



CERN openlab projects with Oracle

Riccardo Castellotti – Viktor Kozlovsky

09/03/2021

Acknowledgments

CERN: Luca Canali, Riccardo Castellotti, Ignacio Coterillo Coz, Eva Dafonte Perez, Lukas Gedvilas, Eric Grancher, Jakub Grazieczny, Alina Grigore, Arash Khodabandeh, Viktor Kozlovsky, Manuel Martin Marquez, Sebastien Masson, Antonio Nappi, Nemanja Nedic, Ioannis Panagiotidis, Luis Rodriguez Fernandez , Aimilios Tsouvelekakis, Artur Wiecek (and many others!)

Oracle: Cemil Alper, Giuseppe Calabrese, Michael Connaughton, Dmitrij Dolgušin, David Ebert, Brent Eyler, Maciej Gruzka, Sevgi Guzzella, Simone Indelicato, Gavin Larson, Vincent Leocorbo, Will Lyons, Pauline Mahrer, Marc Meignier, Çetin Özbütün, Oguz Pastirmaci, Cristobal Pedregal-Martin, Arun Ramakrishnan, Alexandre Reigada, Monica Riccelli, Patrice Scattolin, Engin Senel, Garret Swart, Gabriela Stanescu, Peter Szegedi, Thomas Teske, Reiner Zimmermann

Thanks Oracle and everyone for the support and the work done together

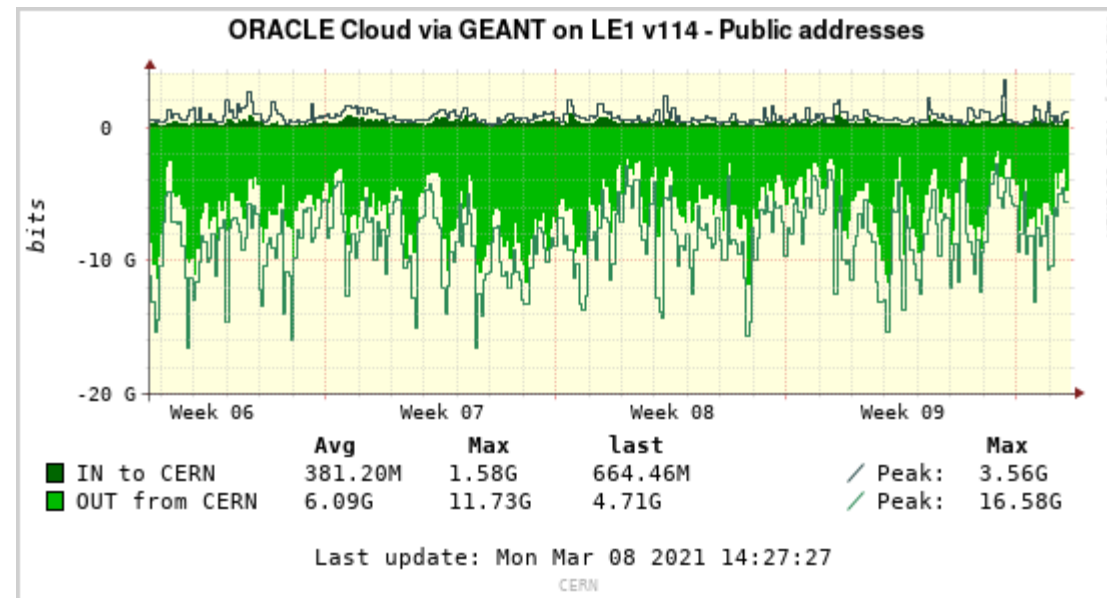
Motivations and Vision

- Partnership to explore new technologies in cooperation with Oracle
 - During this openlab phase, the focus was mostly on cloud services
- Prove that it is possible for an organization with existing infrastructure to profit from public clouds
 - Test new solutions with Oracle development teams
 - Move existing use cases to **cloud native** solutions
 - New solutions: Disaster Recovery

Dedicated connection to OCI Frankfurt

Use OCI as an extension of CERN data center

- Thanks to FastConnect and GÉANT, there is a private connection between CERN and OCI in Frankfurt
- Nodes connected through this link have dedicated bandwidth and are routable from inside CERN

