



EOS Durability Summary

**Manuel Reis
(IT-ST-PDS)**

Overview

EOS Durability

- Concept
- How to apply it n
- Overview on eos-ops-durability

Current Status

- Look back to 2021

Wrap up

- What to improve



EOS Durability

EOS is a reliable filesystem!

How many files are suffering problems overtime? (Data risk, or corruption...)

The **sheer amount** of **use cases/workflows**, demands to keep track on these numbers!

FSCK, an **EOS component** that offers an extra layer of **data integrity**. Introduced @ [EOS 2020 WS](#)

Aside the **embedded functionalities**, we use a set of tools to **monitor** and **repair** this **data**.

- eos-ops-durability – miscellaneous tools
- eos-ns-inspect – namespace explorer

How to apply it

Input sources:

- Namespace scan (time consuming)
- FSCK report / Disk scanning (distributed)
- Failing or draining filesystems (sporadic & passive)

Process:

- Classify and group files per problem/symptom

Outcomes

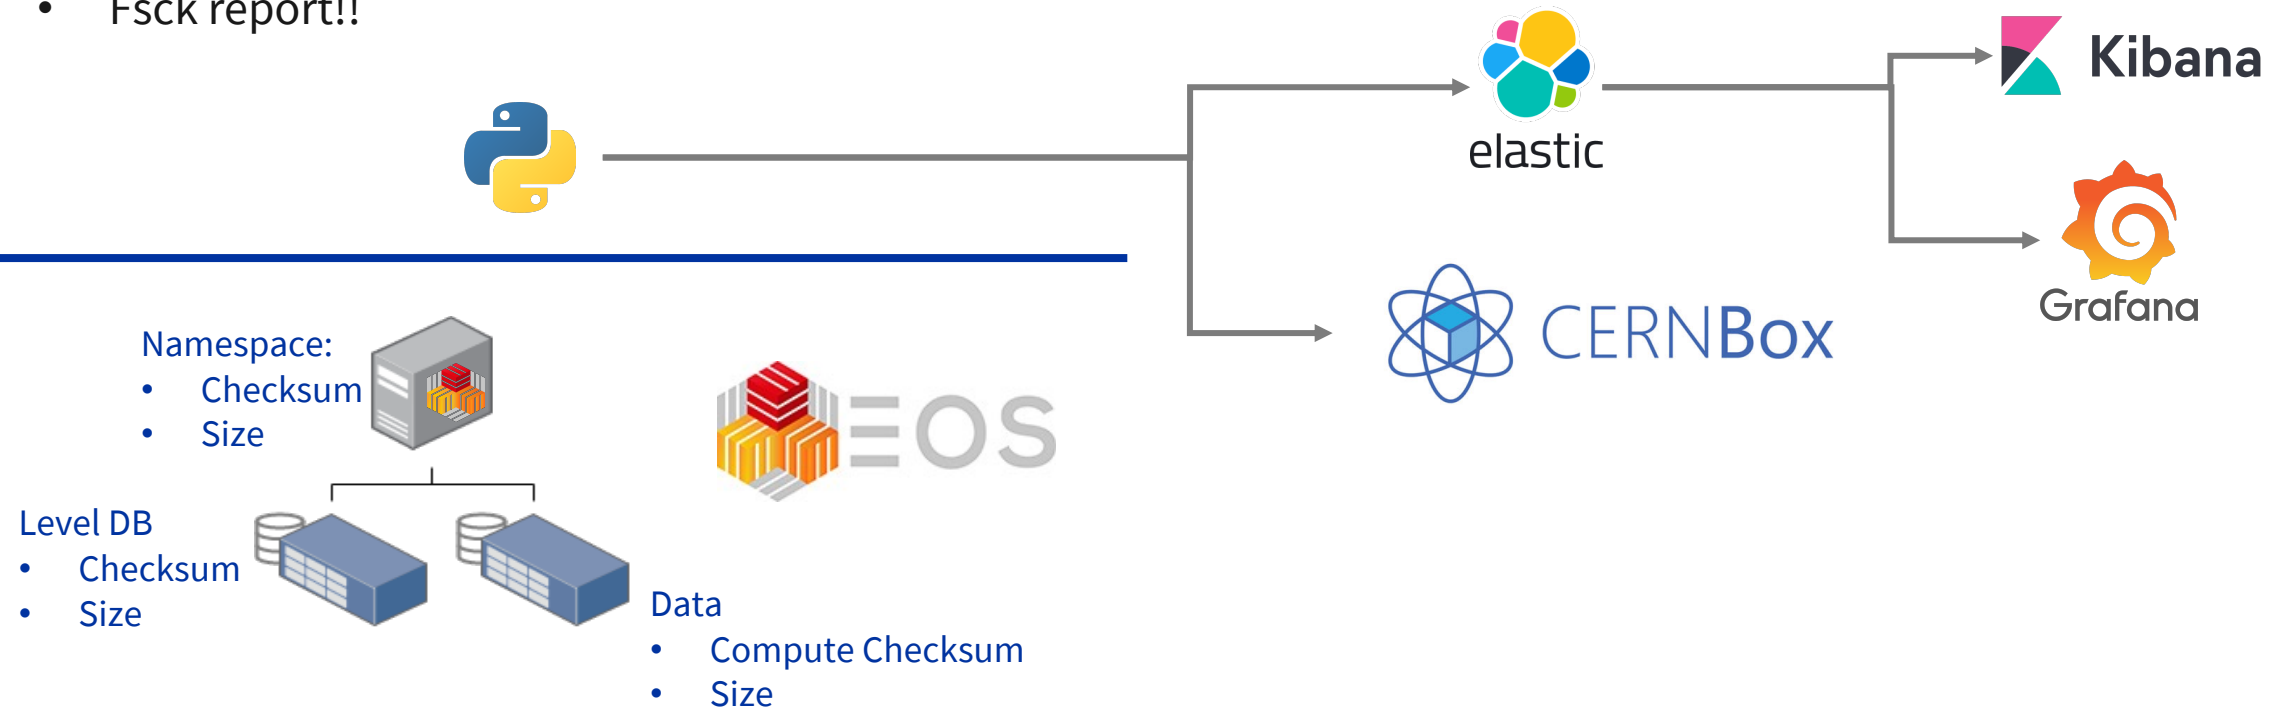
- Spot bugs
- Fixed files
- Monitor status



