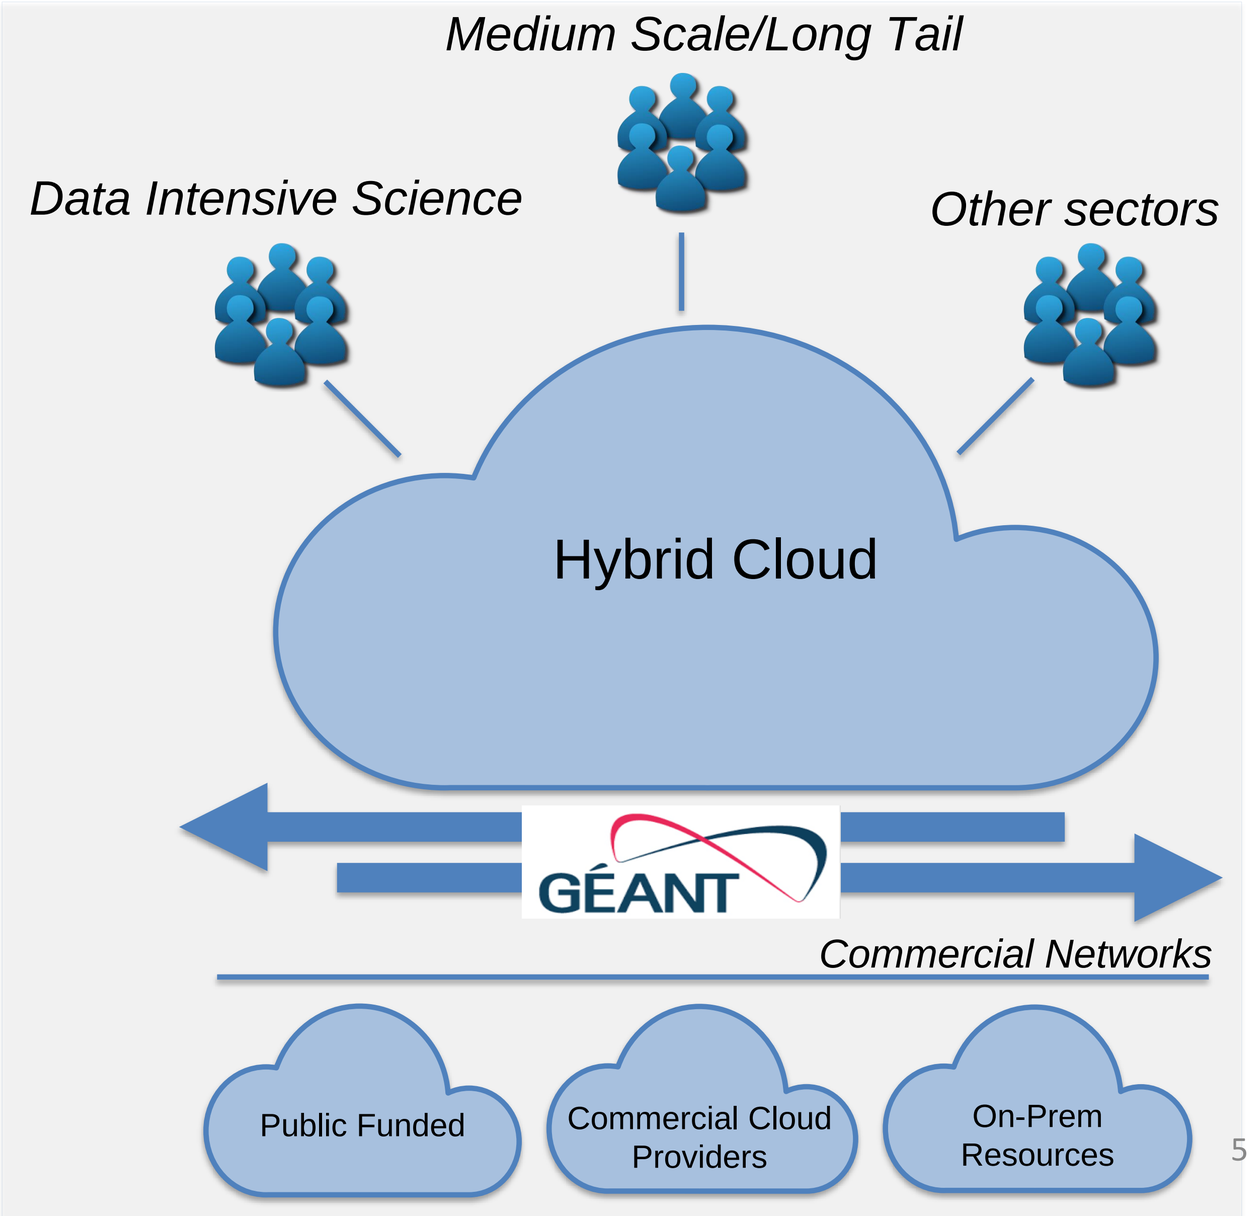

Helix Nebula Science Cloud Joint Pre-Commercial Procurement

Procurers: **CERN, CNRS, DESY, EMBL-EBI, ESRF, IFAE, INFN, KIT, STFC, SURFSara**

Experts: Trust-IT & EGI.eu

Result ...rch communities

eduGAIN

elixir AAI

GÉANT

Deployed in a hybrid cloud mode:

