SOSC 2018, Perugia

MORE SCIENCE, MORE CLOUD - COMMERCIAL CLOUD PERSPECTIVES

T · · Systems ·

SPEAKER

JURRY DE LA MAR

Account Director Research and Aerospace
Ph.D. Physics, Free University Amsterdam

AGENDA

01 INTRODUCTION

02 DIGITIZATION, CLOUDS AND T-SYSTEMS

03 RELEVANCE OF CLOUDS IN SCIENCE

04 OPEN TELEKOM CLOUD – SCIENCE USE CASES AND ACHIEVEMENTS

05 LESSONS LEARNED

T-SYSTEMS – A LEADING ICT PROVIDER THROUGH GROWTH

T-SYSTEMS

FOUNDED BY TWO STRONG PARTNERS

GOAL

Deliver one-stop, best-in-class ICT solutions

Extensive industry- and company-specific expertise due to taking on employees from a variety of organizations

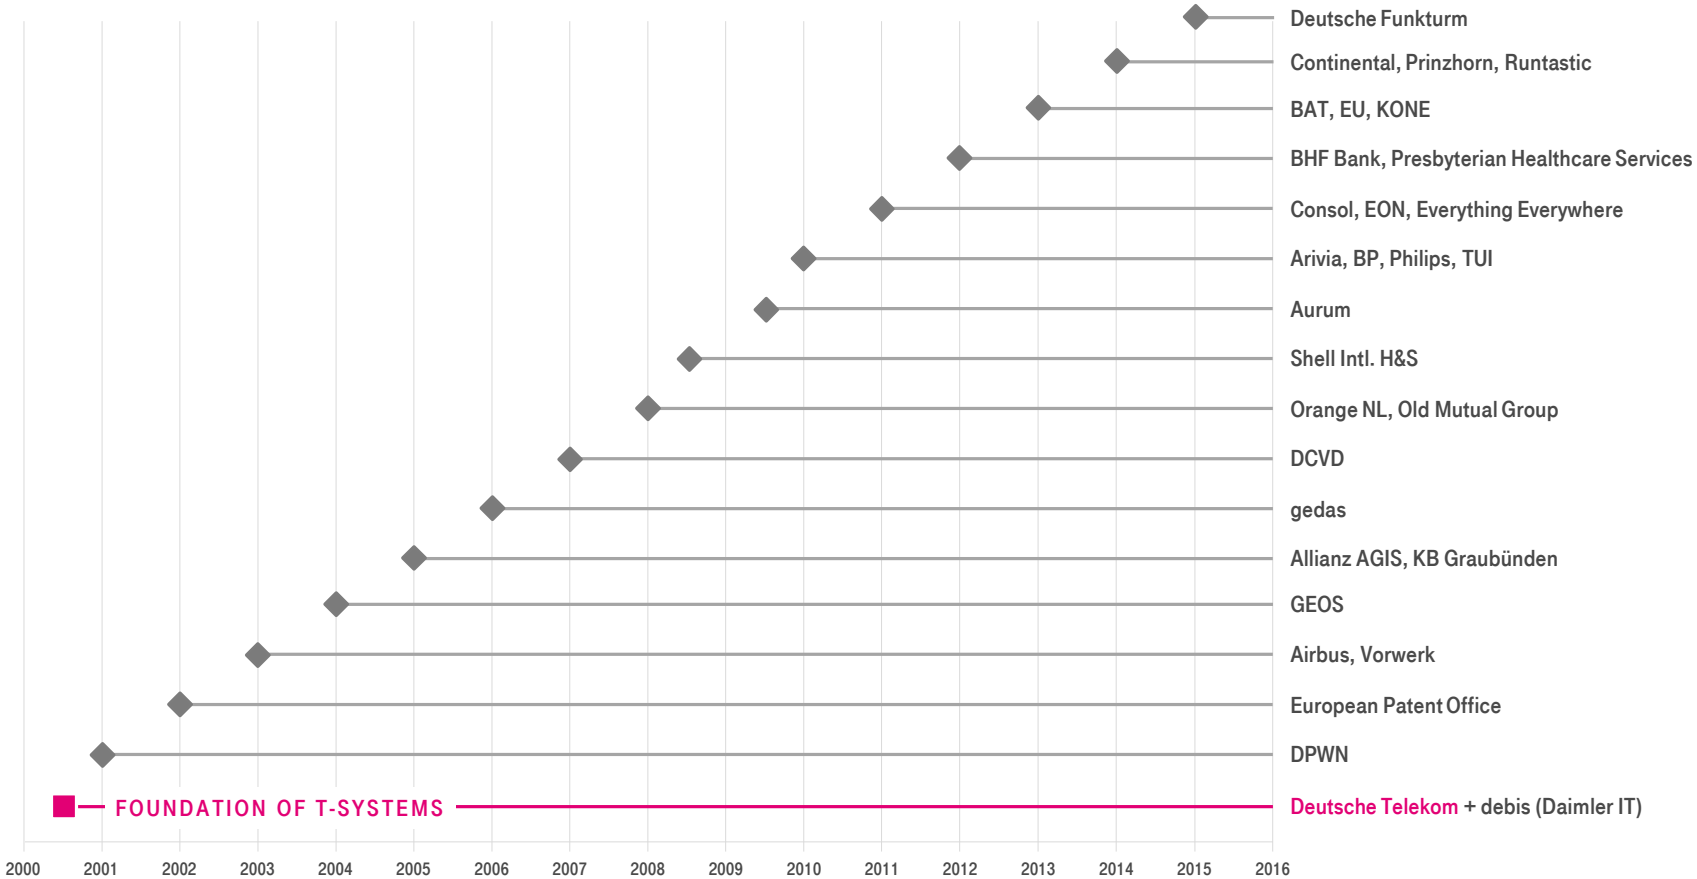

T-Systems

T-SYSTEMS – DEUTSCHE TELEKOM'S SUBSIDIARY FOR LARGE CORPORATIONS AND PUBLIC SECTOR

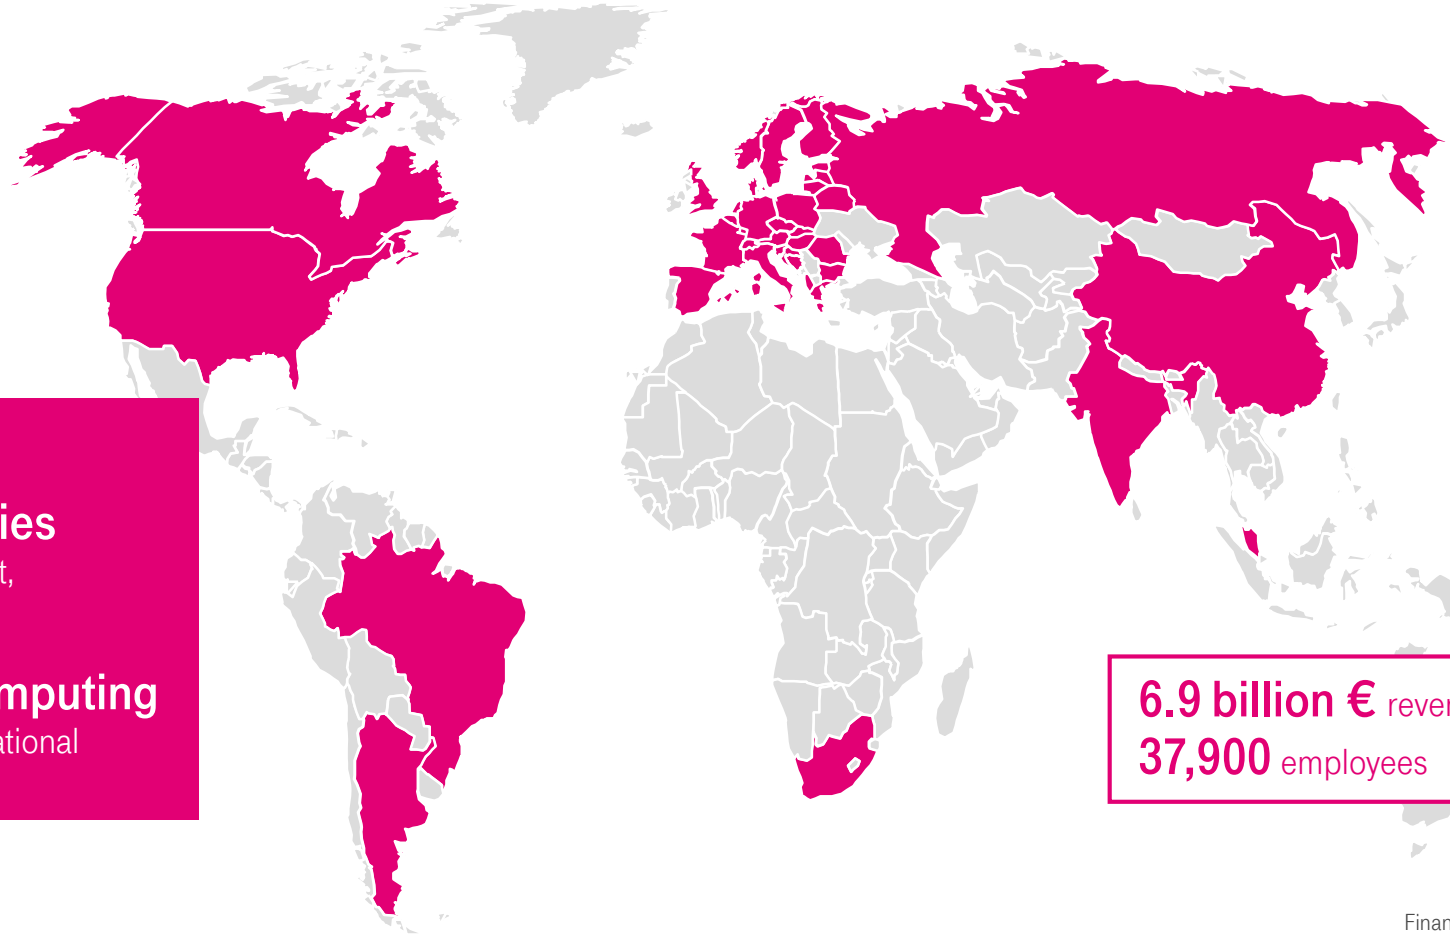

Information technology
and telecommunications

Services for industries
automotive, finance, transport,
retail & public sector

Pioneer in **Cloud Computing**
Corporate customers, multinational
corporations & public sector

6.9 billion € revenue
37,900 employees

Financial figures taken from DT's 2017 annual report

AGENDA

01 INTRODUCTION

02 DIGITIZATION, CLOUDS AND T-SYSTEMS

03 RELEVANCE OF CLOUDS IN SCIENCE

04 OPEN TELEKOM CLOUD – SCIENCE USE CASES AND ACHIEVEMENTS

05 LESSONS LEARNED

DIGITIZATION TODAY...

UBER

WORTH ~ €41 BN

... despite owning no taxis or cars.

airbnb

1.5 M ROOMS

... despite owning no hotels or apartments.

...CREATES NEW BUSINESS MODELS

€1.9 TN

ECONOMIC VALUE-ADD ACROSS ALL INDUSTRIES BY 2020

...CREATES GROWTH OPPORTUNITIES

T · · Systems ·

...AND CHANGES CUSTOMER BEHAVIOR AND REQUIREMENTS

> **"I, ALWAYS, EVERYWHERE AND NOW"**
> Karl-Heinz Land, Digital Evangelist

adidas

Adidas
FUTURECRAFT 3D

3D-PRINTED RUNNING SHOE
midsole tailored to cushioning needs

WeChat

SHOPPING, GAMES AND PAYMENT,
branded accounts

697 MILLION active **USERS WORLDWIDE**
(100 million outside China)

App focuses on **PROVIDING LAST-MILE LOGISTICS** services – worth **$1 BILLION**

800,000
registered & verified delivery
WORKERS

600,000 online
ORDERS EVERY DAY

T · · Systems ·

MULTI-CLOUD UNIVERSE
FIVE BASIC CLOUD ECOSYSTEMS

vmware DSI vCloud

Microsoft Azure Data Trustee

MODE 1

openstack Open Telekom Cloud

MODE 2

T · · Systems ·

OPEN TELEKOM CLOUD
PUBLIC IAAS FOR EUROPEAN STANDARDS

MARKET AND CUSTOMER EXPECTATIONS ARE CHANGING:

Demand for scalable, dynamic IT resources is growing.

