





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*

The group of procurers committed procurement funds, manpower for testing/evaluation, use-cases with applications & data, and in-house IT resources

Resulting made available to end-users from many research communities



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>

☛ *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



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



Cloud Providers

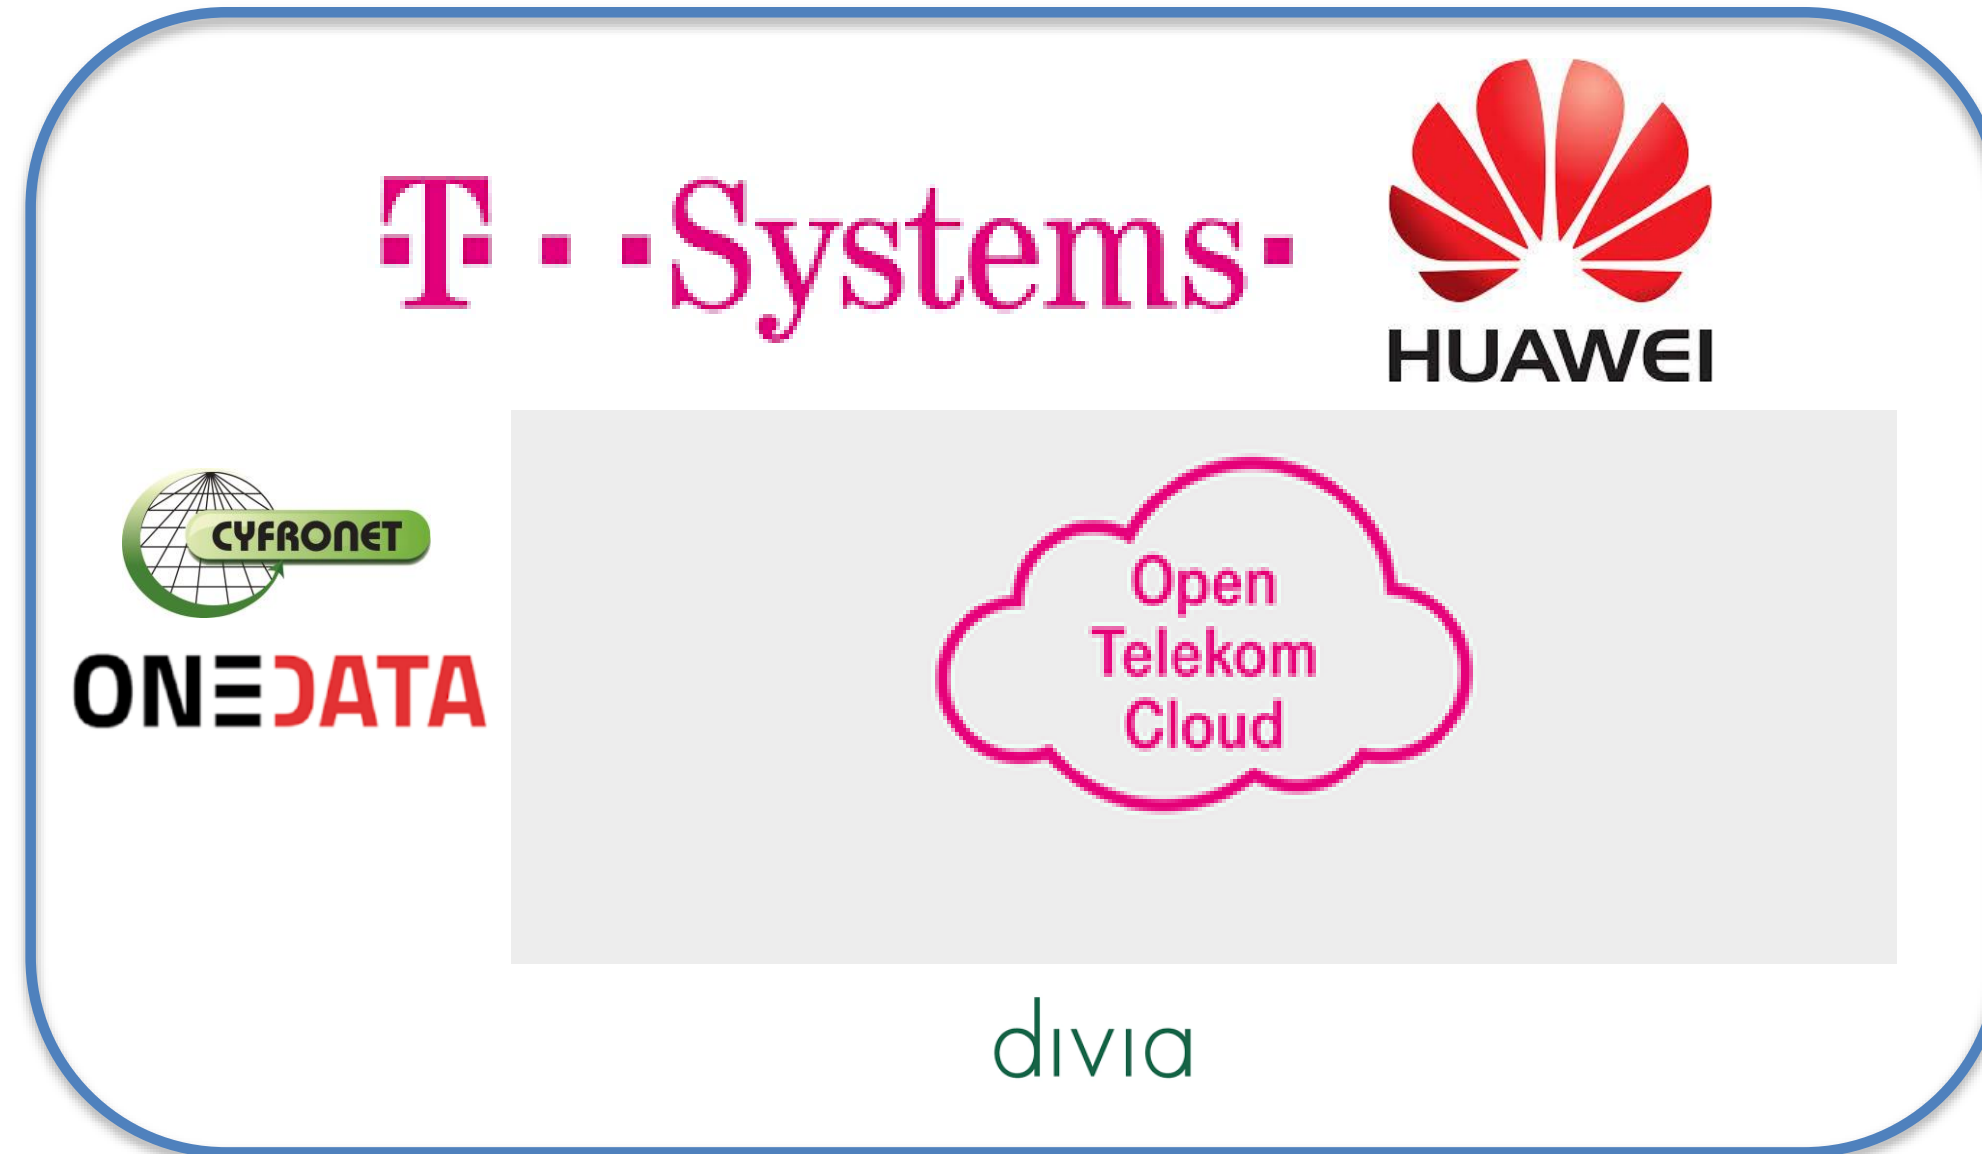
- **T-Systems**

- *IaaS based on OTC*



- **RHEA**

- *IaaS provided by Exoscale*





Initial Deployments



Buyers



Buyers	Tested features			Fields of Research					
	AAI Federated Identity	Storage blocks/objects	Computing capacity (VMs)	Photon / Neutron Science	High Energy Physic	Astronomy	Life sciences		
				FDMNES	CrystFEL	DODAS	LOFAR	PanCancer	WeNMR/HADDOCK
CERN	● ● ●	● ●	● ● ●						
CIFS	● ● ●	● ● ●	● ●						
DESY	● ● ●	● ●	●						
EMBL	● ● ●	● ●	● ●						
ESRF	● ● ●	●	●						
IFAE	● ● ●	●	●						
INFN	●	●	● ● ●						
KIT	● ●	●	●						
SURF SARA	● ● ●	● ●	●						