- procurers data centres
- commercial cloud service providers
- GEANT network, EduGAIN and ELIXIR Federated Identity Management

Co-funded via H2020 Grant Agreement 687614

Total procurement budget >5.3M€

Challenges

Innovative IaaS cloud services integrated with procurers in-house resources to support a range of scientific workloads

- ***Compute and Storage***
  - Support a range of architectures, virtual machine and container configurations including HPCaaS, working with datasets in the petabyte range with transparent data access

- ***Network Connectivity and Federated Identity Management***
  - Provide high-end network capacity via GEANT for the whole platform with common federated identity and access management
    - AAI activities have been described as a ‘pilot’ use-case in a AARC2 project:
      - https://aarc-project.eu/wp-content/uploads/2018/06/DSA1.1-v1.1FINAL.pdf
    - More in Hannah Short (CERN) talk, Track 3:
      - https://indico.cern.ch/event/587955/contributions/2936916/

- ***Service Payment Models***
  - *Explore* a range of purchasing options to determine those most appropriate for the scientific application workloads, including **vouchers** or other means of easy integration in the organisations procurement models and production of a **TCO study** ready by end of 2018

HNSciCloud project phases

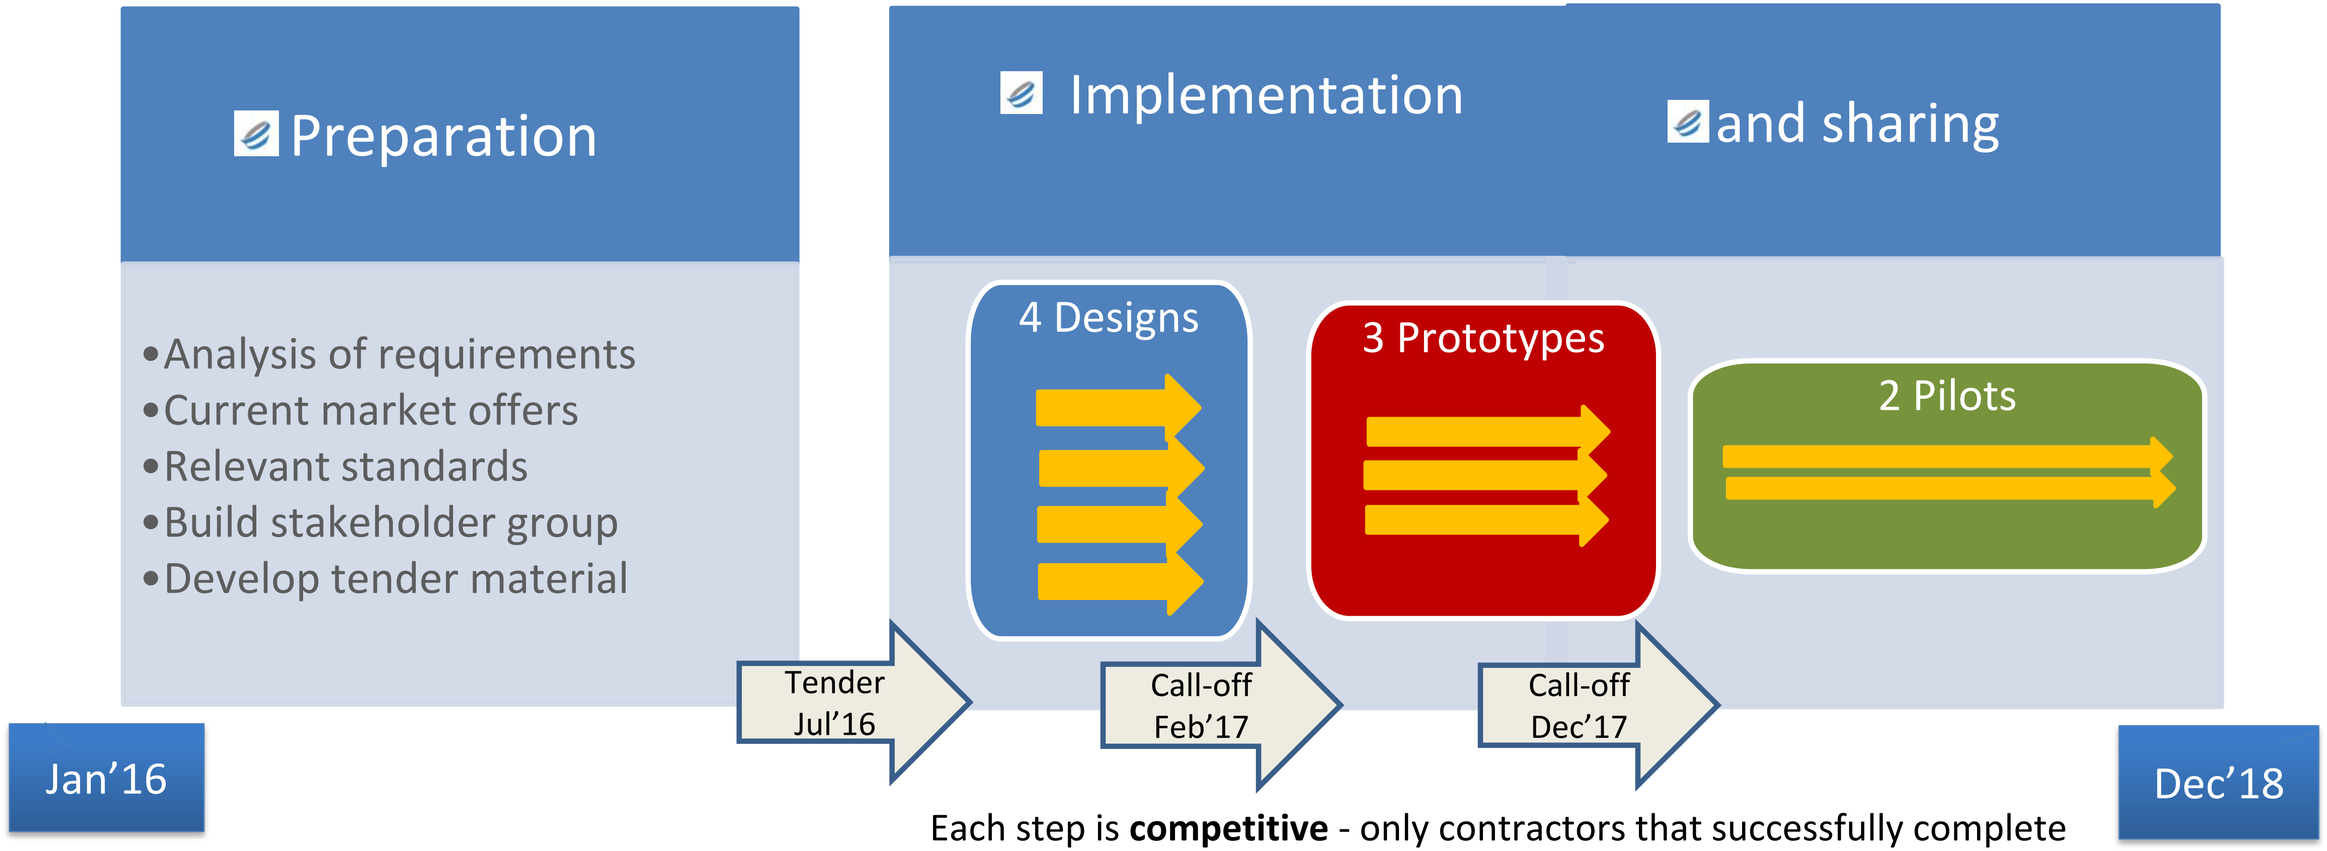

Each step is **competitive** - only contractors that successfully complete the previous step can bid in the next

Phases of the tender are defined by the Horizon 2020 Pre-Commercial Procurement financial instrument

