



---

## HNSciCloud PILOT PHASE AWARD

Alastair Pidgeon, RHEA

February 6<sup>th</sup> 2018

## OUR VISION

### Building an Open and European platform for science data exploitation

Only multi-cloud hybrid strategy can deliver:

- ✓ at scale
- ✓ in a cost effective manner
- ✓ diversity & choice

**Phase 3:**  
10,000 Cores/1PB  
**Designed for:**  
400,000+ cores /100PB+

Built on pure IaaS to **ensure no-lock-in** and **market pricing**, and with **minimum intrusion** into the Buyers' infrastructures

### Laying the foundations of the European Open Science Cloud (EOSC)

# OVERVIEW

## Established Partnership

**RHEA:** Prime contractor, proven management of large contracts for institutional customers, cyber security experts

**T-Systems:** Leading European Public Cloud provider working with CERN and other institutions, GÉANT connection

**Exoscale:** Dynamic, innovative cloud provider worked with CERN, GÉANT connection

**SixSq:** SlipStream/Nuvla; brokerage and automated multi-cloud application management

**Cyfronet:** Providing IT services to research teams and educational institutions and Onedata for Data Management

**Advania:** a High Performance Computing Cloud provider that offers flexible full HPC stack services and solutions.



---

# WORKING TOGETHER

## Network-aware, multi-cloud, agnostic environment

- ✓ We start from a production service:
  - ✓ innovate -> demonstrate -> collect feedback -> improve
- ✓ We build on established, market-leading European cloud services
- ✓ You enhance your community or private cloud with access to public cloud services (hybrid cloud)
- ✓ You create cloud native applications to take full advantage of the cloud, or
- ✓ You drop-in your application as is

... Any App to Any Cloud!

---

# CHALLENGES

## Deliver end-to-end optimised platform and data throughput

- ✓ Deliver throughput from **home institute to the cloud**
- ✓ Deliver throughput from **data source to the VM**
- ✓ Provide a **TRUE hybrid cloud** able to integrate Buyers Group private clouds or academic clouds (e.g. EBI Embassy, EGI federation)
- ✓ Support a **wide range** of applications and use cases
- ✓ Support **different data layout and access patterns**
- ✓ Integrate **containers and their management solutions** with ease
- ✓ Provide seamless user access and security through **single sign-on**

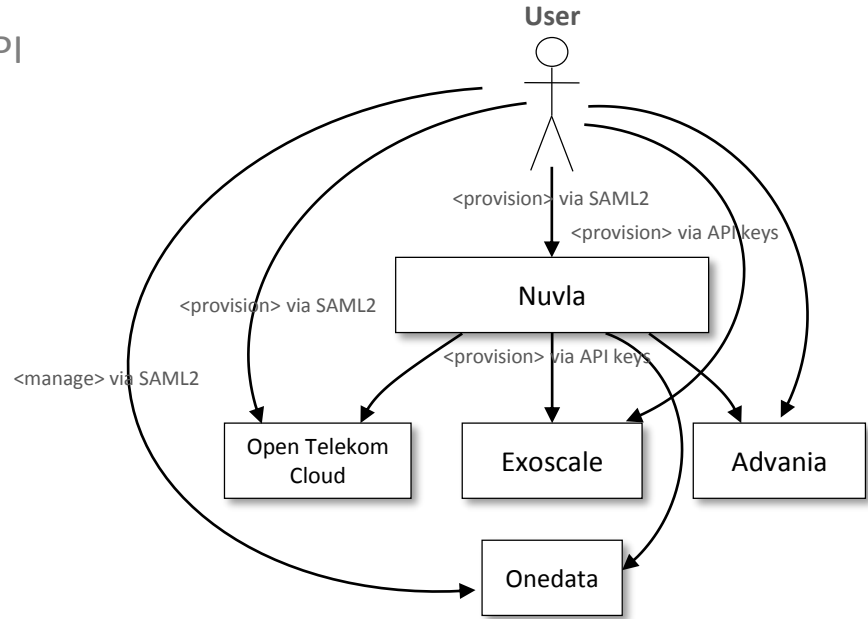
... with **minimum complexity and intrusion!**



# PUTTING ALL TOGETHER

## Complementary partners, coherent solution, flexible implementation...

- Support all connection means
  - Direct, libcloud, jcloud, terraform, native API
- Reporting, Accounting & Management
  - By User, Project, Organisation
- Service Level Agreement Monitoring
  - Define SLA inspired by SLALOM
- Identity & Access Management
  - eduGAIN & Elixir integration
  - GÉANT Data Protection Code of Conduct
  - Browser, command line and API access
- Unified Data Management



# IAAS CLOUD SERVICES & DATA MANAGEMENT

## Proven, innovative cloud service providers

### Exoscale

- Apache CloudStack
- Data centers in Switzerland, Germany, Austria
- Optimized for cloud-native applications



### Open Telekom Cloud from T-Systems

- OpenStack and open standard APIs for all functions
- Public Cloud, data centers in Germany, certified and audited
- Established, high-bandwidth network connectivity throughout Europe
- Scientific use tested and verified by CERN



### IronCloud by Advania Data Centers

- OpenStack
- Data centers in Iceland, powered by green energy
- HPC purpose-built bare metal cloud with high speed interconnect and high performance guarantees.
- Full HPC stack support from seasoned HPC specialists.