Public IaaS is the answer (compute, storage, network, management)

SECURE
- Data protection in compliance with German legislation
- Meeting enterprise needs

SIMPLE
- Rapid access
- Support/help getting started
- Easy to use

OPEN
- OpenStack API
- No vendor lock-in
- Simple to integrate

An open technology platform with built-in compliance, and ease-of-use, for businesses of any size, in any industry – on demand

OPEN TELEKOM CLOUD

T · · Systems ·

OPEN TELEKOM CLOUD
FAST GROWING PORTFOLIO

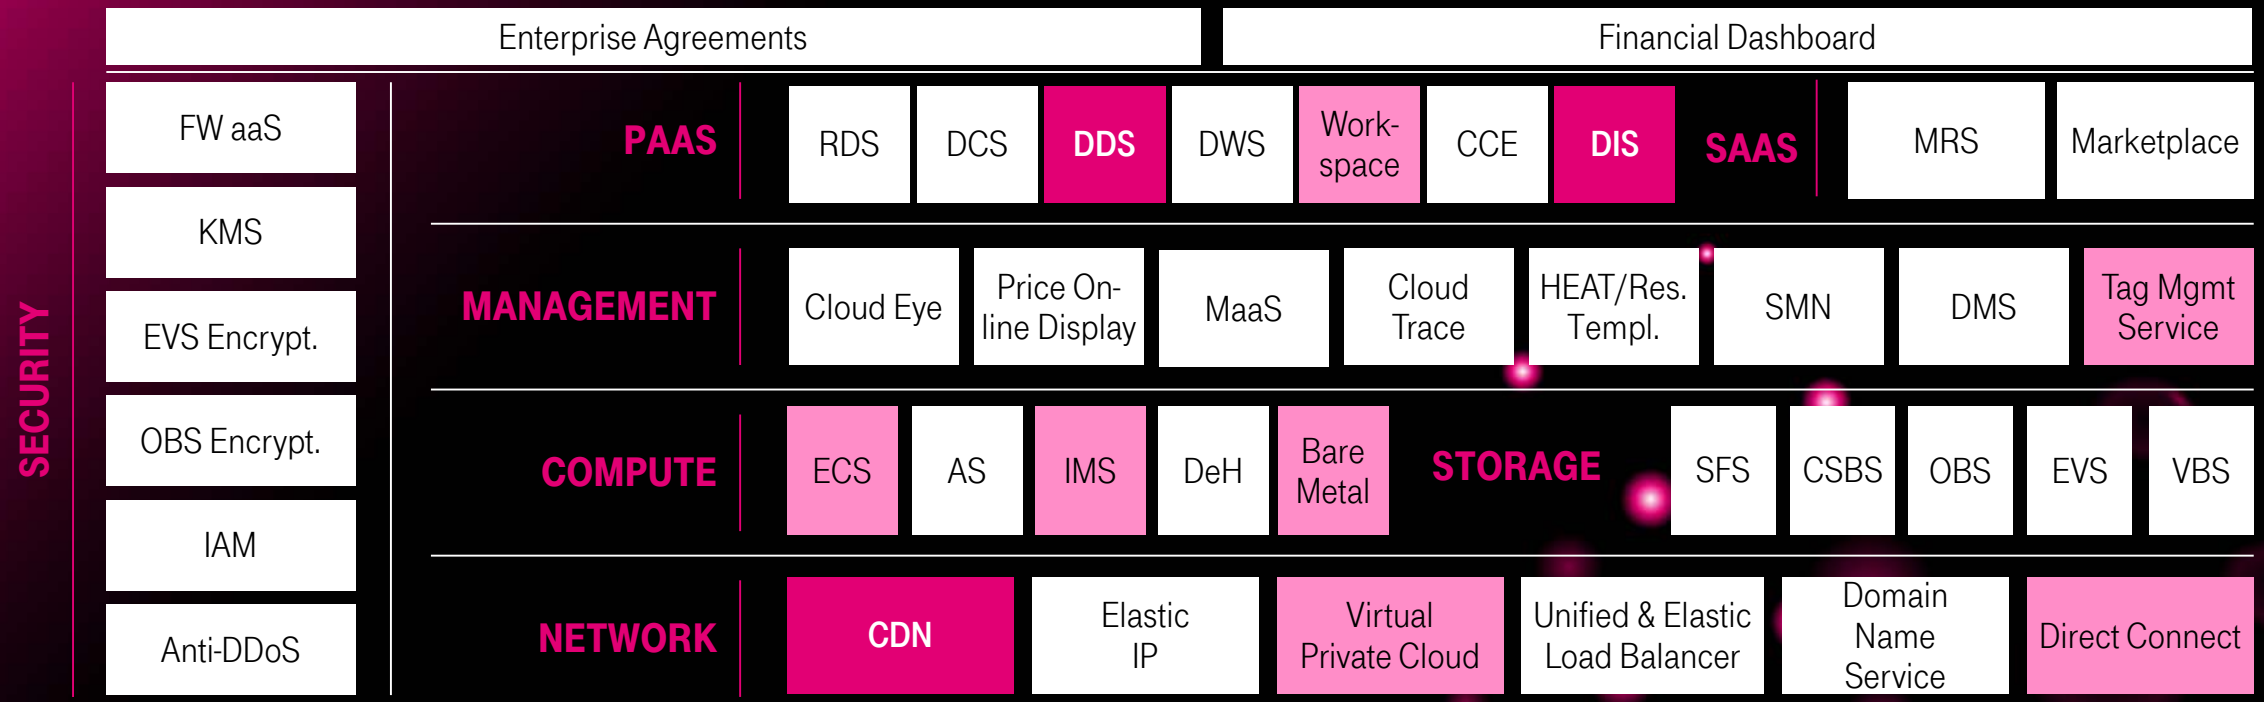

FW = Firewall, KMS = Key Mgmt System, IAM = Identity & Access Mgmt, RDS = Relational Database Service, DCS = Distributed Cache Service, Document DB Service, DWS = Data Warehouse Service, CCE = Cloud Container Engine
MaaS = Migration aaS, SMN = Simple Message Notification, DMS = Distributed Message Service, ECS = Elastic Cloud Server, AS = Auto Scaling, IMS = Image Mgmt System, DeH = Dedicated Host
SFS = Scalable File Service, CSBS = Cloud Server Backup Service, OBS = Object Storage. EVS = Elastic Volume Storage, VBS = Volume Backup Service, CDN = Content Delivery Network

ENHANCED 03/2018 **NEW 03/2018**

OPEN TELEKOM CLOUD
CLOUD CONTAINER ENGINE

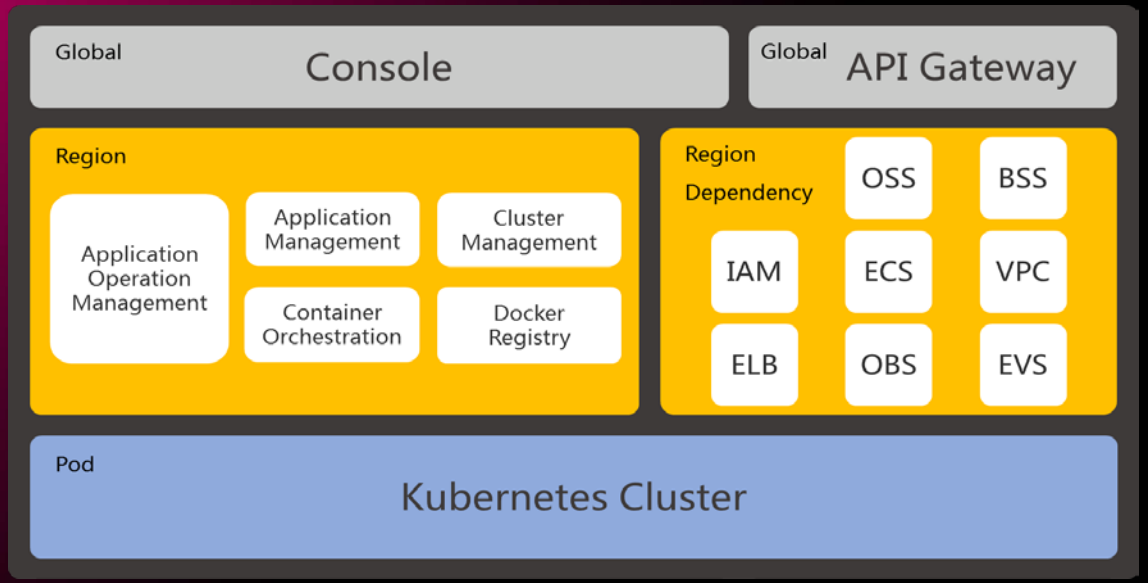

Support of Kubernetes Cluster per Tenant:

#of Clusters	max. #of nodes	max Container Instances per node	max. #Container Instances
5	5000	100	500,000

Note: Bare Metal Cluster support available with OTC V3.2 (Nov 2018)

Feature	CCE V1	CCE V2
Kubernetes Version	1.5	1.9
Docker Version	1.11.2	17.06
Autoscaling for cluster node	Not supported	Supported
Container Storage	EVS	EVS/SFS, Support dynamic provision
Stateful workload(Console)	Not supported	Supported
Job workload(Console)	Not supported	Supported
Configuration Center(Console)	Not supported	Supported
Container network	Overlay	Overlay & Underlay
Application template	Tosca	Kubernetes charts

COST TRANSPARENCY – WHERE ARE THE BENEFITS?

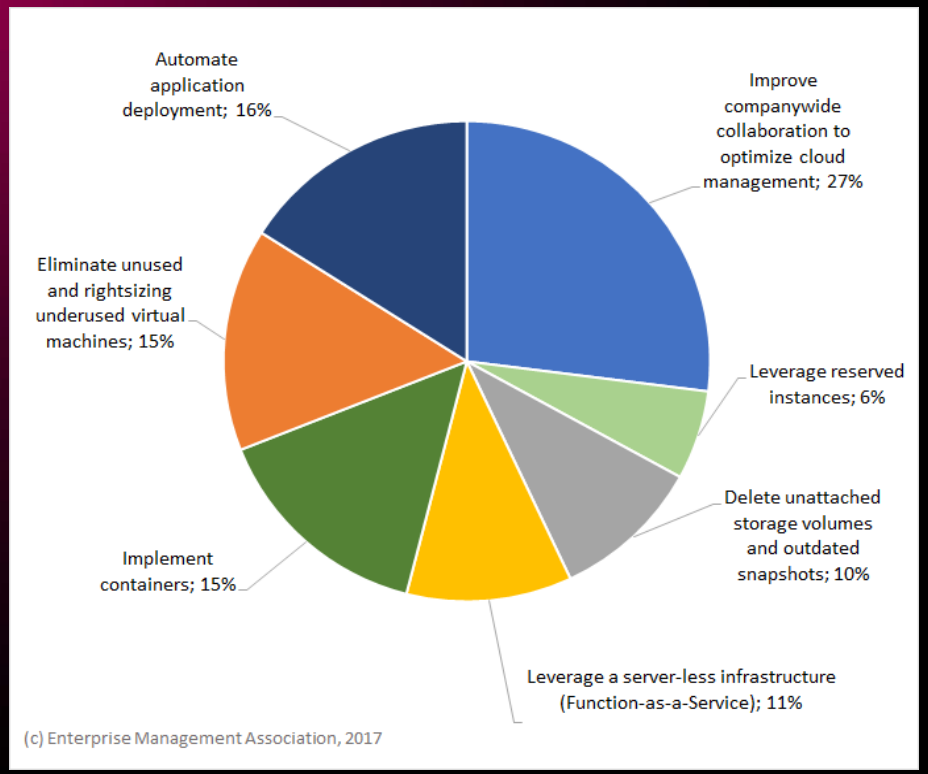

New Financial Dashboard

AGENDA

01 INTRODUCTION

02 DIGITIZATION, CLOUDS AND T-SYSTEMS

03 RELEVANCE OF CLOUDS IN SCIENCE

04 OPEN TELEKOM CLOUD – SCIENCE USE CASES AND ACHIEVEMENTS

05 LESSONS LEARNED

MORE SCIENCE, MORE CLOUD – BUILDING THE EUROPEAN OPEN SCIENCE CLOUD

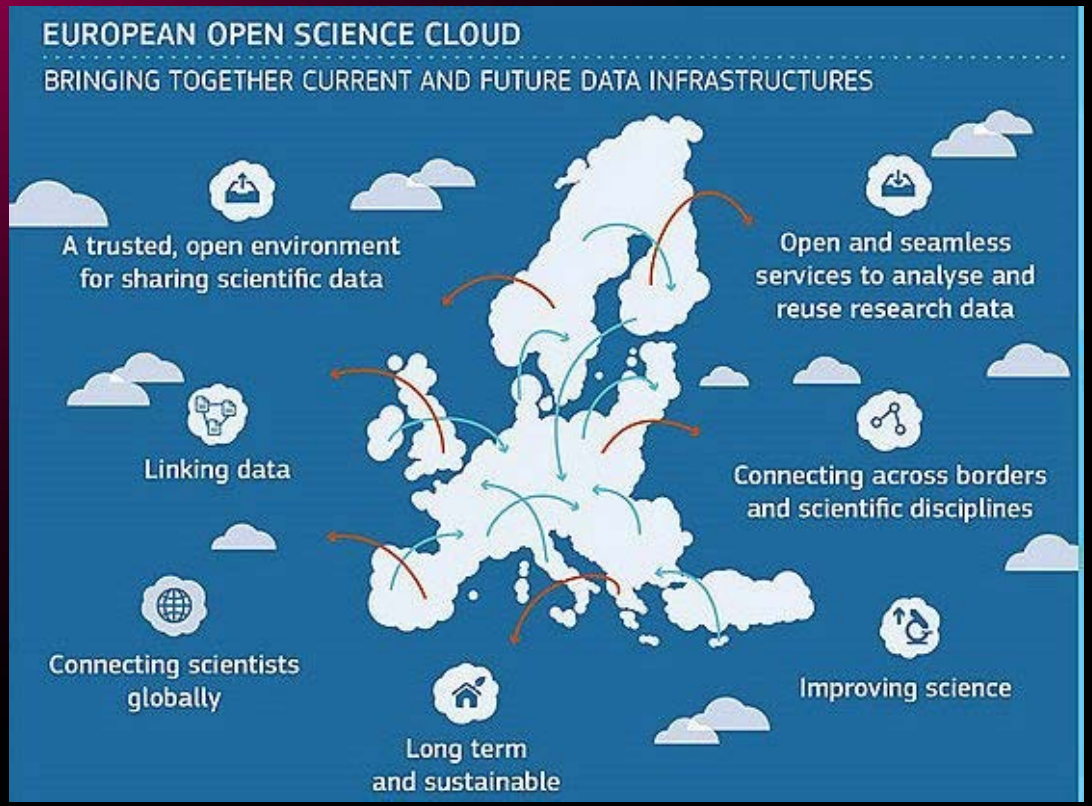

HELIX NEBULA THESCIENCECLOUD

CERN

Joint Pre-Commercial Procurement

High Energy Physics | Astronomy | Life Sciences | Photon/Neutron Sciences | Long Tail of Science

