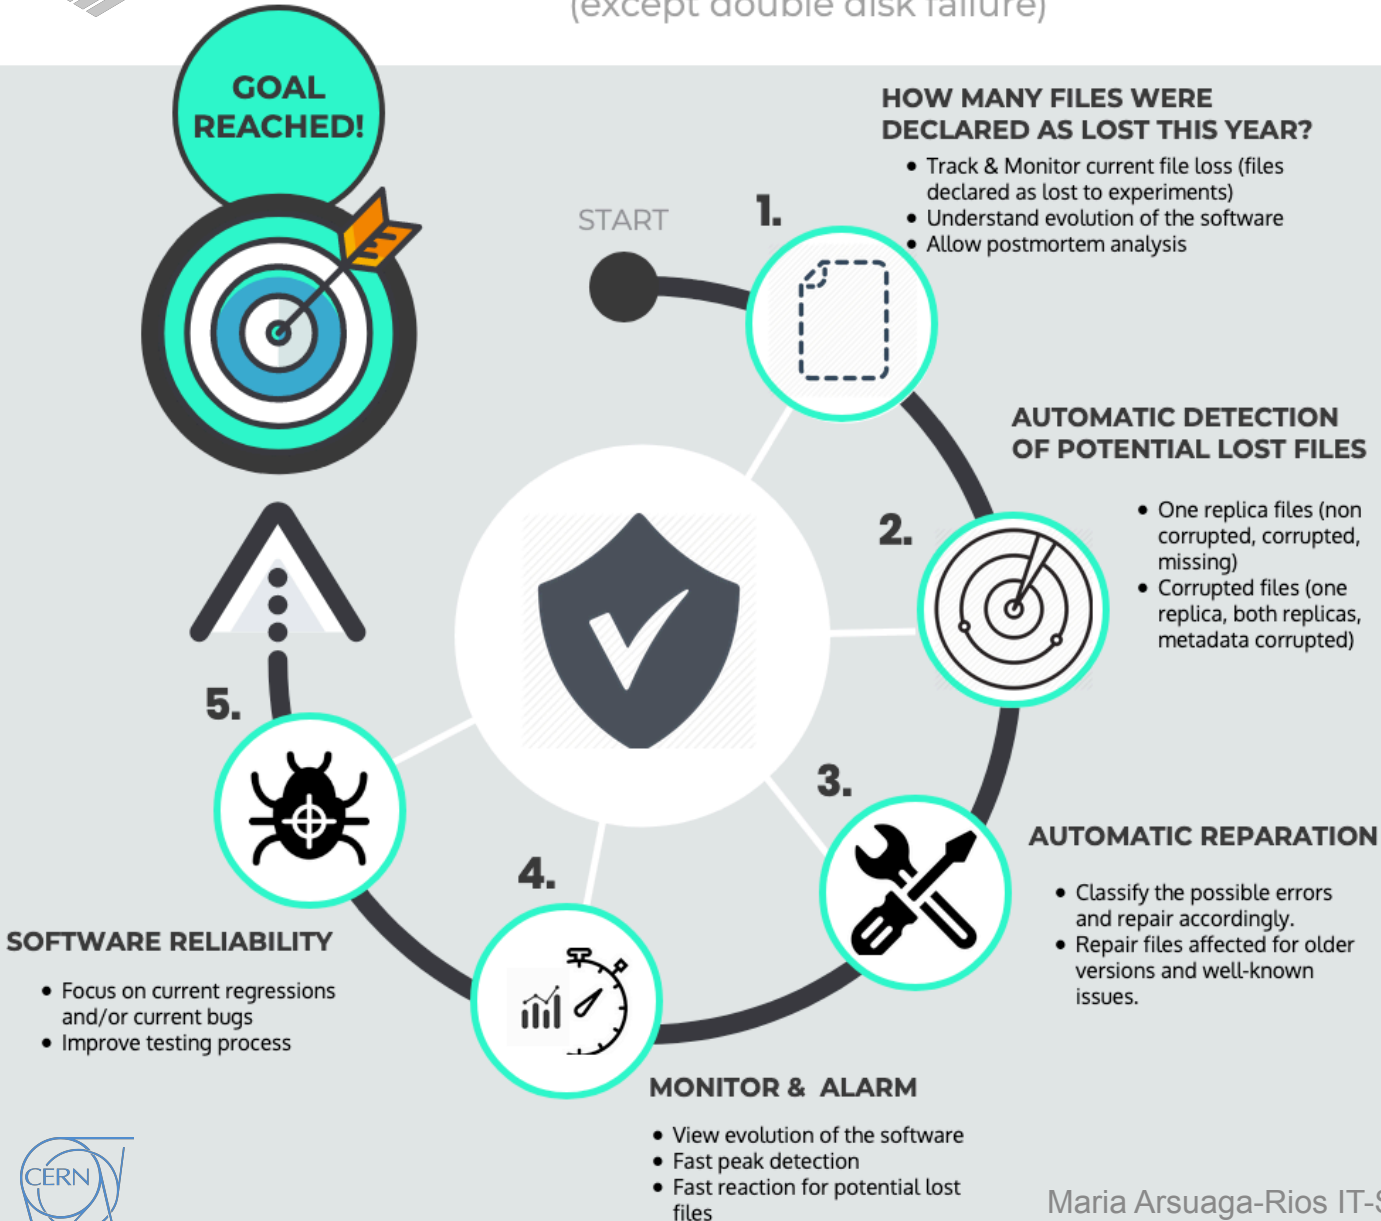


DATA DURABILITY

Data is never lost or compromised
(except double disk failure)

How Data Durability tools use eos-ns-inspect?





- standalone tool which connects directly to QuarkDB
- full view into the namespace contents, no MGM needed.