### Cyfronet

- Onedata integration and Data Management



# HPCAAS

## Cloud and dedicated HPC options

- ✓ Providing optimised performance in virtualised environments
- ✓ IaaS based HPCaaS with:
  - ✓ Bare metal HPC nodes with high speed interconnect
  - ✓ “Super VMs” from clusters of VMs with 24 vCPU (e.g. more than 256 vCPUs, more than 1TB RAM in a cluster)
  - ✓ Infiniband
  - ✓ GPU access
- ✓ Access through GÉANT





---

# MULTI-CLOUD AND HYBRID CLOUD MANAGEMENT

## Fulfilment mechanism to facilitate and organise cloud consumption:

- ✓ **Choice of cloud based on price**, performance, service guarantees, etc. **without vendor lock-in**
- ✓ **Advanced quota management** over multi-tenant user organisations
- ✓ **Transparent interoperability** with your own community and private clouds
- ✓ Fast deployment of new features and fixes supporting **rapid innovation cycle**
- ✓ Optional **Automated deployment** and dynamic management of applications to **minimize costs and maximize performance**
- ✓ **AppStore** for curated applications and user applications




# LARGE DATA SETS

## Provide access while protecting data...

Deliver POSIX and Object Store access to wide range of data sets

- ✓ Understand scalability, performance
- ✓ Integrate existing tools/services
- ✓ Optional: Direct Object Store access
- ✓ Meet data protection norms
- ✓ GDPR compliance
- ✓ ISO 27xxx series security



---

# GÉANT NETWORK CONNECTIVITY

## Efficient data access relies on fast networking...

- ↪ T-Systems connected via DFN NREN (Frankfurt) and SURFNET (Amsterdam) (20x2 Gbps)
- ↪ Exoscale connected via NRENs in Switzerland, Germany and Austria (20/10/10 Gbps)
- ↪ Advania will be connected to Géant via Icelandic NREN Rhnet (10 Gbps)
- ↪ Aggregated bandwidth delivered is above 40 Gbps

# CONTAINER SUPPORT


## Reducing complexity and overheads

- ✓ Recipes to facilitate scalable deployments of standard container management tools
  - ✓ Mesos, Kubernetes, Docker Compose, ...
  - ✓ Supported in the AppStore
- ✓ Provide transparent access to virtualization technologies
- ✓ Additional tools possible

**... your feedback needed!**



MESOS



docker

---

# SECURITY

## Analysing (and targeting) ISO 27000 security compliance for the pilot platform

- ↪ Cloud resource providers need to own various security certifications (from Cloud Security Alliance, CSA)
- ↪ Security needs to be assessed also for the overall pilot (including Nuvla and Onedata):
  - ↪ ISO 27001 chosen source of our pilot platform security baseline
  - ↪ Security Risk Assessment to identify applicable security control points
  - ↪ Definition of routine and emergency procedures to decrease vulnerability
  - ↪ Definition of a security organization for responding to incidents
  - ↪ Security auditing to assess the pilot platform compliance and identify critical areas

---

# BUSINESS MODEL & COMMERCIALISATION

## A Proven Track Record

- ↪ Previous R&D developments commercialised: Software Services & Products
- ↪ Widest possible cloud choices, **avoid vendor lock-in**
- ↪ Evolution of Helix Nebula Marketplace, incorporating lessons learned
  - ↪ Buyers group need to **consider new procurement models** to realise the full benefits of **“as a Service” cloud delivery** –largest discounts from **reserved VM instances, upfront payments**
  - ↪ Support **“Bring your own contract”** concept to coordinate cloud resource consumption
- ↪ Exploring managed services – e.g. Onedata, SLURM batch cluster, Docker Swarm, Kubernetes
- ↪ **Ready to offer this service to science communities on a commercial basis during Phase 3**

---

## GOING FORWARD

### We deliver a flexible and commercial platform for science data exploitation

- ✓ Solution brings together **production products and services** into coherent analysis platform
  - ✓ Established players in cloud ecosystem – open to included others in the future
  - ✓ Proven delivery within Helix Nebula
- ✓ **Extend existing capabilities** to meet the needs of large-scale, scientific data analysis
- ✓ Provided “as a Service”
- ✓ Consortium **led by SMEs delivering innovation at scale**
- ✓ Establishing an **open foundation for the European Open Science Cloud**