



EOS for CERNBox Report

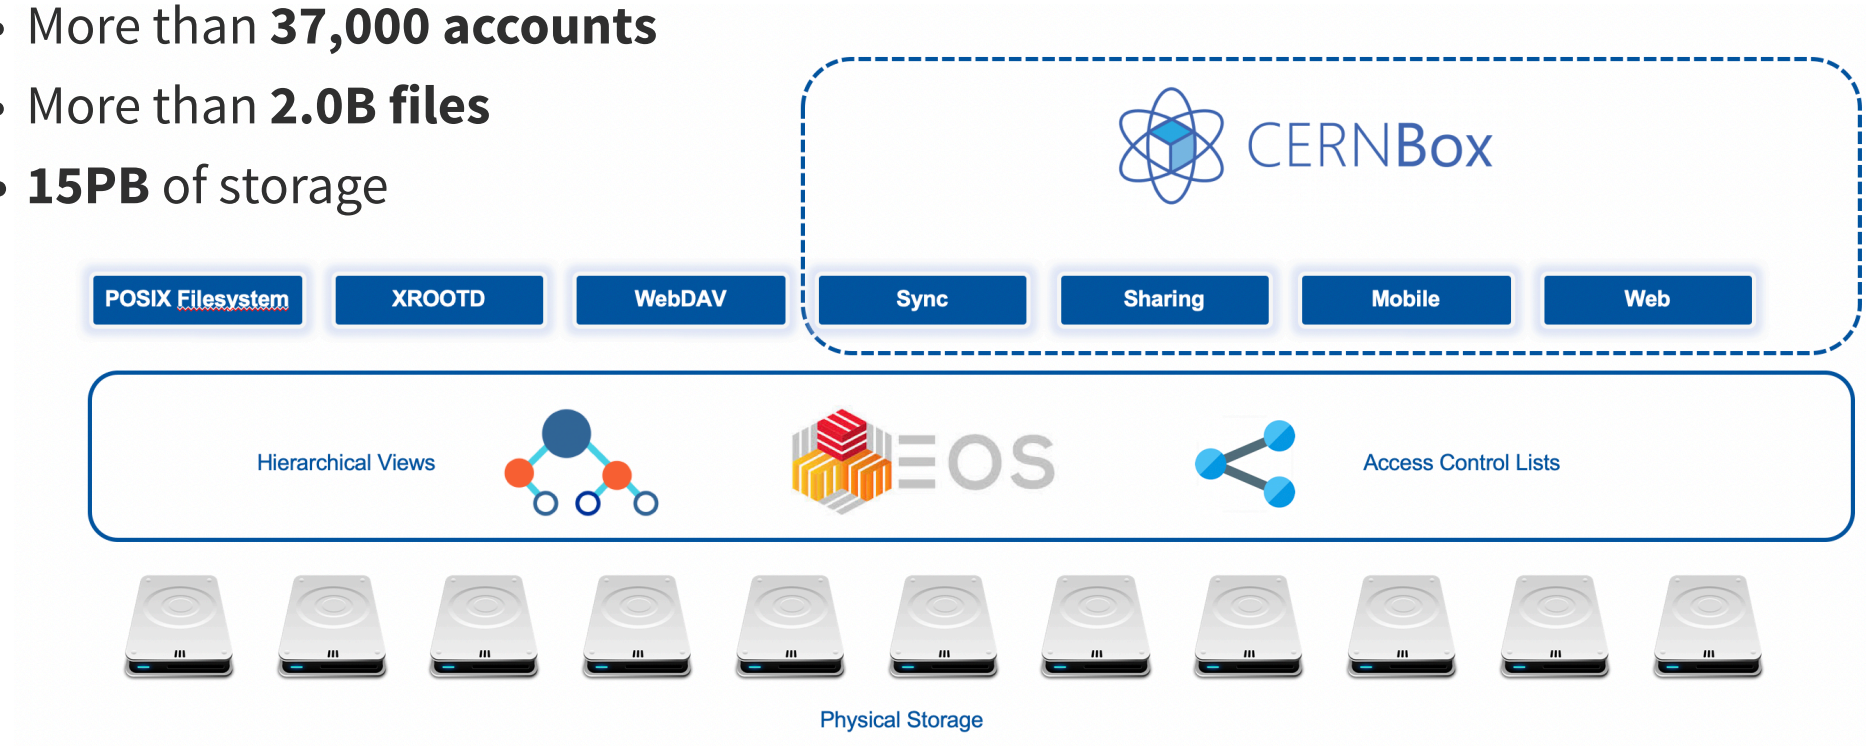
Roberto VALVERDE, in behalf of EOS CERNBox Team (IT-ST, CERN)

08/03/2022

CERNBox: Powered by EOS

CERNBox is the cloud sync and share service used at CERN.