Pilot Phase Timeline

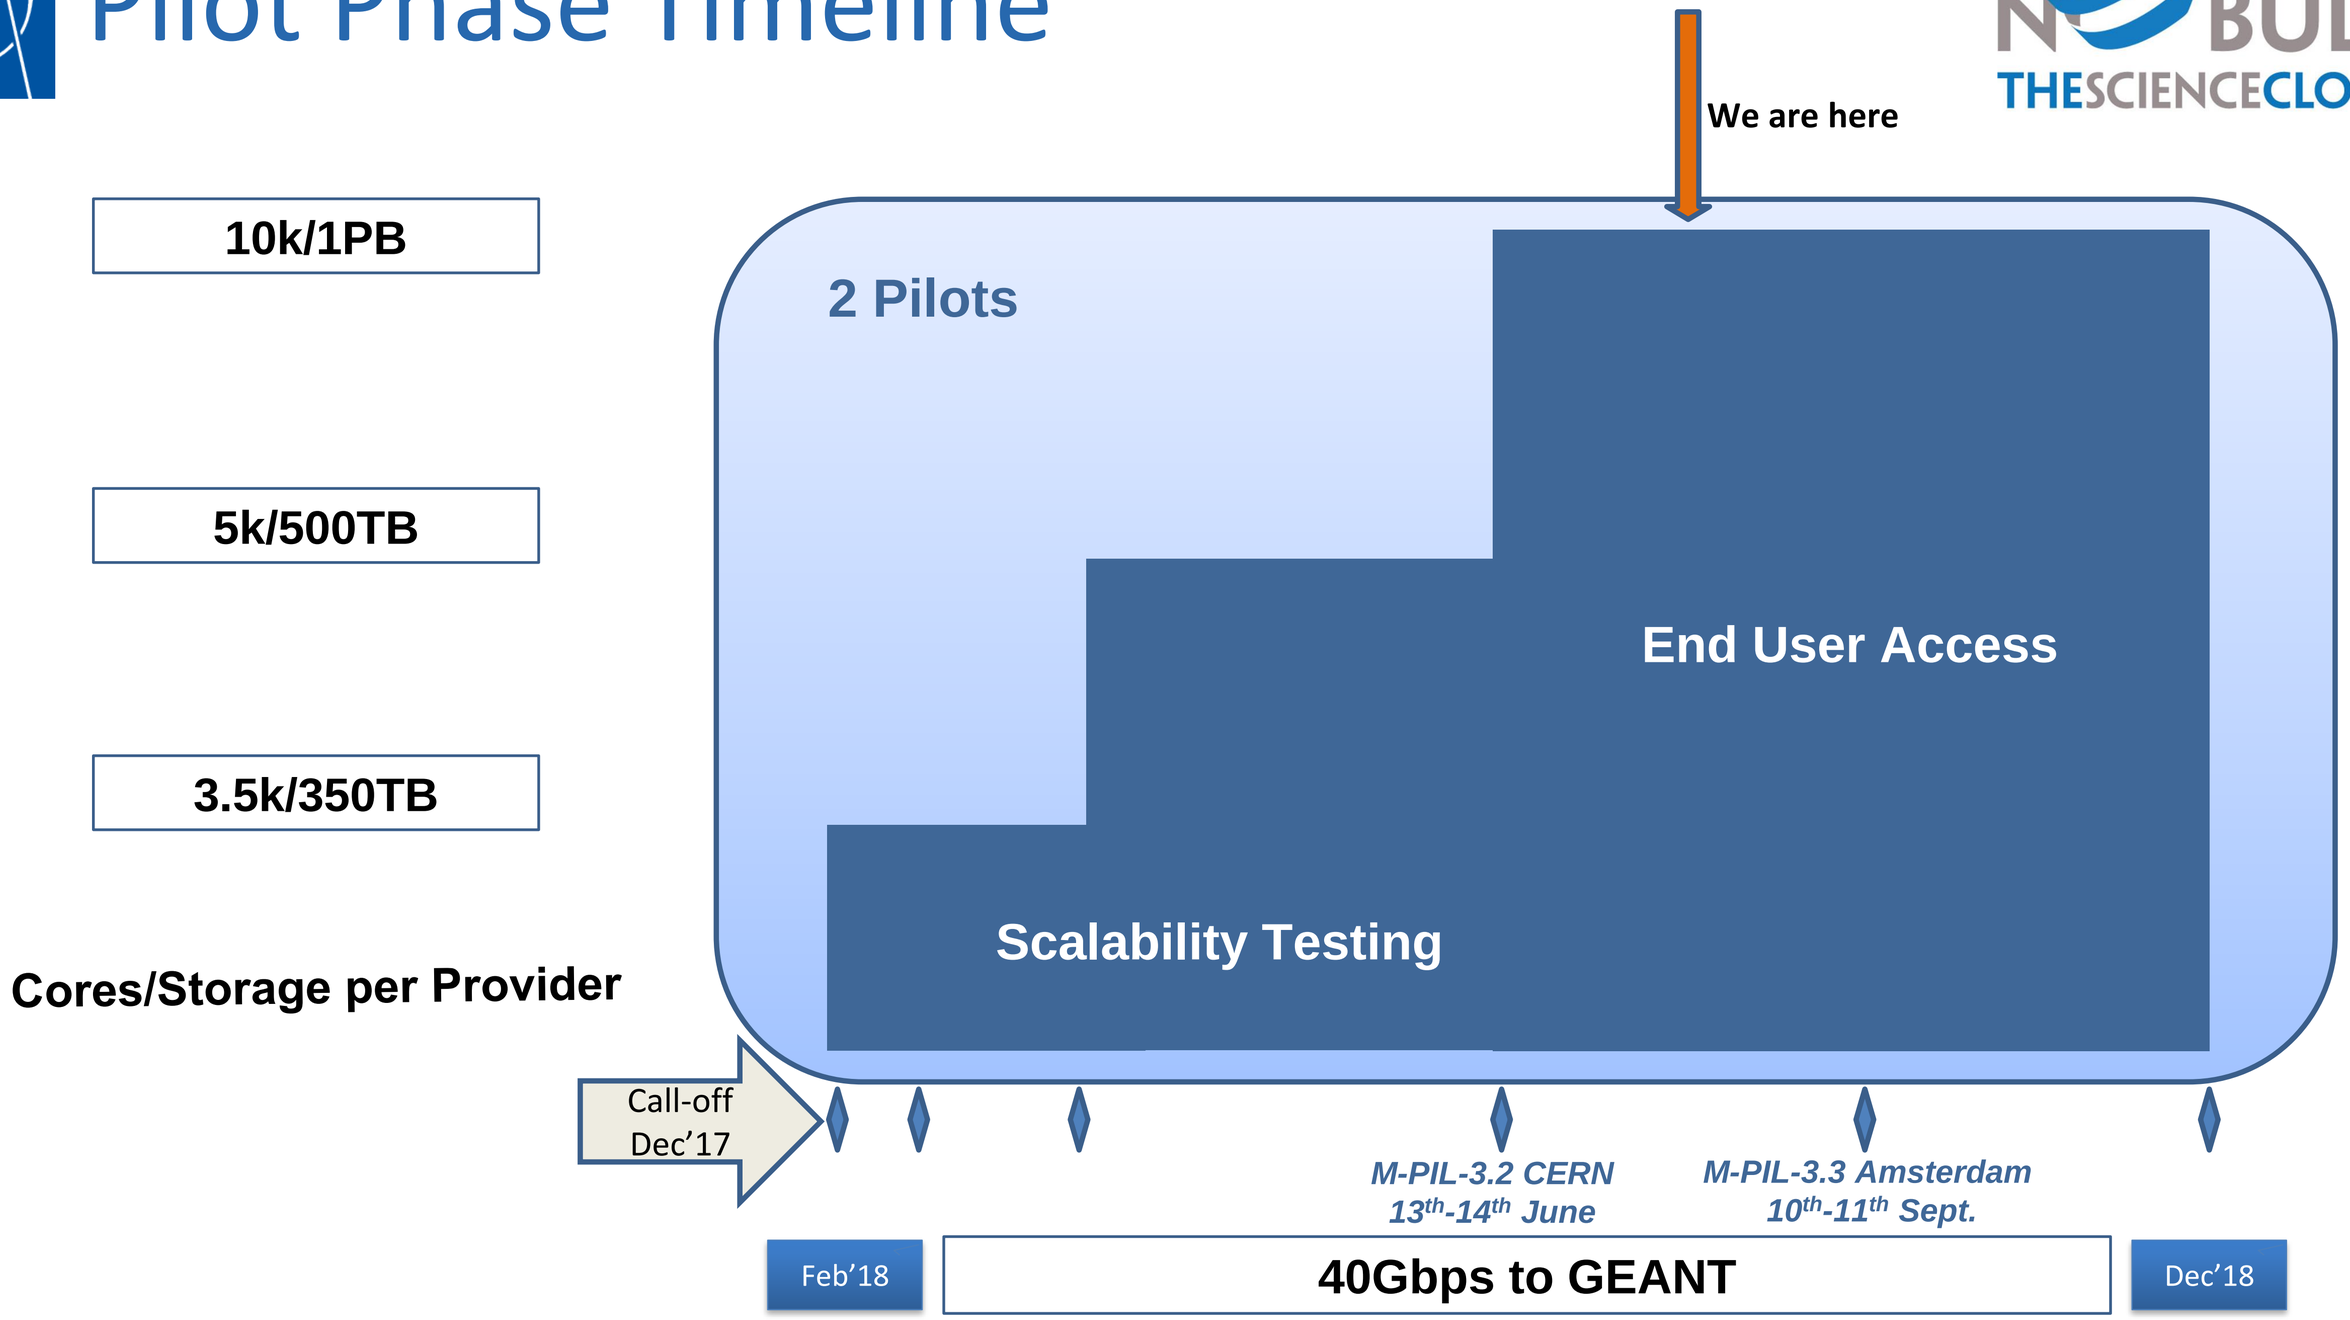

What HNSciCloud is *not*

- ***Not a zero cost research grant***
  - Contractual Relationship: Important to understand how to integrate commercial cloud services into scientific activities
- ***Not a one-off test***
  - *A* predictable production-quality contribution to the scientific programme of the *o*rganisations involved
- ***Not just a technical evaluation***
  - **Legal**: Where is the data stored? Under what jurisdiction? What are the contractual terms and conditions?
  - **Commercial**: what are the commercialisation plans after the project? How much will it cost to use them in production? What purchase models are available?
- ***Not a walk in the park***
  - 2.5 years of intense collaboration for procurers and companies to get here
    - From common requirements specification for workloads of multiple scientific disciplines to tender launch, design and prototype phases
    - Required significant effort by the procuring organisations to follow the process, provide all the necessary material and performing tests on the developed services.

