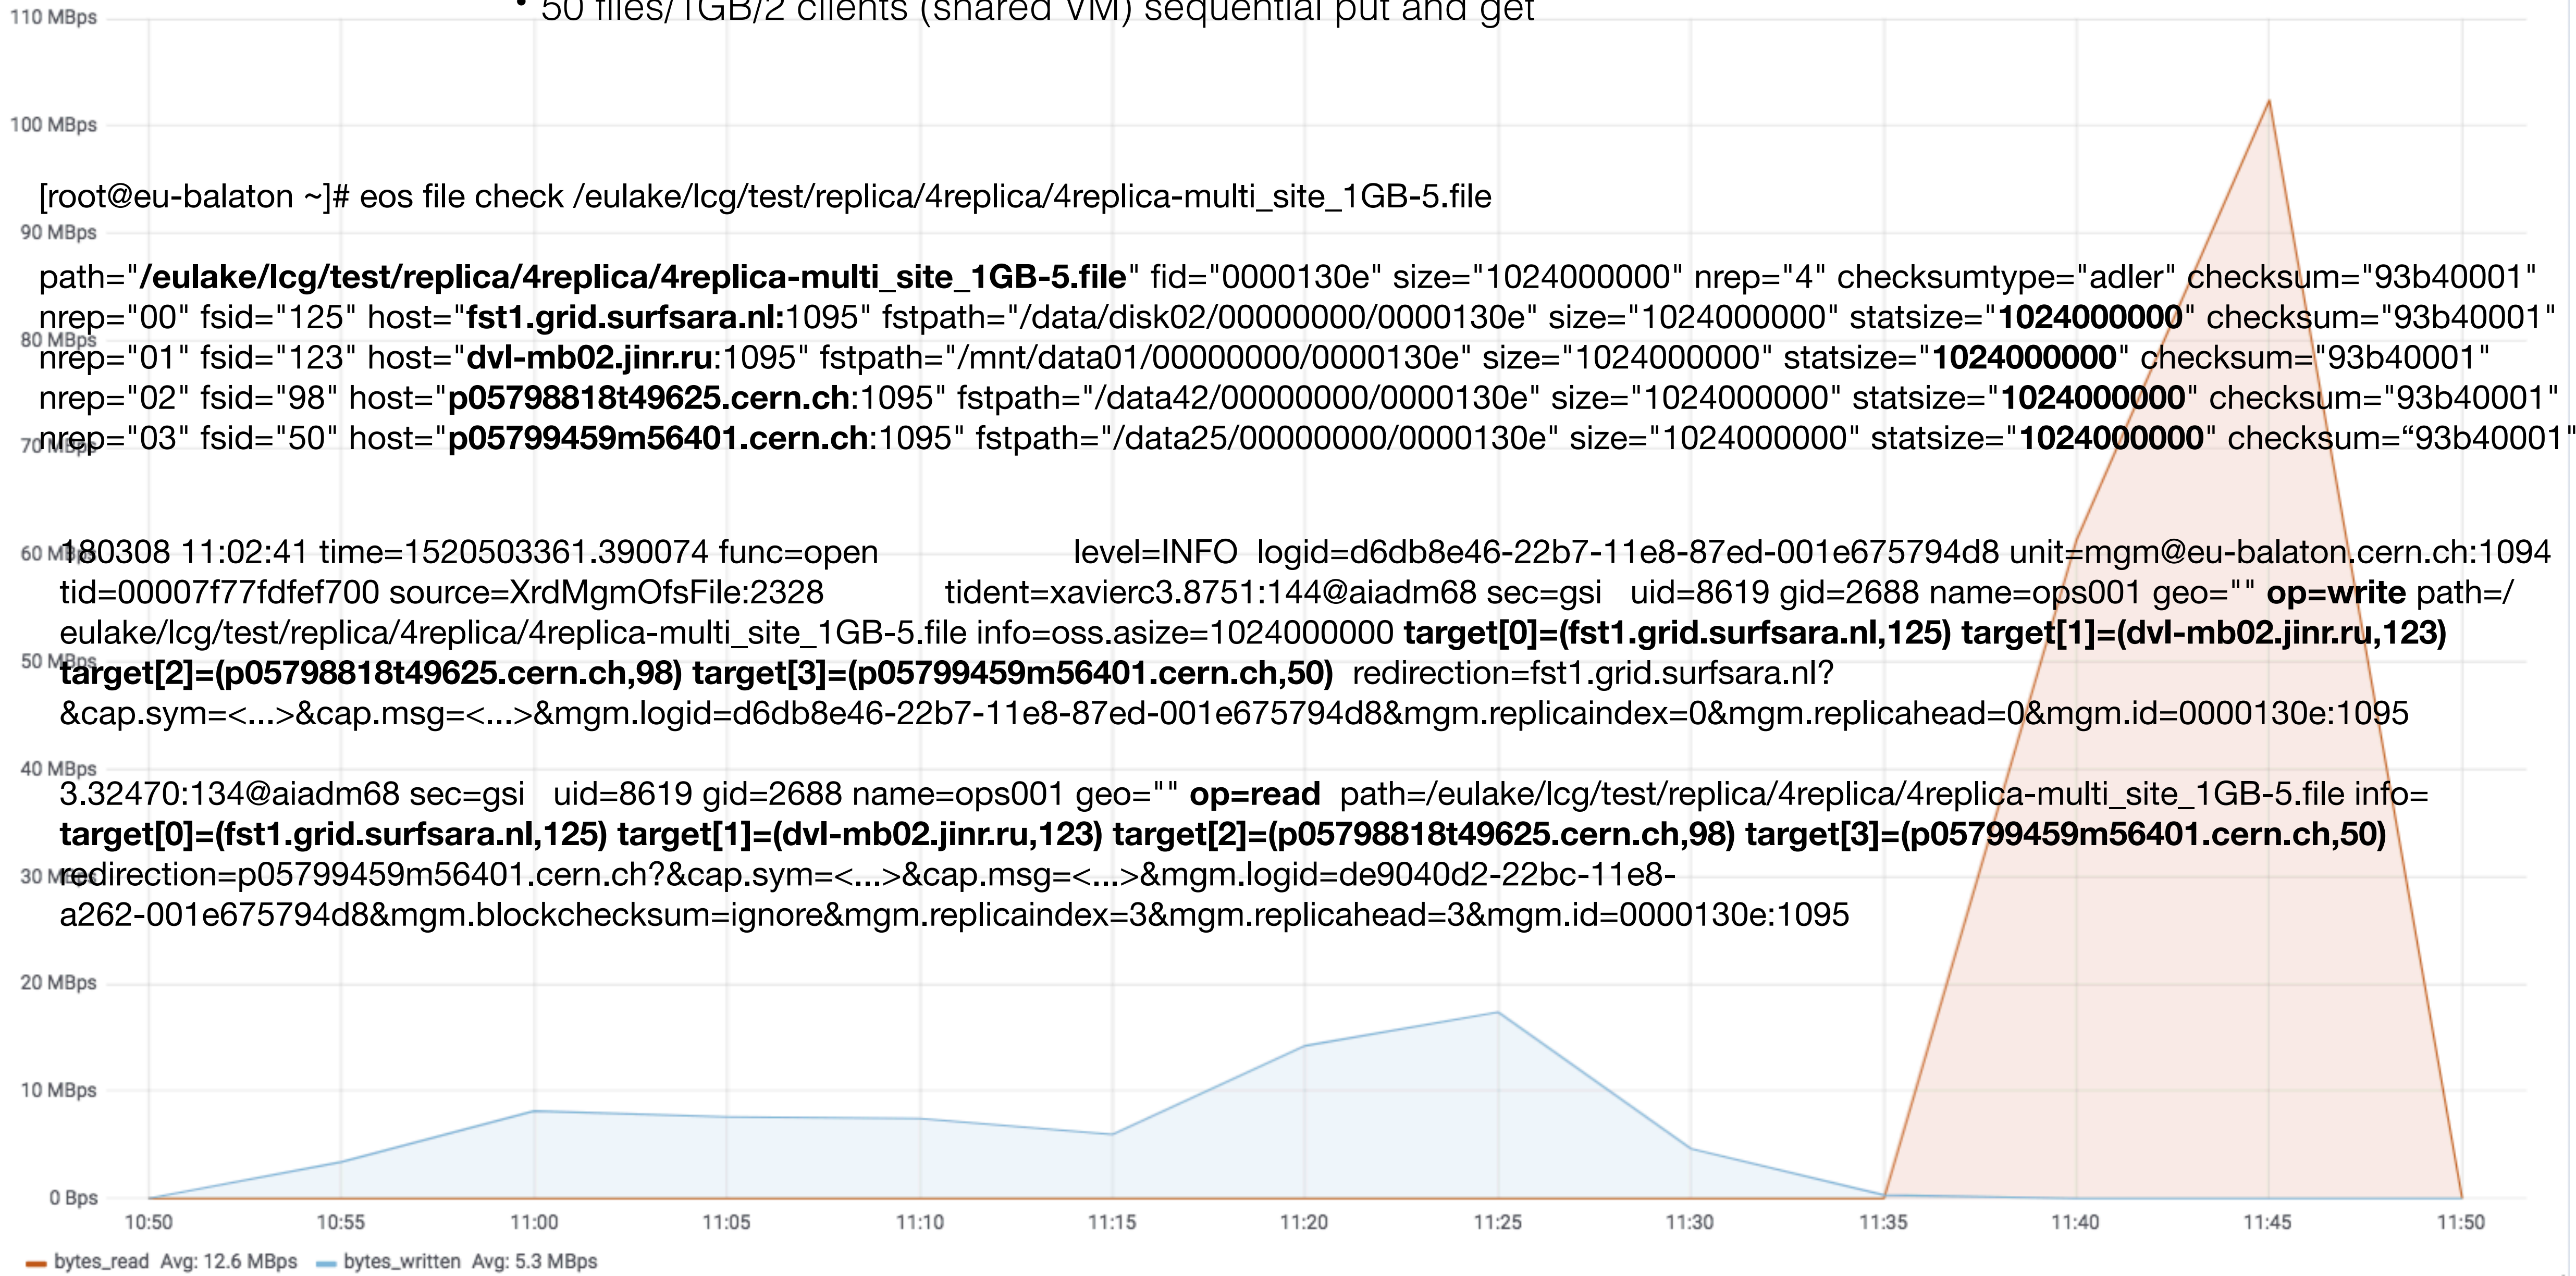


Distributed replicas (4 replicas in 4 sites): /eulake/lcg/test/replica/4replica/

- 50 files/1GB/2 clients (shared VM) sequential put and get

EOS Total IO



```
[root@eu-balaton ~]# eos file check /eulake/lcg/test/replica/4replica/4replica-multi_site_1GB-5.file
```

```
path="/eulake/lcg/test/replica/4replica/4replica-multi_site_1GB-5.file" fid="0000130e" size="1024000000" nrep="4" checksumtype="adler" checksum="93b40001"
nrep="00" fsid="125" host="fst1.grid.surfsara.nl:1095" fstpath="/data/disk02/00000000/0000130e" size="1024000000" statsize="1024000000" checksum="93b40001"
nrep="01" fsid="123" host="dvl-mb02.jinr.ru:1095" fstpath="/mnt/data01/00000000/0000130e" size="1024000000" statsize="1024000000" checksum="93b40001"