- More than **37,000 accounts**
- More than **2.0B files**
- **15PB** of storage



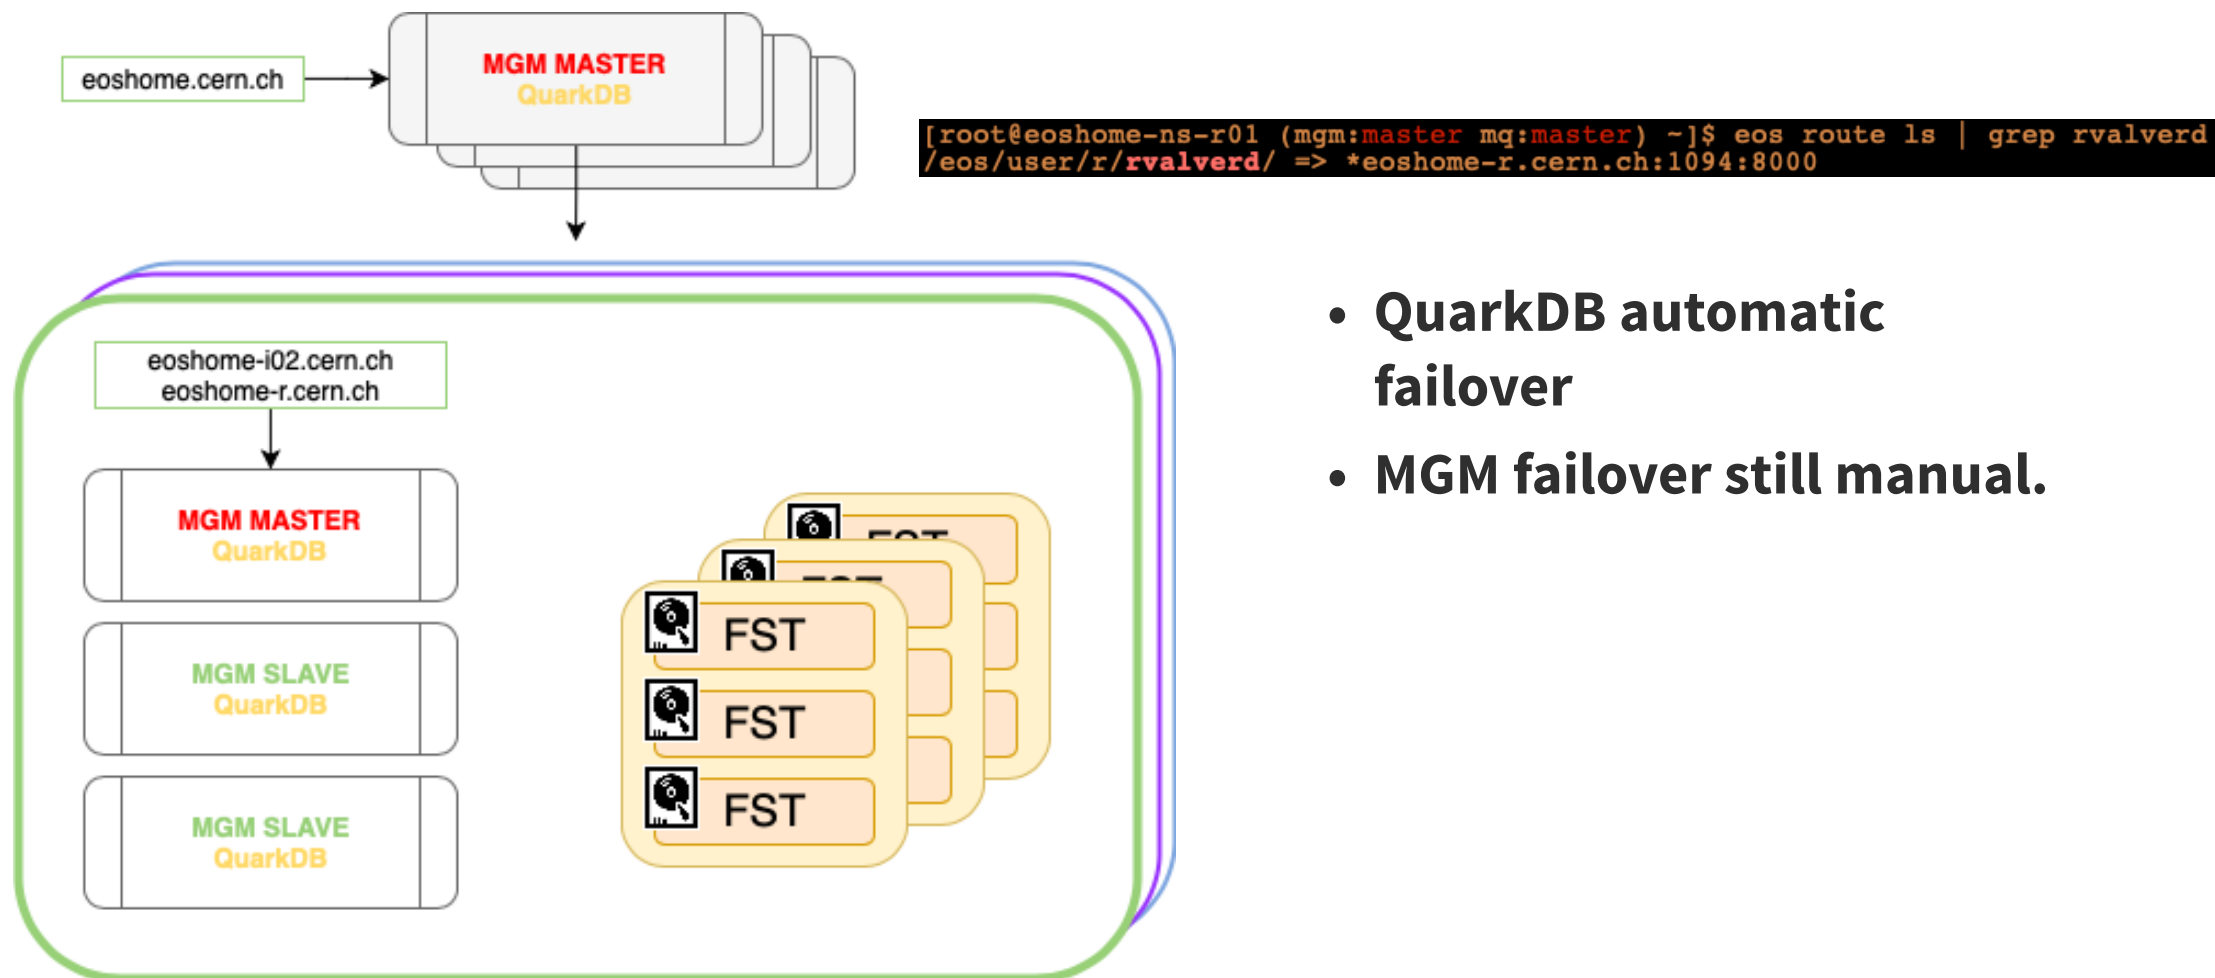
Some numbers...