SaaS

PaaS

IaaS

Key Challenges:

- Compute
- Large-scale Data
- Network
- Federated Identity
- Service Payment

CERN

ORGANIZATIONAL CHALLENGES AND SOLUTION

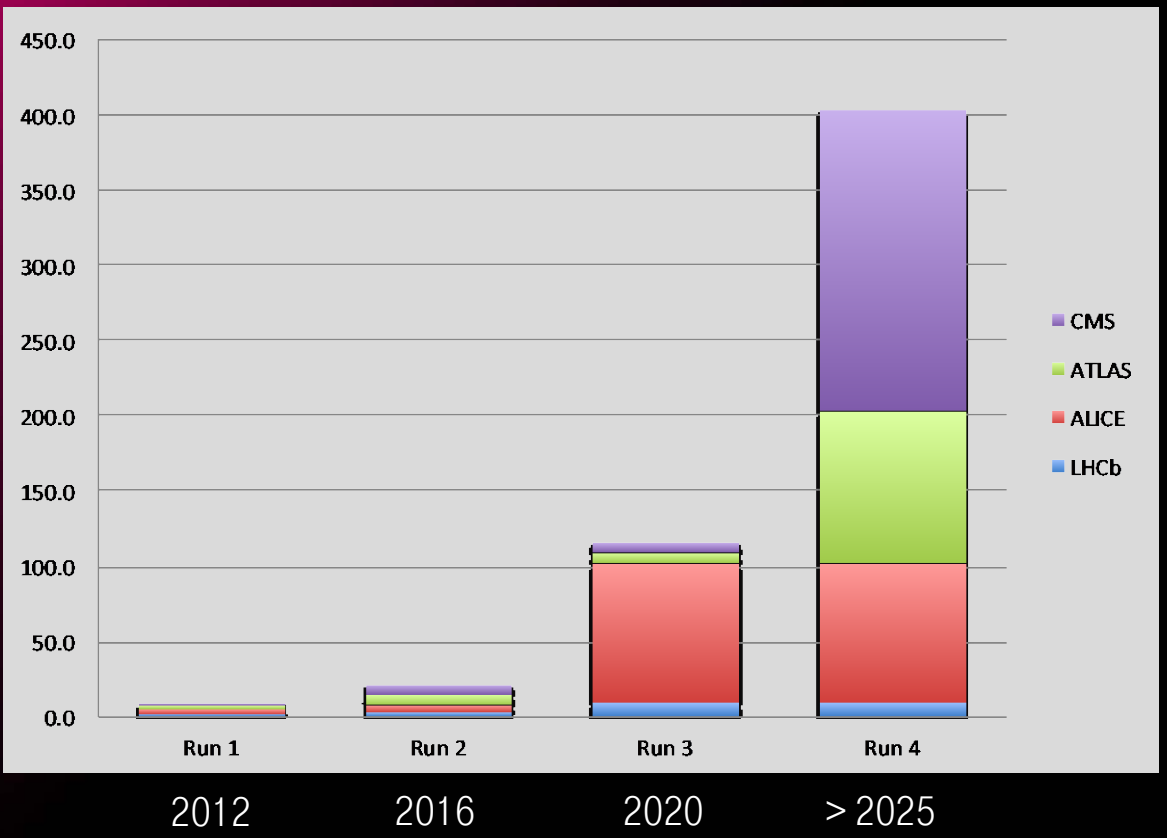

Challenges

- Manage the foreseen data explosion
- Create a more elastic IT infrastructure
- Within a fairly stable budget

Solution

- Moore's law helps, but not sufficient
- Large effort spent to improve software efficiency
- Exploit multi-threading, new instruction sets, ...
- Still need factor 2 in terms of cores, storage etc.

T · · Systems ·

AGENDA

01 INTRODUCTION

02 DIGITIZATION, CLOUDS AND T-SYSTEMS

03 RELEVANCE OF CLOUDS IN SCIENCE

04 OPEN TELEKOM CLOUD – SCIENCE USE CASES AND ACHIEVEMENTS

05 OUTLOOK

THE CERN / HELIX NEBULA APPROACH

Transforming In-House Resources

- 2012: Start with private cloud under OpenStack
- Physical and virtual servers, remote machines
- Responsibilities by layers of service deployment
- Scaling to large numbers
  (> 15'000 physical, several 100'000s virtual)
- Support for dynamic host creation/deletion
  - Deploy new services/servers in hours rather than weeks/months
  - Optimise operational and resource efficiency

Evaluating Public Clouds

- Series of pilot projects increasing in size and complexity

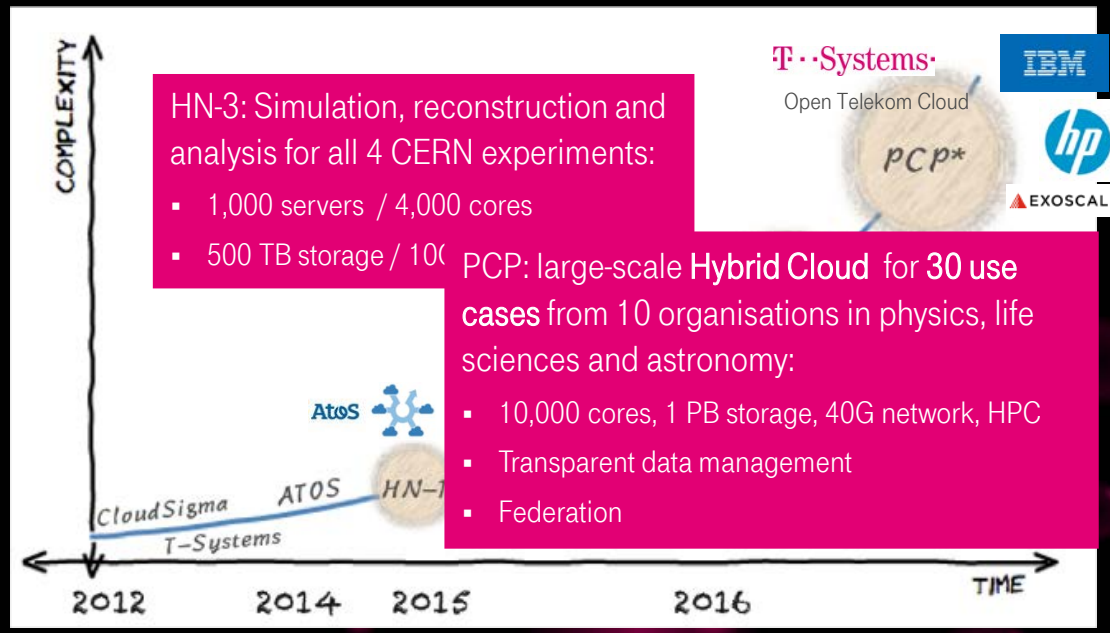

HELIX NEBULA – MORE SCIENCE THROUGH HYBRID CLOUDS

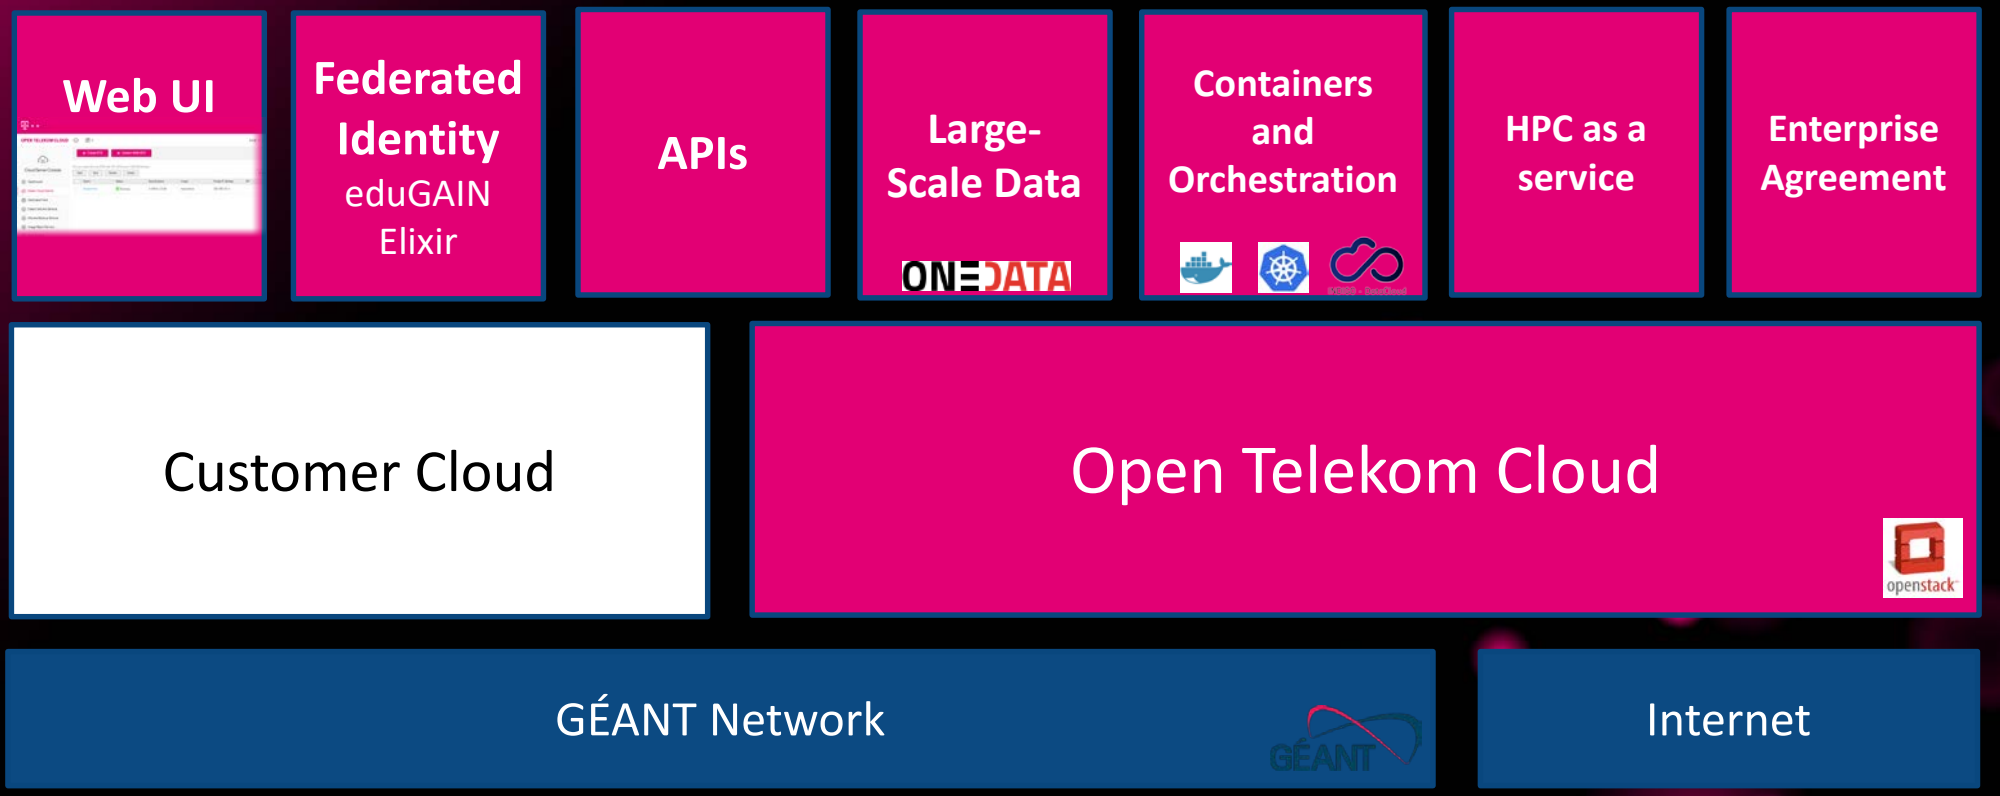

HELIX NEBULA 2018 – USE CASES TESTED

TEST/Deployment NAME	PROCURER	STATUS	Expected Results	Summary of results obtained
Batch Service Deployments	CERN	On-going		
TOTEM Deployment test		On-going	Fully functional data analysis software stack. Scaling out the installation and testing a complex analysis example.	Functional deploymment OK. Now scaling out.
Security test challenge		On-going		
PerfSONAR tests		On-going	Avg Latency (in/out):	
Deep Learning with GPUs		On-going	Parallel training of Generative Adversarial Networks. Scaling experiments and comparison to HPC clusters. Expect linear (strong) scaling.	Initial run on 2 T-Systems GPU.
IAAS access via EduGain and local accounts	CNRS	Completed		in Advania we had issue with the openstack
Onedata Wave 1: HDF5_IO	DESY	Completed		marked as completed but was not able to reach
Onedata Wave 2: Scaling HDF5_IO		On-going	100TB shared data (50MB average file size, min 7MB, max 4GB), more than 100 (parallel) instances (of HDF5_IO)	Currently trying to run tests with multiple files.
SLURM Integration		Completed		Works, still need to check updated LUMA version
Docker SLURM jobs in the Cloud		Completed		works, need to rerun on Exoscale to measure job
CLOUD_BENCHMARK test	EMBL	On-going		Some flavors of VMs are able to match (+/- 20%)
Onedata Wave 1: Data Transparency Layer		Completed	The data transparency layer can be established, sub-tenants and tenants are able to share datasets with other sub-	Data transparency layer cannot be established or
Onedata Wave 2: PanCancer deployment		On-going		
FDMNES: HPC Test	ESRF	On-going	- start a multi-process MPI and/or OpenMP (later also GPU) job on a cluster with a user-configurable number of node	RHEA: no HPC tests yet because no cloud
Cloud Access Interface		On-going	The HPC based use-case, such as the ESRF/FDMNES in the Buyers Group test-suite, must be executable in a script	RHEA: so far, I have not been able to set up a
Performance test		On-going	**Once the first multi-process HPC job has successfully run, it would be interesting to do some tests: how do**	
Onedata Wave 1: DODAS	INFN	Completed		
Onedata Wave 2: Scaling DODAS		On-going		
Batch system extension	KIT	On-going		
CPU benchmarks		On-going		
Non-browser federated authentication with ECP SAML extension		Completed		T-Systems fixed the registration of the SP for the C
Onedata Wave 1: dCache io	SURFsara	On-going	**Objective: a custom tool will be provided to be run in a simple small VM, with enough allocated storage for t**	We have problems with our own dcache to oneDat
Federated AAI for OTC setup		Completed		SAML translation worked
WeNMR / HADDOCK SLURM deployment		On-going		
NIKHEF scale-out		On-going		
LOFAR SLURM deployment		On-going	**Run LOFAR pipelines on a SLURM cluster on both providers**	
Batch-system-extension	STFC	Not Started		
Jobs-on-kubernetes		Not Started		