nrep="02" fsid="98" host="p05798818t49625.cern.ch:1095" fstpath="/data42/00000000/0000130e" size="1024000000" statsize="1024000000" checksum="93b40001"
nrep="03" fsid="50" host="p05799459m56401.cern.ch:1095" fstpath="/data25/00000000/0000130e" size="1024000000" statsize="1024000000" checksum="93b40001"
```

```
180308 11:02:41 time=1520503361.390074 func=open level=INFO logid=d6db8e46-22b7-11e8-87ed-001e675794d8 unit=mgm@eu-balaton.cern.ch:1094
tid=00007f77dfef700 source=XrdMgmOfsFile:2328 tident=xavierc3.8751:144@aiadm68 sec=gsi uid=8619 gid=2688 name=ops001 geo="" op=write path=/
eulake/lcg/test/replica/4replica/4replica-multi_site_1GB-5.file info=oss.asize=1024000000 target[0]=(fst1.grid.surfsara.nl,125) target[1]=(dvl-mb02.jinr.ru,123)
target[2]=(p05798818t49625.cern.ch,98) target[3]=(p05799459m56401.cern.ch,50) redirection=fst1.grid.surfsara.nl?
&cap.sym=<...>&cap.msg=<...>&mgm.logid=d6db8e46-22b7-11e8-87ed-001e675794d8&mgm.replicaindex=0&mgm.replicahead=0&mgm.id=0000130e:1095
3.32470:134@aiadm68 sec=gsi uid=8619 gid=2688 name=ops001 geo="" op=read path=/eulake/lcg/test/replica/4replica/4replica-multi_site_1GB-5.file info=
target[0]=(fst1.grid.surfsara.nl,125) target[1]=(dvl-mb02.jinr.ru,123) target[2]=(p05798818t49625.cern.ch,98) target[3]=(p05799459m56401.cern.ch,50)
redirection=p05799459m56401.cern.ch?&cap.sym=<...>&cap.msg=<...>&mgm.logid=de9040d2-22bc-11e8-
a262-001e675794d8&mgm.blockchecksum=ignore&mgm.replicaindex=3&mgm.replicahead=3&mgm.id=0000130e:1095
```