Detecting inconsistencies

» Detecting under and over replications



Under replication



Over replication

```
# eos-ns-inspect stripediff --members eoslhcb-qdb.cern.ch:7777 --password-file  
<keytab_location>
```

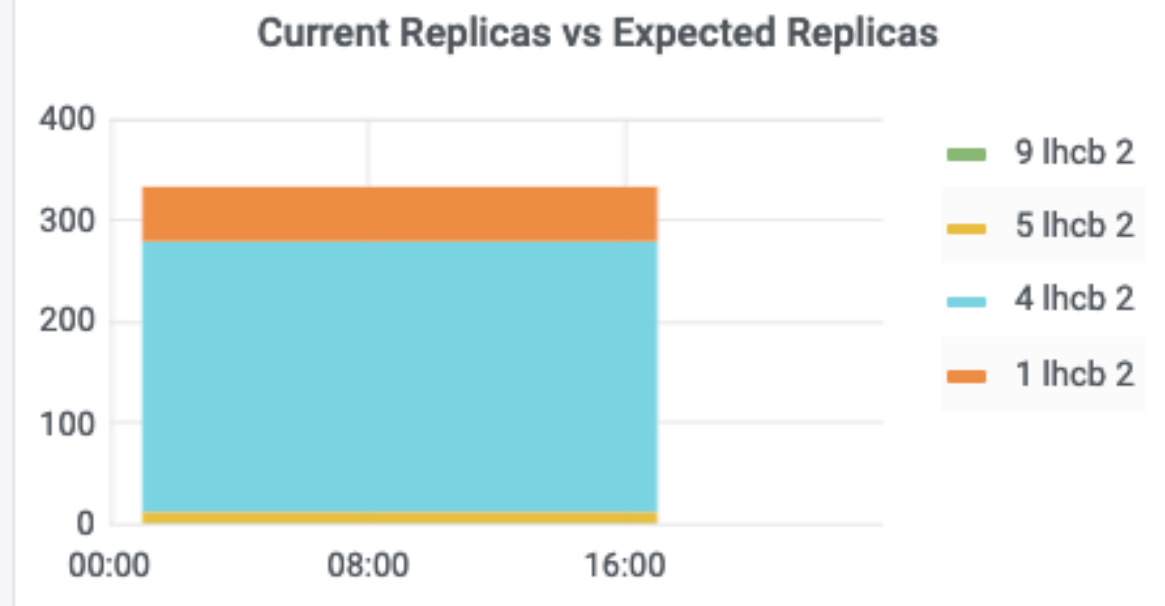
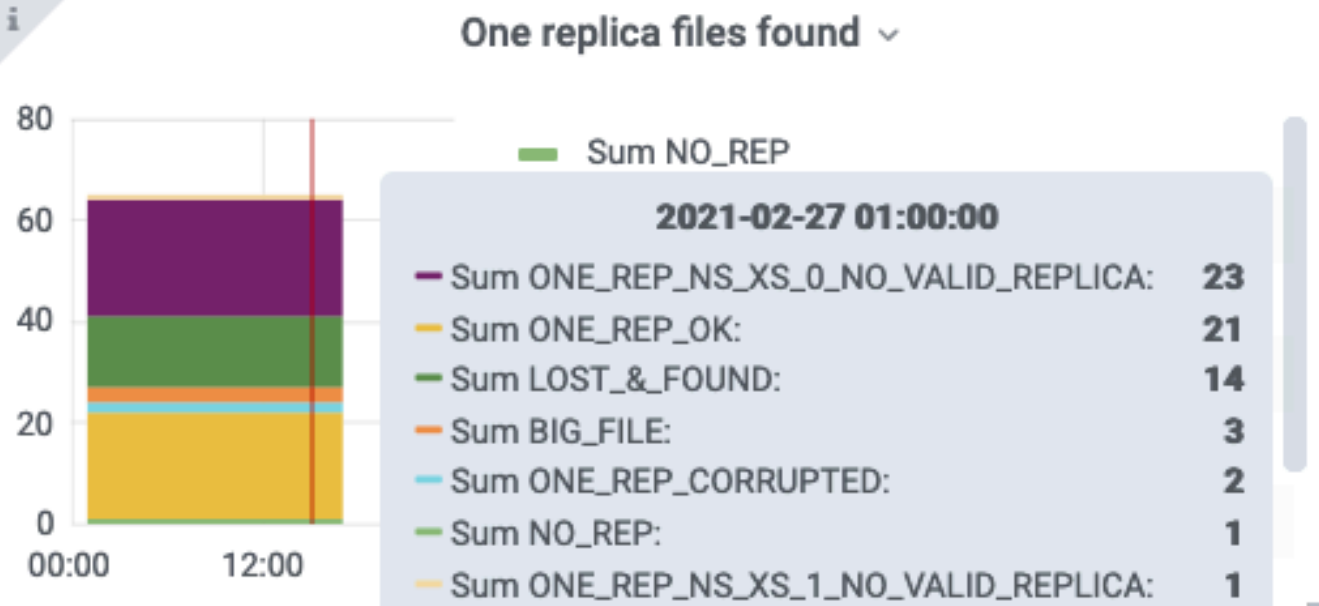