● useful, they can ease the way I do my job ● ● relevant, they can really improve the way I do my job ● ● ● strategic, they are fundamental to execute my job

See [booklet](https://www.hnscicloud.eu) on <https://www.hnscicloud.eu>



Some more deployed use cases



Integration of commercial cloud capacity in batch services for the LHC experiments

On demand computing facilities generation

- CMS Data Reduction Facility – Physics Analysis with Apache Spark
- Hybrid Cloud auto-scaling with Kubernetes
- DODAS/Lightweight WLCG sites deployments
- Interactive user analysis services

Data Management and transparent data access

- Pancancer (EMBL-EBI)
- HDF5 (DESY)
- Grid sites deployment (INFN)

HPCaaS

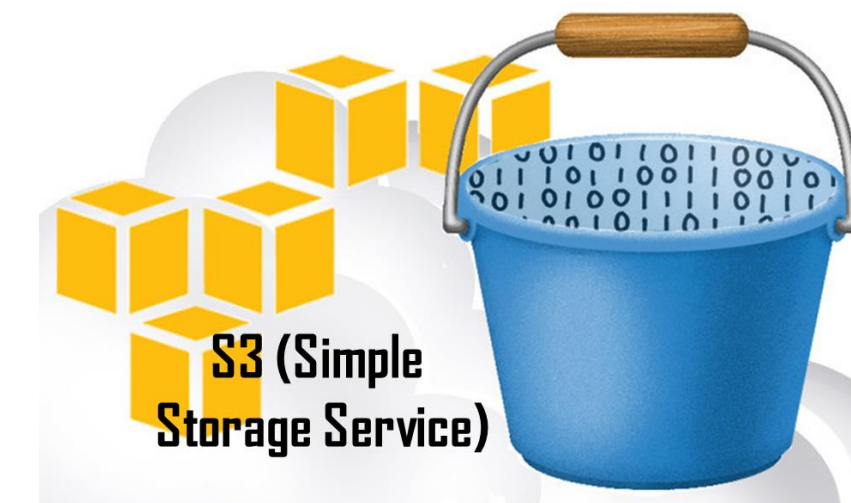

- FDMNES (ESRF)

S3 object stores

- Hybrid S3 services for data replication using Ceph
- Use S3 for preparatory analysis jobs

Machine Learning and Deep Learning for Simulation

- Scale out model training for Neural Network optimization on GPUs
- Extend to other hardware accelerators (FPGAs); Generalise the approach to satellite imagery analysis and medical applications


