

LARGE-SCALE DATA PROCESSING AND HPC FOR SCIENCE WITH T-SYSTEMS HYBRID CLOUD



Open
Telekom
Cloud



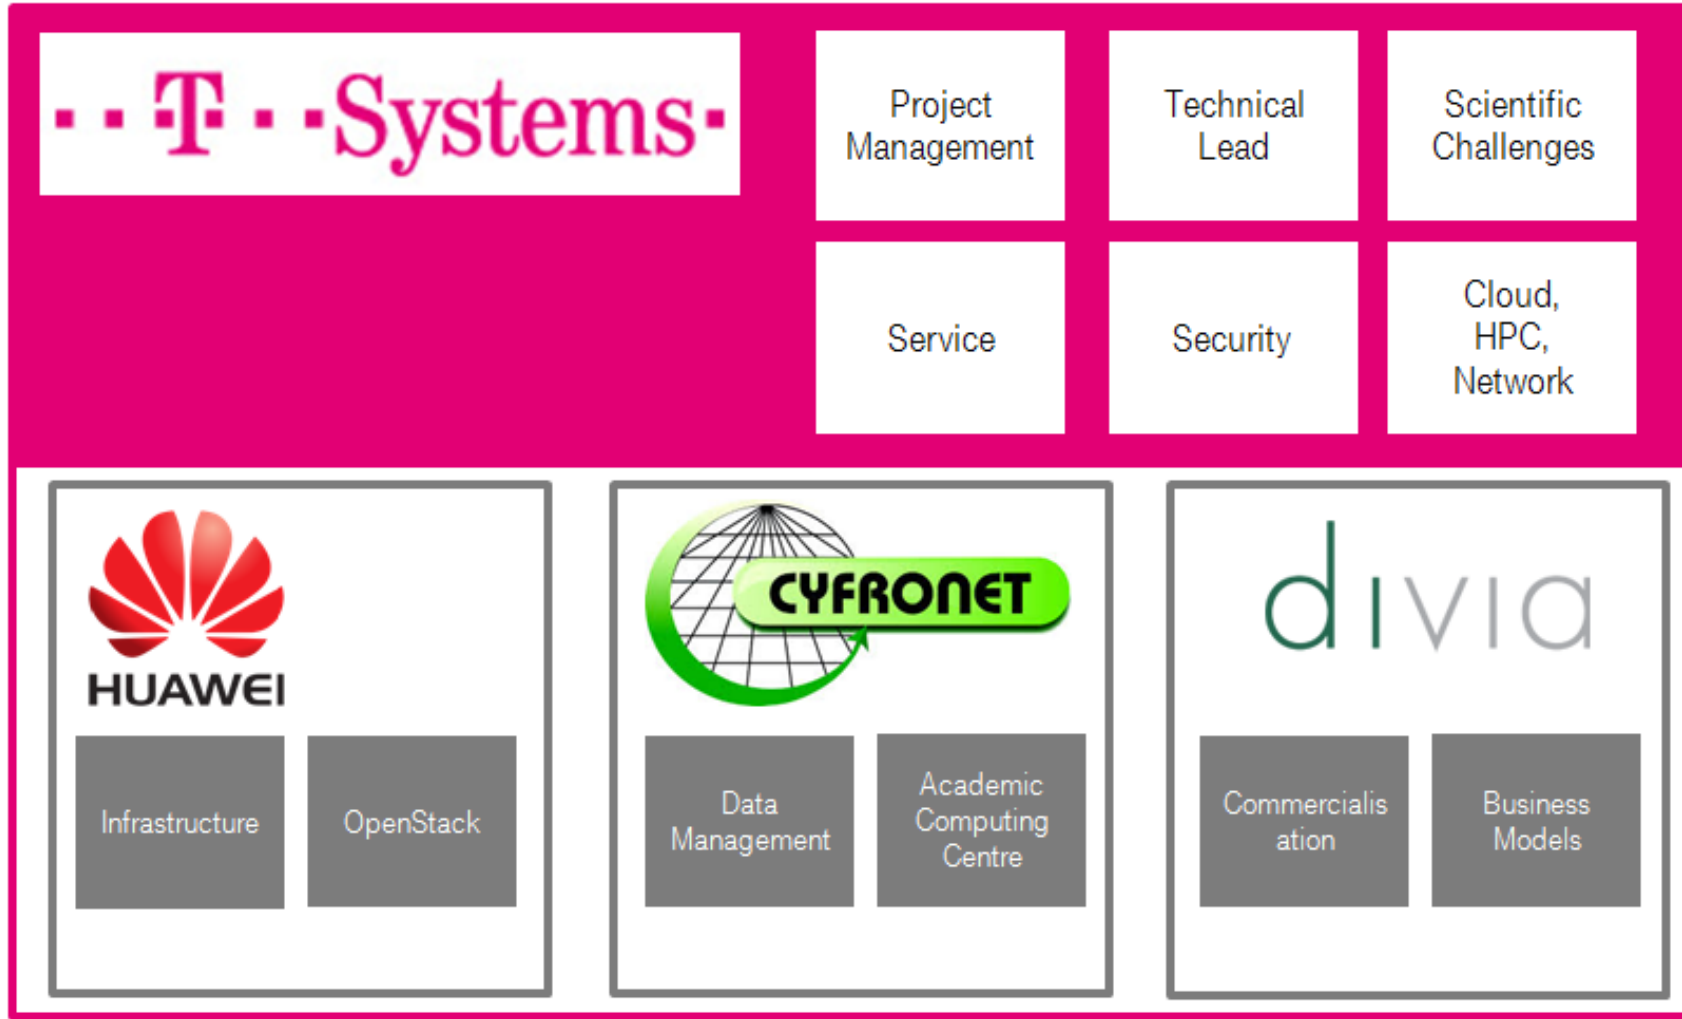
Jurry de la Mar, T-Systems

CHEP 2018, July 9th, Sofia

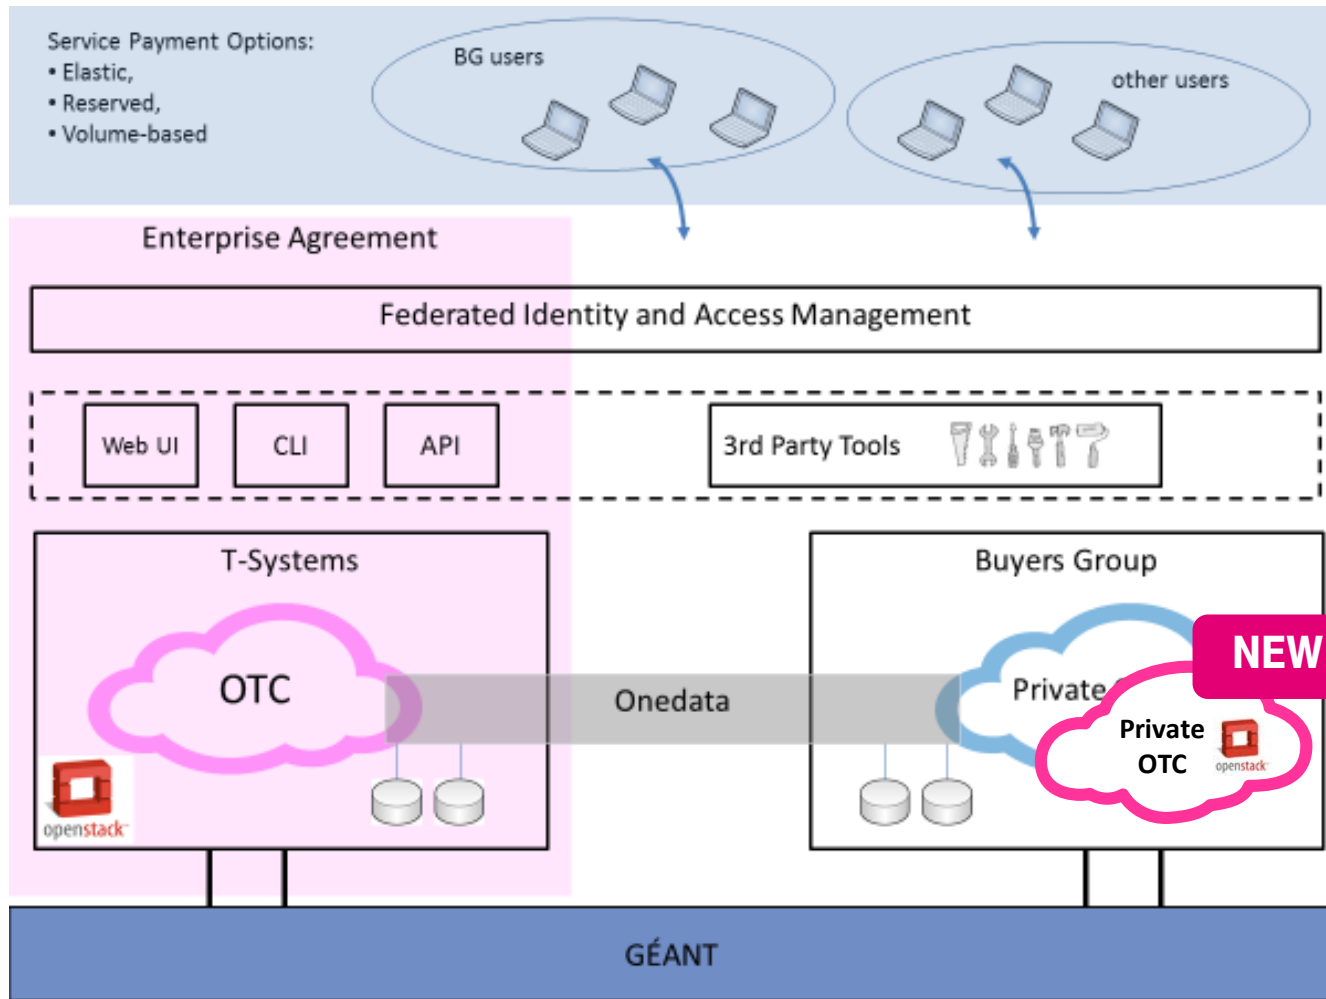
T · · Systems ·



T-Systems HNSCICLOUD TEAM