10 EOS instances

- **5 for HOMES**
 - **3 for PROJECTS**
 - **1 for Audio/Video data (EOSMEDIA)**
 - **1 Pre-Production - EOSHOME**
- Canary**

Instance	Used	Capacity	Disk servers
EOSHOME-00	1.6P	3.3P	9
EOSHOME-01	2.2P	3.8P	11
EOSHOME-02	1.4P	2.9P	9
EOSHOME-03	1.5P	3.2P	9
EOSHOME-04	1.4P	2.7P	7
EOSPROJECT-00	1.4P	2.9P	8
EOSPROJECT-01	1.1P	2.3P	7
EOSPROJECT-02	1.2P	2.3P	6
EOSMEDIA	2.8P	4.5P	12
EOSHOMECANARY	2.6T	618T	5
TOTAL	15.0P	29.0P	78

EOSHOME: Architecture



- QuarkDB automatic failover
- MGM failover still manual.

EOSHOME: Hardware

MGM + QuarkDB (x3)

- CC7
- Intel Xeon, 32 cores.
- 192G/384G RAM
- 2 x 2T SSD (1 for QuarkDB)
- One node per switch
- 10G Network

FST

- CC7
- Intel Xeon/AMD Epyc, 20/32 cores
- 128G/256G RAM
- 48x6T/14T JBOD, 1 filesystem / disk.
XFS
- 10G/25G Network


Service availability



- **99.999%** availability in last year
 - Stable upgrade cycles
 - Pre-production testing (EOSHOME Canary)
 - Software improvements (Kudos EOS dev Team!)
 - Better handling of batch overloads.


Service growth

- **817TB** written in last year (+6%)



Service growth (II)

- **2 Billion** files reached (+317M, +22%).
- **200 Million** directories reached (+49M, +29%)