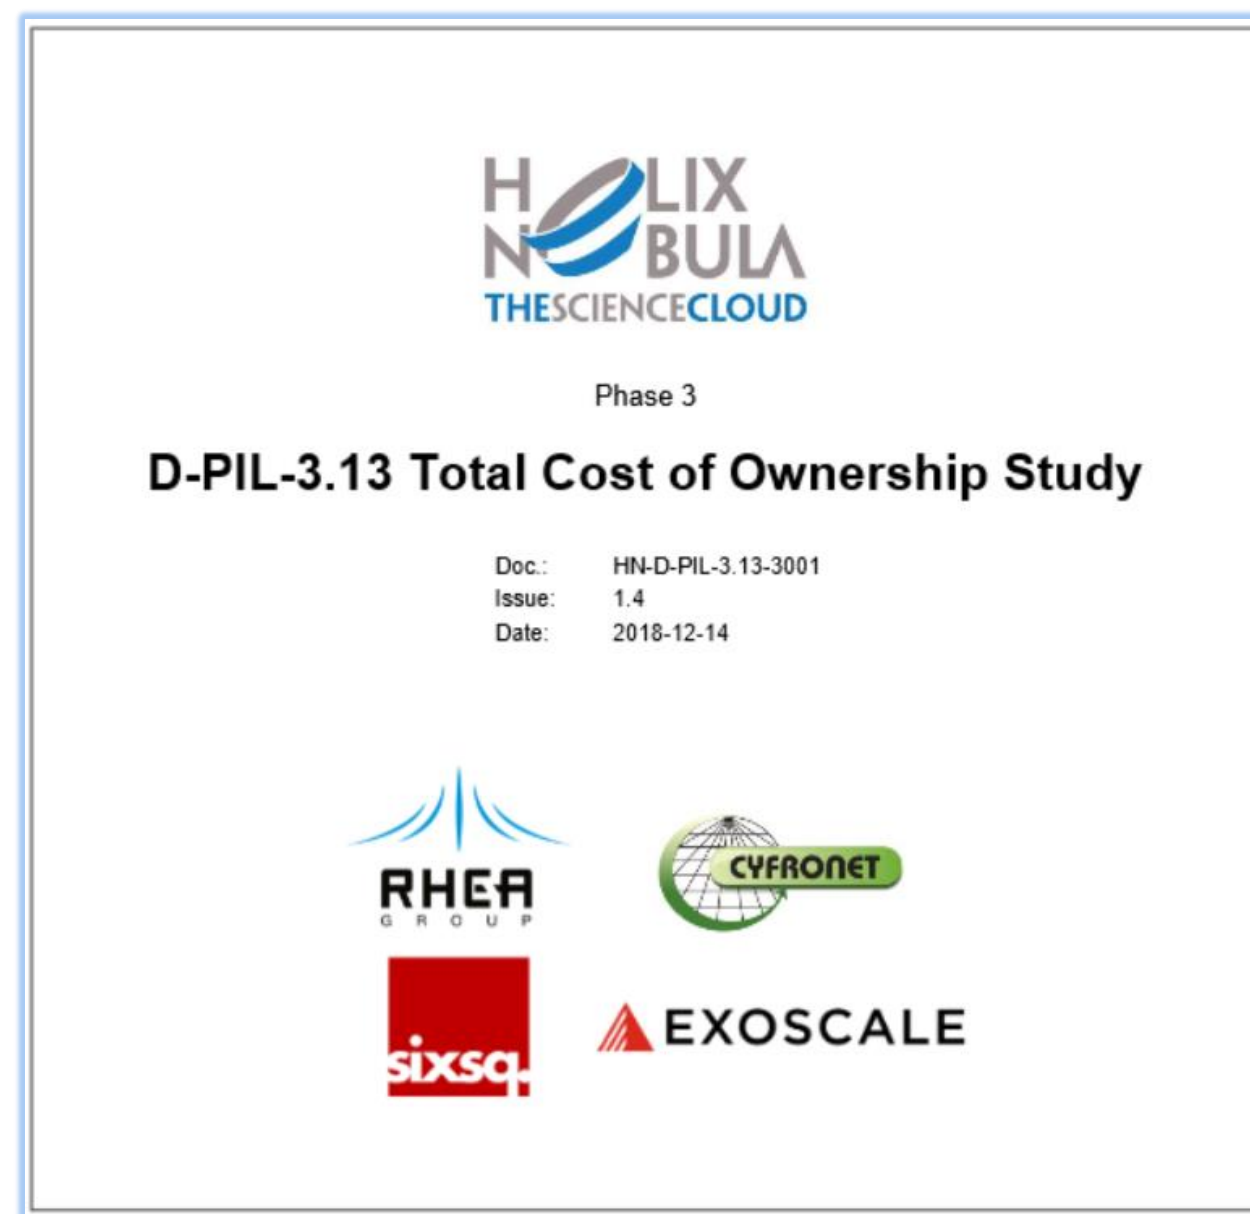




Total Cost of Ownership study



- Understand the costs of using commercial cloud services as part of a hybrid cloud model
- 2 use-cases selected with different requirements
 - **ALICE**: single core jobs, up to 50.000 at any time (monte-carlo, reconstruction, analysis)
 - **PANCANCER**: burst pattern, with minimal resources constantly used (few VMs) and periods of ramp-up (up to 400 VMs)