HNSCICLOUD PILOT – T-SYSTEMS SOLUTION



Highlights

- OpenStack Cloud and APIs
- eduGAIN and ELIXIR AAI access
- Extensive 3rd party tools support
- HPC and Big Data options
- GÉANT and Internet networking
- Help Center & Enterprise Agreement (SLA, custom support)

Features

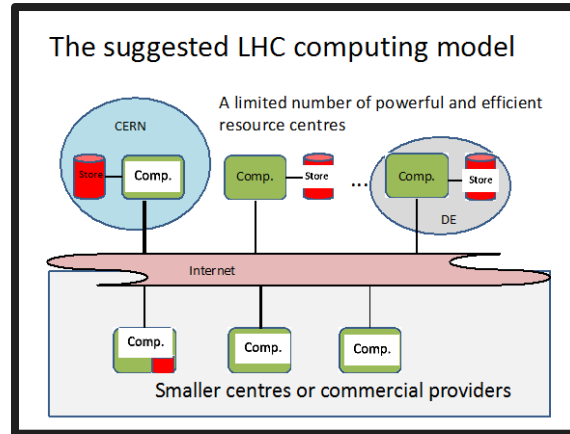
Benefits

- Transparent processing and data mgmt. between institute and public cloud
- On-demand burst and scale-out
- Flexible service payment
- Committed services
- Reduced costs for science experiments

HYBRID CLOUD LIGHTHOUSE USE CASES



CERN
FEDERATED KUBERNETES CLUSTER
 (CLENIMAR FILEMON,
 RICARDO ROCHA)




