Helix Nebula Science Cloud

GridKa - School of Computing 2018

Karlsruher Institut für Technologie

Paulo Canilho · CERN · IT Department

Helix Nebula Science Cloud

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

What?

- What?
- How?

- What?
- How?
- Use-case?

- What?
- How?
- Use-case?
- What is next?

What is the HNSciCloud?

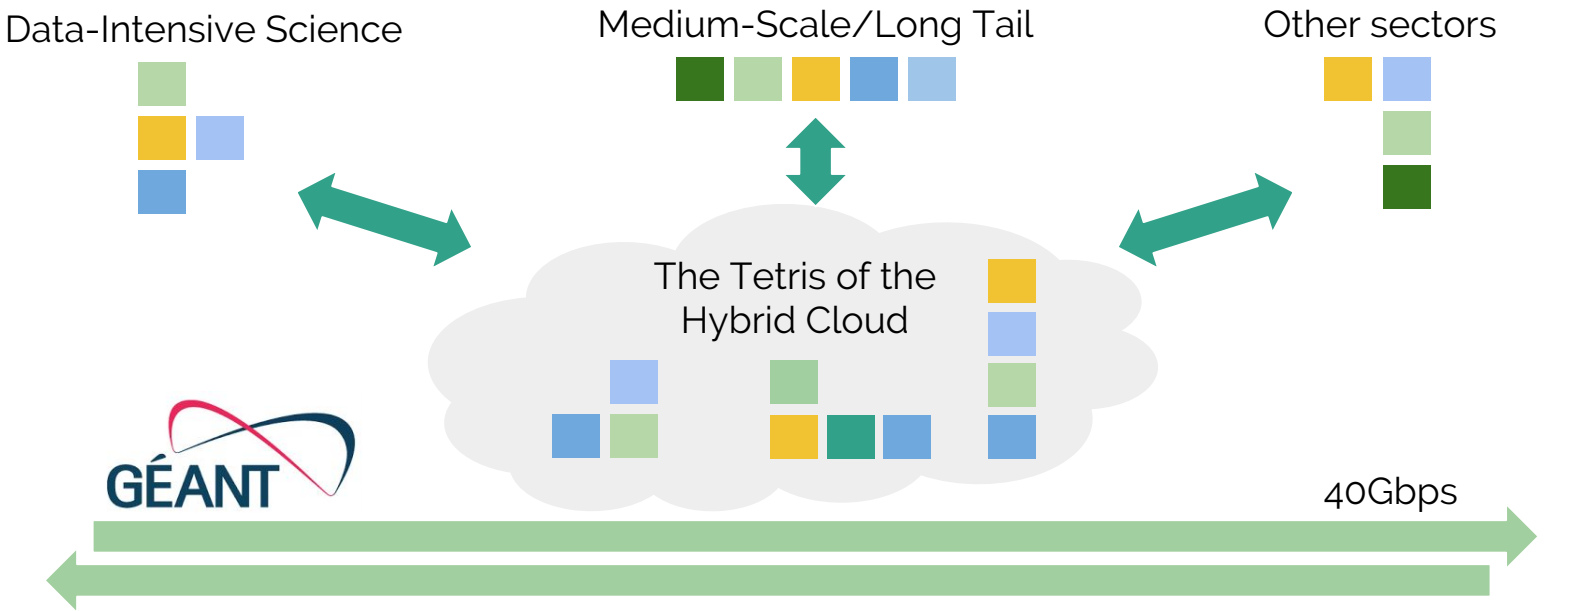

What is the HNSciCloud?

Desy

SURFsara

STFC

EMBL

KIT

CERN

CNRS

ESRF

IFAE

ESRF

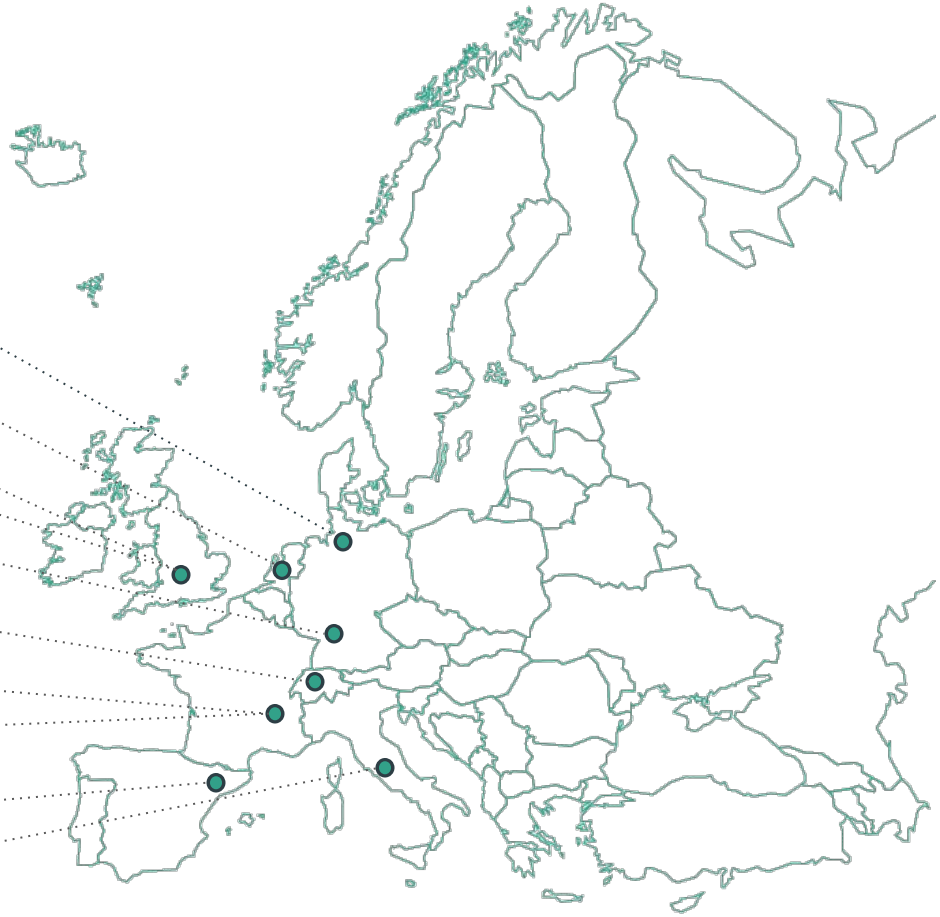

Procurers Who is involved?