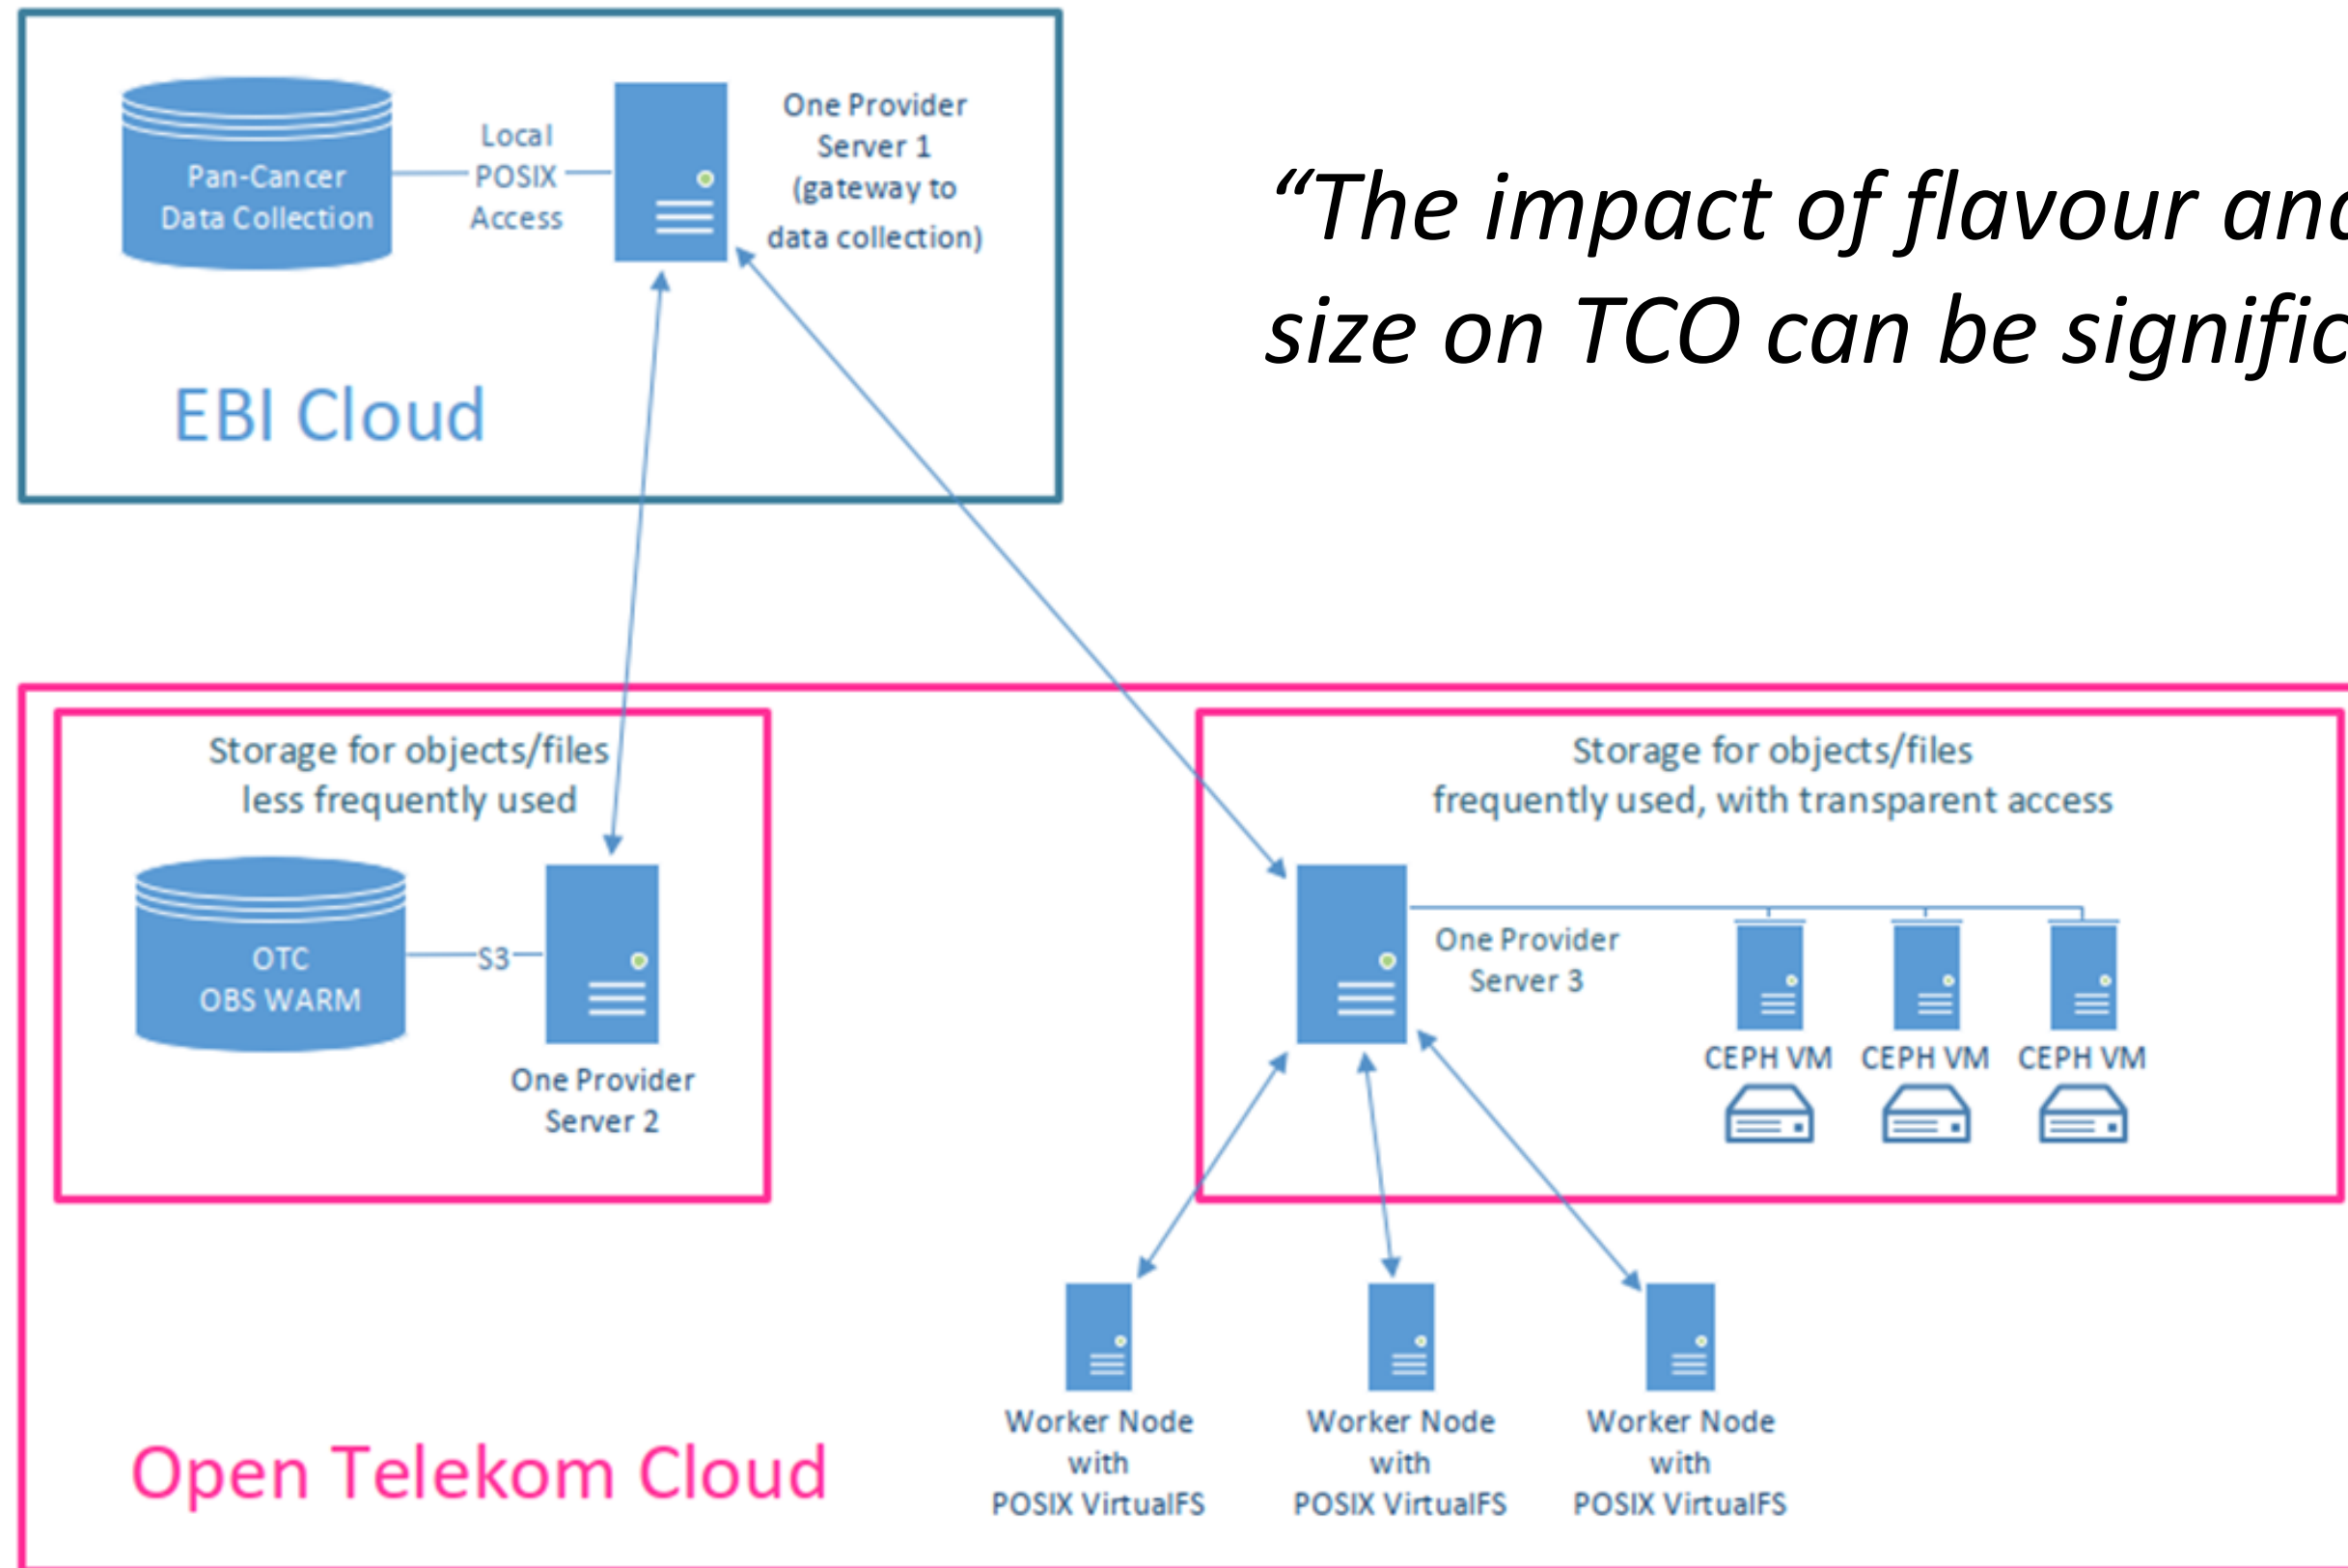


<https://doi.org/10.5281/zenodo.2605187>

<https://doi.org/10.5281/zenodo.2605176>



TCO results: T-Systems



“The impact of flavour and VM size on TCO can be significant.”

Figure 1: Pan-Cancer Data Management scenario

Table 5: ALICE job costs compared between CERN cloud and OTC public cloud

Job Type	TCO Benefit OTC Public Cloud
Monte Carlo	
Reconstruction	
Analysis	

Table 4: Pan-Cancer qualitative factors

Factor	Importance (H, M, L)	Effectiveness	Notes	Additional Information
Agility	H	+1	The solution provides easy and quick deployment and changes to the processing and storage functionality	
Contract Review and Negotiation	M	+1	Can be implemented with minimum review	
Elasticity and Scalability	H	+1	Solution facilitates easy and quick expansion of available processing and/or storage capacity	
Regulatory and Policy Requirements	H	+1	Solution adequately enables compliance with external regulations	User data can be stored encrypted in the public cloud. Through Onedata, data privacy can be maintained between public, restricted and confidential data. Confidential data would only be hosted on-premise.
Security	H	+1	Solution provides effective mechanisms by which constantly escalating security threats are prevented and security events or breaches are constantly monitored	Cloud service is certified e.g. for relevant ISO standards, EU GDPR and German national regulations.

Factor	Importance (H, M, L)	Effectiveness	Notes	Additional Information
Service Levels	M	+1	Service level targets were not provided. The solution would be able to match or improve the current SLA based on on-premise services. Especially with regards to performance, the cloud offers more granularity and diversity to tune performance to requirements when compared to on-premise resources.	Service Availability of 99,95% is achievable when using 2 availability zones and load balancing. Support is provided 24/7.