INFN
DODAS SERVICE FOR HEP
 (DANIELE SPIGA)

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

CERN OPENLAB
DEEP LEARNING FOR FAST SIMULATION
 (SOFIA VALLECORSA)

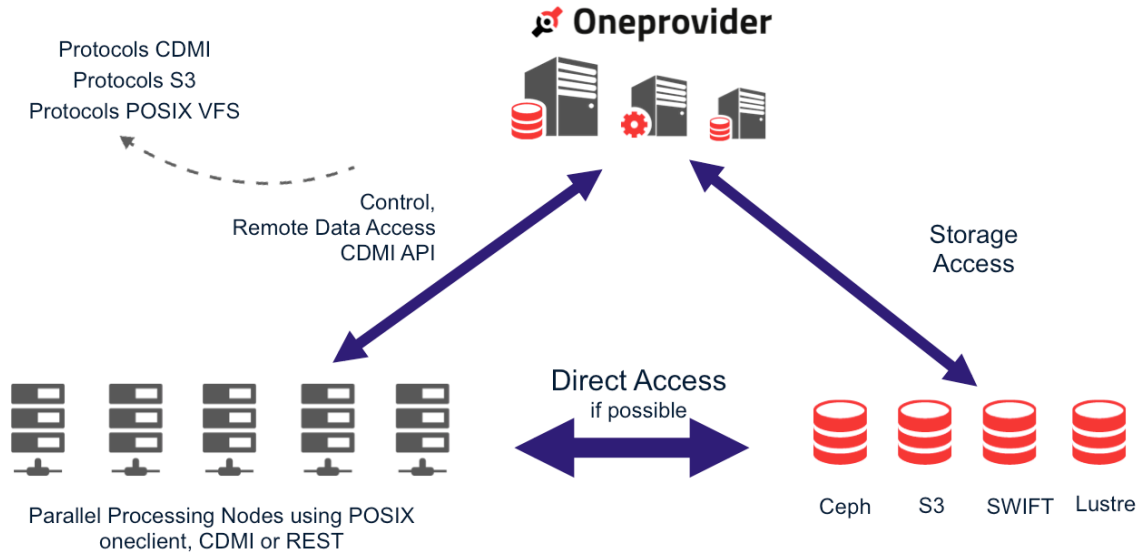
DATA MANAGEMENT CHALLENGES



HETEROGENEITY

DATA LAKES

TIME TO WISDOM



STORAGE FEDERATION AND MULTI-TENANCY

- Ease of access to distributed data in hybrid cloud
- Secure data sharing between users

