
Status of Tape-less Archive Storage project @ KISTI-GSDC

AHN SANG-UN @ WLCG GDB, 9 DECEMBER 2020

Outline

- Introduction
- System Architecture
- QRAIN Layout
- Current Status
- Monitoring
- Plan
- Summary


Introduction

- Replacing tape library (+3PB) with disk-only storage for archiving @ KISTI for ALICE experiment
 - Simpler architecture, less operational efforts, cost-effective comparable to tape
- Found domestic suppliers of high-density(> 60 disks/box) JBOD models
- Relying on EOS erasure coding implementation (RAIN layout) for data protection
- About 1M CHF budget (2019) included
 - 18 High-density JBOD boxes (84 disks/box \simeq 18PB raw capacity)
 - 9 Servers for EOS front-end nodes (12Gbps SAS HBAs, 40Gbps uplinks + switches)
- Providing production service to ALICE before the start of RUN3 (by June 2021) ← **POSTPONED**


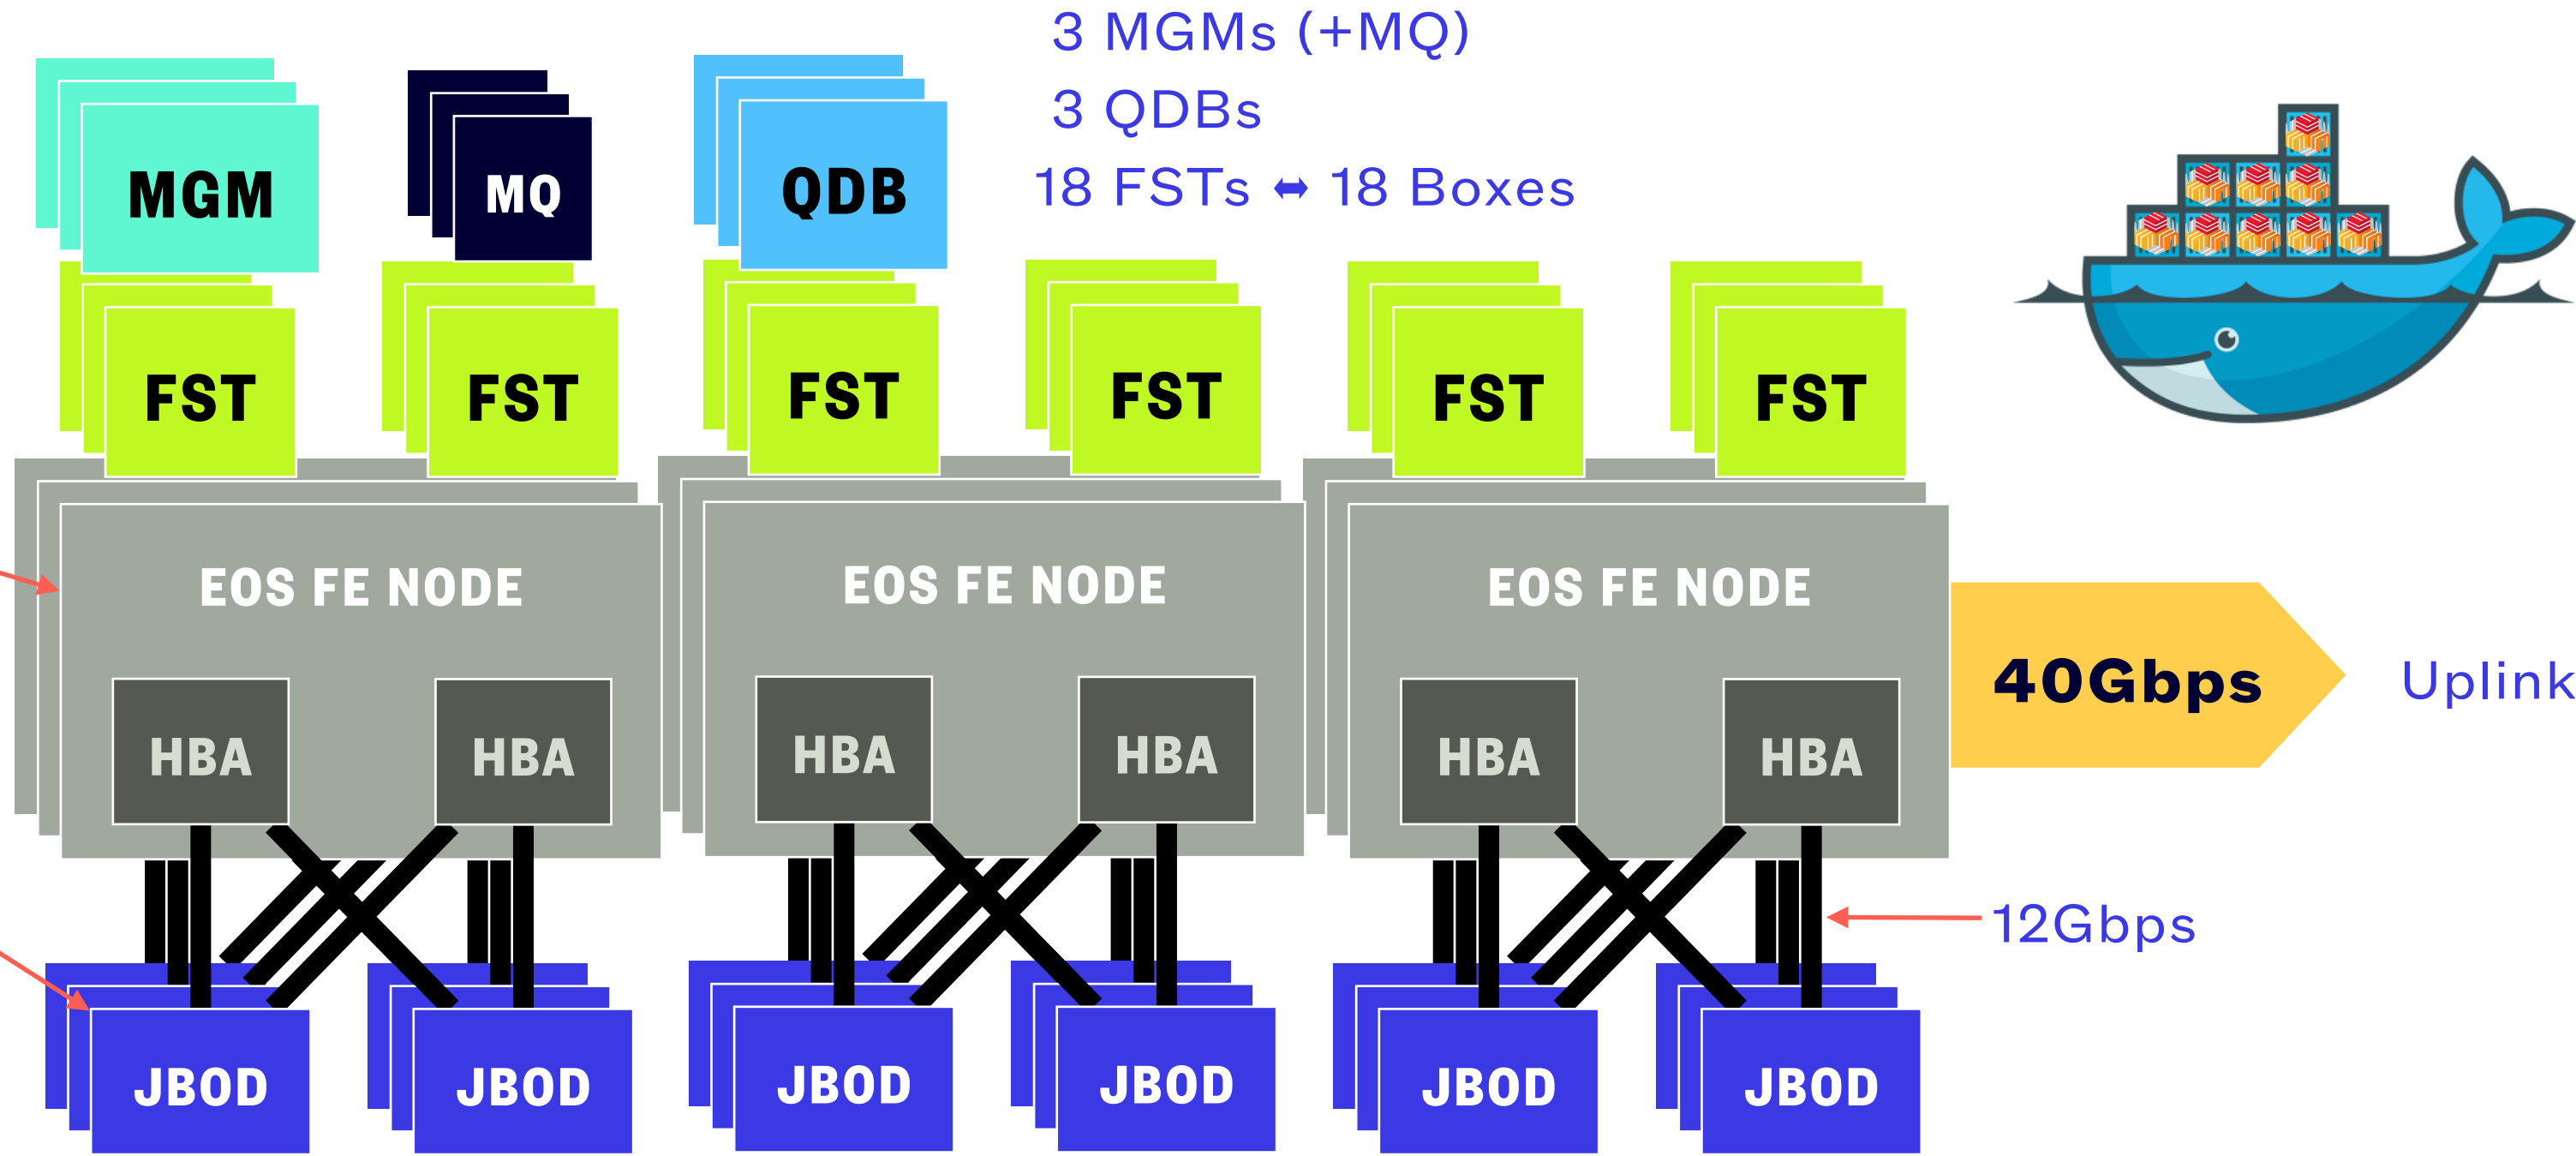
System Architecture



9 servers
18 boxes



84 DISKS
IN ONE BOX



- Total raw capacity = 18,144TB (= 12TB * 84 disks * 18 boxes)
- EOS version = 4.8.25 (released on 2020.11.10)
- EOS components are running on containers (a fork of EOS-Docker project)
 - Ansible playbook available at <https://github.com/jeongheon81/gsd-eos-docker>


Recap CHEP2019

<https://doi.org/10.1051/epjconf/202024504001>

- EOS Erasure Coding implementation => RAIN layout
 - 2 FSTs (data server) on a single FE node to maximize usable space (~70%) out of raw capacity
- Upper cap on total throughput limit by PCI-e 3.0 (~6GB/s)
- Power consumption ~ 1.75W/TB (JBOD+Server+S/W)
 - Enterprise Disk storage: 5 ~ 9W/TB
 - Tape library ~ 0.5W/TB


I/O Test: Read/Write

- XFS read/write performance (simultaneous read and/or write from 70 disks)
 - **VDBench** shows full read/write transfer performance @ transfer size >= 2048k (6GB/s)
 - **IOZone** shows full read/write transfer performance @ transfer size ~ 2048k (6GB/s)