TCO results: RHEA



Table 15: TCO Cost Summary for Each Use Case

Use-Case	Exoscale	AWS	CloudFerro
PANCANCER (no storage)	123,170	147,575	155,247
PANCANCER (with storage)	239,562	-	-
ALICE (all three use-cases) – 50000 jobs*	6,216,252	7,907,449	9,282,265
ALICE Monte Carlo – 1 job**	0.082	0.106	0.125
ALICE Raw Data Reconstruction – 1 job**	0.096	0.138	0.159
ALICE Analysis Trains – 1 job**	0.038	0.034	0.042

“The PANCANCER and ALICE use-cases each have their own specific requirements for cloud deployment. Both use-cases have more than one job type requiring different VM flavours and/or pricing options. Also, use-cases can be supported without having to store large volumes of data in the cloud, minimising storage costs. Therefore, we have not included the cost data management solutions, since our assumptions for use-cases are that the data can be ‘streamed’ to the VMs. “



Voucher Scheme

Experience in HNSciCloud





Experience from HNSciCloud



- Needed a simple & flexible way to distribute part of the procured capacity to end-users

Encourage the uptake of services deployed in HNSciCloud by Long Tail of Science (LToS)

Perform R&D on cloud providers infrastructure & to **validate** resources before wider use

- Explored use of vouchers in the pilot phase

- Voucher providers:  **EXOSCALE**

- Paper with experience gained published in Zenodo: <https://doi.org/10.5281/zenodo.2615456>



The Process



Definition of scheme characteristics

- Face value: 250€
- Validity: 1 Year
- To be redeemed under existing tenant

Tests by the Buyers Group

- Usage monitoring
- Data repatriation
- Access to the resources after credit exhausted with additional vouchers

Distribution to end-users

- To individual researchers
- Selected by external organisation (Eurodoc)
- **Feedback via survey form**

eurodoc
The European Council of Doctoral Candidates and Junior Researchers



Lessons learned

- Based on the feedback received
- **Apply it in OCRE**




Essential Features of the Voucher Scheme




Simplified user interface & Up to date documentation and trainings




Cost Calculators



Compatibility with Data Management Plans







Clear data repatriation policies & Long-term data storage solution





Lessons Learned



- Framework agreements provide a convenient structure for service procurements in the scientific community*
- Volume and requirement aggregation across a group of scientific organisations brings advantages ->  Open Clouds for Research Environments*
- Use a small number of cloud providers in parallel to avoid lock-in and ensure service continuity*
- Need a scientific test validation suite and need to repatriate data at the end of contracts ->  Open Clouds for Research Environments*
- Commercial clouds offer opportunities to rapidly scale cutting edge technology for R&D ->  Open Clouds for Research Environments*
- Vouchers/Credits are a practical means to provide limited-scale access to commercial cloud services for end-users ->  Open Clouds for Research Environments*
- Commercial clouds providers offer services that are certified against international standards (e.g. ISO 27000 family) and consistent with legislation (e.g. GDPR)*

OCRE

Open Clouds
for Research
Environments

Stimulate the adoption of
commercial cloud and Earth
Observation services by the
European Research community

Driven by

GÉANT



RHEA
GROUP

Trust-IT Services
Communicating ICT to markets

OCRE will become the procurement vehicle of the European Open Science Cloud (EOSC) and provide opportunities for research communities to adopt commercial cloud services as part of their IT computing strategy

Challenge

Basic Scientific Data
Archiving & Bit
Preservation



Data volumes and complexity growing; Communities Growing; Needs of advanced functionality increase

Objective: perform R&D to demonstrate long-term preservation and archiving services for scientific data in the PB range under F.A.I.R. principles, using commons solutions, while research groups keep total stewardship of their data sets

- ☁️ OAI, FAIR and European Regulation (GDPR)
- ☁️ European Open Science Cloud (EOSC) integration
 - ☁️ Reduce fragmentation expanding existing solutions to several scientific domains
 - ☁️ Access through GÉANT network; Federated AAI support
- ☁️ Scalability: PB data volume of multi domain scientific data; High sustained ingest rates (1-10 Gbps)
- ☁️ Cost-effective and transparent business models: range of purchasing and deployment options, stimulating uptake by the wider research community



Accelerating Science