Experts

- Trust-IT
- EGI.eu

Total procurement budget

> 5.3 million euros

Budget

Who is involved?

Astronomy

Photon/Neutron Sciences

WLCG

Belle II

Compass

ILC

High-Energy Physics

CTA

MAGIC

Pierre Auger

ELIXIR

Euro-BioImaging

PCAWG

BBMRI-ERIC

WE-NMR

ADDOCK

Life Sciences

PETRA III

3DIX

OCEAN

OSIRIS

.

.

.

.

Long Tail of Science

Use-Cases Who is involved?

Compute & Storage

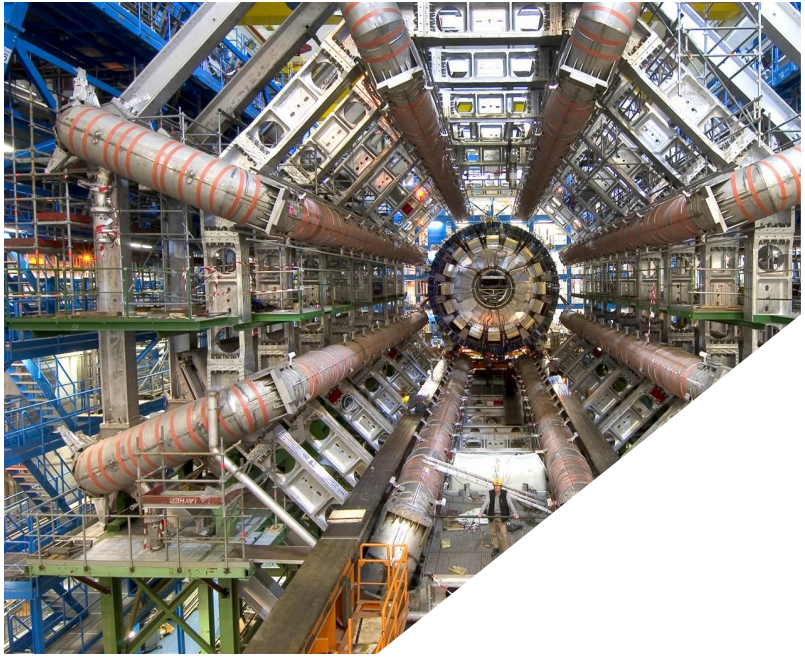

Challenges

Networking & FIM*

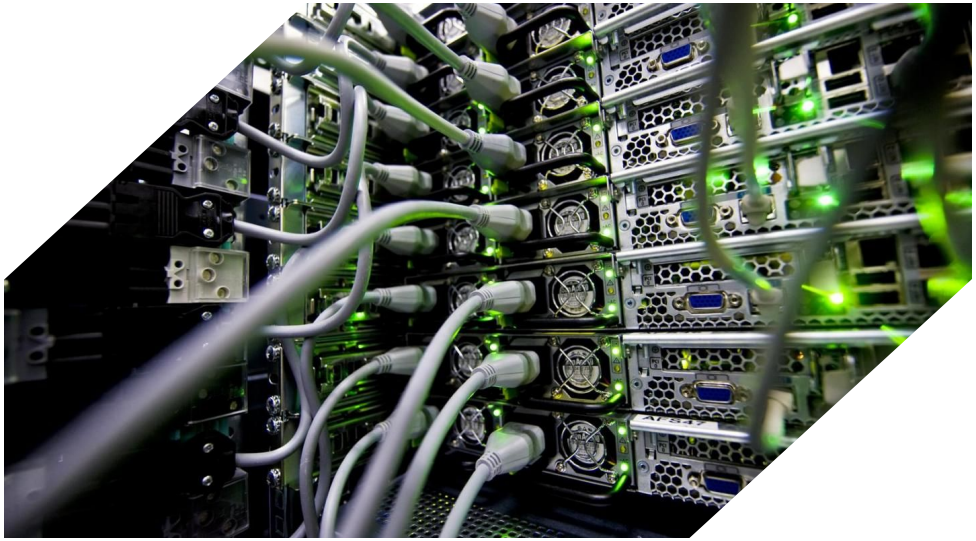

* - *Federated Identity Management*

Challenges

Service Payment Models

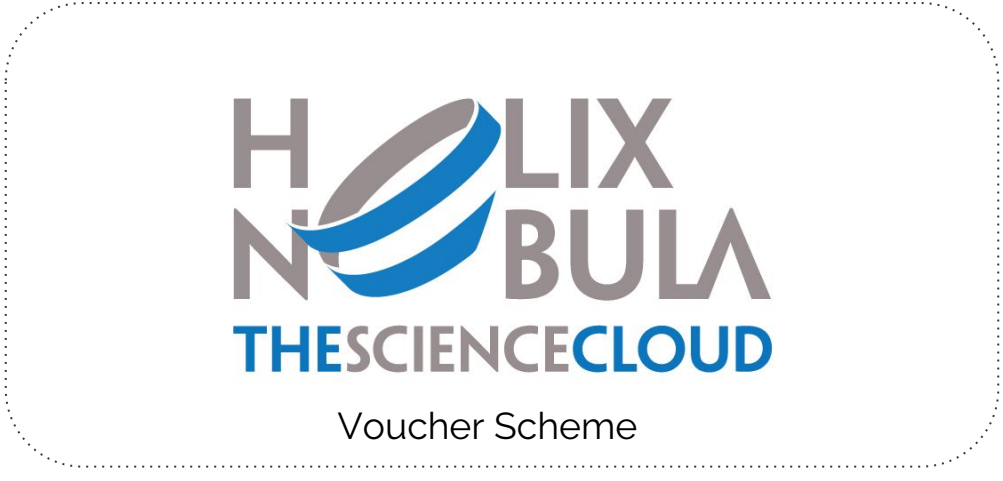

Example | Challenges

Preparation

- Analysis of Requirements
- Evaluation of current market offers
- Identification of relevant standards
- Formation of the stakeholder group
- Development of tender material