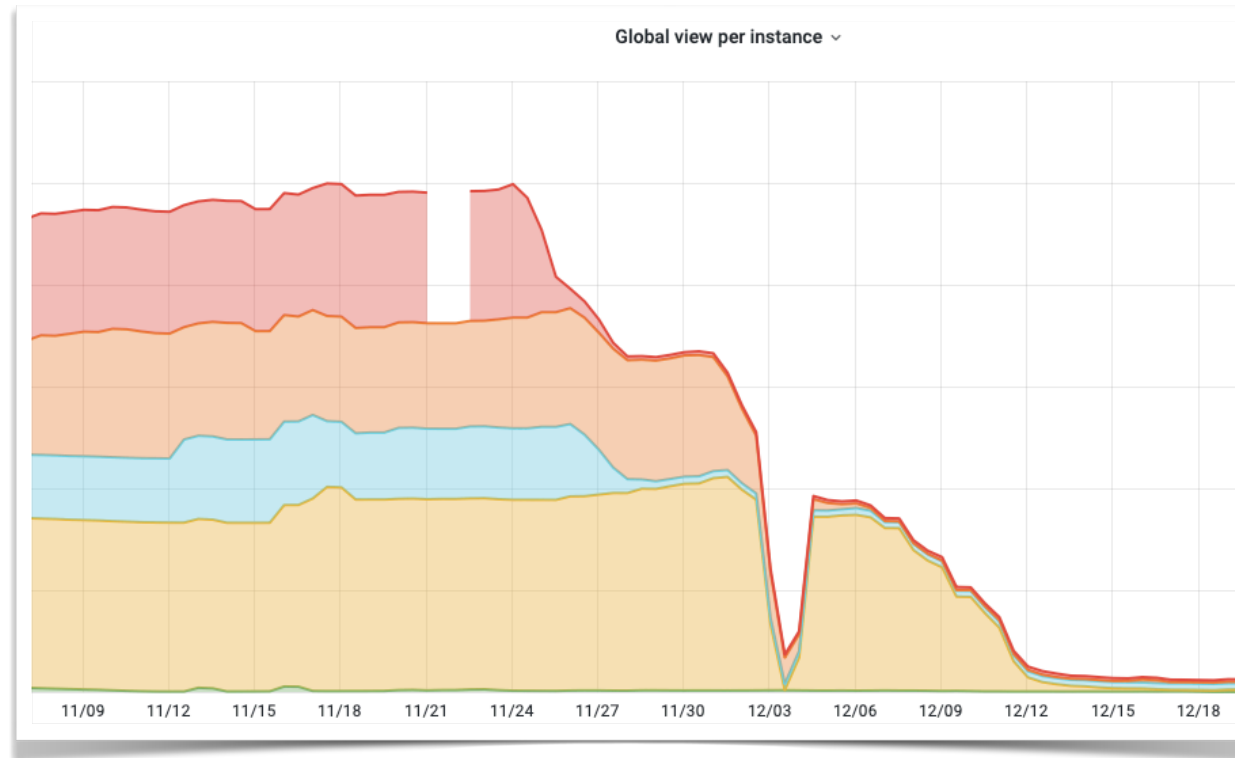
2021 Highlights

- **EOS on CephFS** Evaluation
 - FST on CephFS instead of directly on hard disks
 - Fully virtualised EOS
 - Streamlined operations in ST-GSS
- Check my presentation later today [\[Link\]](#)



2021 Highlights

- **EOS FSCK** in automatic repair enabled