Cloud Providers

- *T-Systems*
  - *IaaS based on OTC*

- *RHEA*
  - *IaaS provided by Exoscale*

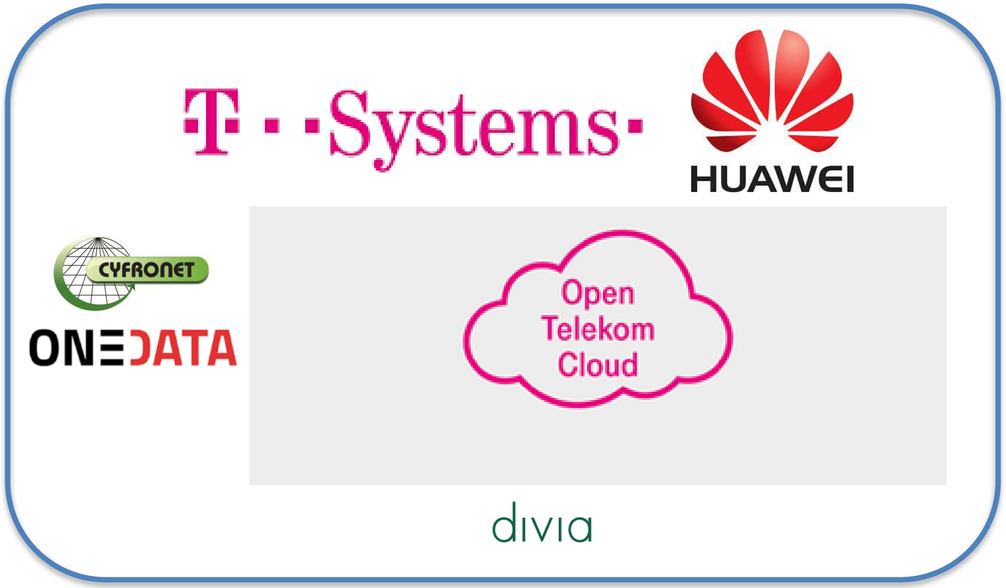

HEP Flagship Deployments

- WLCG
  - ALICE, ATLAS, CMS and LHCb
    - Daniele Spiga (INFN) talk about DODAS, Track 7:
      - https://indico.cern.ch/event/587955/contributions/2937198/
    - Matthias Schnepf (KIT) talk about Dynamic Integration of resources, Track 8:
      - https://indico.cern.ch/event/587955/contributions/2937900/

- CERN Batch Service
- Container Federation
  - Openstack Summit, reference talk:
    - https://www.openstack.org/summit/vancouver-2018/summit-schedule/events/20768/cern-experiences-with-multi-cloud-federated-kubernetes

- Belle II
  - Silvio Pardi (INFN) poster about experience of Belle II with commercial clouds, Track 7:
    - https://indico.cern.ch/event/587955/contributions/2937060/

- Interactive Analysis for End Users for TOTEM
  - https://indico.cern.ch/event/727193/contributions/3039091/attachments/1667076/2674030/TotemTest_HNSciCloud.pdf

- Machine Learning/Deep Learning for Fast Detector Simulation using GPUs
  - Sofia Vallecorsa (CERN openlab), Track 2 and Jean-Roch Vlimant (CMS), Track 6:
    - https://indico.cern.ch/event/587955/contributions/2937595/
    - https://indico.cern.ch/event/587955/contributions/2937513/

WLCG Cloud Consolidation

- 8 of the 10 members of the Buyers Group are actively supporting the LHC programme
  - Operating sites in WLCG
- Agreement to consolidate to a shared WLCG tenant
  - Reduce effort at sites to support WLCG workloads
  - Reduce network traffic across each site and commercial cloud provider data centres
- Hybrid Model bringing together WLCG sites and commercial cloud providers data centres linked via GEANT or/and NRENs
- More WLCG procuring organisations can participate via the early adopter programme:

  https://www.hnscicloud.eu/the-hnscicloud-adopter-group

Unconsolidated

Slides from Ben Jones (CERN IT-CM)

- In this setup, an experiment could have to define 8 additional sites to send jobs to the same resources