Project phases

Preparation

- Analysis of Requirements
- Evaluation of current market offers
- Identification of relevant standards
- Formation of the stakeholder group
- Development of tender material

Competitive Implementation & Sharing

4 designs

T*-Jul'16

* - *Tender*

Project phases

Preparation

- Analysis of Requirements
- Evaluation of current market offers
- Identification of relevant standards
- Formation of the stakeholder group
- Development of tender material

Competitive Implementation & Sharing

4 designs

3 prototypes

T*-Jul'16 CO**-Feb'17

Project phases

* - *Tender*

** - *Call-off*

Project phases

Preparation

- Analysis of Requirements
- Evaluation of current market offers
- Identification of relevant standards
- Formation of the stakeholder group
- Development of tender material

Competitive Implementation & Sharing

T*-Jul'16	CO**-Feb'17	CO**-Dec'17
4 designs	**3** prototypes	**2** pilots

* - *Tender*

** - *Call-off*

Scalability Testing

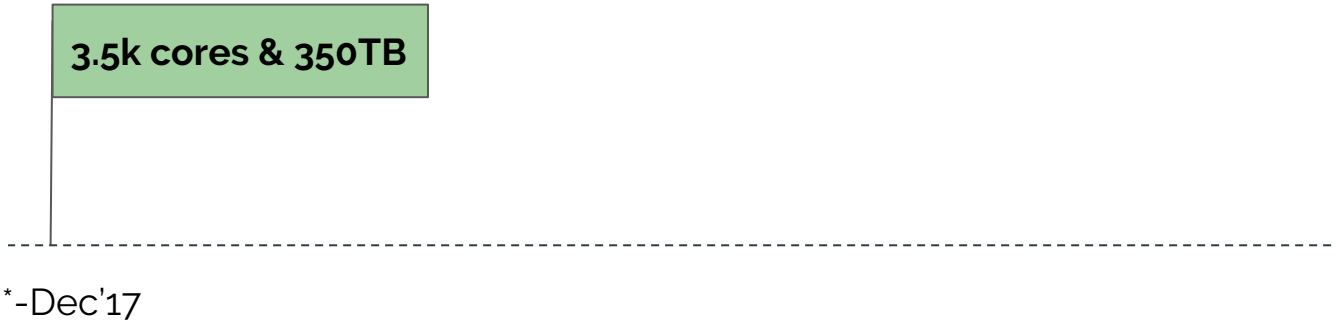

** - *Call-off*

Timeline

Per provider capacity

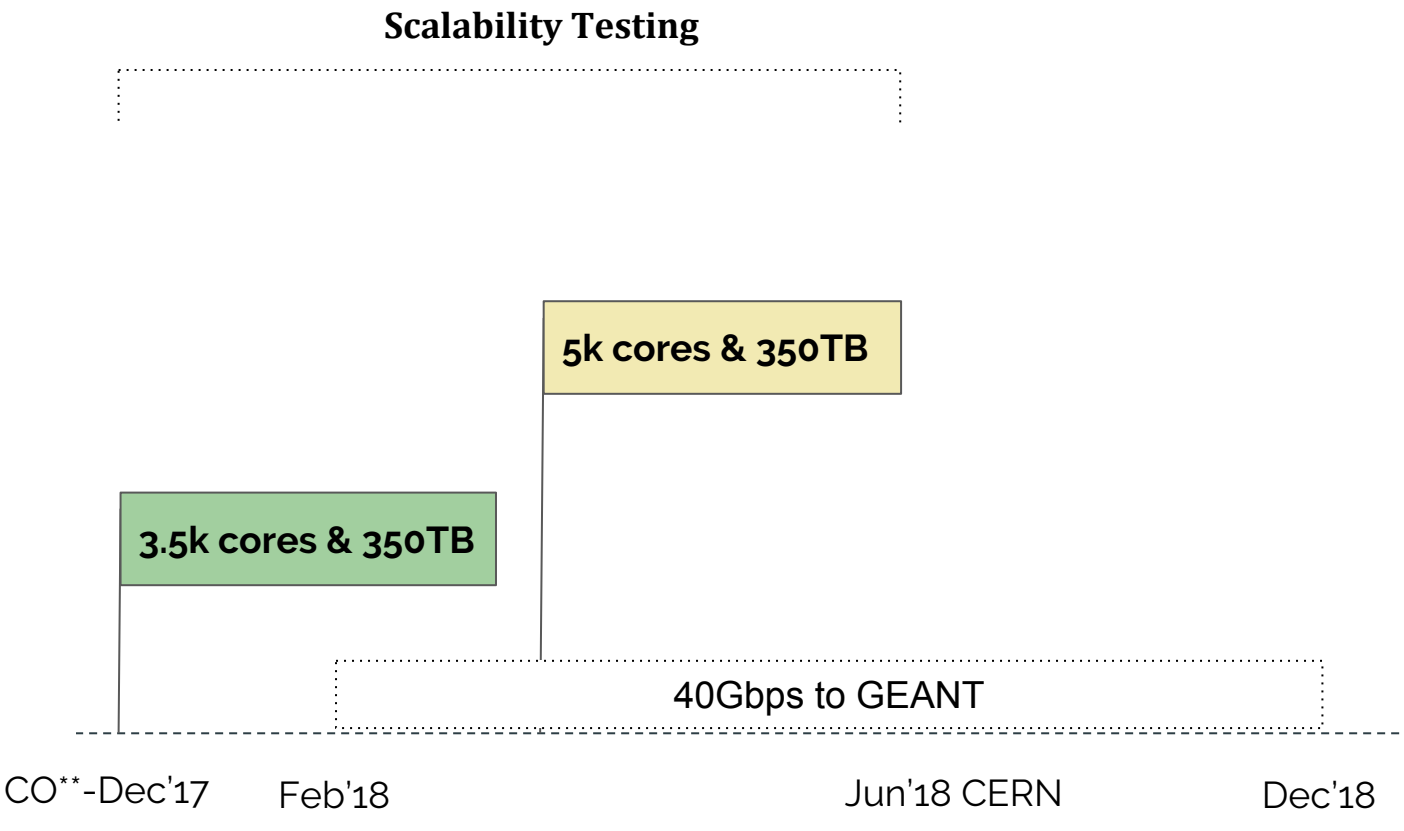

** - *Call-off*

Timeline

Per provider capacity

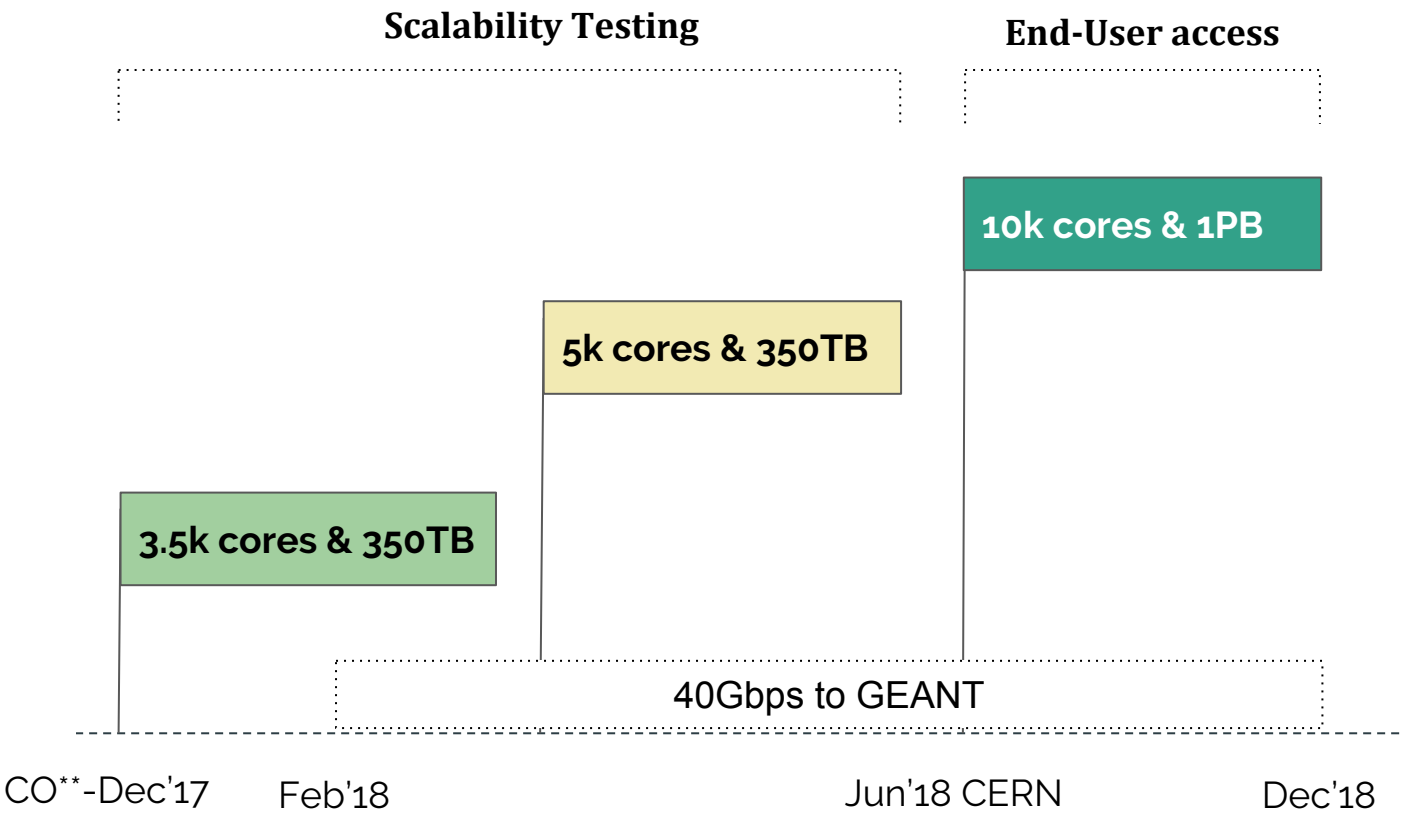

** - *Call-off*

Timeline

Per provider capacity

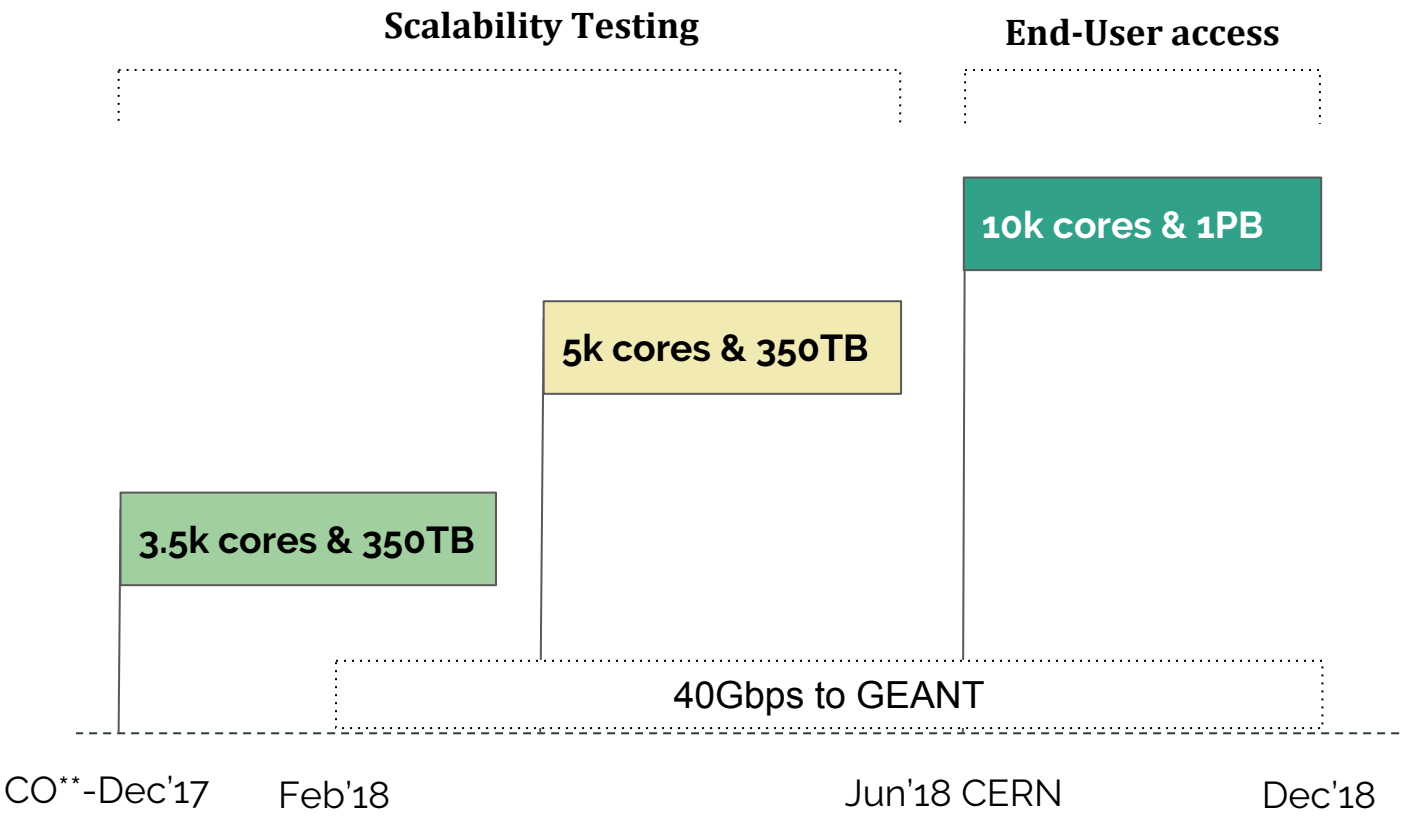

Future events

- Sep'18 SURFsara

** - *Call-off*

Timeline

Per provider capacity

What HNSciCloud is ***not***

- Zero-cost research grant

What HNSciCloud is ***not***

- Zero-cost research grant
- One-off test

What HNSciCloud is ***not***

- Zero-cost research grant
- One-off test
- Just a technical evaluation

What HNSciCloud is ***not***

- Zero-cost research grant
- One-off test
- Just a technical evaluation
- A “*walk in the park*”

What HNSciCloud is ***not***

Who are the commercial cloud providers?

Providers - IaaS

CERN use-case

Provisioning

Cloud Provider

openstack.

