



# EOS site report of the Joint Research Centre

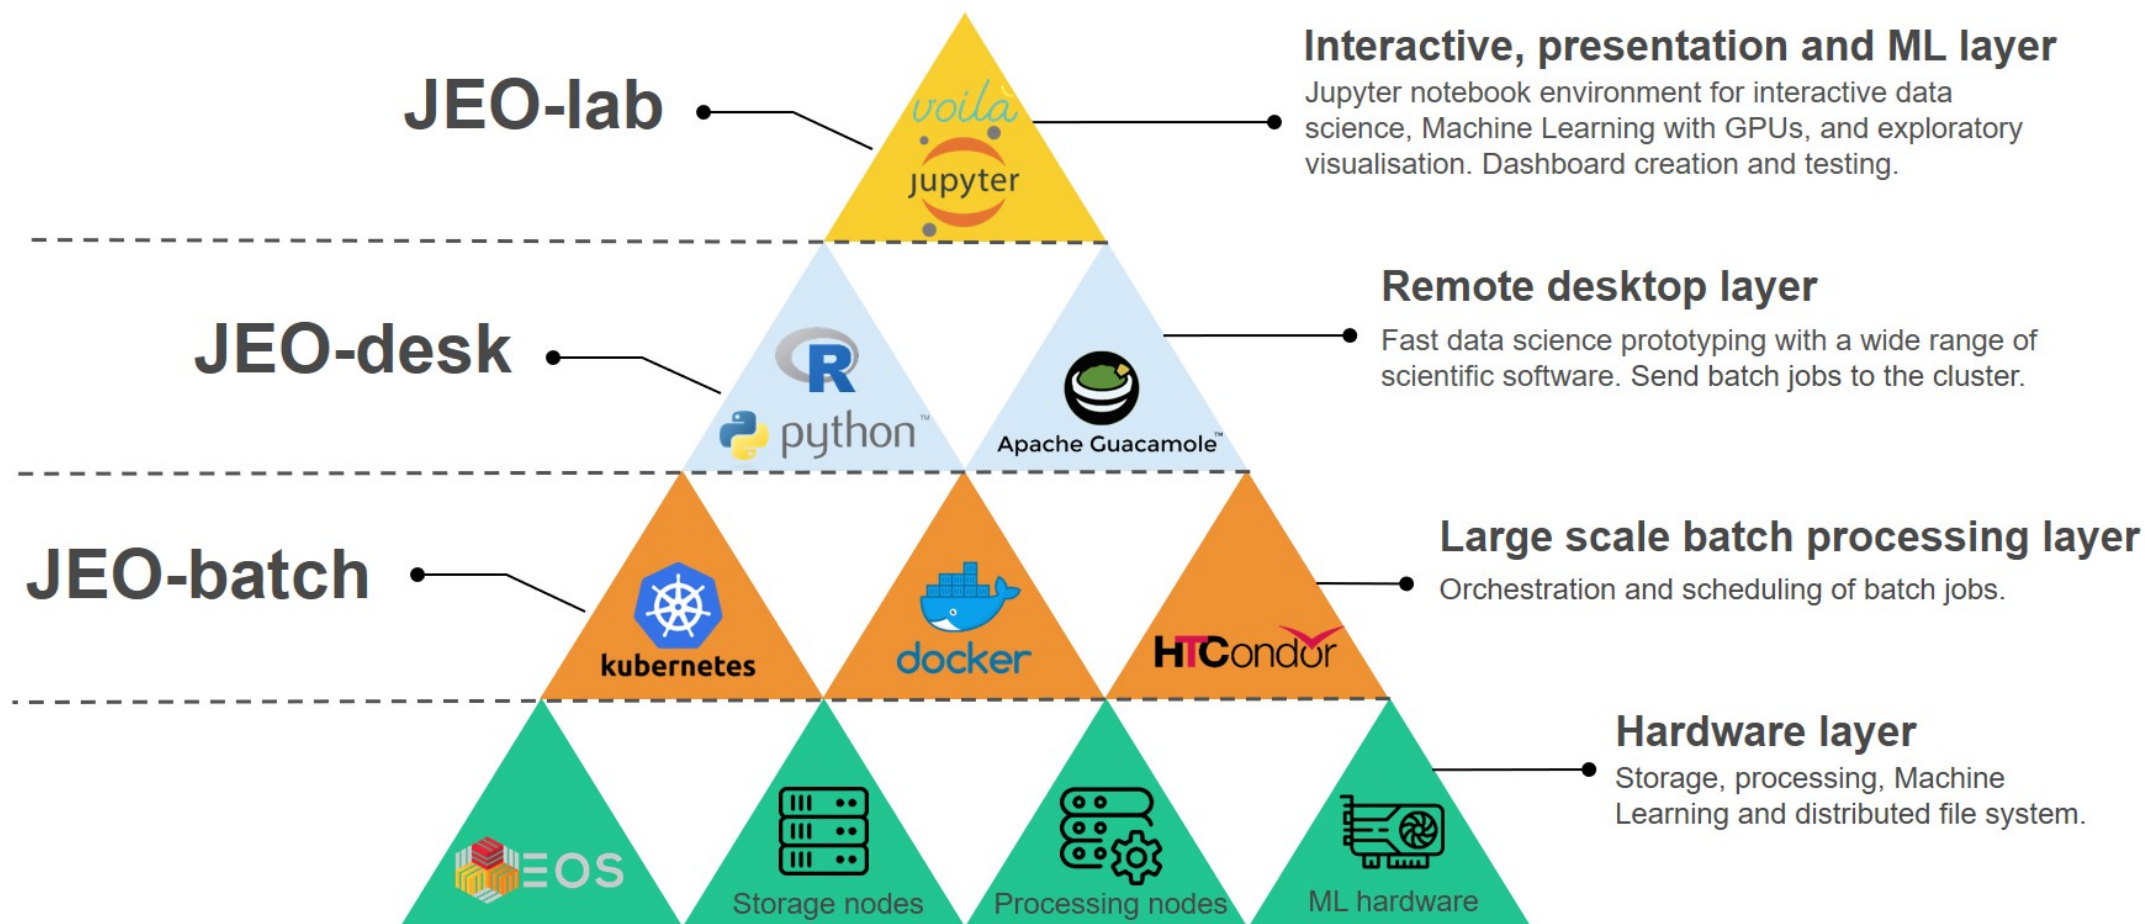
*EOS workshop March 2022*

*Armin Burger, Franck Eyraud, Marco Scavazzon*

# EOS usage at the Joint Research Centre

- Storage backend of the JRC Big Data Analytics Platform (BDAP)
- Used as file storage by >50 JRC projects
- Main data collections are geo-spatial and meteo data
  - Copernicus programme, NASA/USGS, ECMWF (meteo), etc.
- Growing use for other data domains
  - social media, economy, health data, any type of text data, etc.


# Big Data Analytics Platform – Services



# Big Data Analytics Platform – Software



PostgreSQL



# EOS set-up at JRC – Overview



- EOS in production since 2016
- 1 production instance, 23.6 PB current gross capacity
- 1 small test instance on VM's
- MGM using QuarkDB namespace
- Access via FUSE(X) client only
- Replica 2 used throughout

# EOS instance at JRC – Software

- EOS versions
  - MGM: 4.5.15
  - FST: 4.5.17
  - Processing & service nodes: 4.8.x
- QuarkDB : 0.4.2
- EOS FUSE client:
  - FUSEX: all processing and service nodes
  - FUSE: Samba gateway, remote desktop services, a few legacy services
- Server OS: CentOS 7

# EOS instance at JRC – Hardware

- Heterogeneous hardware
- JBODs
  - Xyratex, Lenovo, Supermicro, Quanta, Western Digital, Seagate
  - Disks per JBOD: 12, 24, 60, 84
  - Disk sizes: 6, 10, 14, 16 TB
- FST servers
  - IBM/Lenovo, HPE
  - 12-20 cores, 128-256 GB RAM, 1-2 TB space for logs
  - 10 Gbps Ethernet bonded network connection



# EOS instance at JRC – Current status

- 58 configured FSTs with 1 or 2 JBODs connected
- 23.6 PB available gross space, soon extended by 8 PB
- 21.3 PB used, 2.3 PB free space
- 1.35 G files (2.1 G total created files)
- 242 M directories (262 M total created directories)
- More than 1 year uptime