Overview eos-ops-durability

Detect, Classify, Repair & Monitor

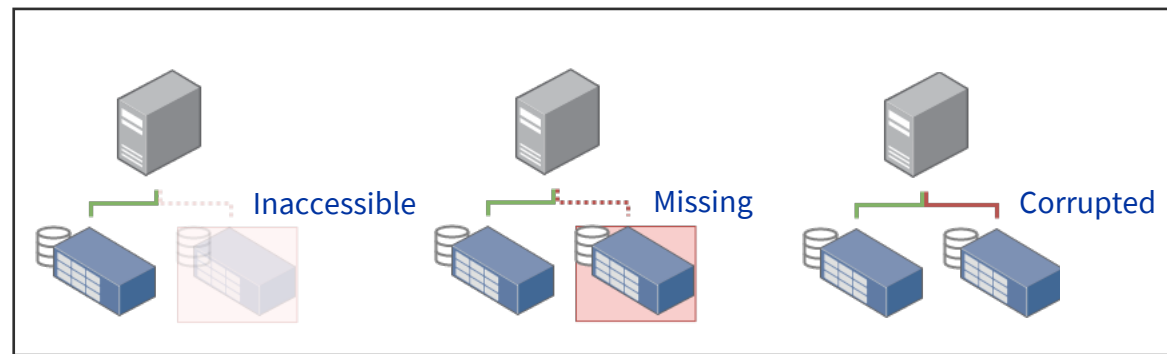
- `eos-ns-inspect` - one replica and stripe diff files
- Drain failed filesystems
- Fsck report!!




Classification overview – branching

The concepts to have in mind:


- Under replicated data!!!
 - Misconfiguration (unlikely)
 - ENOENT, IO Error, etc... in one replica
- Corruption!!
 - Mismatch of checksum or size
- Over replicated data!?




One Replica OK




Zero Replicas




Mismatching
Replicas




Single
Mismatching
Replica



Missing
Replicas



Mismatching
pair of
Replicas*




All Replicas
Ok

Current status

What happened during 2021

- `eos fsck repair` was enabled (in all instances)
 - Works on a complete scope of data (NS + disk)
 - Cleanup the working bench
 - Room to investigation (search root causes)
 - Bug fix ([EOS-4782](#), [EOS-4828](#), [EOS-5052](#), ...)
- eos-ops-durability - main highlights
 - Added `eos fsck report` as supported input
 - Now ported to Python3

FSCK report dashboard (non unique numbers)



0.29% of total files
(from which, most are **not** @ risk)

Wrap up

We have the resources to protect endangered data, but there's risk while actions aren't taken.

- Prioritize recovery of data at risk
 - Files with one replica only!!
 - Set alarms to trigger upon sudden growths of such cases
- Invest more on root causes investigations

Ensuring durability is a continuous work (as new problems can be found)


Having the system in good shape has given the opportunity to work on preserving it!



<https://storage-ci.web.cern.ch/storage-ci/eos-ops-durability/tag/>



<https://gitlab.cern.ch/eos/eos-ops-durability/>



home.cern