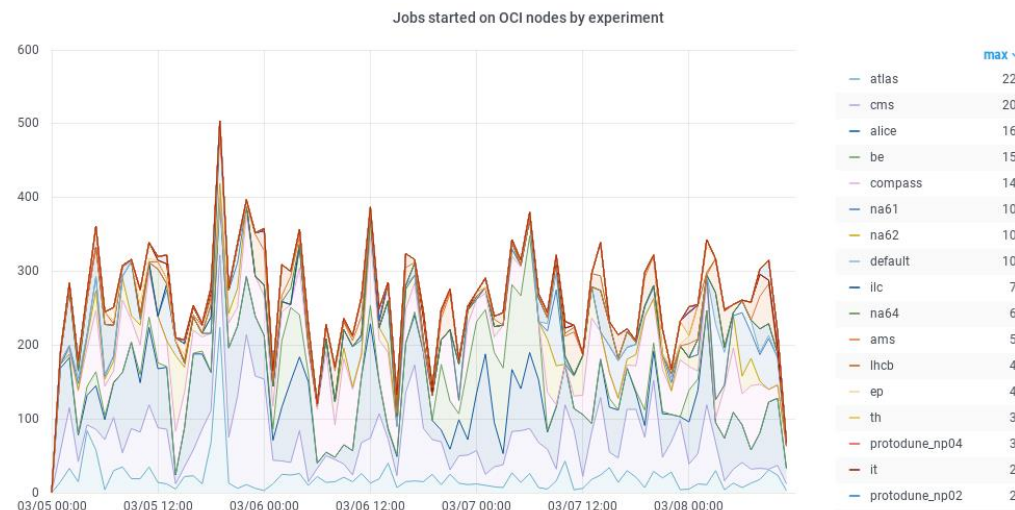


From <https://netstat.cern.ch>

Running CERN Tools in OCI

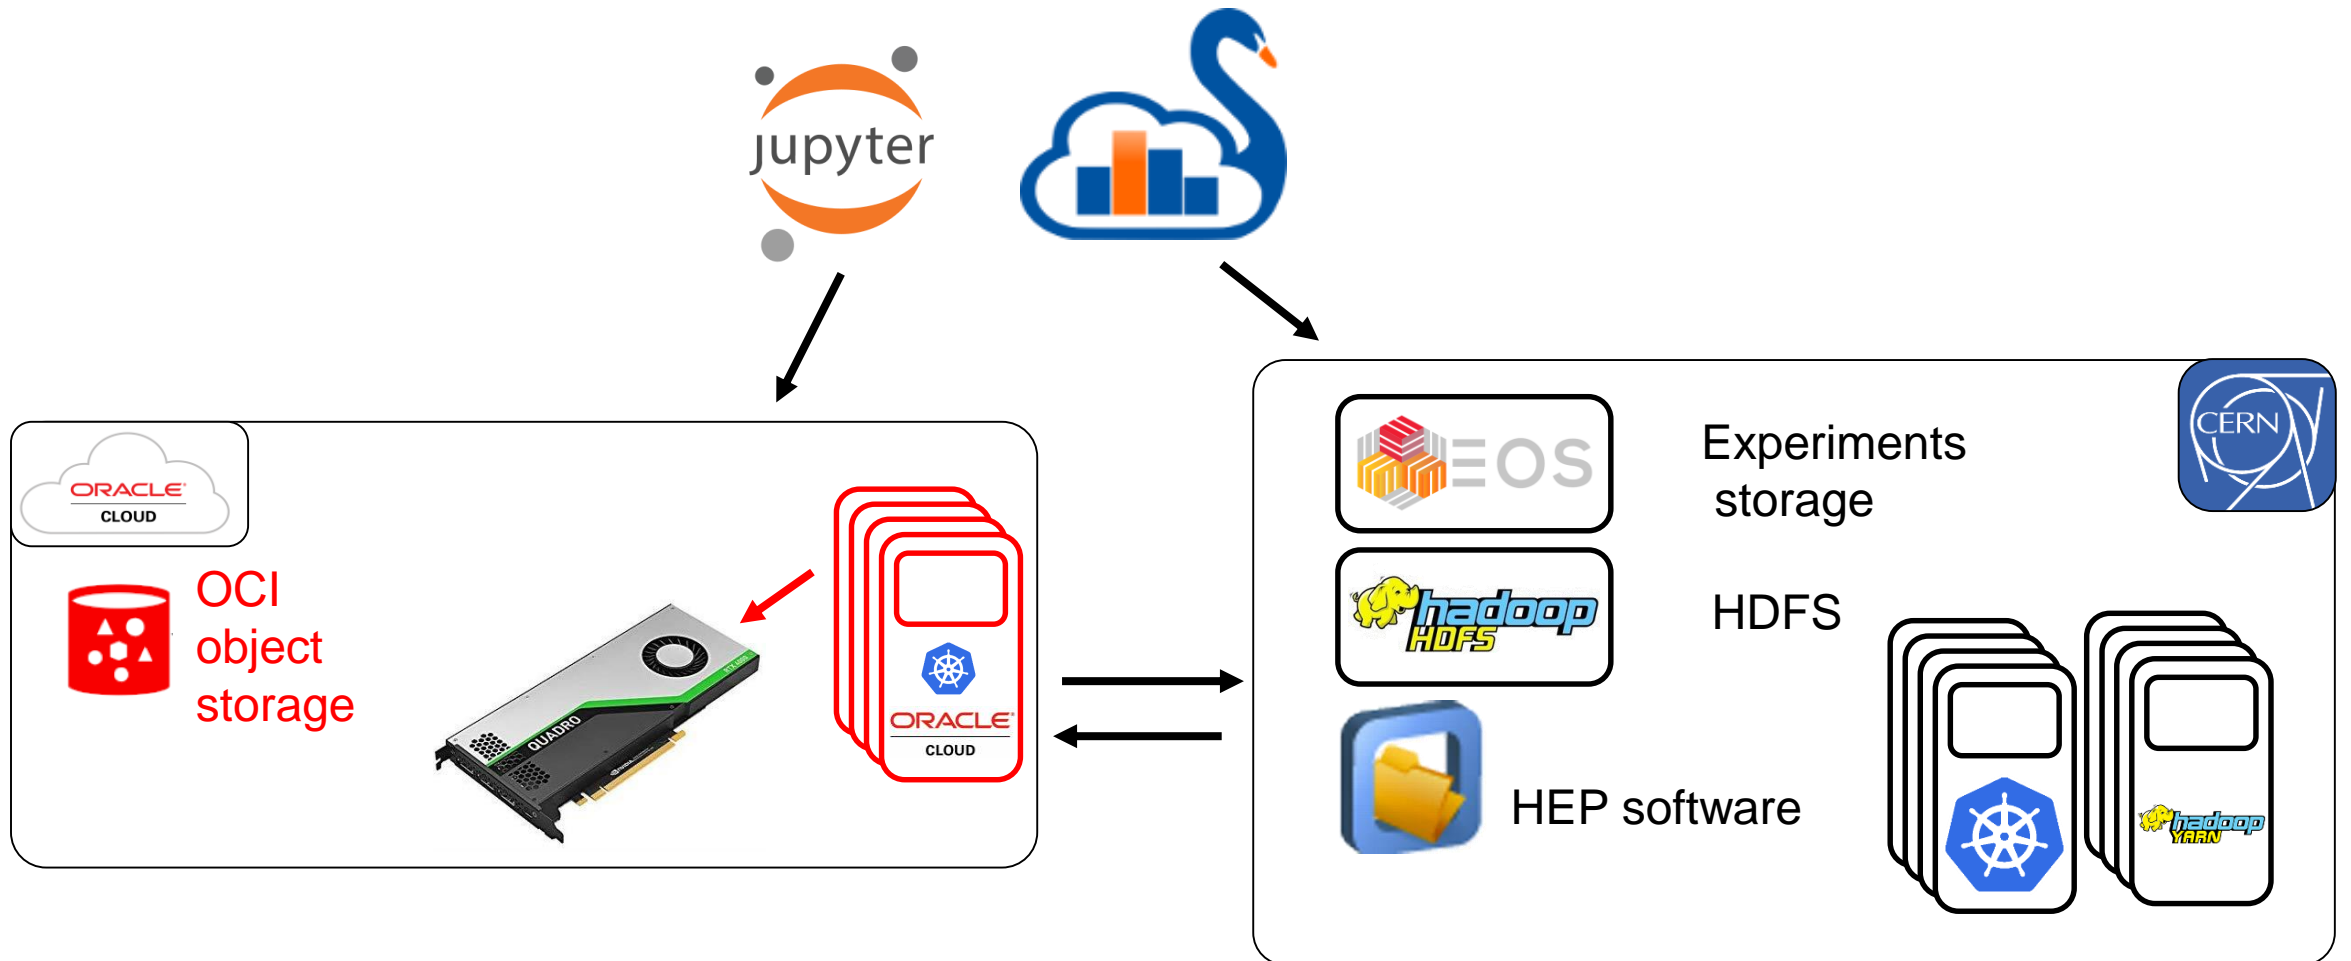
Integrate nodes in OCI with existing CERN services

- Implemented `oci-bs` :
 - Creates a VM in the OCI with the OCI SDK
 - Registers the node in CERN network database
 - Runs internal tool to register the node in other CERN services
- Deployed 128 batch compute nodes (collaboration with Batch service)