Power Consumption

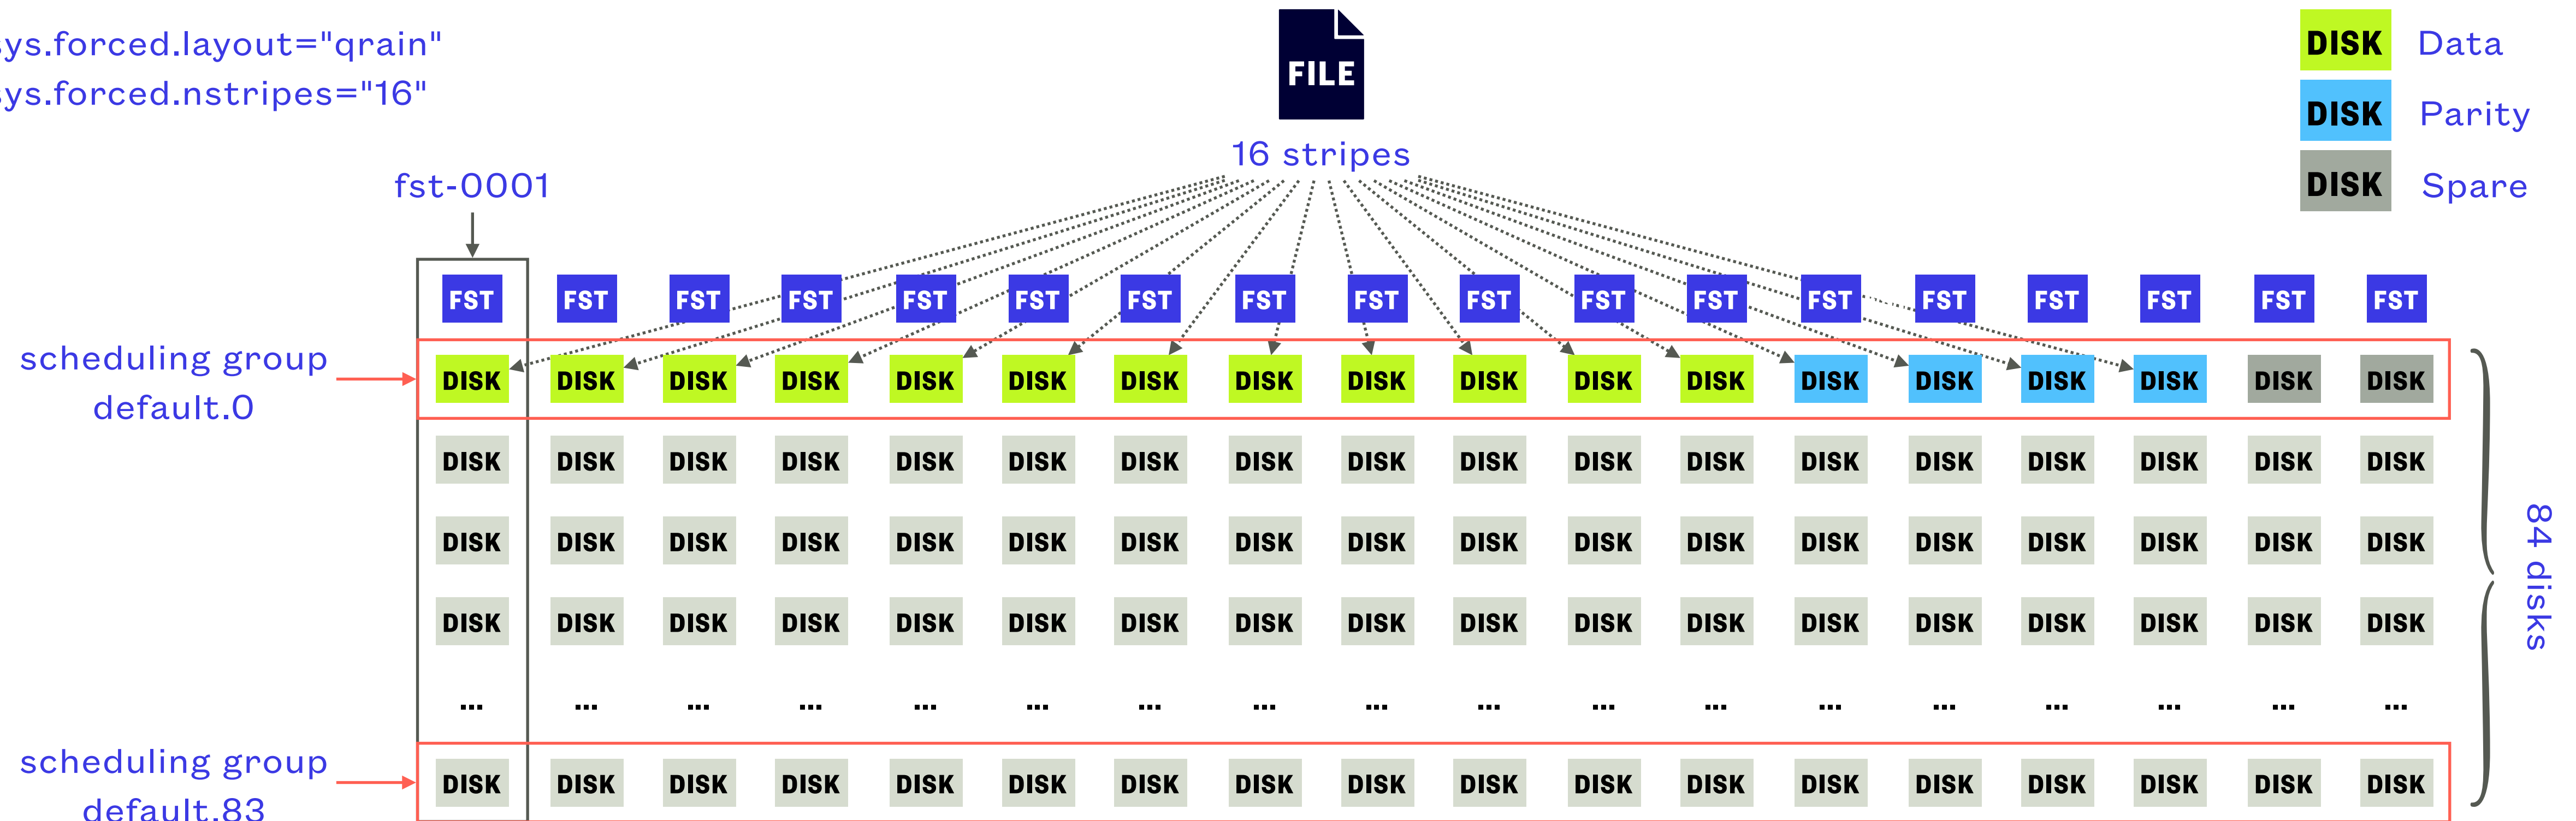
- JBOD Test Equipment (70 Disks)
 - JBOD (DELL ME484): idle = 830W; load = 860W (Max 960) (**1.12W/TB**)
 - Server: idle = 200W; load = 270W
 - Switch: idle = 246W; load = 246W
 - **1.75W/TB** including JBOD, Server and Switch
- Disk Storages (Full Load)
 - DellEMC SC7020, 2.5PB - 12,120W (**4.8W/TB**)
 - EMC Isilon, 16 Nodes, 2.95 PB- 13,730W (**4.6W/TB**)
 - EMC VNX, 12 Nodes, 2.36 PB - 5,100W (**2.2W/TB**)
 - HITACHI VSP, 2 PB - 18,300W (**9.15W/TB**)
 - EMC Isilon, 15 Nodes, 1.43 PB - 12,880W (**9W/TB**)
 - EMC CX4-960, 1.5PB - 14,900W (**9.9W/TB**)
- Tape Library (Full Load)
 - IBM TS3500 5-Frame (3.2PB) - 1,600W (**0.5W/TB**)