HPCAAS: FDMNES USE CASE (ESRF)

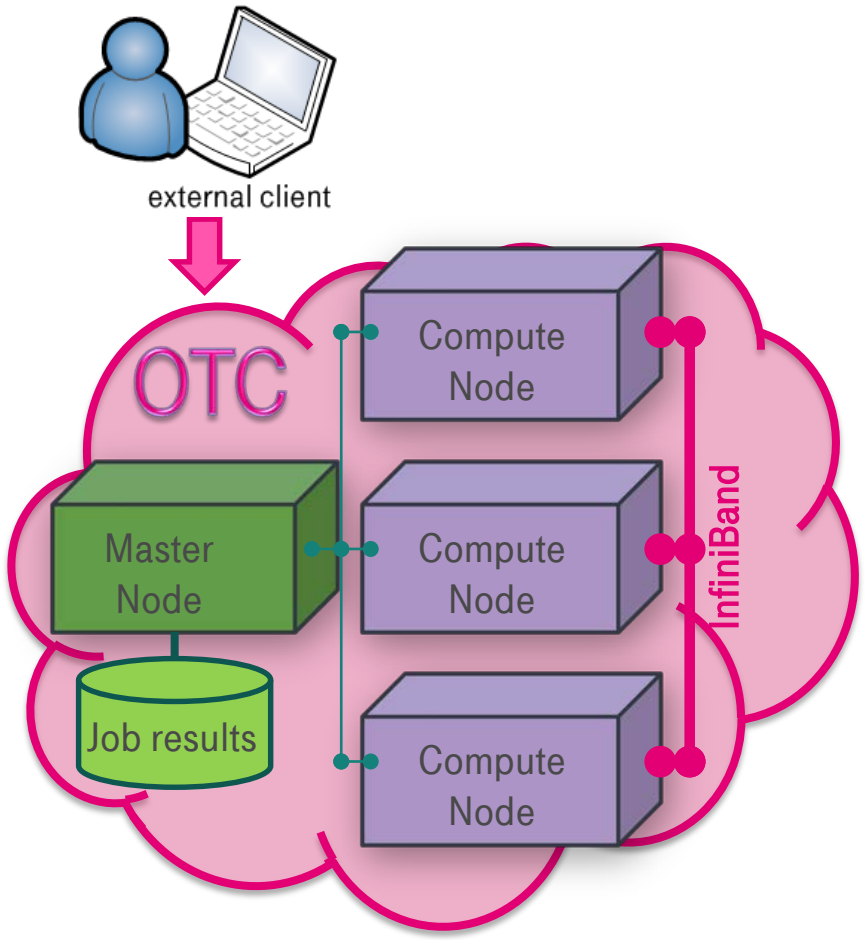

- "CFN" User Menu (CloudFormatioN)
- Workflow (high level):
  1. Create cluster (on local pc)
  2. Enter master node
  3. Submit job (on master node)
  4. Execute job
  5. Auto-delete compute nodes after job termination
  6. Keep results

- HPCAAS menu's prepared to also integrate e.g., SLURM workload manager

Example master flavor: s2.medium.1 (cheaper, exists longer)
Example compute flavor : h2.3xlarge.10 (more expensive, auto-deleted after job-end)

HELIX NEBULA – DODAS USE CASE

DODAS in a nutshell

- **D**ynamic **O**n **D**emand **A**nalysis **S**ervice:
  DODAS
  - Platform as a Service tool which aim is to guarantee deployment of complex and intricate setup on "any cloud provider" with almost zero effort.
  - Allows to instantiate on-demand container based cluster to execute software applications

DODAS is a Thematic Service under EOSC-hub Horizon 2020 EU project.

- **Initially developed as an INDIGO-DataCloud use cases based on CMS workflow.**

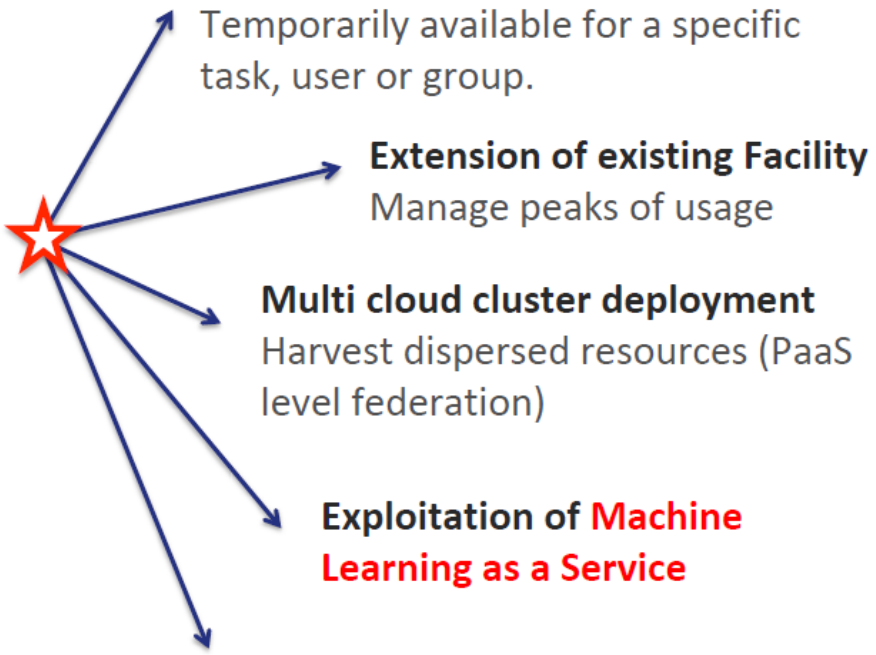

HELIX NEBULA – DODAS USE CASE

EOSC-hub

Resources Management with DODAS

T··Systems·

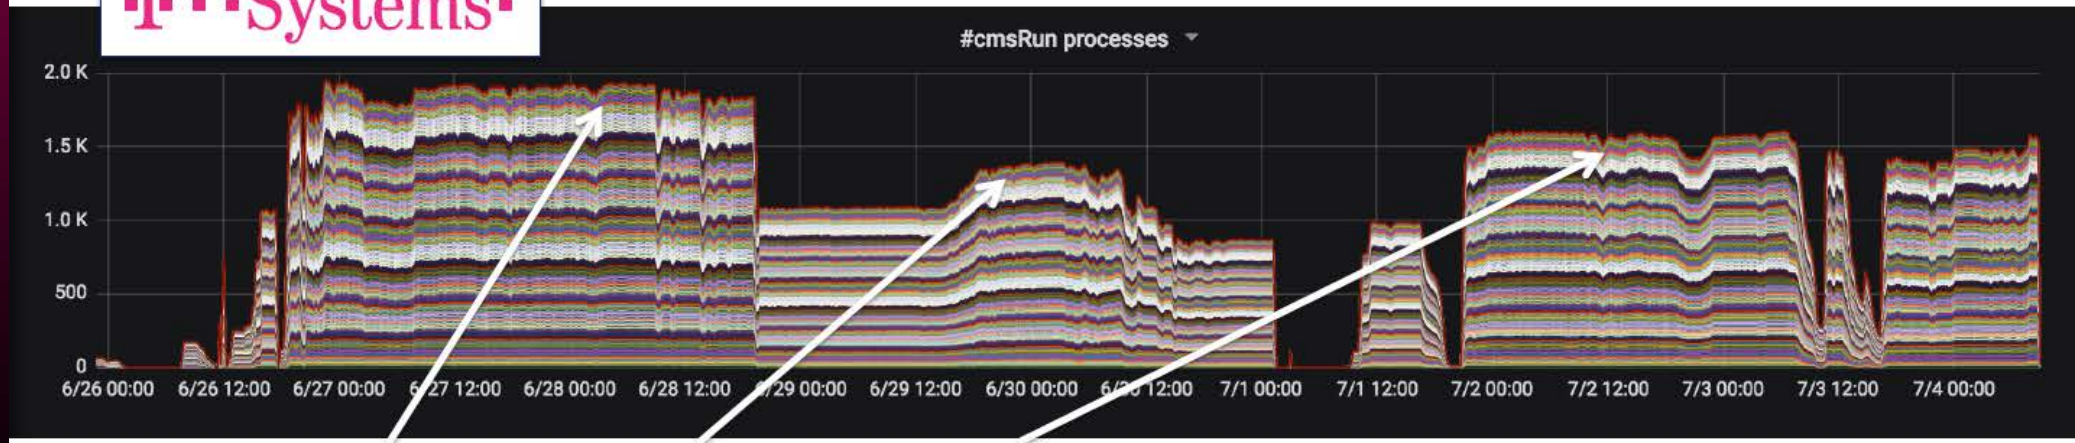

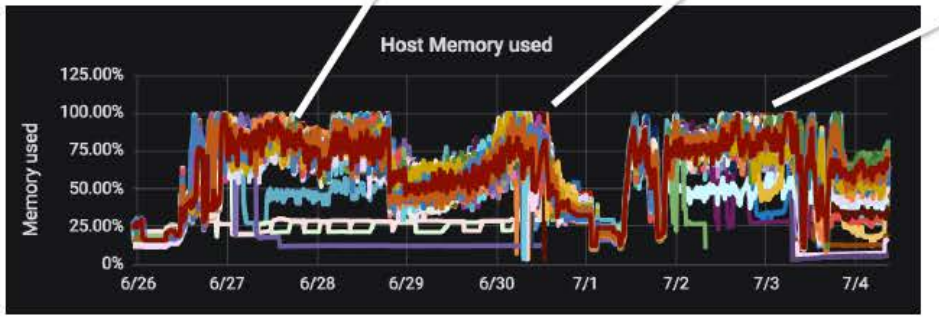

- Elasticity and self-healing
- Handling "special requirements" high memory jobs
- Stability over days/weeks (120k jobs)

→ Work done in conjunction with
Thanks to Andrea Chierici

HELIX NEBULA
THESCIENCECLOUD

Monte Carlo Simulation: Why

Detailed simulation of subatomic particles is essential for data analysis, detector design

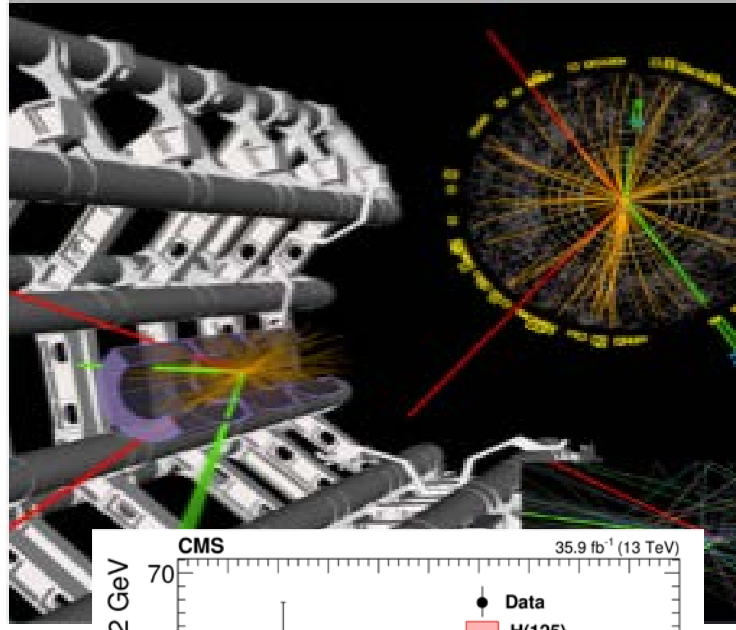

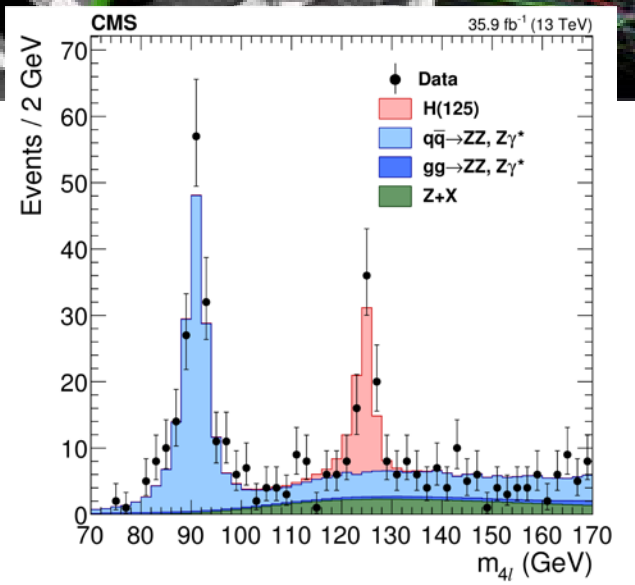

Understand how detector design affect measurements and physics

Correct for inefficiencies, inaccuracies, unknowns.