# EOS and user interaction

- Typical user expectation: use EOS like a local disk space
- Requires some “education” of users:
  - Read documentation and be aware of some limitations
  - Write processing output to local scratch space and move to EOS afterwards
- Mitigate better possible crashes of FUSE client, e.g. via dedicated containers or side-car containers
- Put user home directories on NAS storage with NFSv4 mount

# EOS activities and issues in 2021

- Completion of network migration: move EOS instance to a new private network
- Crash in the fusex client when recursive rm block is enable (rm-rf-protect-levels) fixed in the code
- Issue with draining an offline node, then putting it back online
- Heterogeneous nodes: added nodes with more disks than previous ones, prefer to disable small scheduling groups until more disks than 3 are available
- The fact that we have a lots of files (1M per disk) is a potential cause of difficulties (long listing time, difficulty with maintenance, etc...)

# EOS planned activities in 2022

- Upgrade EOS to latest 4.8.x version; later upgrade to 5.x TBD
- Extend capacity with additional 60-disk JBODs
- Decommission oldest FSTs/JBODs (6+ years)
- Test of RAIN layout for data collections with large data files
- Test OAuth2 & token-based authentication
- Test HTTP/S3 interface for web application clients
- If feasible migrate all FUSE clients to FUSEX

# Thanks to

## EOS development and management team for the support



© **European Union** 2022

Unless otherwise noted the reuse of this presentation is authorised under the [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/) license. For any use or reproduction of elements that are not owned by the EU, permission may need to be sought directly from the respective right holders.