QRAIN Layout



sys.forced.layout="grain"
sys.forced.nstripes="16"



- Thanks to spare FSTs,
 - Data are still accessible if 6 FSTs are offline
 - Data can be written if 2 FSTs are offline
 - One node (= 2 FSTs) can be turned off for maintenance at any time
- Data loss rate in a year is $\approx 8.6 \times 10^{-5}\%$, where 5 disks are failed simultaneously, considering 1.17% of AFR in practice
cf. vendor published AFR is 0.35% (AFR = Annualized Failure Rate)

QRAIN Configuration

- 'eos attr' command
 - One can have different layouts on different directories in an EOS instance
 - Available layouts = plain (default, 1 single copy); replica (2 copies); raid6, raiddp (2 parities); archive (3 parities); qrain (4 parities), ...

```
eos attr -r set default=raid6 /eos/gsdctestarea/raid6
eos attr -r set default=archive /eos/gsdctestarea/archive
eos attr -r set default=replica /eos/gsdctestarea/replica
eos attr -r set default=qrain /eos/gsdctestarea/rain12
```

- # of stripes can be changed, e.g. 3 copies, 16 stripes...

```
eos attr -r set default=qrain /eos/gsdctestarea/rain16
eos attr -r set sys.forced.nstripes=16 /eos/gsdctestarea/rain16
```

```
sh-4.2# eos attr ls /eos/gsdctestarea/rain16
sys.eos.btime="1605069261.927407367"
sys.forced.blockchecksum="crc32c"
sys.forced.blocksize="1M"
sys.forced.checksum="adler"
sys.forced.layout="qrain"
sys.forced.nstripes="16"
sys.forced.space="default"
sys.recycle="/eos/gsdctestarea/recycle/"
```


Fileinfo



EOS fileinfo command

```
sh-4.2# eos fileinfo /eos/gsdctestarea/rain16/testfile.10G
```

```
File: '/eos/gsdctestarea/rain16/testfile.10G' Flags: 0640
Size: 10485760000
Modify: Thu Oct 22 00:01:35 2020 Timestamp: 1603324895.724750000
Change: Thu Oct 22 00:00:51 2020 Timestamp: 1603324851.619542497
Birth: Thu Oct 22 00:00:51 2020 Timestamp: 1603324851.619542497
CUid: 0 CGid: 0 Fxid: 0000159b Fid: 5531 Pid: 40 Pxid: 00000028
XStype: adler XS: a1 1c 00 01 ETAGs: "1484716507136:a11c0001"
```

Layout type
of stripes
of replica

```
Layout: qrain Stripes: 16 Blocksize: 1M LayoutId: 40640f52 Redundancy: d5::t0
#Rep: 16
```

File chunk location
Scheduling group
Filesystem status

no.	fs-id	host	schedgroup	path	boot	configstatus	drain	active	geotag
0	995	jbod-mgmt-06.sdfarm.kr	default.70	/jbod/box_12_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02
1	1499	jbod-mgmt-09.sdfarm.kr	default.70	/jbod/box_18_disk_070	booted	rw	nodrain	online	kisti::gsdc::g03
2	659	jbod-mgmt-04.sdfarm.kr	default.70	/jbod/box_08_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02
3	407	jbod-mgmt-03.sdfarm.kr	default.70	/jbod/box_05_disk_070	booted	rw	nodrain	online	kisti::gsdc::g01
4	827	jbod-mgmt-05.sdfarm.kr	default.70	/jbod/box_10_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02
5	491	jbod-mgmt-03.sdfarm.kr	default.70	/jbod/box_06_disk_070	booted	rw	nodrain	online	kisti::gsdc::g01
6	1079	jbod-mgmt-07.sdfarm.kr	default.70	/jbod/box_13_disk_070	booted	rw	nodrain	online	kisti::gsdc::g03
7	71	jbod-mgmt-01.sdfarm.kr	default.70	/jbod/box_01_disk_070	booted	rw	nodrain	online	kisti::gsdc::g01
8	743	jbod-mgmt-05.sdfarm.kr	default.70	/jbod/box_09_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02
9	1247	jbod-mgmt-08.sdfarm.kr	default.70	/jbod/box_15_disk_070	booted	rw	nodrain	online	kisti::gsdc::g03
10	155	jbod-mgmt-01.sdfarm.kr	default.70	/jbod/box_02_disk_070	booted	rw	nodrain	online	kisti::gsdc::g01
11	1415	jbod-mgmt-09.sdfarm.kr	default.70	/jbod/box_17_disk_070	booted	rw	nodrain	online	kisti::gsdc::g03
12	911	jbod-mgmt-06.sdfarm.kr	default.70	/jbod/box_11_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02
13	1331	jbod-mgmt-08.sdfarm.kr	default.70	/jbod/box_16_disk_070	booted	rw	nodrain	online	kisti::gsdc::g03
14	239	jbod-mgmt-02.sdfarm.kr	default.70	/jbod/box_03_disk_070	booted	rw	nodrain	online	kisti::gsdc::g01
15	575	jbod-mgmt-04.sdfarm.kr	default.70	/jbod/box_07_disk_070	booted	rw	nodrain	online	kisti::gsdc::g02

```
*****
```