Theoretical model to compare data against.

>50% of WLCG power for simulations

Current code cannot cope with 2025 LHC upgrade

CERN openlab

WLCG
Worldwide LHC Computing Grid

200 Computing centers in 20 countries: > 600k cores

@CERN (20% WLCG): 65k cores; 30PB disk + >35PB tape storage

Generated images

Interpret detector output as a 3D image

3D convolutional GAN generate realistic detector output

Customized architecture (includes auxiliary regression tasks)

Agreement to standard Monte Carlo in terms of physics is remarkable!

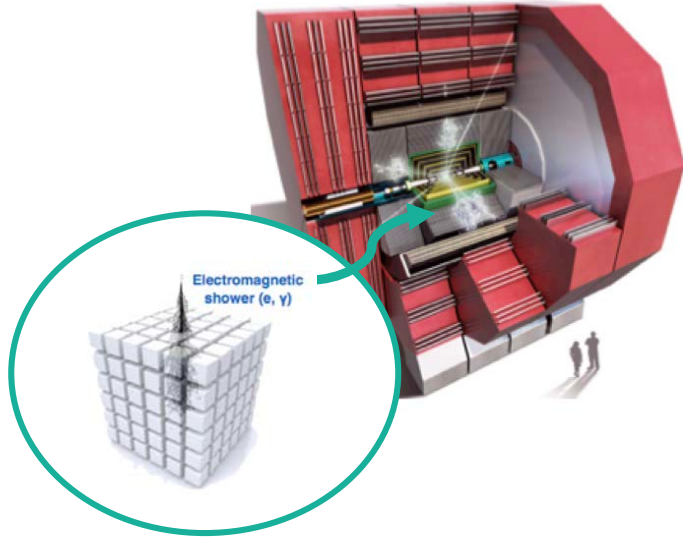

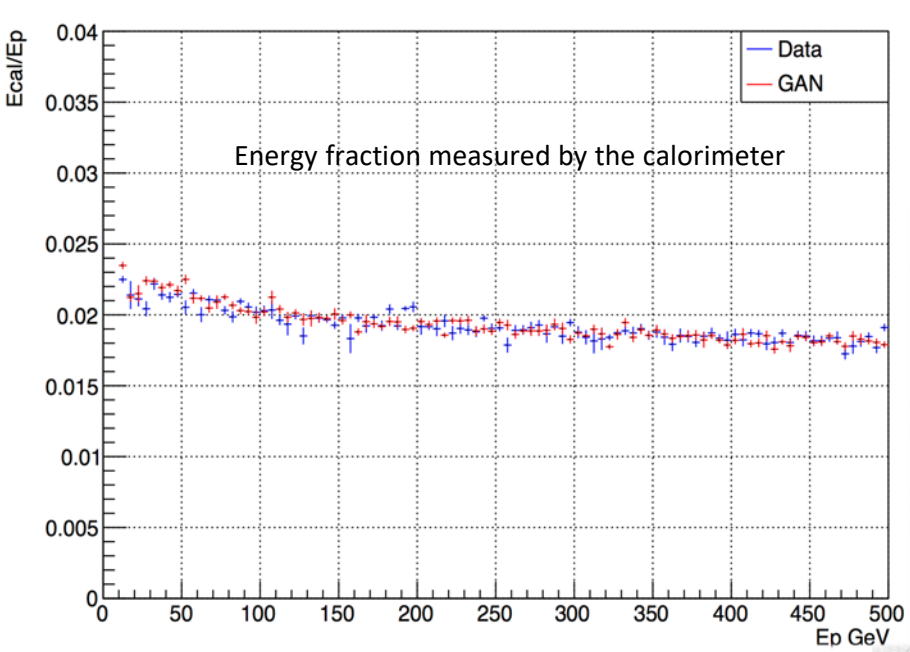

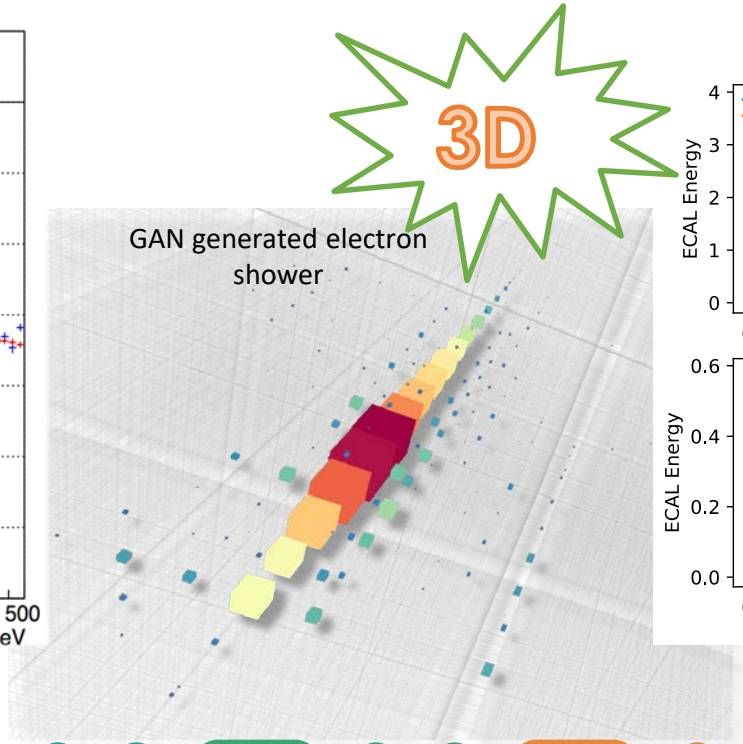

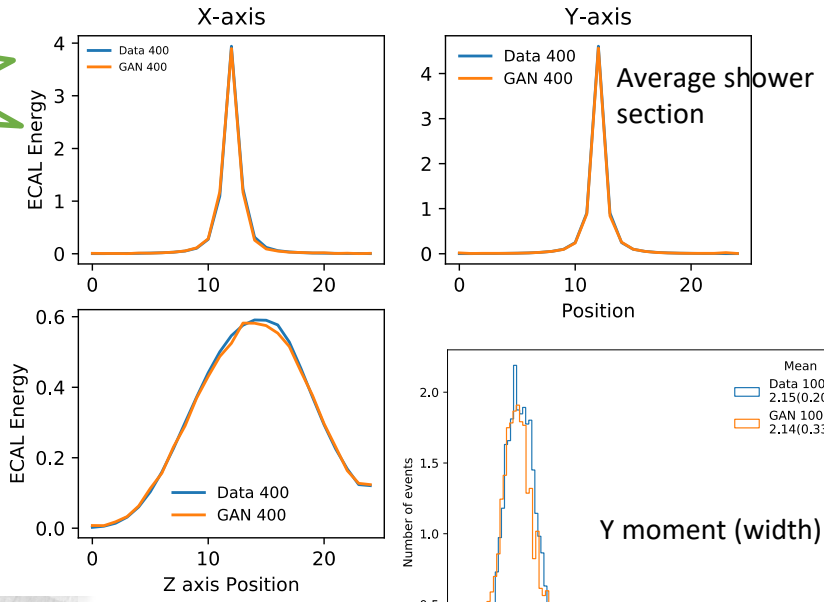

Distributed training is needed

Optimised training and access to cloud facilities enables a larger range of applications

Inference: Monte Carlo: 17 s/particle vs 3DGAN: 7 ms/particle

➔ speedup factor > 2500!!

Training: 1 day on a NVIDIA GTX1080

- **Introduce mpi based data parallel training**
- First results using cloud GPUs via Helix Nebula Science Cloud (T-System)
- Test on 8xV100 BMS on OTC

Time to create an electron shower		
Method	Machine	Time/Shower (msec)
Full Simulation (geant4)	Intel Xeon Platinum 8180	17000
3d GAN (batch size 128)	**Intel Xeon Platinum 8180**	**7**
3d GAN (batchsize 128)	GeForce GTX 1080	0.04

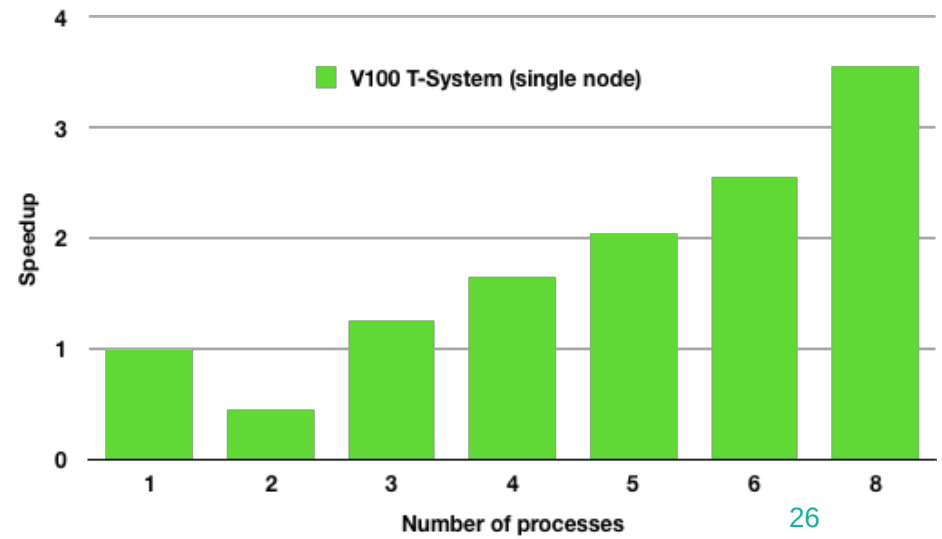

LIFE SCIENCES - PAN-CANCER USE CASE

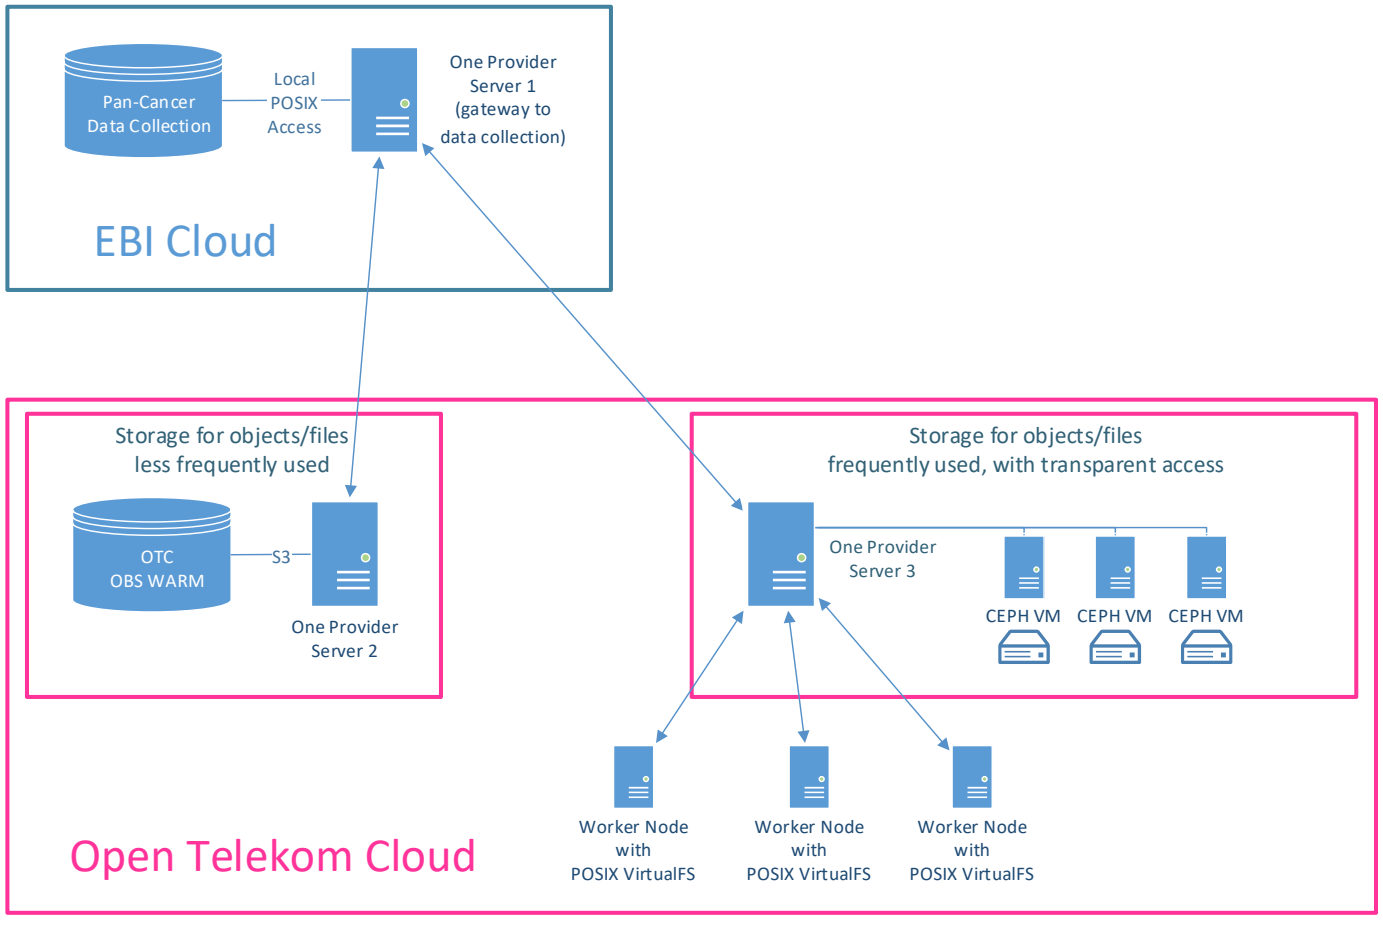

- Use case with strong burst character:
  Continuous use: 7 VMs
  Burst use (2.5 days p.m.): up to 400 VMs
- Storage infrastructure with High I/O and infrequent usage patterns.
- Users require transparent data management and “simple” parallel processing.
- Total-Cost-of-Ownership study comparing on-premise processing with Hybrid Cloud model.