EOS Detecting under and over replications

```
# eos-ns-inspect stripediff --members eoslhcb-qdb.cern.ch:7777 --password-file  
<keytab_location>
```

instance

lhcb ▾





EOS Detecting under and over replications

```
# eos-ns-inspect stripediff --members eos<instance>-qdb.cern.ch:7777 --password-file <keytab_location>
```

```
[root@eos-pilot-ns-00 (mgm:master mq:master) ~]$ eos-ns-inspect stripediff --members eos-pilot-qdb.cern.ch:7777 --password-file /etc/eos.keytab
```

```
id=337 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=338 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10  
id=375 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10  
id=374 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=412 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10  
id=411 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=596 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=819 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10  
id=781 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=818 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=855 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=892 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=929 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9  
id=1219 container=152 size=37 actual-stripes=1 expected-stripes=2 unlinked-stripes=0 locations=2 unlinked-locations=
```

Under replication!!



EOS Detecting under and over replications

```
# eos-ns-inspect stripediff --members eoslhcb-qdb.cern.ch:7777 --password-  
file <keytab_location>
```

Under replication

```
id=1219 container=152 size=37 actual-stripes=1 expected-stripes=2 unlinked-stripes=0 locations=2 unlinked-locations=
```

Solution for under and over replications



```
#eos file adjustreplica fid:1219
```



Detecting unlinked stripes

```
# eos-ns-inspect stripediff --members eoslhcb-qdb.cern.ch:7777 --password-file <keytab_location>
```

```
[root@eos-pilot-ns-00 (mgm:master mq:master) ~]$ eos-ns-inspect stripediff --members eos-pilot-qdb.cern.ch:7777 --password-file /etc/eos.keytab
```

```
id=337 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=338 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10
id=375 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10
id=374 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=412 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10
id=411 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=596 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=819 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10
id=781 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=818 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=855 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=892 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=929 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=9
id=1219 container=152 size=37 actual-stripes=1 expected-stripes=2 unlinked-stripes=0 locations=2 unlinked-locations=
```

Unlinked stripes that should be removed!