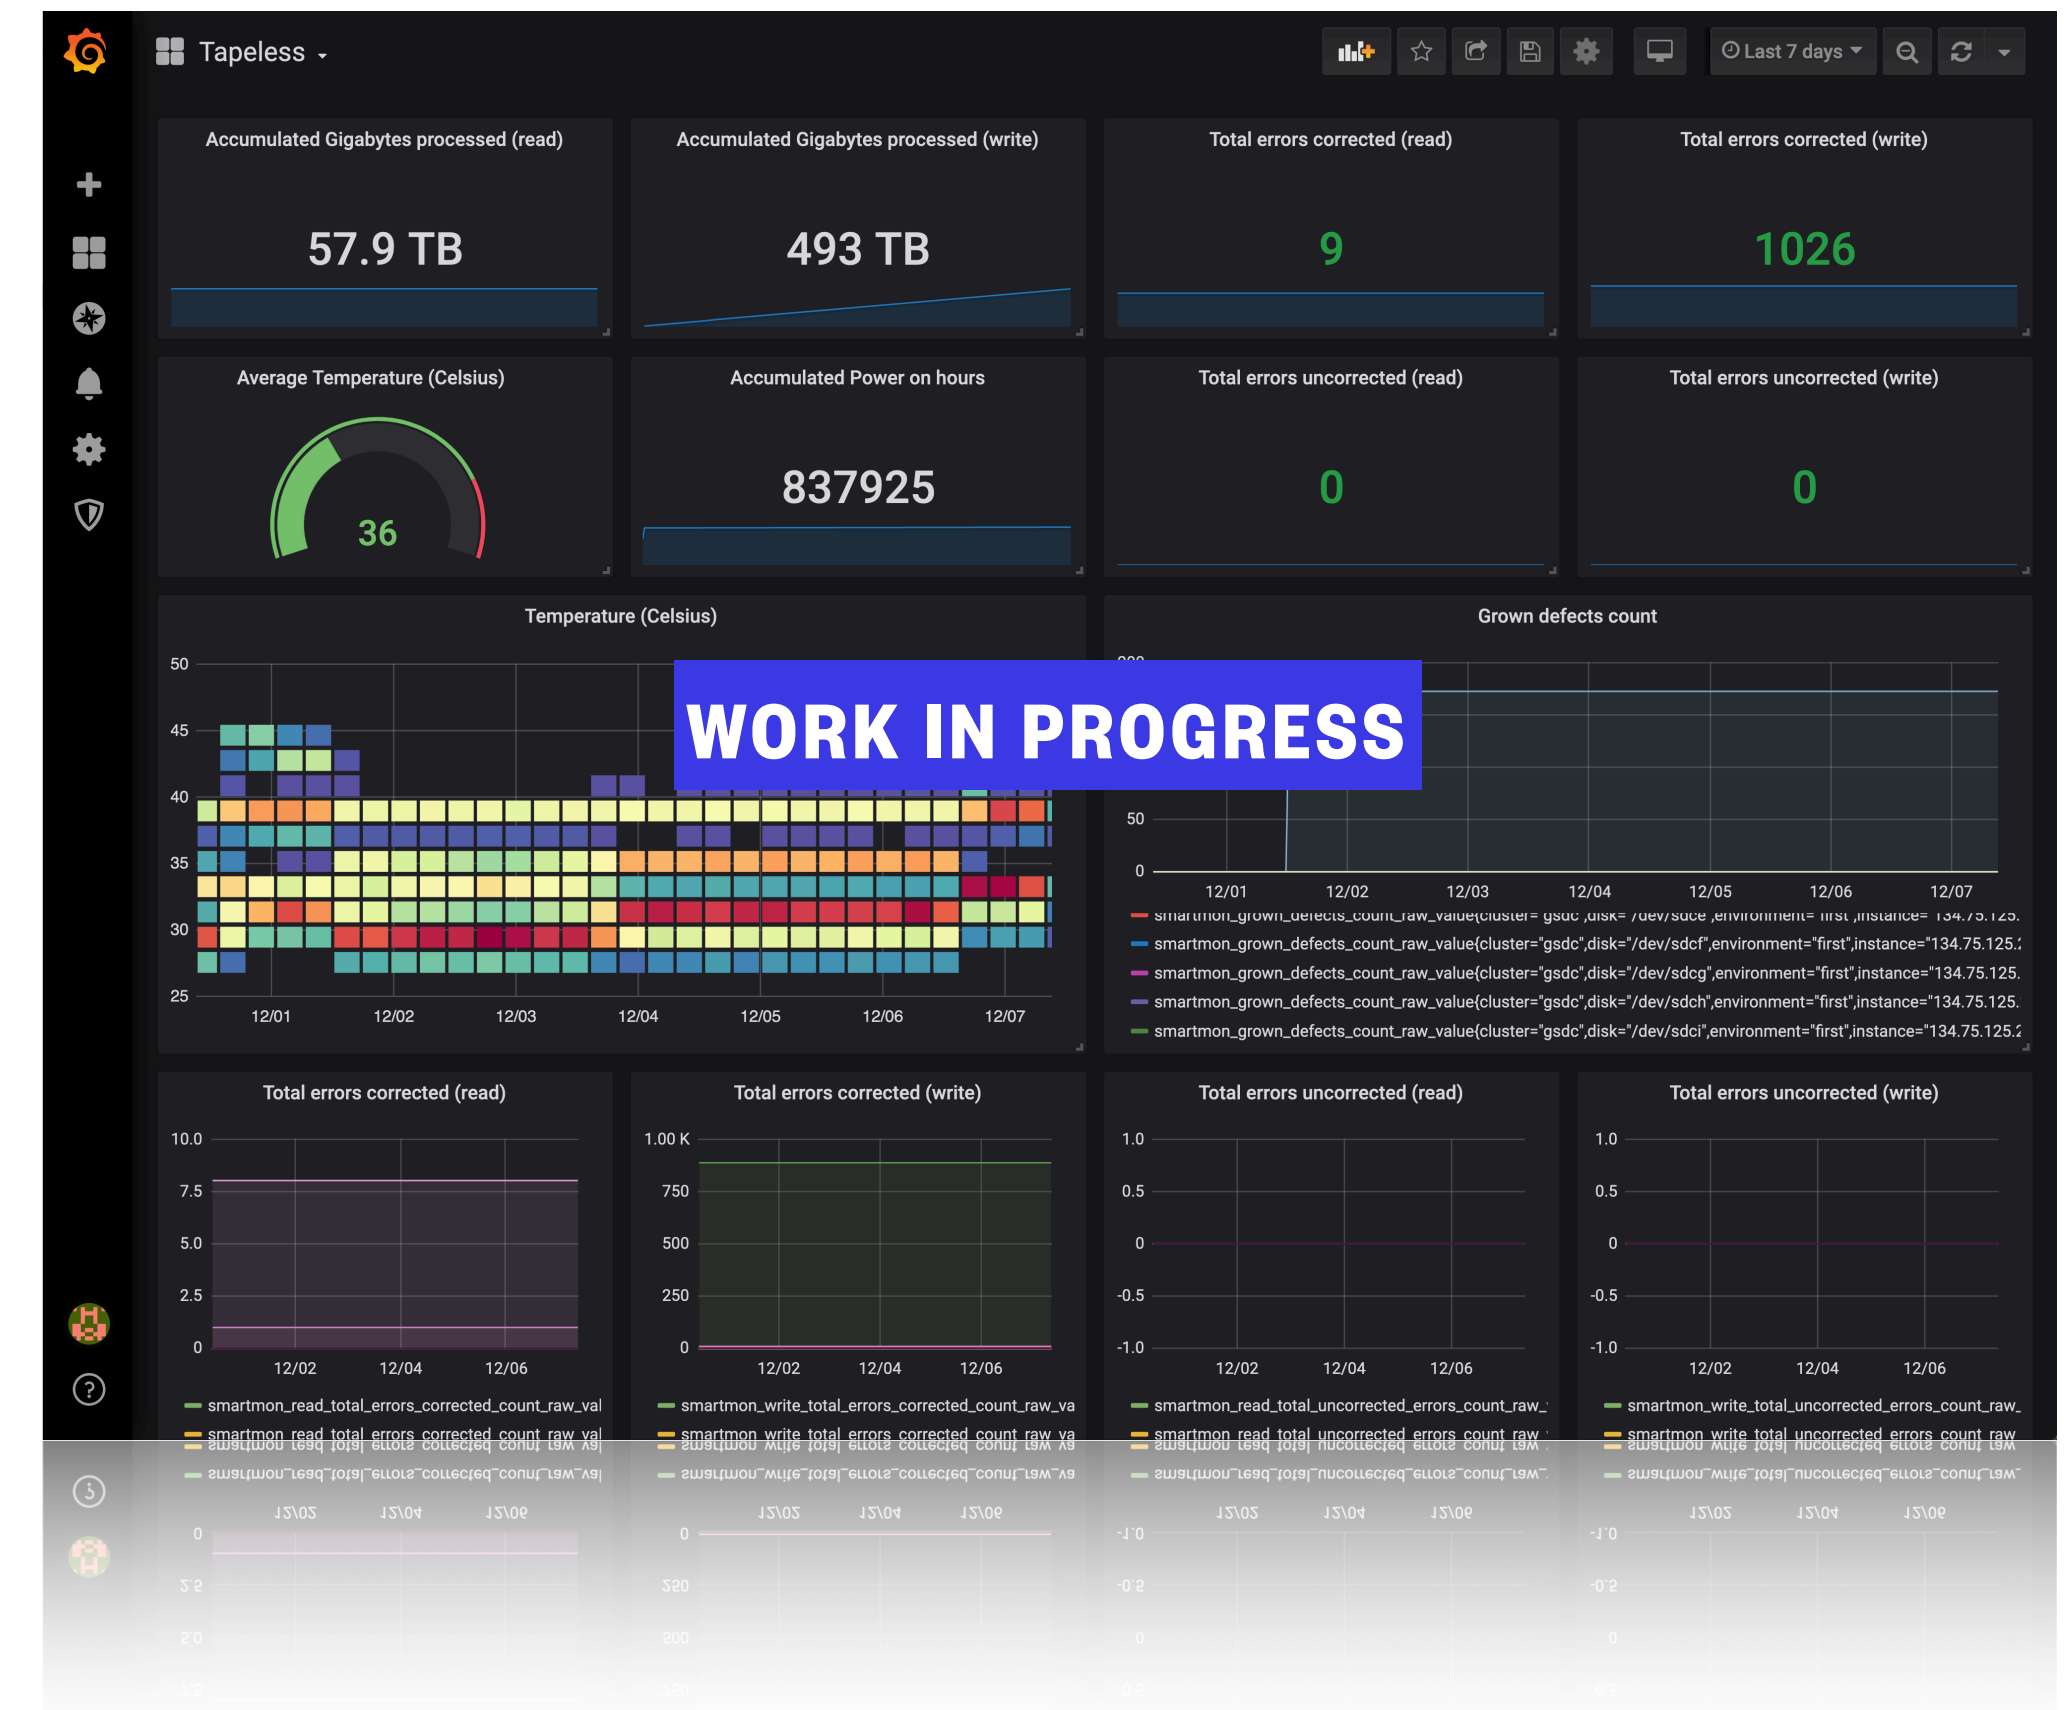
Current Status

- EOS update: 4.8.12 → 4.8.25 (2020.11.10 release)
- MGM load-balancing and fail-over tests confirming that
 - Xrdcp write/read requests through MGM secondaries work well
 - Automatic fail-over happens within a few seconds (~10s) when primary MGM is unresponsive
- Public DNS alias pointing to 3 MGMs
- IPv4/IPv6 dual stack configured

Monitoring




- Prometheus node_exporter + Grafana dashboard
 - Hardware level monitoring using *smartmon* tools
 - Corrected/uncorrected error counts
 - Temperature, defects count, etc.
 - Learning metrics to identifying and predicting disk failures
- Custom script parsing JBOD enclosures status from *sg_ses* command output



Debugging



- EOS services log dump using *loki*, *promtail*
 - Easy access to container level logs (*MGMs*, *FSTs*, ...) by name