 - Make sure you use **EOS v4.8.76** or newer to use this feature.



2021 Highlights

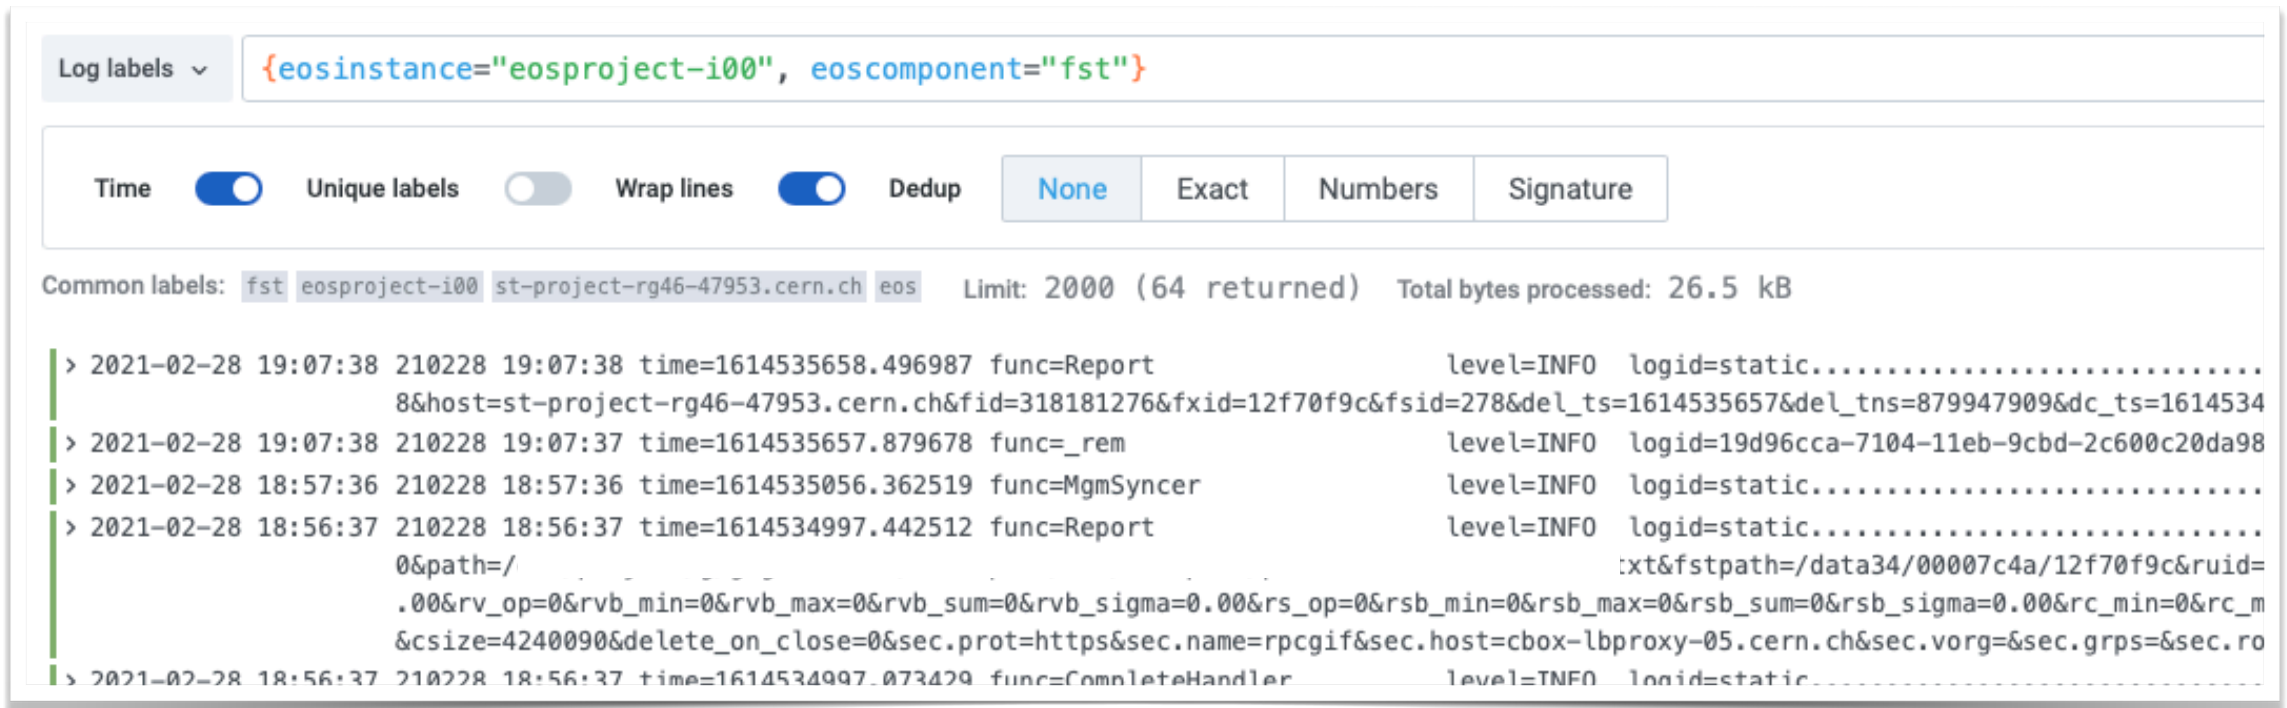
- **CERNBox** backup Evolution
 - More than 40,000 users and projects spaces are now backed up by our **cback** backup orchestrator.
 - Check Gianmaria's presentation [\[Link\]](#)