Under replication!!



Detecting unlinked stripes

```
# eos-ns-inspect stripediff --members eoslhcb-qdb.cern.ch:7777 --password-  
file <keytab_location>
```

Unlinked stripes that should be removed!

```
id=338 container=0 size=1024 actual-stripes=0 expected-stripes=2 unlinked-stripes=1 locations= unlinked-locations=10
```


Solution



```
#eos-ns-inspect drop-from-deathrow --fid 338 --members eospilot-qdb.cern.ch:7777 --  
password-file /etc/eos.keytab --no-dry-run
```



Detecting inconsistencies



» Detecting under and over replications

» Detecting cursed names from older EOS versions

Files named as: “.“, “..“ and “/“




EOS


Detecting cursed name files

```
# eos-ns-inspect check-cursed-names --members eos<instance>-qdb.cern.ch:7777 --password-file  
<keytab_location>
```

Files with cursed names (., .., slashes..)



Created by users and they cannot be removed with the standard eos rm



```
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 32906608, 'instance': 'atlas'}  
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 49607212, 'instance': 'atlas'}  
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 22160498, 'instance': 'atlas'}  
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 46710679, 'instance': 'atlas'}  
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 43970140, 'instance': 'atlas'}  
{'date': '2020-10-29', 'cursed-name': '.', 'fid': 11102599, 'instance': 'atlas'}
```



EOS

Detecting cursed name files

```
#eos-ns-inspect check-cursed-names --members eos<instance>-qdb.cern.ch:7777 --password-file <keytab_location>
```

Cursed name files

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos file info fid:32906608
File: '/eos/atlas/user/b/berge/MadGraphGen/.' Flags: 0000 Clock: 1667e410495c17bd
Size: 4640202
Modify: Tue Aug 5 12:18:26 2014 Timestamp: 1407233906.412797000
Change: Wed Feb 22 15:18:17 2012 Timestamp: 1329920297.264076000
Birth: Thu Jan 1 01:00:00 1970 Timestamp: 0.000000000
CUid: 1724 CGid: 1307 Fxid: 01f61d70 Fid: 32906608 Pid: 116071 Pxid: 0001c567
XStype: adler XS: 45 3d d1 89 ETAGs: "8833300323893248:453dd189"
Layout: replica Stripes: 2 Blocksize: 4k LayoutId: 00600112 Redundancy: d2::t0
#Rep: 2
```

no.	fs_id	host	schedgroup	path	boot	configstatus	drain	active	geotag
0	33206	st_192_22b367fd.cern.ch	default.84	/data112	booted	rw	nodrain	online	0513::R::0050::CB17::69131943
1	36649	st_096_dd9061e6.cern.ch	default.84	/data85	booted	rw	nodrain	online	0513::S::0034::SM08::09718fc3





Detecting cursed name files

```
#eos-ns-inspect check-cursed-names --members eos<instance>-qdb.cern.ch:7777 --password-file  
                    <keytab_location>
```

Cursed name files

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos rm fid:32906608  
# pre-configuring default route to /eos/user/r/root/  
# -use $EOSHOME variable to override  
error: unable to remove file/directory '/eos/atlas/user/b/berge/MadGraphGen/.'
```



Solution

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos file info eos/recovered/2020-10-29  
Directory: '/eos/recovered/2020-10-29' Treesize: 0  
Container: 0 Files: 0 Flags: 40750  
Modify: Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453  
Change: Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453  
Sync : Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453  
Birth : Fri Feb 21 11:23:14 2020 Timestamp: 1582280594.674495679  
CUid: 0 CGid: 0 Fxid: 09882355 Fid: 159916885 Pid: 146696311 Pxid: 08be6877  
ETAG: 9882355:1614510016.280
```



EOS

Detecting cursed name files

```
#eos-ns-inspect check-cursed-names --members eos<instance>-qdb.cern.ch:7777 --password-file <keytab_location>
```

Cursed name files

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos file info eos/recovered/2020-10-29
Directory: '/eos/recovered/2020-10-29' Treesize: 0
Container: 0 Files: 0 Flags: 40750
Modify: Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453
Change: Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453
Sync : Sun Feb 28 12:00:16 2021 Timestamp: 1614510016.280054453
Birth : Fri Feb 21 11:23:14 2020 Timestamp: 1582280594.674495679
CUID: 0 CGid: 0 Fxid: 09882355 Fid: 159916885 Pid: 146696311 Pxid: 08be6877
ETAG: 9882355:1614510016.280
```

eos/recovered/2020-10-29

Located in

Renamed!