LIFE SCIENCES – DATA MANAGEMENT

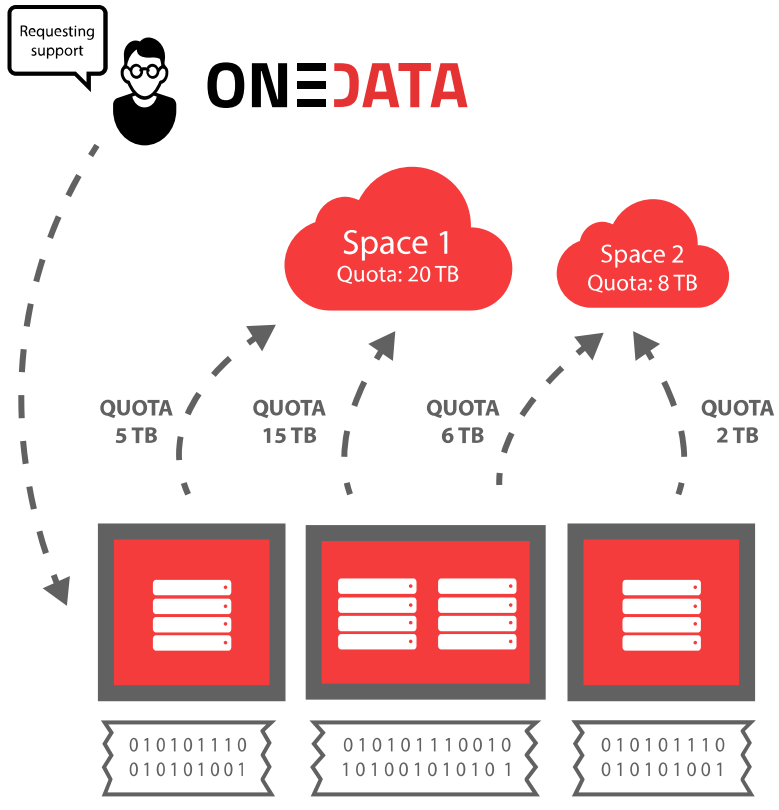

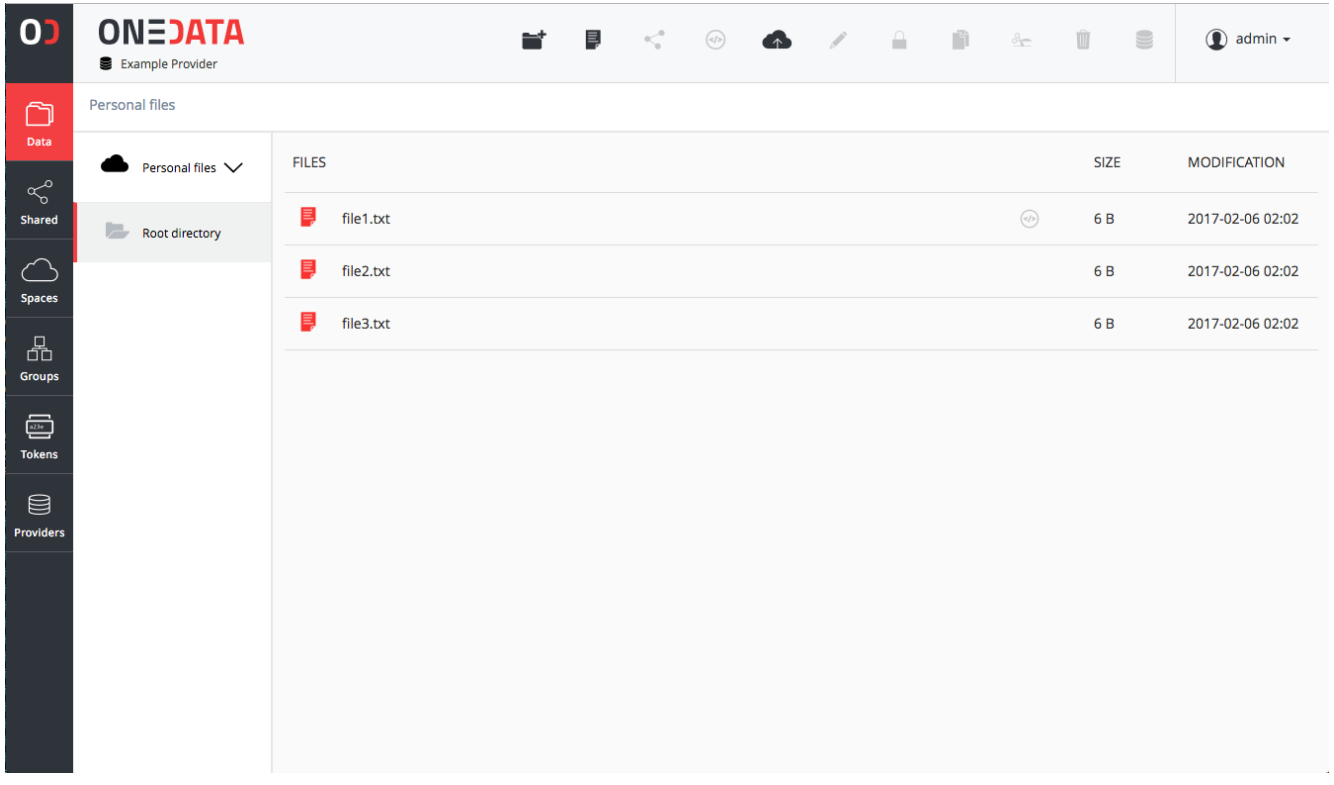

ONEDATA – High-Throughput Data Processing

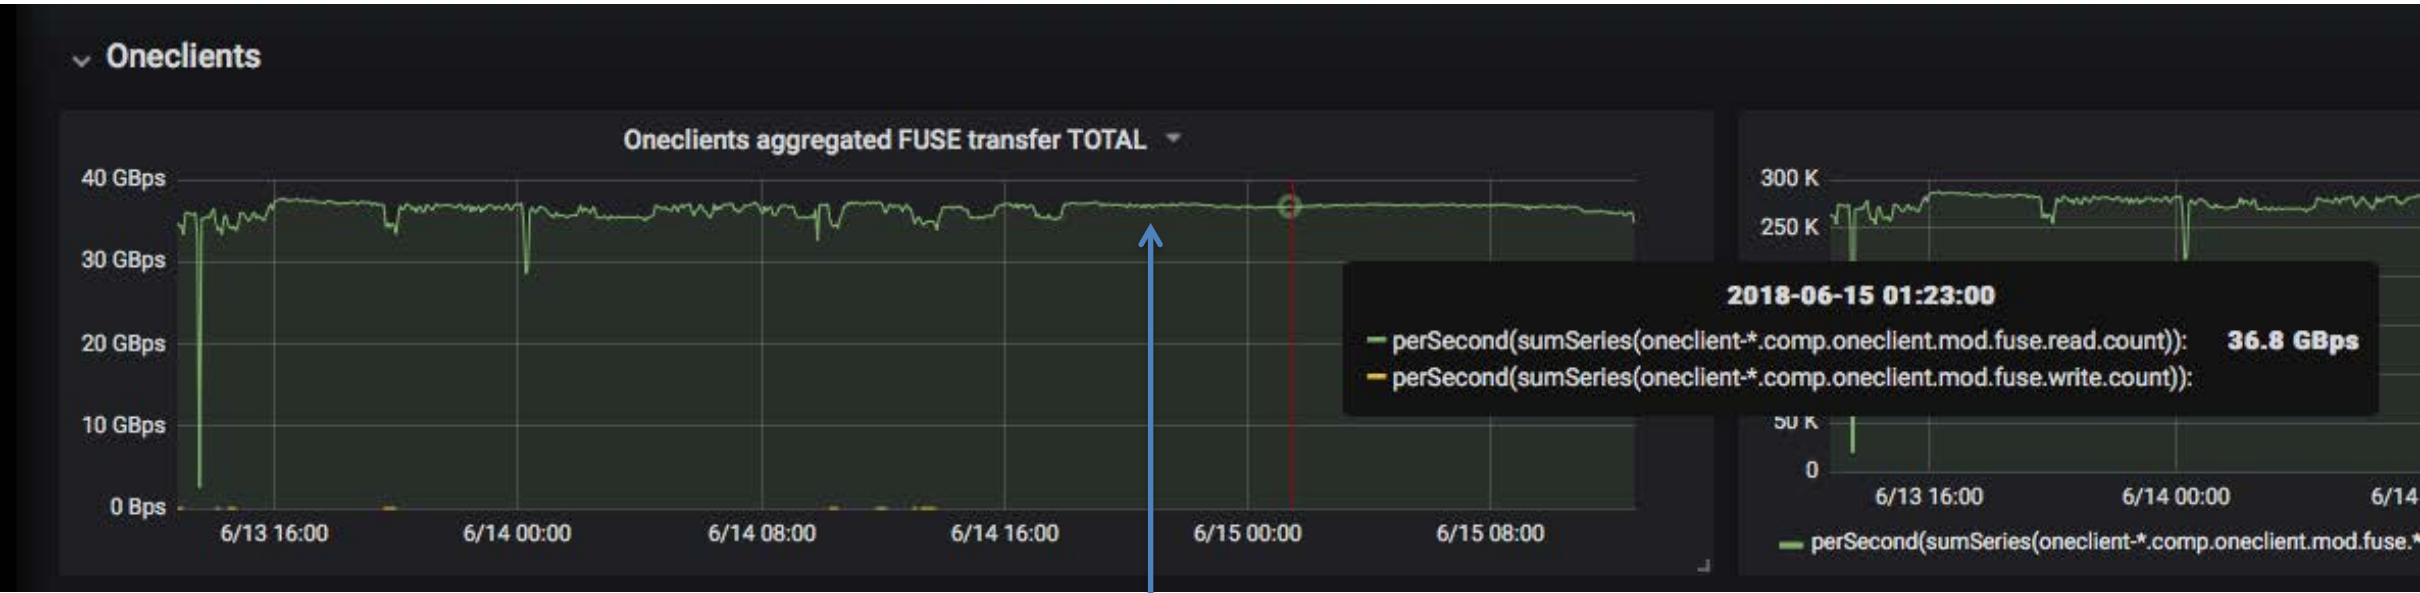

Onedata Transparent POSIX File System
Processing transparently cached data - 37GBytes/sec