cloudstack

Provisioning stack

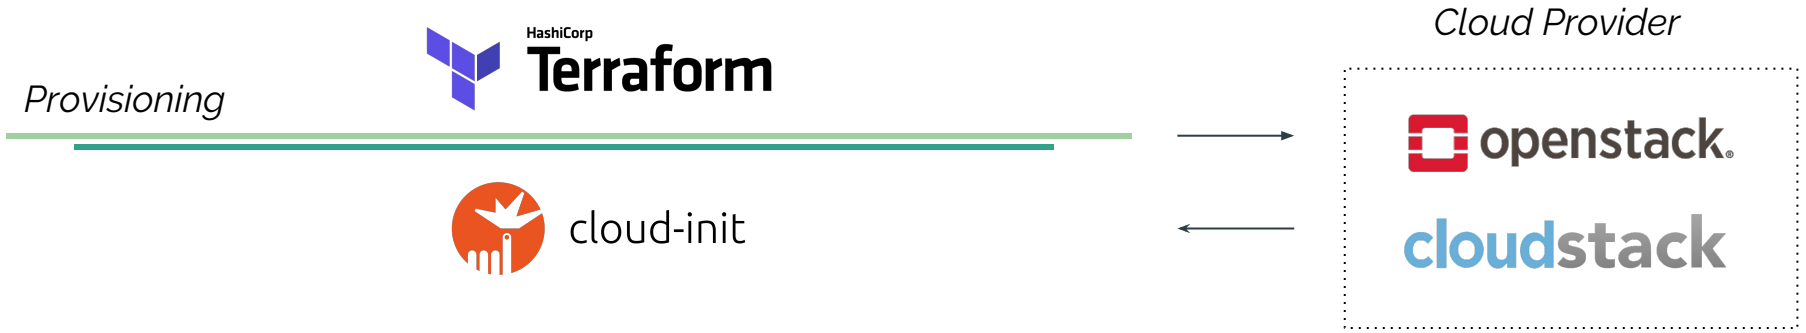

Provisioning stack

Provisioning stack

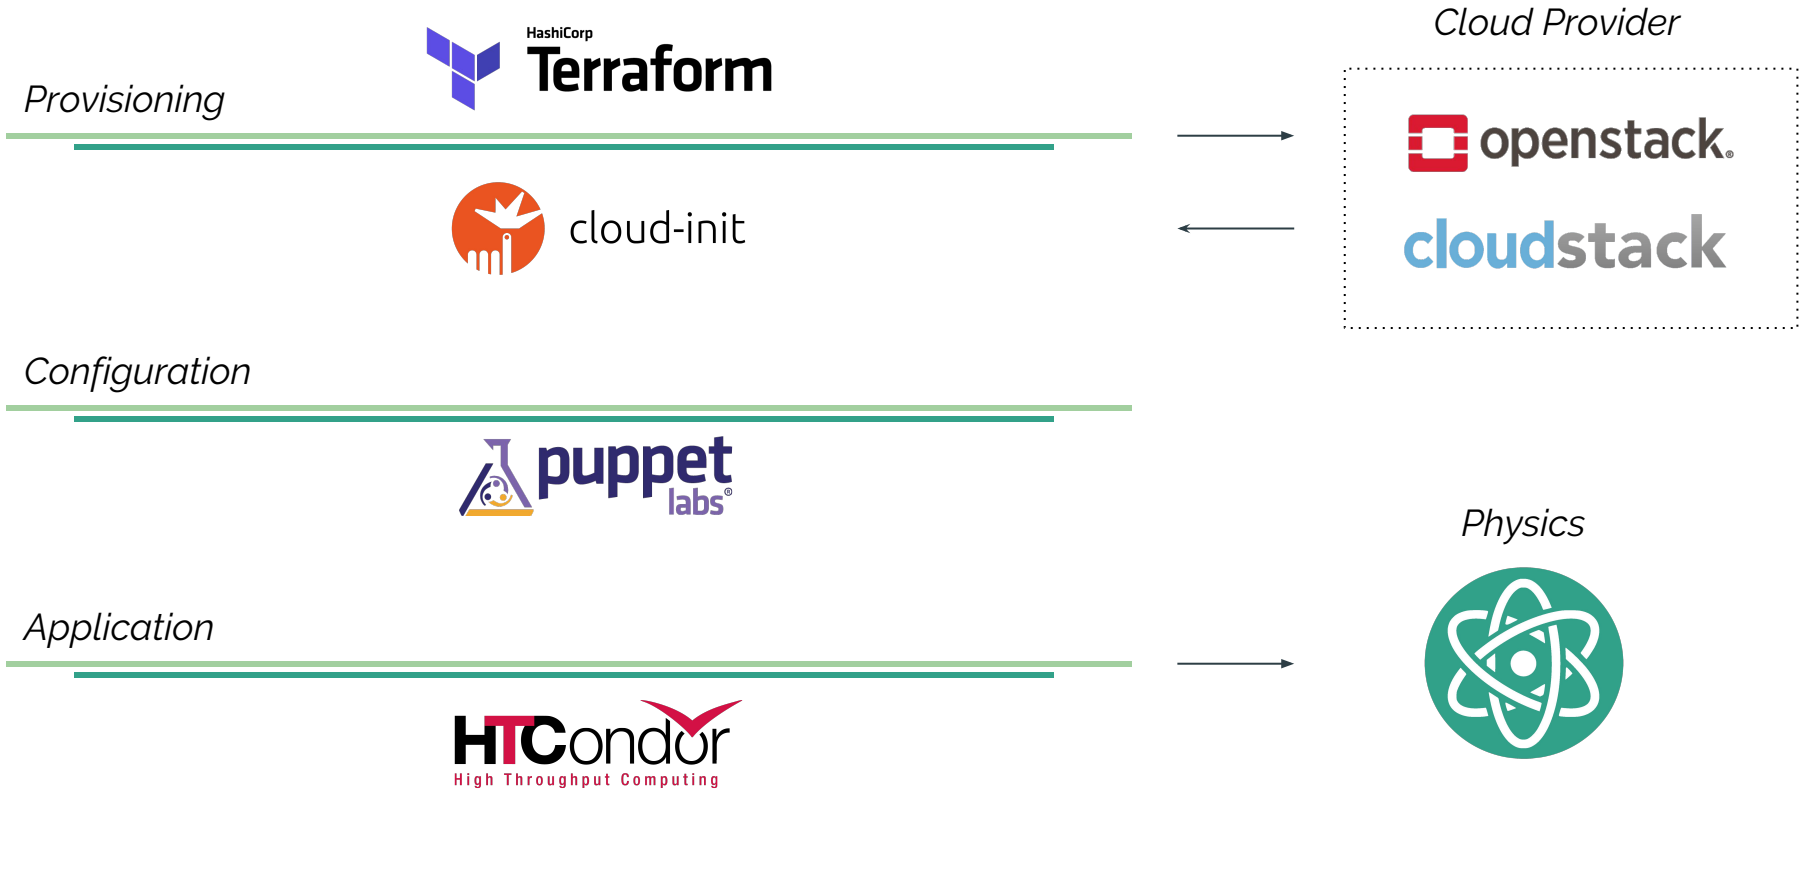

Provisioning stack

```
resource "cloudstack_instance" "resource_reference_name" {
  name                = "instance_name"
  service_offering    = "${var.flavor_name}"
  template            = "${var.image_id}"
  zone                = "${(count.index % var.az_factor) ? var.zone_1 : var.zone_2}"
  # Bootstrap
  user_data           = "${file(var.bootstrap_file)}"
  keypair             = "${cloudstack_ssh_keypair.keypair.name}"
  # Security
  security_group_ids  = ["${cloudstack_security_group.secgrp_xbatch.id}"]
}
```

Example Provisioning stack

Provisioning

HashiCorp Terraform

openstack.

```
resource "openstack_compute_instance_v2" "resource_reference_name" {
  name              = "instance_name"
  image_id          = "${var.image_id}"
  flavor_name       = "${var.flavor_name}"
  key_pair          = "${openstack_compute_keypair_v2.keypair.name}"
  user_data         = "${file(var.bootstrap_file)}"
  # Bootstrap
  metadata {
     ssl_pw         = "${data.external.data.result.ssl_pw}"
     certname       = "hostname_reference"
  }
  network {
    port            = "${openstack_networking_port_v2.port.id}"
    access_network = true
  }
}
```

Security is done on the "*port*" resource type

Example Provisioning stack

Future CERN use-case

kubernetes

“all things”

Evaluation Batch in Kubernetes

kubernetes

“all things”

- What to use for configuration management?

kubernetes

"all things"

- What to use for configuration management?
- How to handle live-updates?

kubernetes

"all things"

- What to use for configuration management?
- How to handle live-updates?
- HTCondor authentication per node, hypervisor, etc?