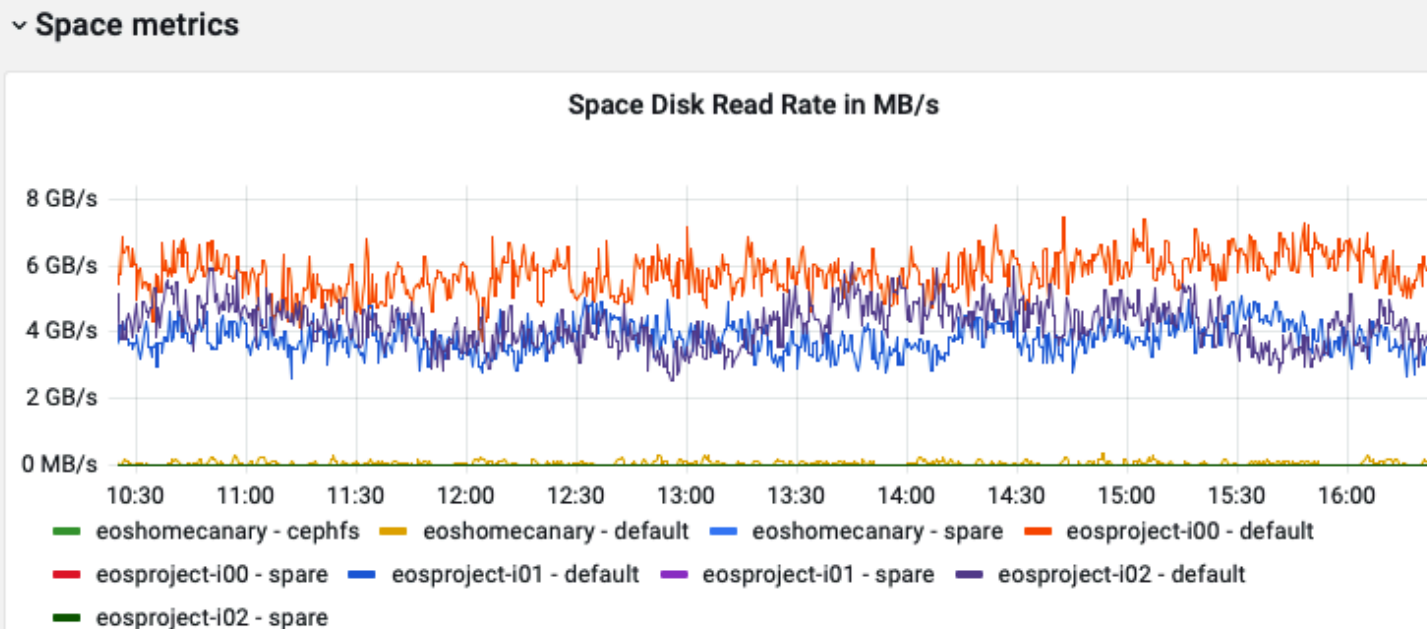
2021: Monitoring / Logging



- Centralised EOS logs with **Grafana Loki**
 - Easier tracing of operations across MGM/FST
 - Check Sami's presentation [\[Link\]](#)