ONEDATA – High-Throughput Data Migration

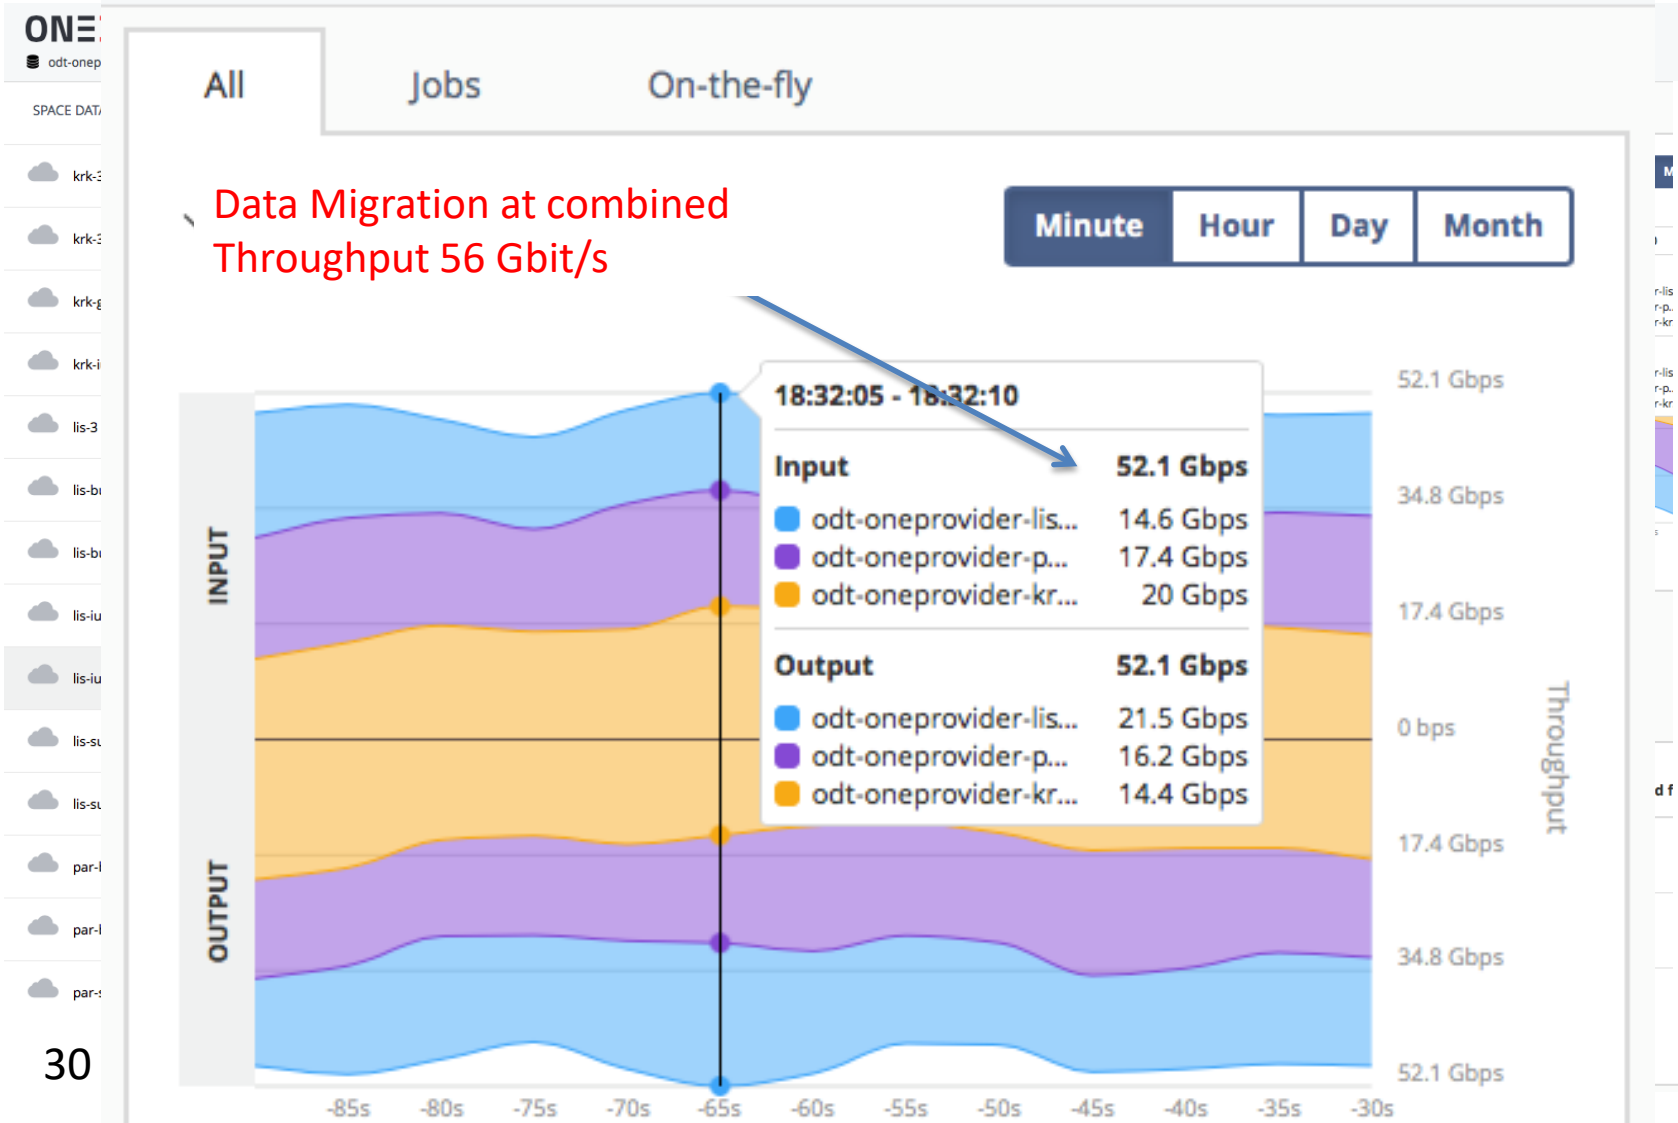

- Data Transfer Mesh
- 3 Oneproviders connected by 20+Gbit/s links
- Transfer data between all them
- Single VM Node per Provider
- Linear scalability

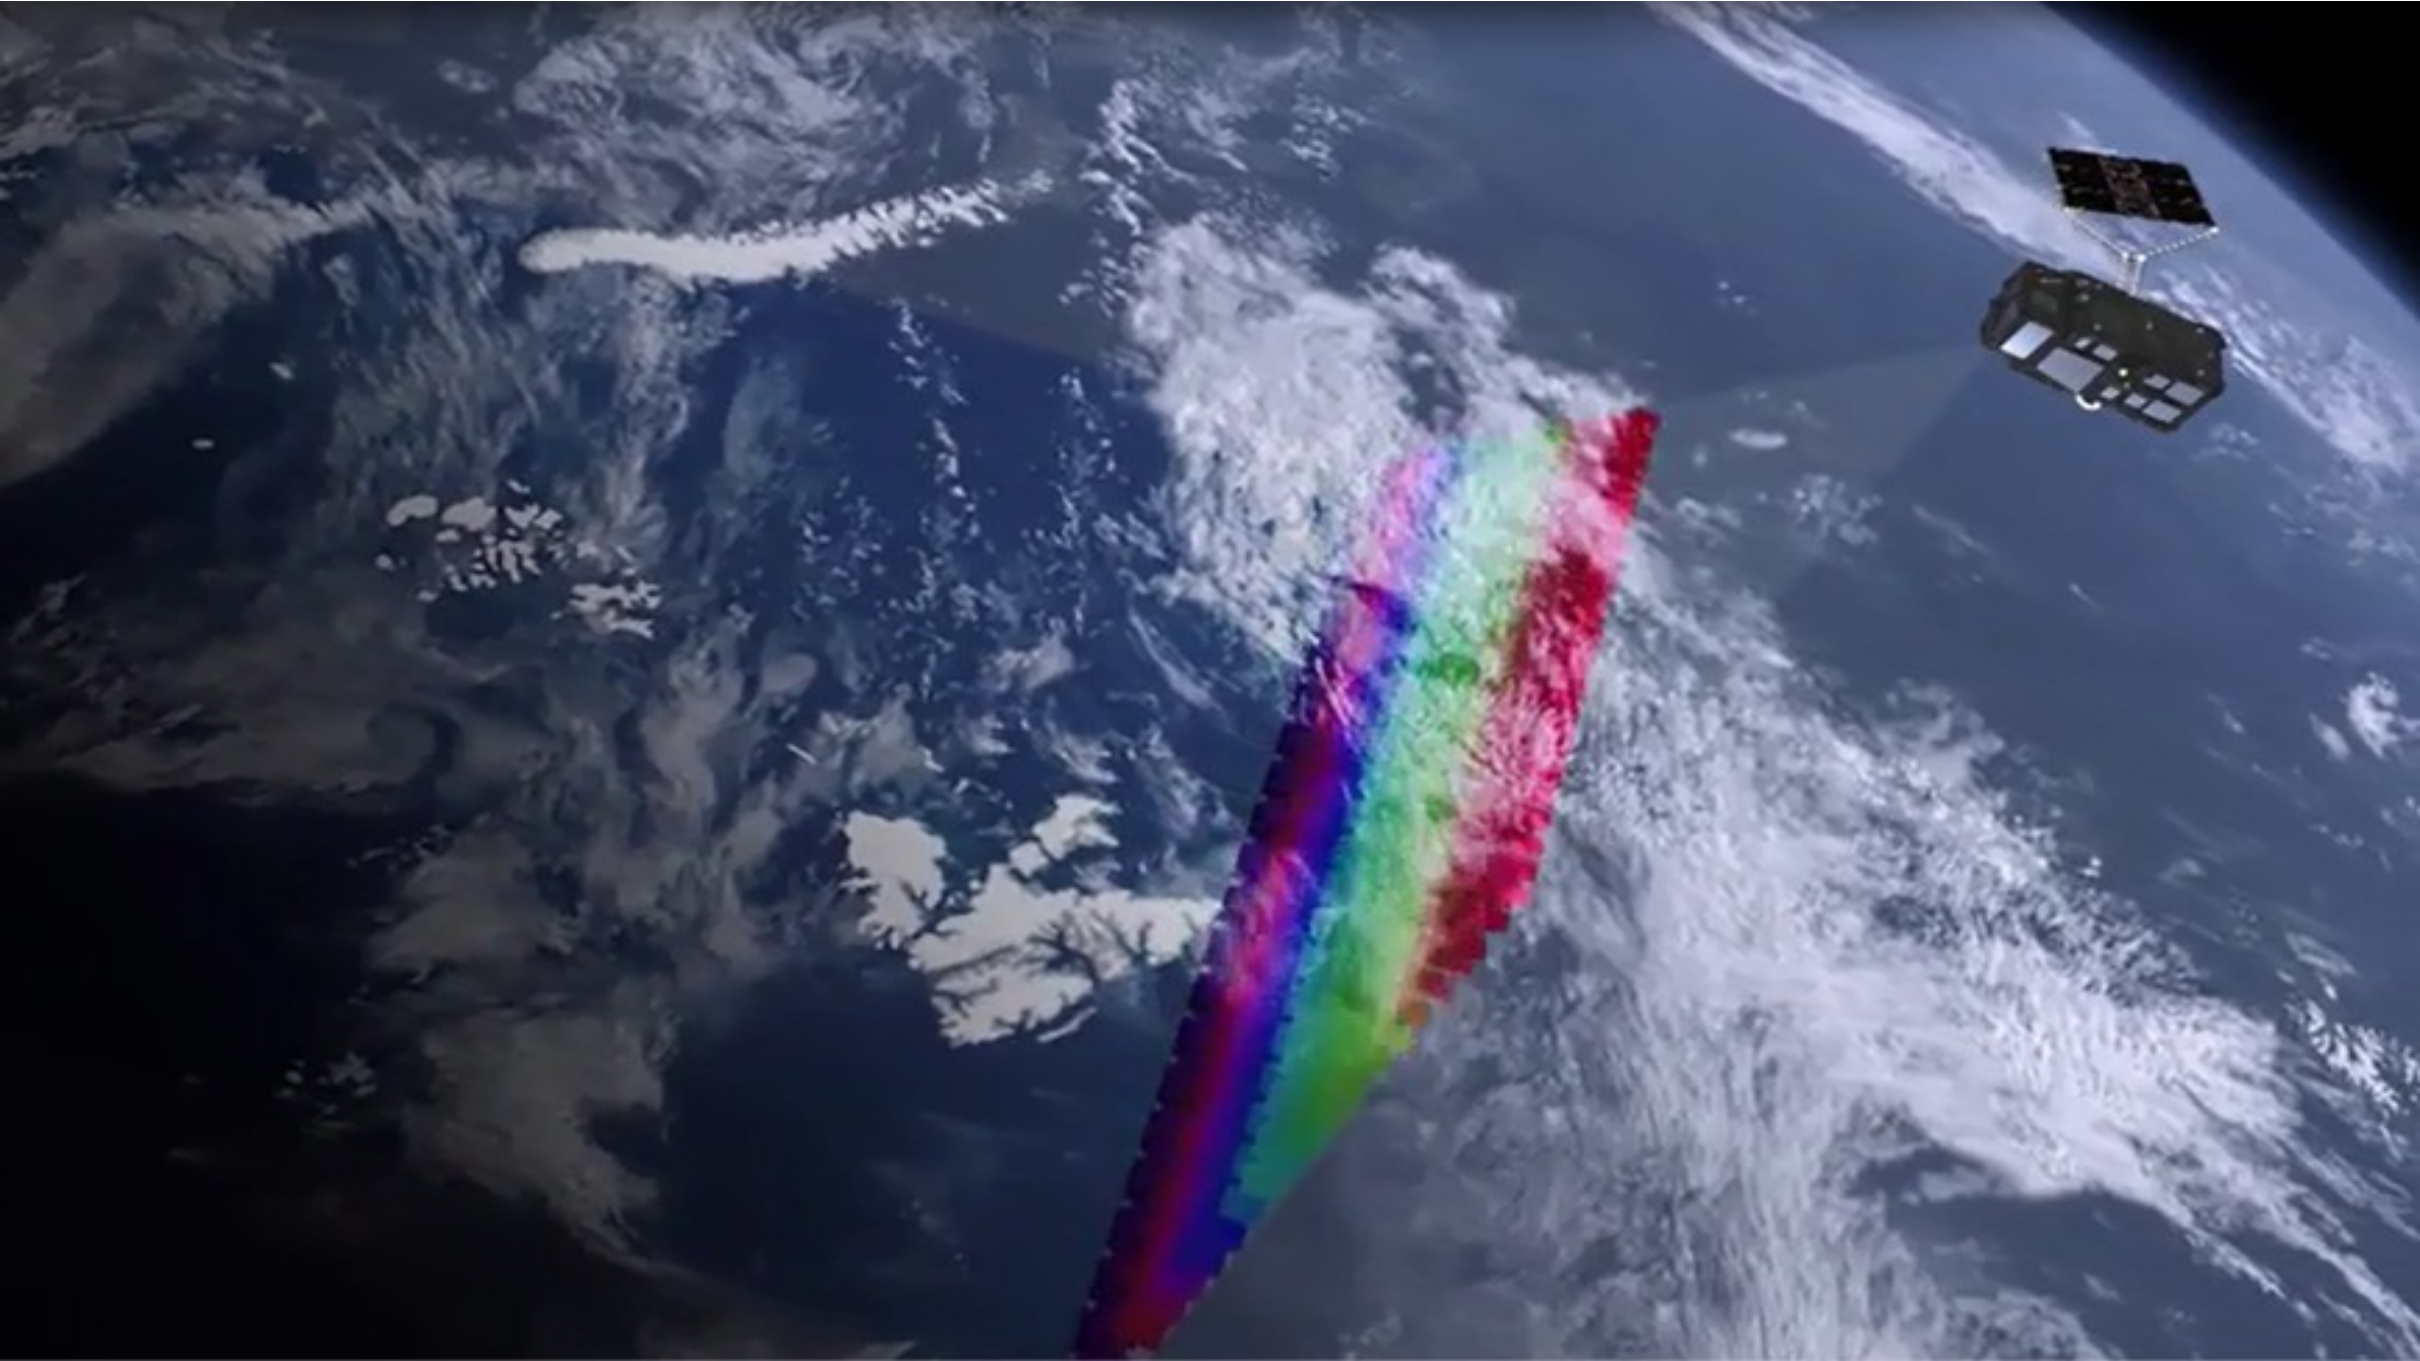

INFORMATION SERVICES FOR COPERNICUS DIAS (ESA/EU)