UNIFIED TRANSPARENT ACCESS TO HYBRID-CLOUD STORAGE RESOURCES


- Transparent data access across hybrid-cloud deployment based on POSIX and CDMI

<https://onedata.org>

<https://onedata.hnsc.otc-service.com>

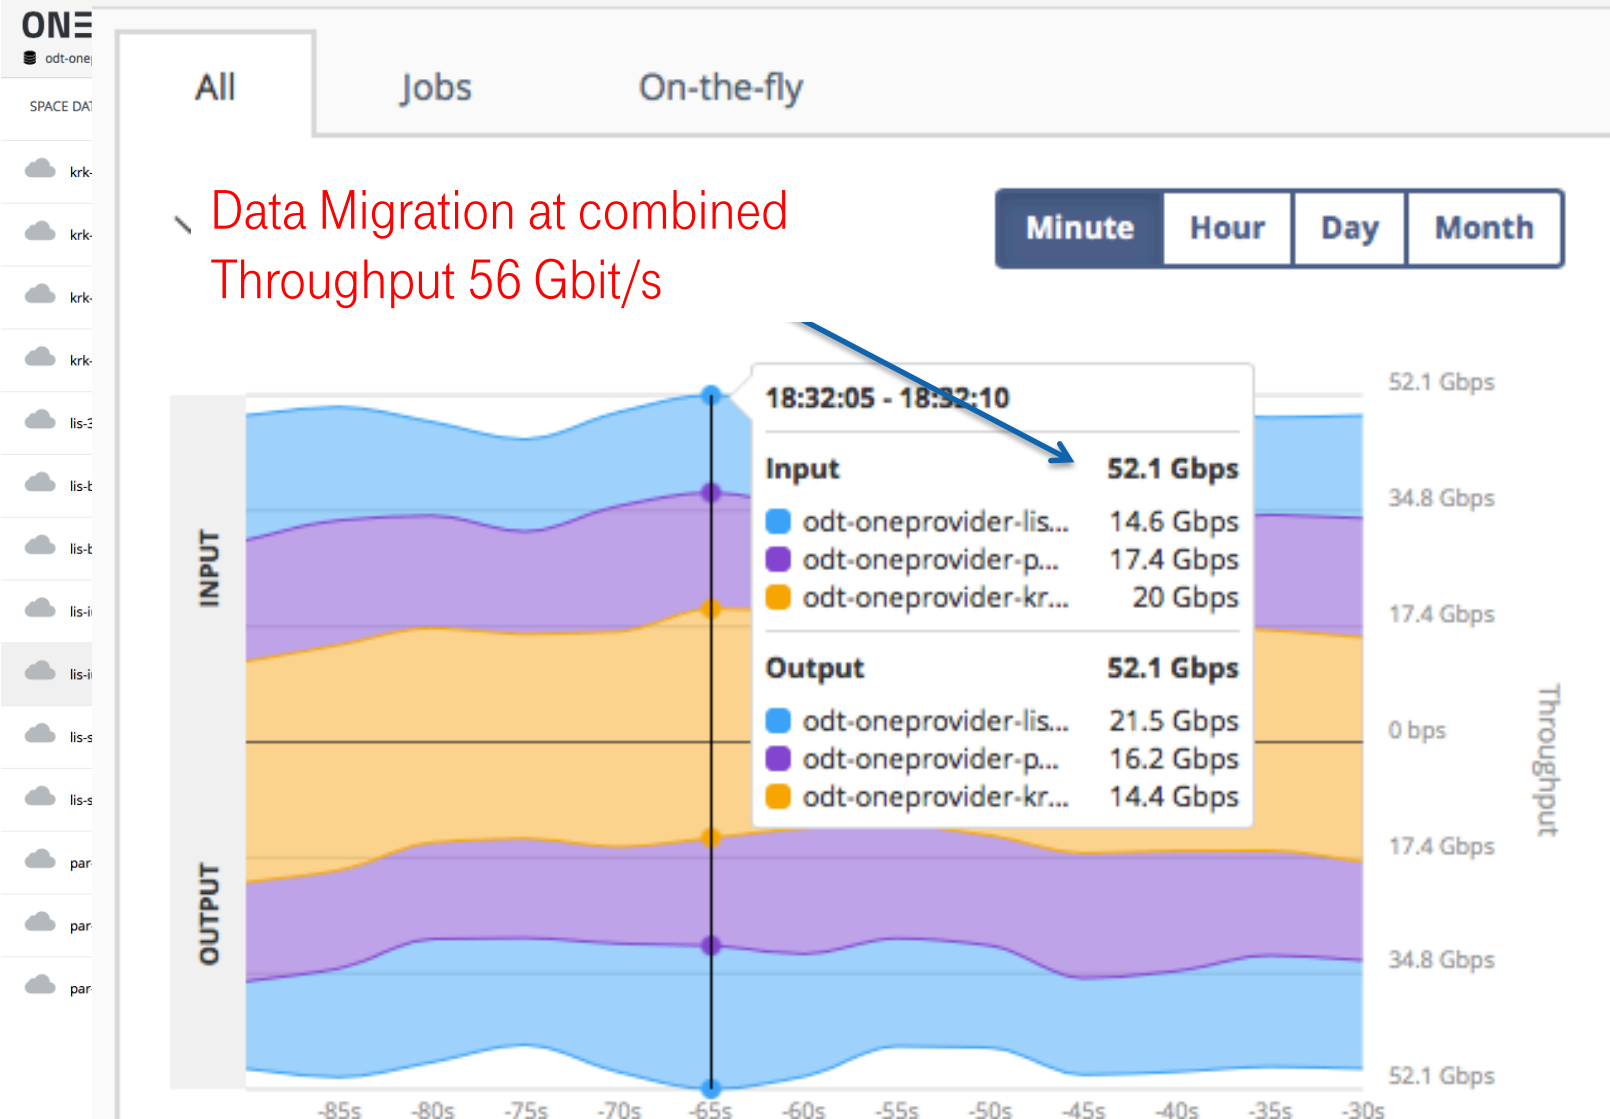
ONEDATA HIGH-THROUGHPUT DATA PROCESSING ON OTC

Oneclients



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

HIGH-THROUGHPUT DATA MIGRATION IN PEER TO PEER MESH



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

ONEDATA FILE POPULARITY AND SMART CACHING

The screenshot displays the ONEDATA PROVIDER PANEL interface. The left sidebar shows navigation options: CLUSTERS, Nodes, Provider, Storages, and Spaces. The main content area shows the configuration for a cluster named 'krk-n-par-3'. The 'Providers support' section shows a total of 1.9 GiB, with a donut chart indicating 953.7 MiB from 'release17060-rc8...' and 953.7 MiB from 'Provider#VIF-F-T'. The 'AUTO CLEANING' section is active, showing 'CLEAN REPLICATED FILES' and 'CLEANING BOUNDARIES' settings. A progress bar indicates 'Total space: 953.7 MiB', 'Used space: 384.1 MiB', 'Free space: 569.6 MiB', and 'To release: 179 MiB'. A 'CLEANING REPORTS' table is visible at the bottom.

ONEDATA PROVIDER PANEL

CLUSTERS

Search...

release17060-rc8-oneprovider-paris

Nodes

Provider

Storages

Spaces

krk-n-par-3

Name: krk-n-par-3

Id: IHa7SuFLUjm-3XEAPGogxWSQR_tixPIMwqMuf950rn4

Mount in root:

This provider storage: s3

Providers support (total 1.9 GiB)

Chart Table

1.9 GiB

release17060-rc8... 953.7 MiB

Provider#VIF-F-T 953.7 MiB

Storage synchronization Files popularity Auto cleaning

AUTO CLEANING

CLEAN REPLICATED FILES Saving...

Lower size limit: 200 MIB

Upper size limit: MIB

Not opened for: Hours

CLEANING BOUNDARIES

Total space: 953.7 MiB Used space: 384.1 MiB Free space: 569.6 MiB To release: 179 MiB