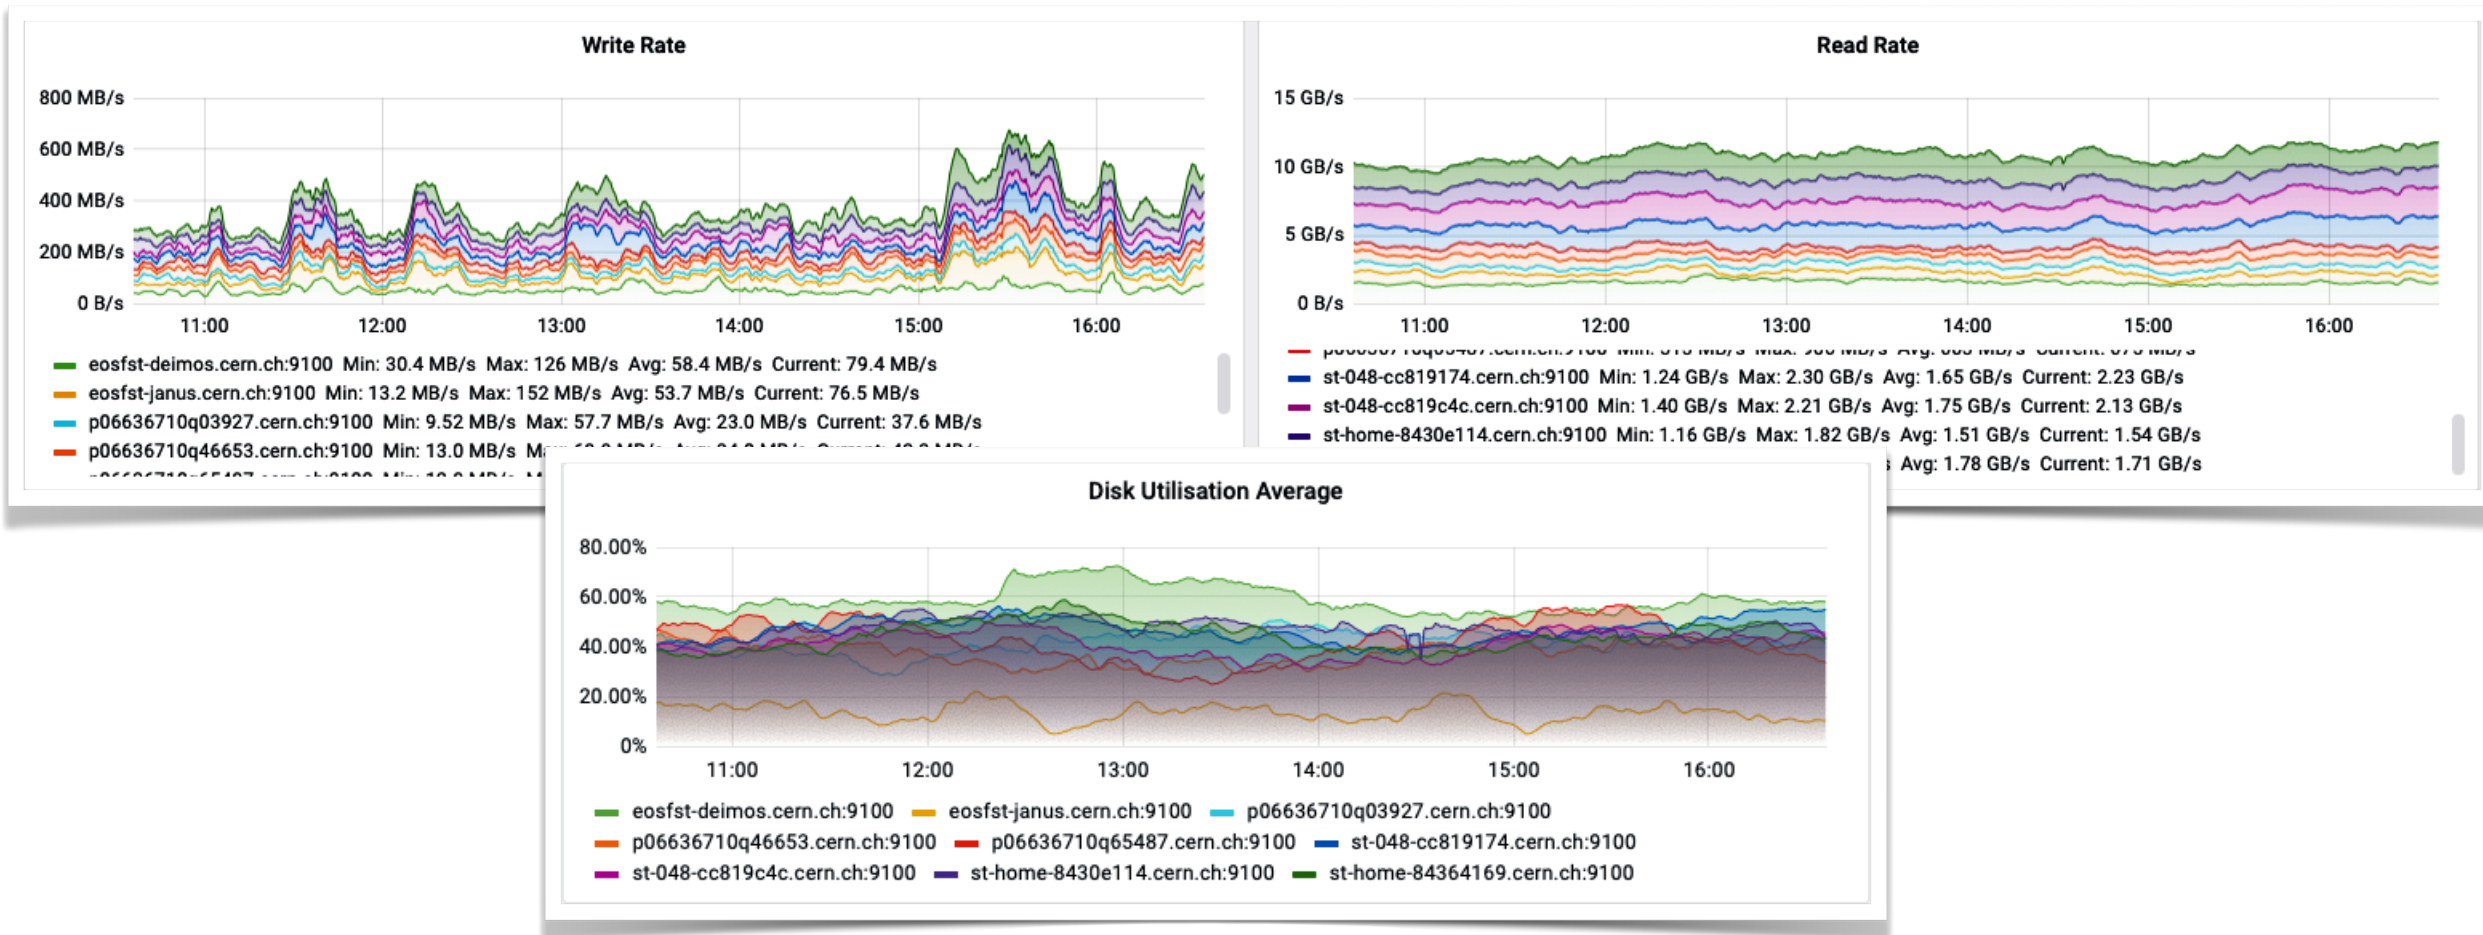
- ✓ Space metrics



2021: Monitoring / Logging (III)