↓ Solution

```
#eos-ns-inspect rename-fid --fid 32906608 --destination-cid 159916885 --new-name dot-fid-32906608-cid-116071
--members eosatlas-qdb.cern.ch:7777 --password-file /etc/eos.keytab --no-dry-run
```

dot-fid-32906608-cid-116071

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos file info fid:32906608 -m
keylength.file=42 file=/eos/recovered/dot-fid-32906608-cid-116071 size=4640202 mtime=1407233906.412797000
ctime=1329920297.264076000 btime=0.0 clock=1614510399046335575 mode=0000 uid=1724 gid=1307 fxid=01f61d70
fid=32906608 ino=8833300323893248 pid=146696311 pxid=08be6877 xstype=adler xs=453dd189
etag="8833300323893248:453dd189" detached=0 layout=replica nstripes=2 lid=00600112 nrep=2 fsid=33206 fsid=36649
```





Detecting cursed name files

```
#eos-ns-inspect check-cursed-names --members eos<instance>-qdb.cern.ch:7777 --password-file  
                                <keytab_location>
```


Cursed name files

```
[root@eosatlas-qdb-c49567be6d (mgm:master mq:master) ~]$ eos ls -la eos/recovered/2020-10-29  
drwxr-x--+ 1 root  root   4640202 Feb 28 12:12 .  
drwxr-x--+ 1 root  root     0 Feb 28 12:12 ..  
-rw-r--r-- 2 berge  zp    4640202 Aug 5 2014 dot_fid_32906608_cid_116071
```