205.1 MiB 711 MiB

Start cleaning now

CLEANING REPORTS

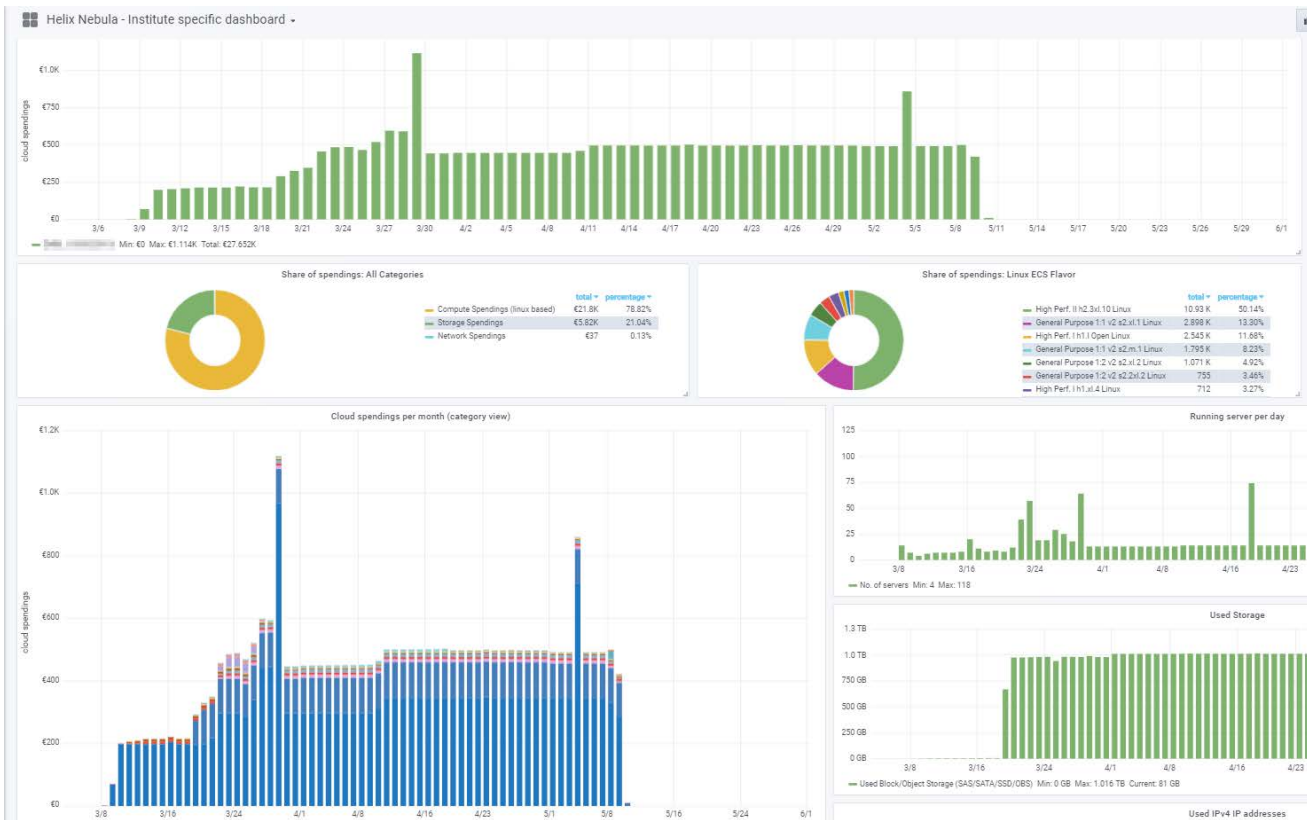
Start	Stop	Released size	Files number	Status
8 Dec 2017 13:32:13	-	0 B (out of 179 MiB)	de la Mar - CHEP 2018	09/07/2018
8 Dec 2017 13:31:53	8 Dec 2017 13:31:55	448 MiB (out of 565.4 MiB)	1	✖

Onedata can evict least important replicas
Based on file popularity



ONEDATA

SERVICE PAYMENT MONITORING & CONTROL



Highlights

Functionalities

- Global view
- One Dashboard per Institute
- Budget alert
- Dashboard can be customized
- Automated data upload
- Responsive UI design

Benefits

- More transparency of costs
- Budget management
- Open source
- User self-service

A TOP PRIORITY – DATA PRIVACY AND SECURITY

<https://cloud.telekom.de/en/infrastructure/open-telekom-cloud/compliance>



CERTIFICATION

ISO 20000, ISO 27001, ISO 27017, ISO 27018, CSA Star Level 2, TISAX/ENX, Trusted Cloud (EU-GDPR) and TÜV Trusted Cloud



DATA STORED IN GERMANY

All data stored in Germany only, meeting EU and German legal requirements



LEGALLY COMPLIANT

In accordance with Sections 3 and 11, German Federal Data Protection Act (BDSG)



PRIVACY

Technical and organizational measures (TOM) for contract data processing according to EU-GDPR



CONTRACTUALLY DEFINED

Information on legislation provided in advance in the service specifications, T&Cs and contract data processing agreement

MORE SCIENCE WITH MORE CLOUD



OPEN TELEKOM CLOUD
PRIVATE, PUBLIC OR HYBRID

<https://cloud.telekom.de/en/infrastructure/open-telekom-cloud>

THANK YOU.