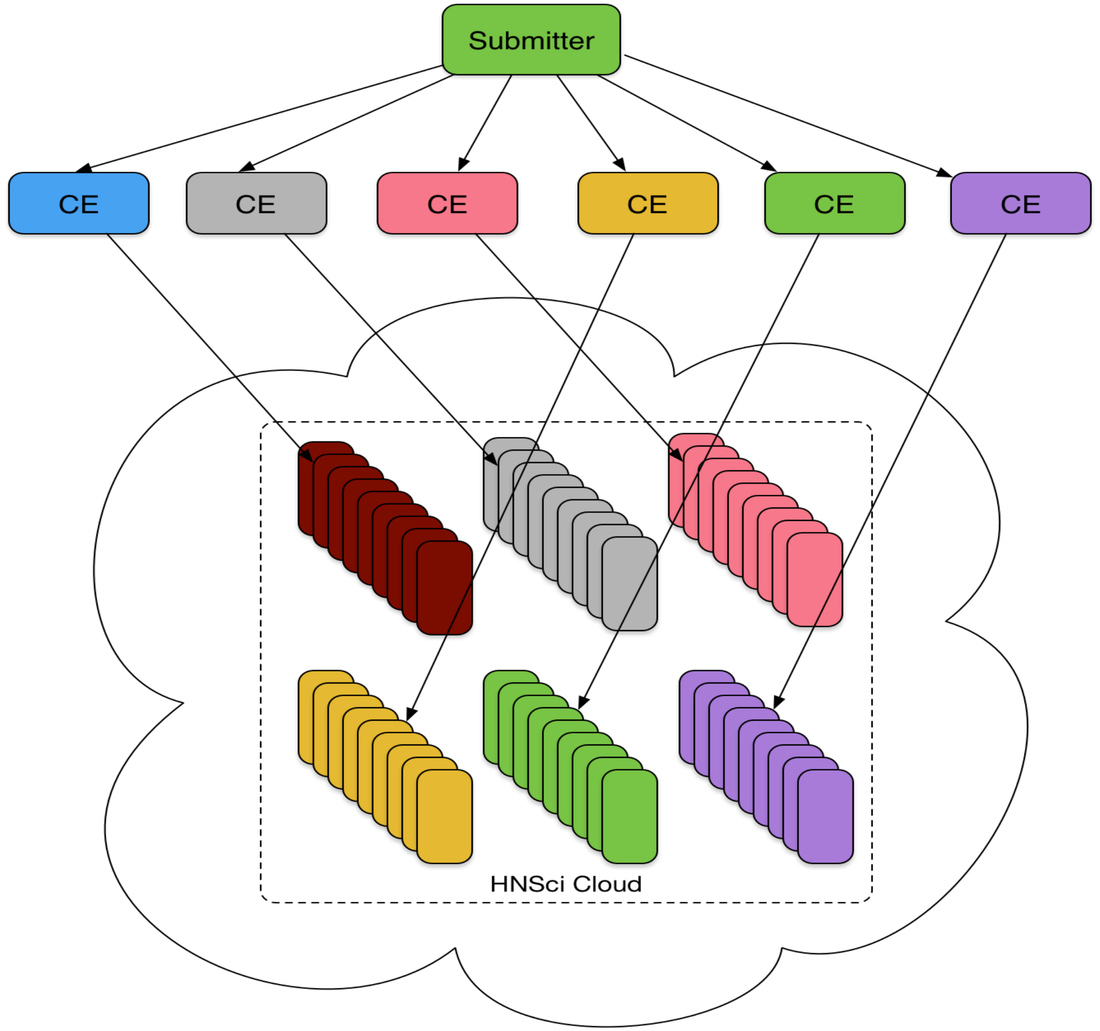

Shared Tenant

Slides from Ben Jones (CERN IT-CM)

- With a shared tenant and single entry point (CE), only one site needs to be setup

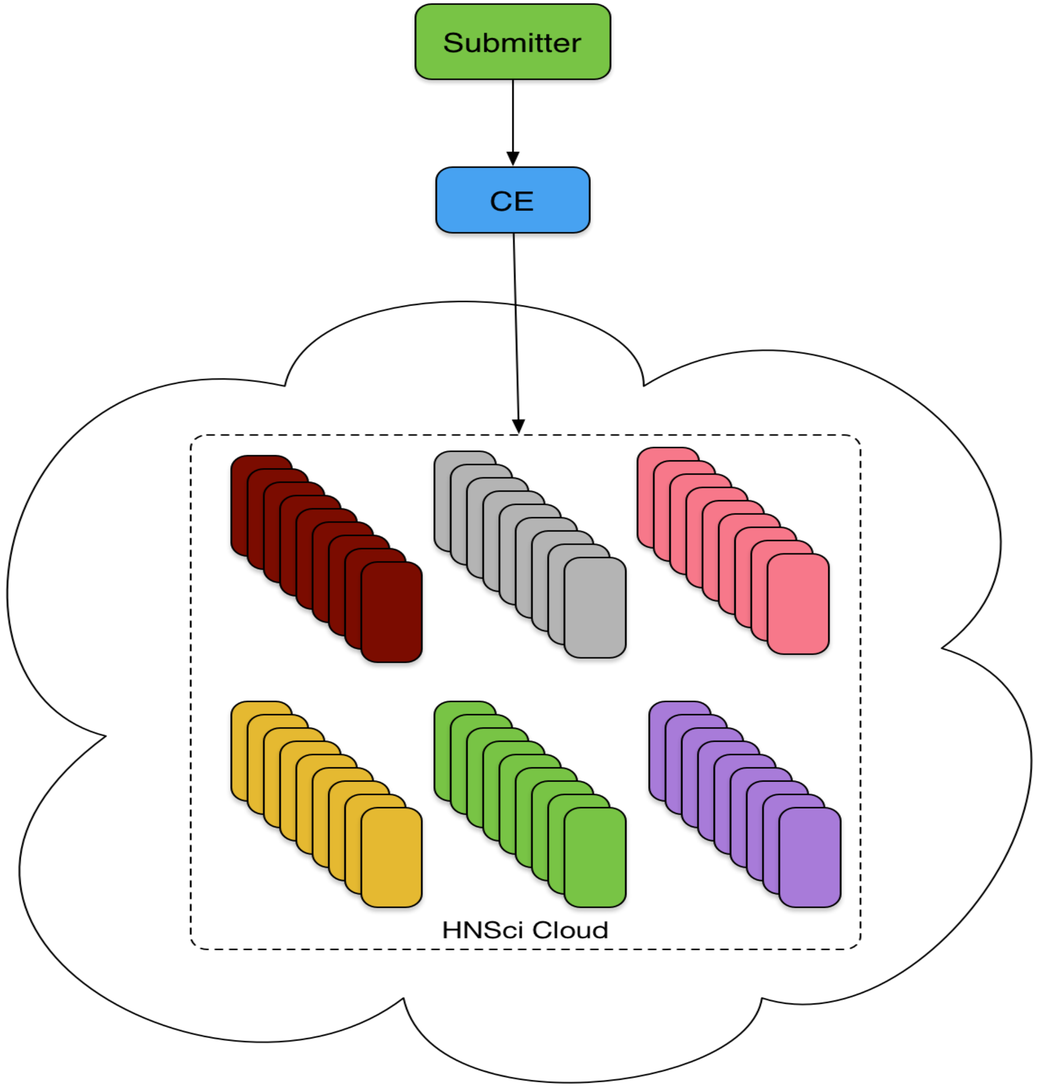

Multiple Clouds

Slides from Ben Jones (CERN IT-CM)