**This file can be removed without problems now,
but we always check with the experiments before if they wish that 😊**



Detecting inconsistencies



» Detecting under and over replications

» Detecting cursed names from older EOS versions

Files named as: “.”, “..” and “/”

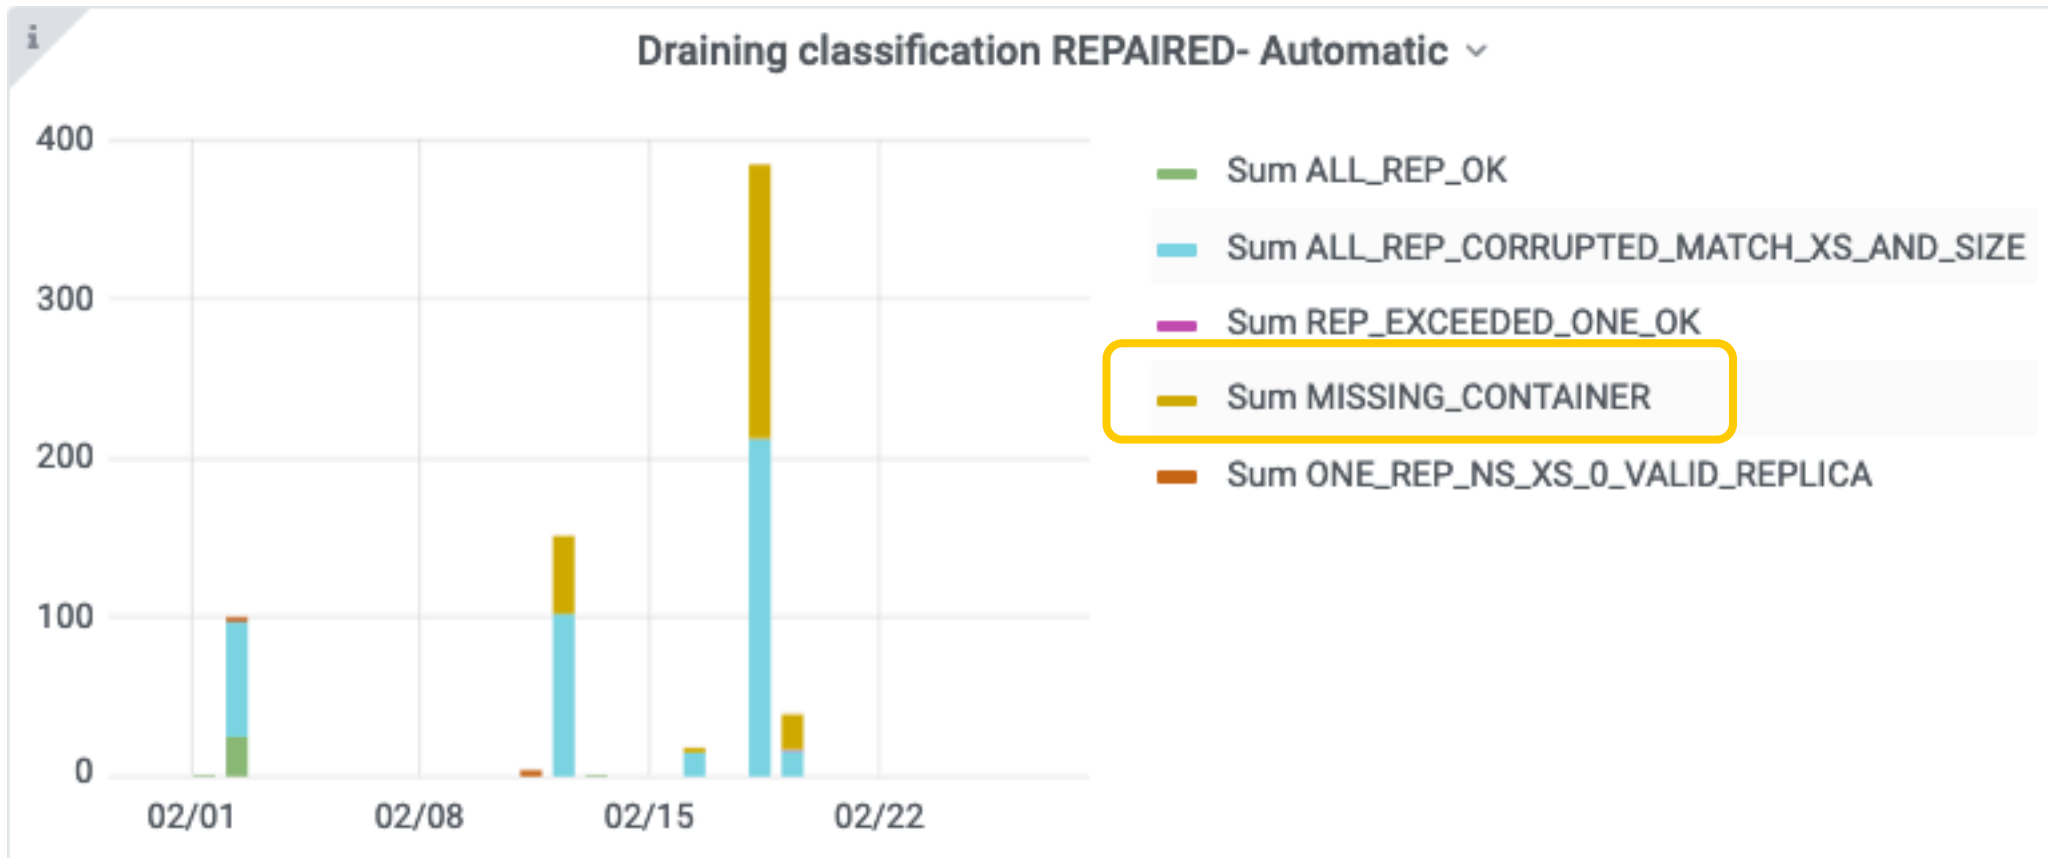
» Repairing missing containers

When files do not have a path associated and instead of getting the path you get the fid.



EOS

Repairing missing containers





EOS

Repairing missing containers

Missing containers

```
[root@eos-pilot-ns-00 (mgm:master mq:master) ~]$ eos file info fid:892 -m  
keylength.file=7 file=fid:892 size=1024 mtime=1605098294.465860000 ctime=1605098294.383594074 btime=1605098294.383594074 ...
```

"Error while fetching ContainerMD"


Solution



```
#eos-ns-inspect fix-detached-parent -fid <fid> --destination-path /eos/recovered/<date> --  
members eos<instance>-qdb:7777 --password-file <keytab_location> --no-dry-run
```




Detecting inconsistencies



» Detecting under and over replications

» Detecting cursed names from older EOS versions

Files named as: “.”, “..” and “/”

» Repairing missing containers

When files do not have a path associated and instead of getting the path you get the fid.