 - Live stream (like *'tail -f /var/log/eos/...xrd.log'*) from multiple nodes
 - Complementary to EOS CLI tools
 - Remote support by EOS developers



Issues Fixed

- When read a file with "eos cp" (or eoscp) from RAIN layout with more than 12 stripes
 - Redirection information over 2kB truncated
 - <https://gitlab.cern.ch/dss/eos/-/commit/4cb0f733650e041a3153da60610a8d5a0e4672f4>
 - <https://gitlab.cern.ch/dss/eos/-/commit/03bcc4b55556bc0d7b3160670a41a98bfa50a941>
- Failed to read a file with "eos cp" through MGM secondaries
 - Quota information propagation corrupting namespace
 - <https://gitlab.cern.ch/dss/eos/-/commit/32b93012459f73409b45fcea3c575cc6c47f421>

SPECIAL THANKS TO ELVIN ALIN SINDRILARU !!!

Plan

- Update to EOS 4.8.31
 - Fix for namespace corruption due to wrong propagation of quota information
- Enabling token based authentication for ALICE
 - Host certificates for MGMs are issued and installed
- Integrating as a new ALICE tape storage element *without tape*
 - Performing periodic functional tests, and monitoring stability and reliability
- Establishing maintenance and operation scheme
 - Disk replacement, JBOD and server maintenance, rolling update of EOS components

Summary

- Working on providing a disk-only archive storage for ALICE experiment with the help of EOS erasure coding implementation for data protection
- Successfully deployed QRAIN layout with 16 stripes including 4 parities and made it working by fixing a couple of issues
- Working on establishing a dedicated monitoring framework for the archive storage
- On track of schedule to provide production service by June 2021, even though the start of RUN3 has been postponed

Thank you