Distributed RAIN (3+2: 5 locations): /eulake/lcg/test/rain/rain32/

EOS Total IO

- 50 files/1GB/2 clients (shared VM) sequential put and get

```
[root@eu-balaton ~]# eos file check /eulake/lcg/test/rain/rain32/rain32-multi_site_1GB-20.file
```

```
path="/eulake/lcg/test/rain/rain32/rain32-multi_site_1GB-20.file" fid="0000138c" size="1024000000" nrep="5" checksumtype="adler" checksum="93b40001"
nrep="00" fsid="125" host="fst1.grid.surfsara.nl:1095" fstpath="/data/disk02/00000000/0000138c" size="1024000000" statsize="341839872" checksum="93b40001"
nrep="01" fsid="50" host="p05799459m56401.cern.ch:1095" fstpath="/data25/00000000/0000138c" size="1024000000" statsize="341839872" checksum="93b40001"
nrep="02" fsid="123" host="dvl-mb02.jinr.ru:1095" fstpath="/mnt/data01/00000000/0000138c" size="1024000000" statsize="341839872" checksum="93b40001"
nrep="03" fsid="98" host="p05798818t49625.cern.ch:1095" fstpath="/data42/00000000/0000138c" size="1024000000" statsize="341839872" checksum="93b40001"
nrep="04" fsid="4" host="p05496644k62259.cern.ch:1095" fstpath="/data04/00000000/0000138c" size="1024000000" statsize="341839872" checksum="93b40001"
```

```
3.20250:144@aiadm69 sec=gsi uid=8619 gid=2688 name=ops001 geo="" op=write path=/eulake/lcg/test/rain/rain32/rain32-multi_site_1GB-20.file
info=oss.asize=1024000000 target[0]=(fst1.grid.surfsara.nl,125) target[1]=(p05799459m56401.cern.ch,50) target[2]=(dvl-mb02.jinr.ru,123)
```


```
target[3]=(p05798818t49625.cern.ch,98) target[4]=(p05496644k62259.cern.ch,4) redirection=fst1.grid.surfsara.nl?
```

```
&cap.sym=<...>&cap.msg=<...>&mgm.logid=84bd740c-22e5-11e8-9375-001e675794d8&mgm.replicaindex=0&mgm.replicahead=0&mgm.id=0000138c:1095
```

```
3.2470:145@aiadm69 sec=gsi uid=8619 gid=2688 name=ops001 geo="" op=read path=/eulake/lcg/test/rain/rain32/rain32-multi_site_1GB-20.file info=
```

```
target[0]=(fst1.grid.surfsara.nl,125) target[1]=(p05799459m56401.cern.ch,50) target[2]=(dvl-mb02.jinr.ru,123) target[3]=(p05798818t49625.cern.ch,98)
```

```
target[4]=(p05496644k62259.cern.ch,4) redirection=p05799459m56401.cern.ch?&cap.sym=<...>&cap.msg=<...>&mgm.logid=244386ee-22e9-11e8-
a262-001e675794d8&mgm.replicaindex=1&mgm.replicahead=1&mgm.id=0000138c:1095
```



Distributed Rain(3+2) vs Multi-replica (4) Read-Write performance

50 files/1GB/2 clients (shared VM) sequential puts and gets.

Rain 3+2 distributed across 5 disk servers: NL, RU, CH and HU. Replicas (4) distributed across four sites: NL, RU, CH and HU.