- Current setup being explored at HNSciCloud with a single entry point

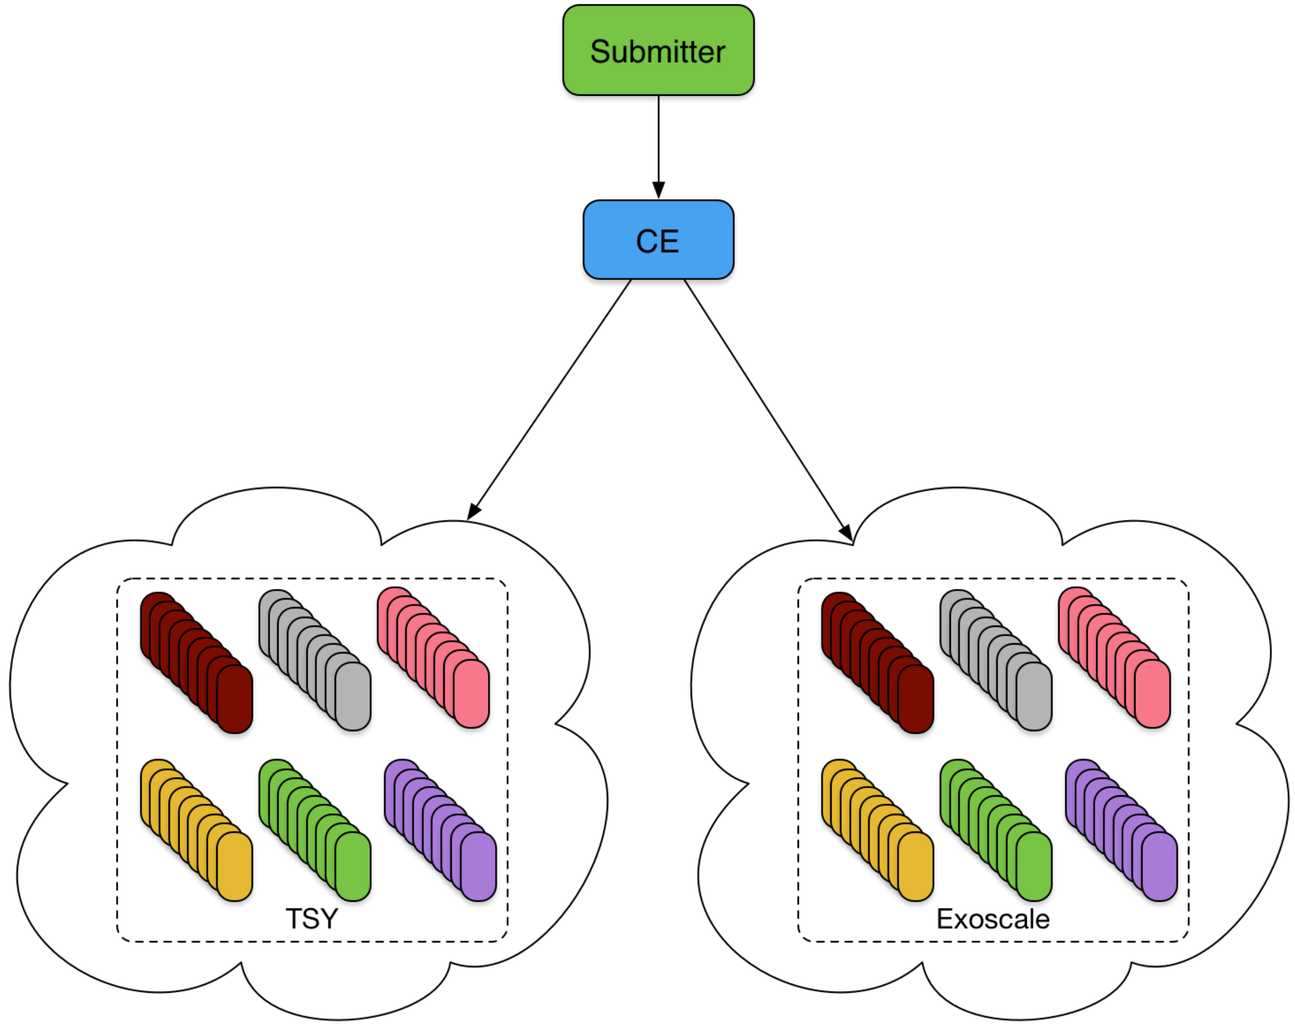

European Open Science Cloud

HNSciCloud promoted as a working example of an Open Science Cloud by the EC High Level Expert Group

Upcoming Events

•August 28: GridKa School, Karlsruhe
Hands-On session organised by KIT

• September 11: HNSciCloud meeting, Amsterdam
Organised by SURFsara

• October 9-11: DI4R 2018, Lisbon

• October 24th: Hamburg
Organised by DESY