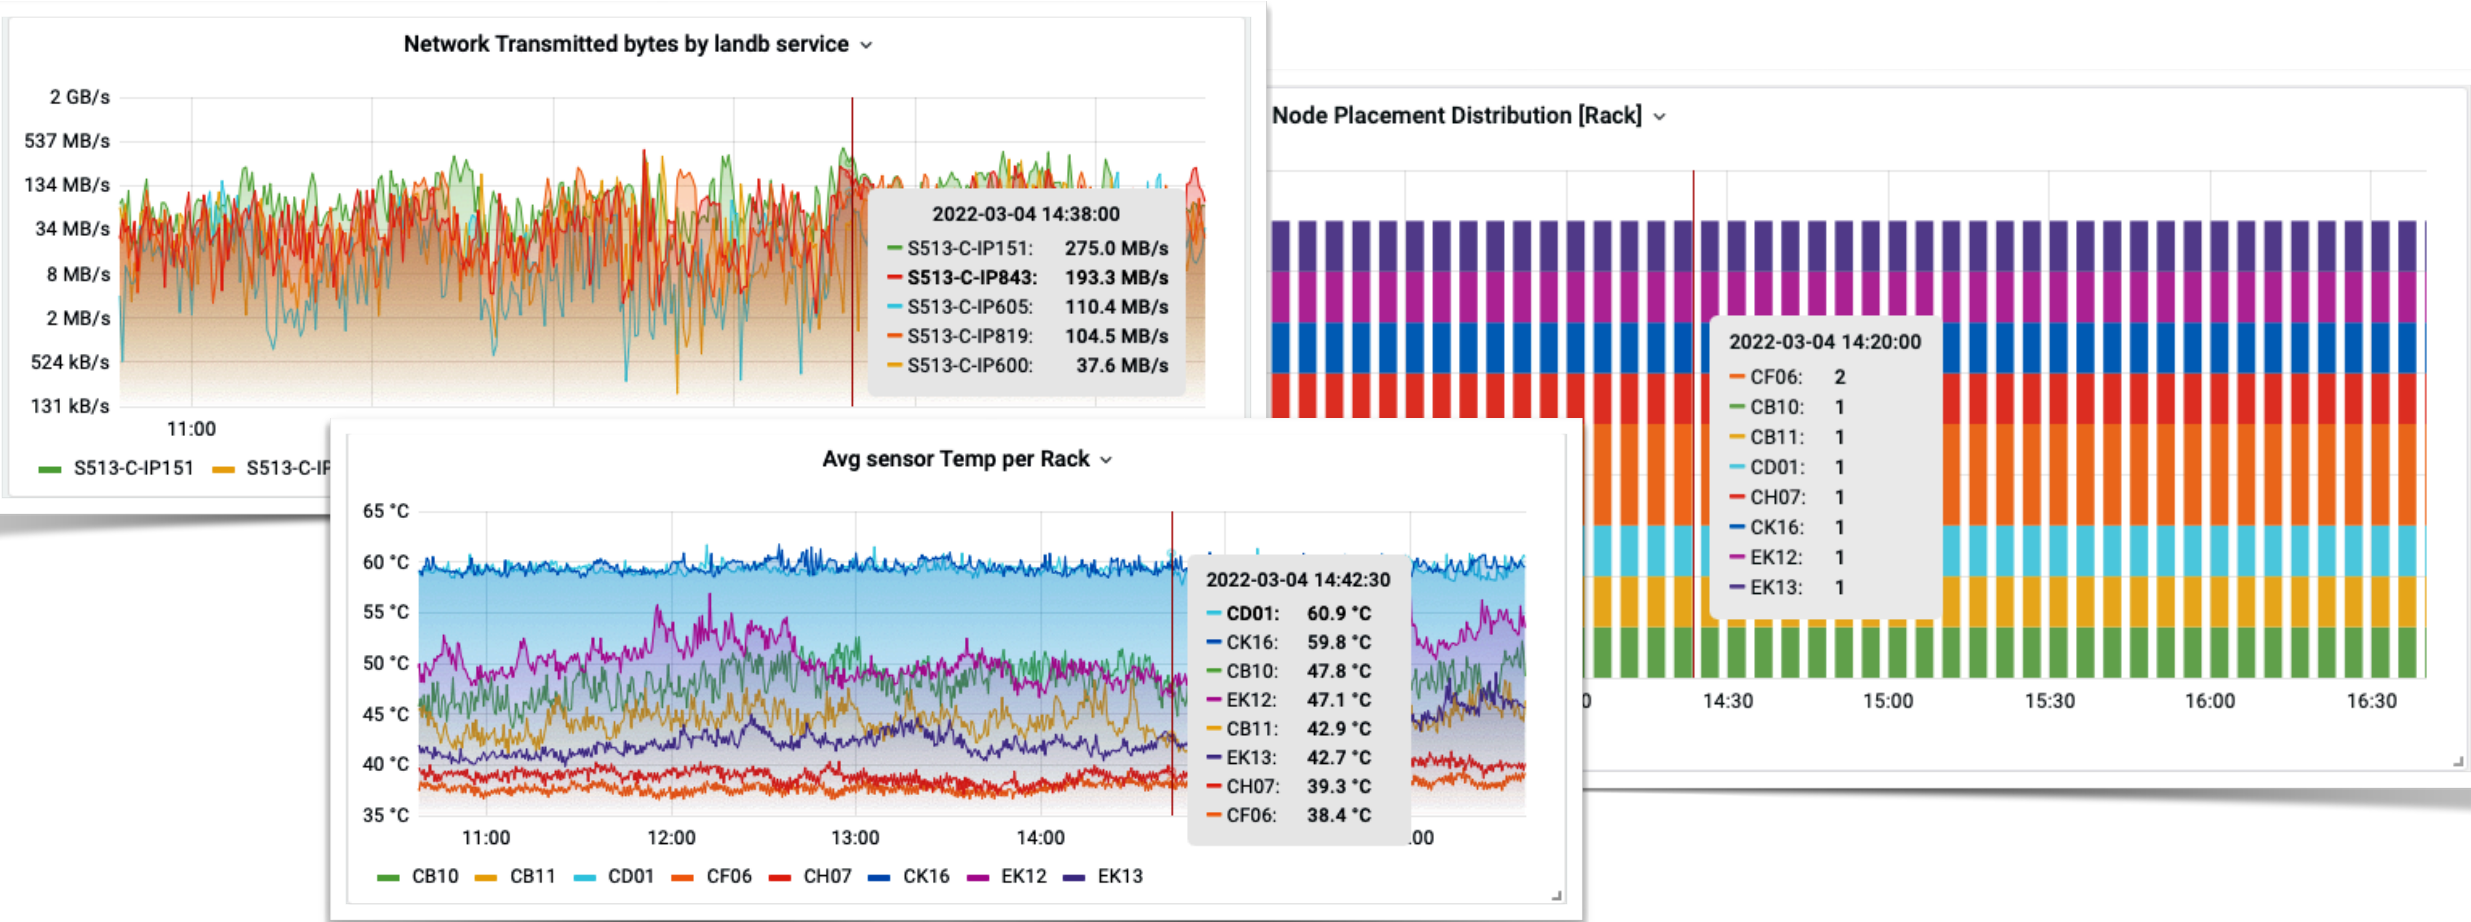
- Improved FST monitoring with Prometheus **node_exporter**



2021: Monitoring / Logging (IV)



- Improved FST monitoring with Prometheus **node_exporter**



2021: Hardware decommission / allocation

- Retirement campaign:
 - 32 disk servers - **6.5P**
 - 24 MGMs
- Hardware allocation:
 - 17 disk servers + **11.4P**
 - 24 MGMs



2022 Roadmap

- **CephFS** as FST backend
- **EOS 5** deployment
- **HA** on EOS MGM
 - Do even more transparent upgrades / interventions
- Provide **self-service backup restore** for CERNBox users.
- Improve **monitor** and **alerting**
 - Durability, backup

Thank you!



home.cern