One potential solution

kubernetes

Configuration

puppet labs → HELM

Updated by

git

Example Batch in Kubernetes

kubernetes

Configuration

puppet labs → HELM

Updated by

git

Provisioning + Application

HTCondor High Throughput Computing → docker + HTCondor High Throughput Computing

Example Batch in Kubernetes

Job submission
WLCG use-case

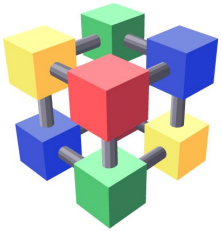

WLCG

Worldwide LHC Computing Grid

42 countries

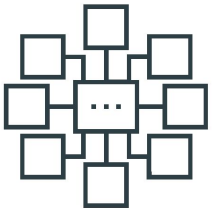

170 computing centres

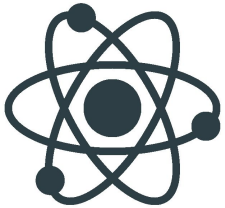

4 LHC experiments

ALICE

LHCb

CMS

ATLAS EXPERIMENT

What is the WLCG?

Submitter

Batch

Unconsolidated queue

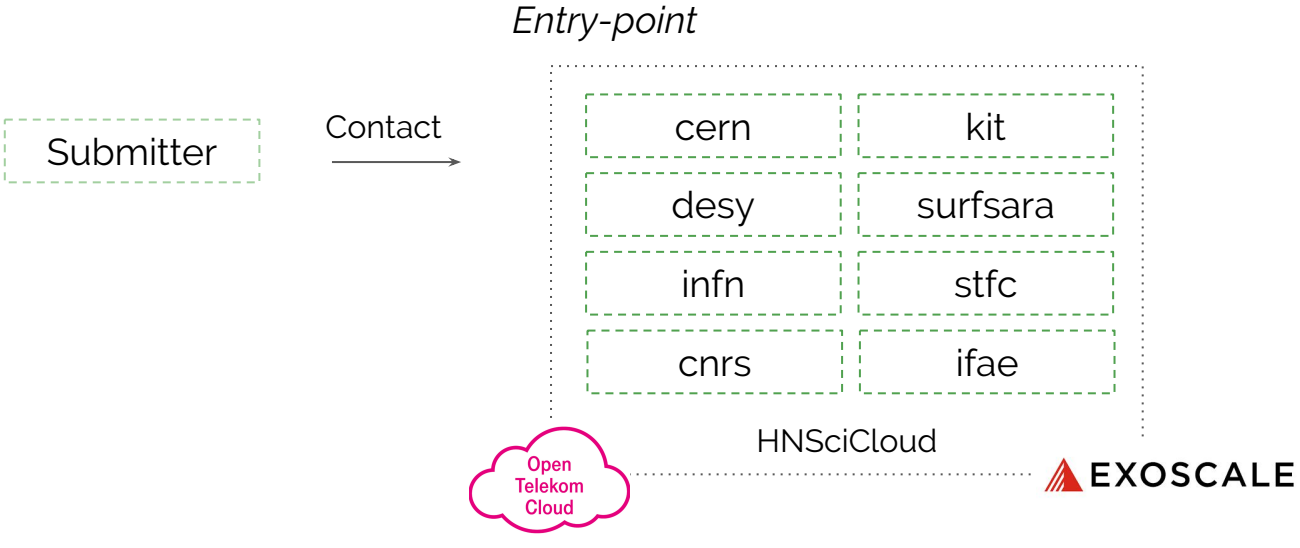

Batch

Unconsolidated queue

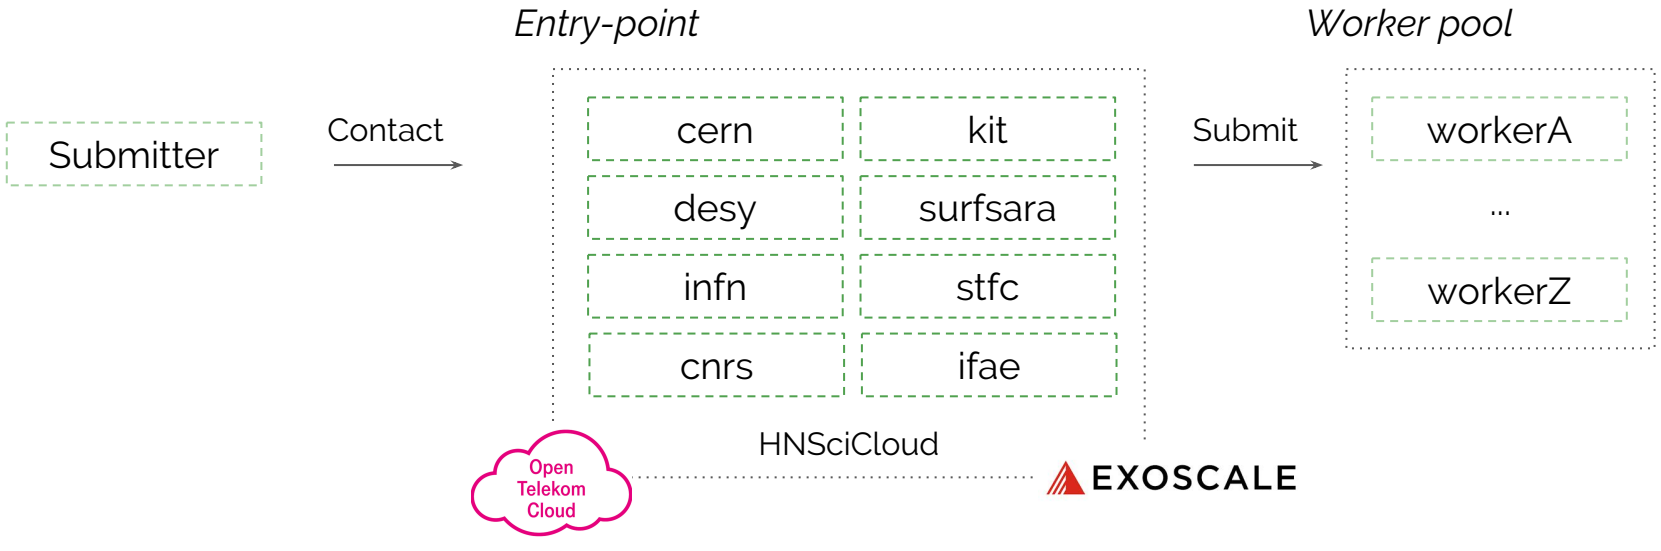

Batch

Unconsolidated queue

- Each submitter would need to define all the 8 entry-points to contact the HNSciCloud resources