Analytics Platform


Compute across OCI and CERN cloud, storage at CERN



ML Tools Scalability

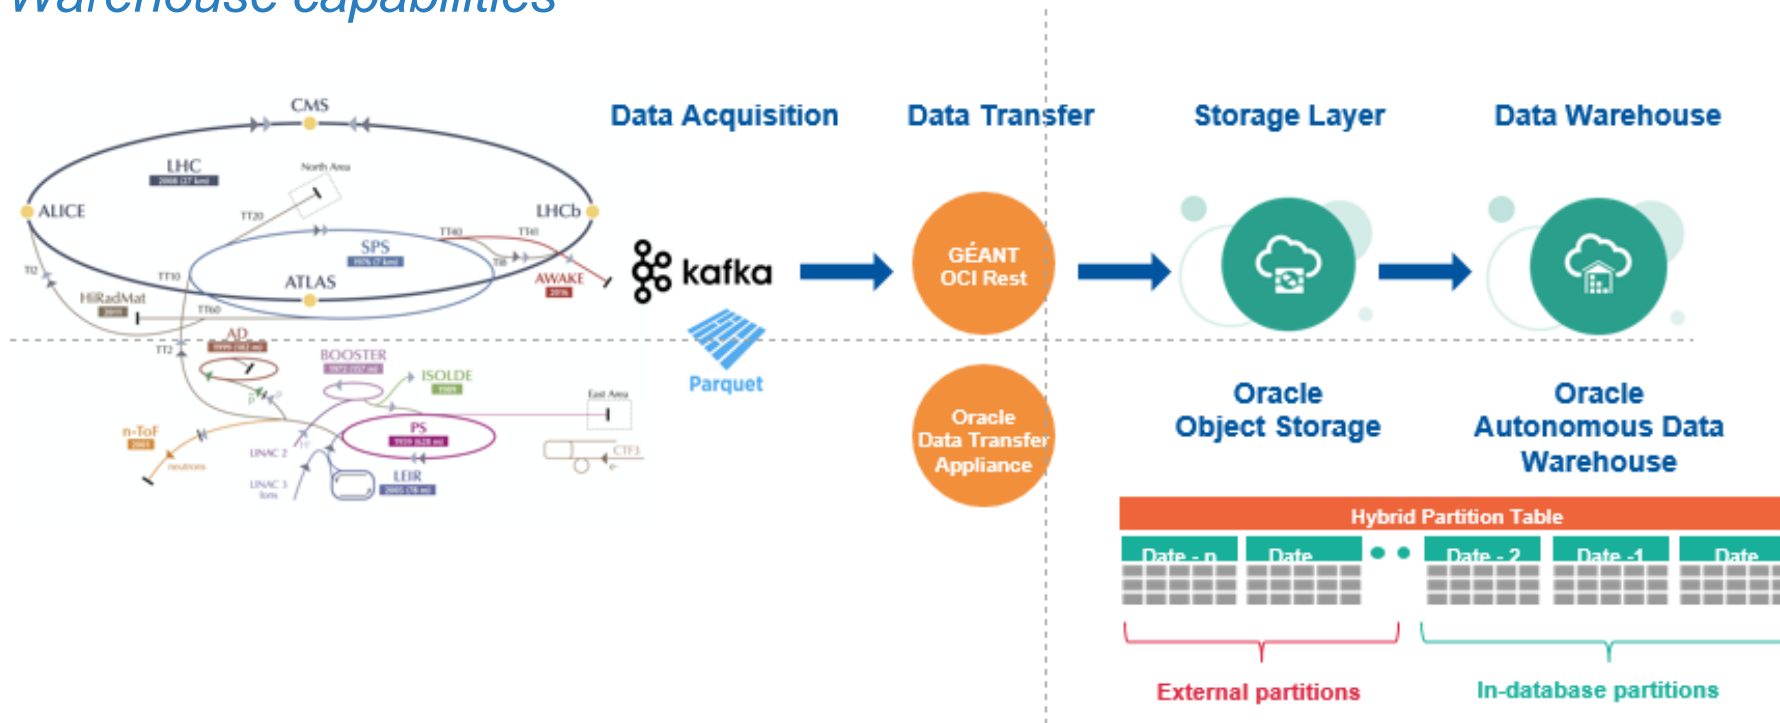
“Machine Learning Pipelines with Modern Big Data Tools for High Energy Physics”,
Comput Softw Big Sci 4, 8 (2020)

- Can we **parallelize** the workload execution **efficiently**?
- We tested up to 480 CPU cores and 10 GPUs



Autonomous Data Warehouse

Use CERN controls system's complexity as an advanced use case to test Autonomous Data Warehouse capabilities



- Used data from CERN controls Middleware (CMW)
- Worked with the development team to support **nested and complex data types**
- **Hybrid** solution based on **in-database** and **external table partitions**



CERN openlab projects with Oracle part II.

Viktor Kozlovsky

Openlab Technical workshop (2021.03.09-11)