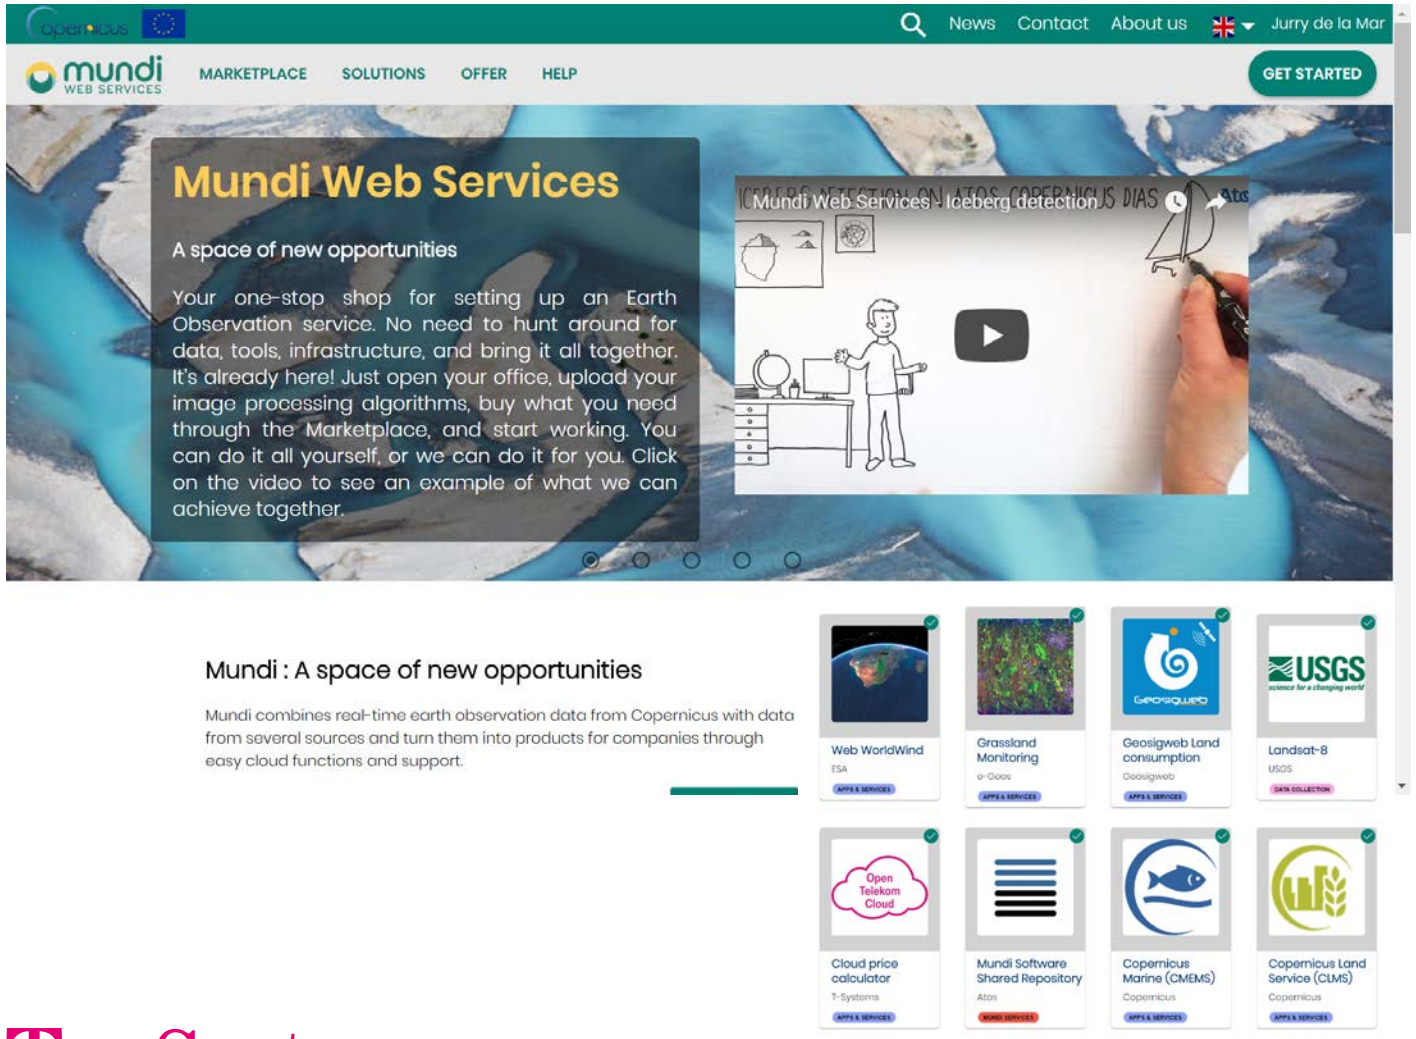

The challenge:

- To provide data access to at least two years of Copernicus Satellite data from the ESA Sentinel Missions.
- To provide geo-analytics tools and services with attractive business models.
- To complement the data with other relevant information for big-data analytics.

The solution:

- Mundi Web Services – SaaS and PaaS services on OTC
- Live start June 2018
- 5 PB data, expected to grow to more than 40 PB by End of 2019
- Sentinel Cloudless, Grasslands Monitoring, Regional Agriculture and Forestry Services established
- Various other use cases under implementation.

THE PARADIGM – ADOPT THE CLOUD NATIVE WORLD

- Consequent application development, taking cloud computing architectures into account
- Use of microservices, that can be run on any combination of infrastructure(s)
- Containers, automation, resilience and scalability as the framework
- New tool ecosystems emerging to support developers and users e.g., HEAT, Terraform, Ansible, Kubernetes, Mesos/Marathon

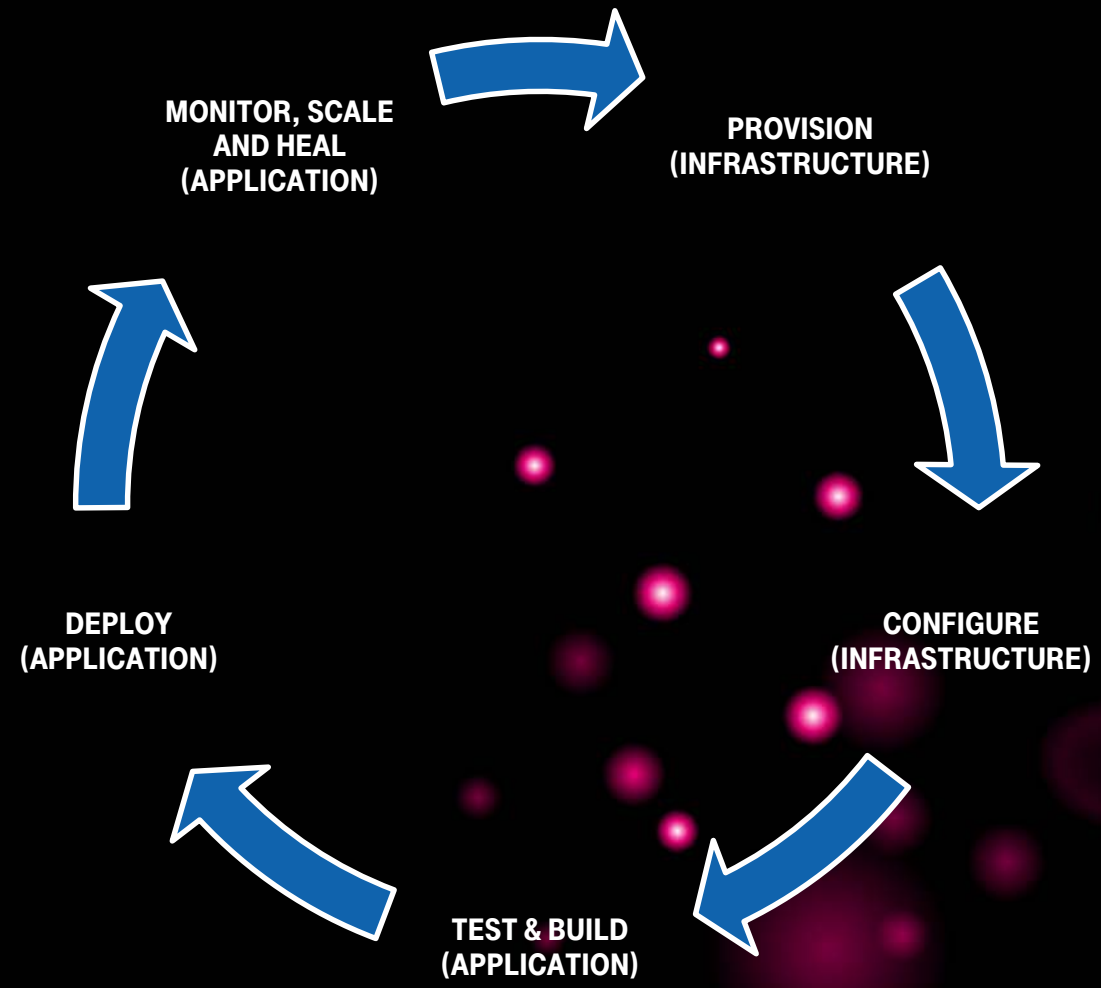

AGENDA

01 INTRODUCTION

02 DIGITIZATION, CLOUDS AND T-SYSTEMS

03 RELEVANCE OF CLOUDS IN SCIENCE

04 OPEN TELEKOM CLOUD – SCIENCE USE CASES AND ACHIEVEMENTS

05 LESSONS LEARNED

LESSONS LEARNED

- Main challenge is to adept to a **cloud-native** world
- **Federation** is key to win user acceptance, while maintaining security
- **Automation and self-healing** will generate totally new dimensions of computing
- **Data Management** challenges often underestimated and require next generation approaches
- Organisations require flexible and highly-granular **cost transparency** and management

QUESTIONS?

What do you consider main challenges for using cloud computing in your environment today?

Mentimeter

53

What functionality would be relevant for your work or environment?

Mentimeter

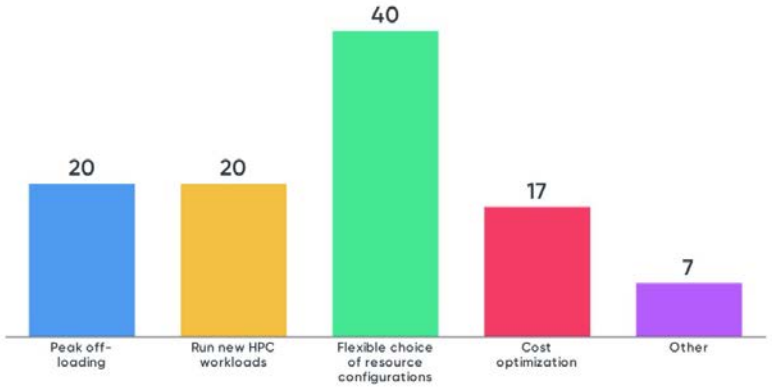

58

What technologies are you using or interested in today?

Mentimeter

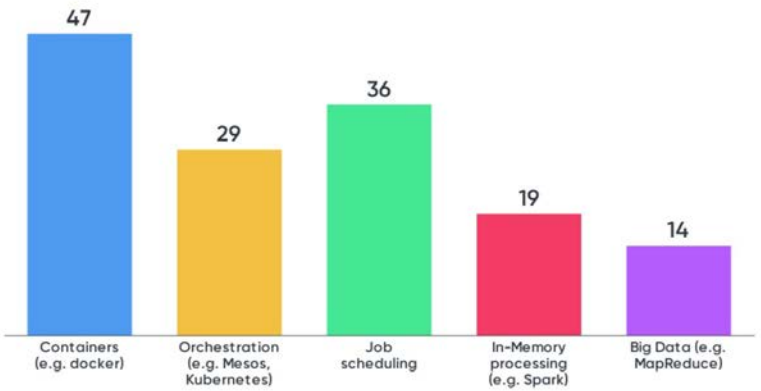

55

What would you need to exploit more benefits from AI?

Mentimeter

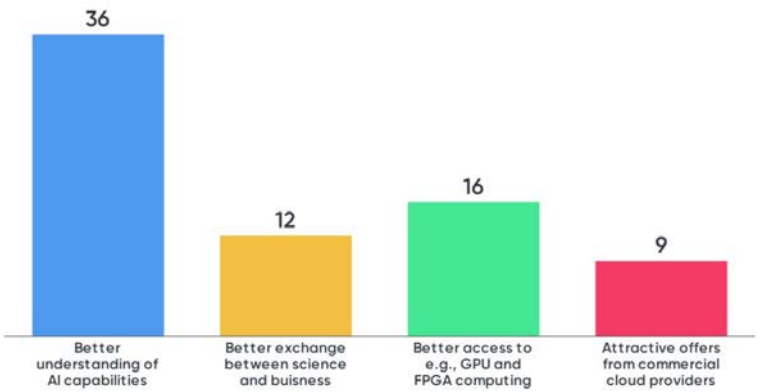

49

KEY CONTACTS

JURRY DE LA MAR

+49 170 9236058
jurry.delamar@t-systems.com

THANK YOU!

T · · Systems ·