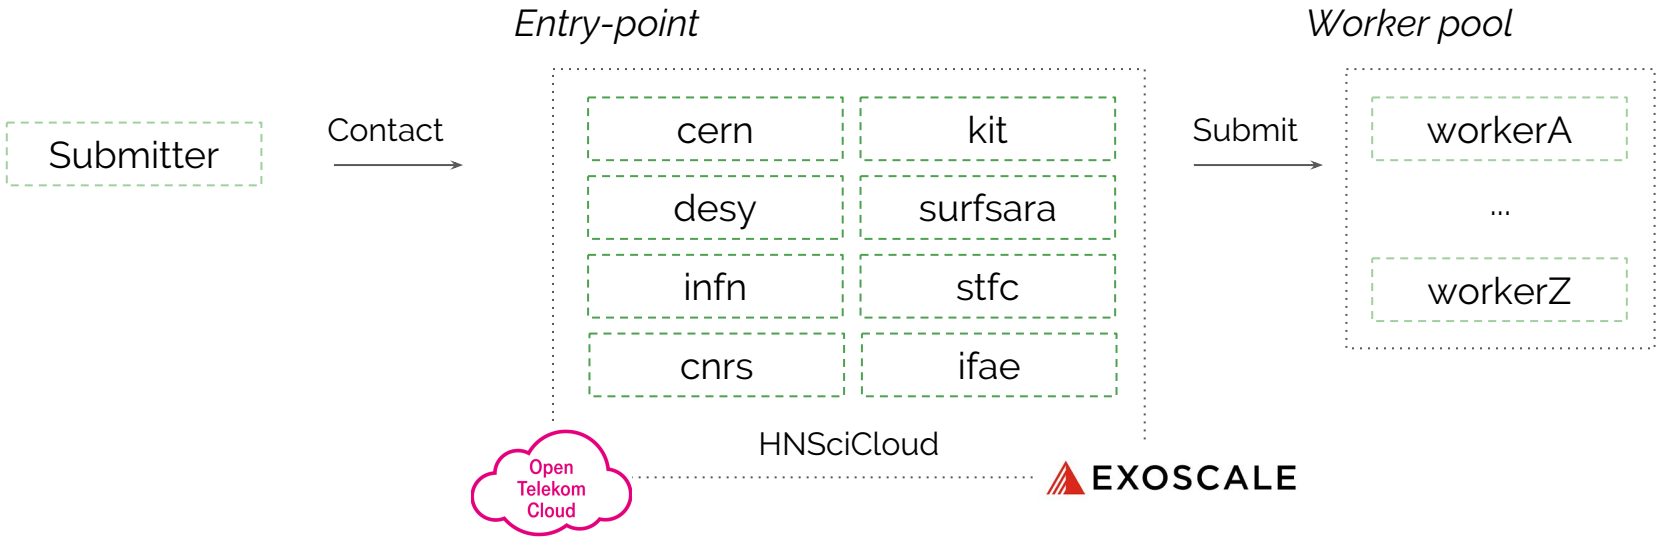

Batch

Unconsolidated queue

Solution?

Entry-point consolidation

- One entry-point can be used by all experiments to contact the HNSciCloud resources, sharing over 7 thousand cores with the 4 experiments across T-Systems and Exoscale

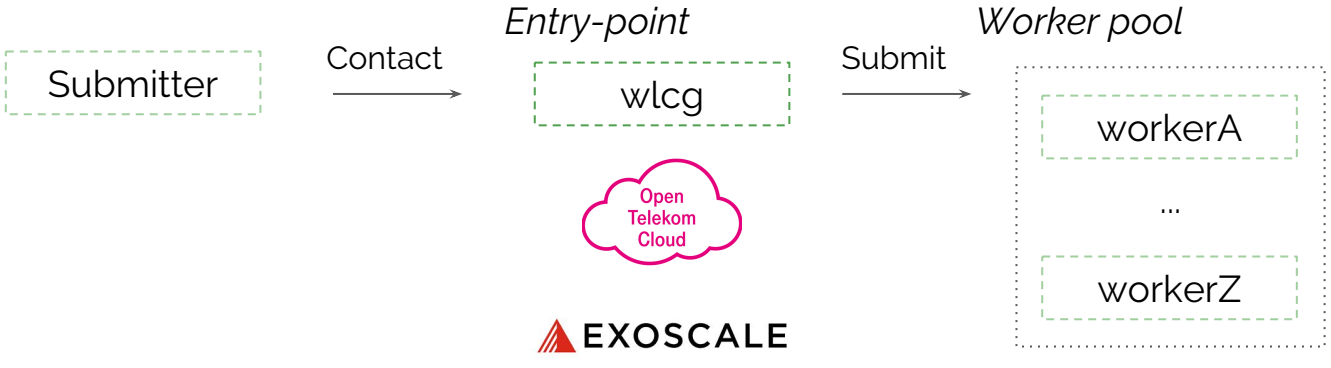

Batch | Consolidated queue

What‘s next?

Oct'24 - Hamburg

Desy

Sep'11 - Amsterdam

SURFsara

What's next?

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

What's next?

Thank you!