

Kolkata Tier-2 Storage Evolution

Grid Computing Facility, VECC, Kolkata, India

SAN Storage

Started with SAN Based Storage system

Procured HP EVA 6100 in 2008

2 Racks
Total 200 Spindles
Each 500GB
Total raw capacity 100 TB
usable capacity 68 TB only.
ALICE::Kolkata::SE

Procured IBM DS5100 in 2011

1 Racks
Total 96 Spindles
Each 2TB
Total raw capacity 192 TB
usable capacity 152 TB only.
Added under ALICE::Kolkata::SE

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Kolkata::SE	OK	223.2 TB	24.36%	168.8 TB	54.36 TB	1.722 M	FILE	OK

Monalisa snapshot during 2012 showing Kolkata::SE

Towards Disk based Storage

First EOS storage system at Kolkata Tier2 during March 2017

Procured in March 2017

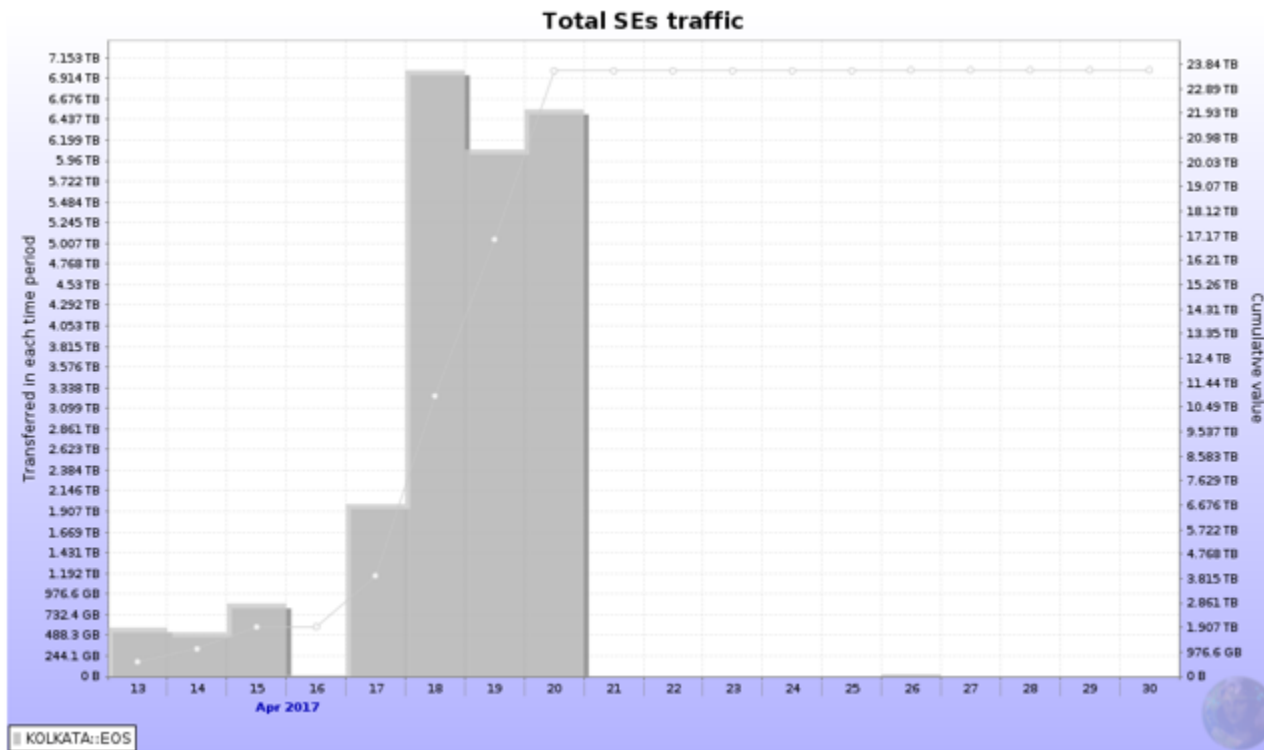
Name of Instance:	ALICE::Kolkata::EOS	EOS version:	EOS AQUAMARINE (0.3.256)
No. of FST Disk Server:	3 Nos.	RAW Space:	144 TB.
Total Physical Space	6 U	No. of Management Server:	1
Each Disk Capacity:	4TB	Disk Type:	NLSAS
Nos. of Disk in each:	12	No. of Disk Group:	6
H/W RAID type:	0 (Zero)	RAIN type (Redundant Array of Independent nodes) :	RAIN6
OS:	Scientific Linux 6	Used as prototype concept.	

```
[root@eos /]# eos -b
EOS Console [root://localhost] |/> space ls
#-----
# type # name # groupsize # groupmod #N(fs) #N(fs-rw) #sum(usedbytes) #sum(capacity) #capacity(rw) #nom.capacity #quota #balancing # threshold # converter # ntx # active #intergroup
#-----
spaceview default 0 24 36 36 52.16 T 141.74 T 141.74 T 0 off on 20 off 2 0 off
```

```
EOS Console [root://localhost] |/> node ls
#-----
# type # hostport # geotag # status # status # txgw #gw-queued # gw-ntx #gw-rate # heartbeatdelta #nofs
#-----
nodesview eos01.tier2-kol.res.in:1095 online on off 0 10 120 2 12
nodesview eos02.tier2-kol.res.in:1095 online on off 0 10 120 2 12
nodesview eos03.tier2-kol.res.in:1095 online on off 0 10 120 2 12
EOS Console [root://localhost] |/>
```

```
[root@eos ~]# eos -b group ls
#-----
# type # name # status #nofs #dev(filled) #avg(filled) #sig(filled) #balancing # bal-shd #drain-shd
#-----
groupview default.0 on 6 0.08 94.88 0.05 idle 0 0
groupview default.1 on 6 0.12 94.87 0.07 idle 0 0
groupview default.2 on 6 0.13 94.87 0.07 idle 0 0
groupview default.3 on 6 0.15 94.86 0.08 idle 0 0
groupview default.4 on 6 23.82 58.87 11.22 balancing 0 0
groupview default.5 on 6 0.09 94.88 0.06 idle 0 0
```

Data transfers during April 2017



- ✓ 25 TB of files,
- ✓ 1Gbps Network,
- ✓ Helped to upgrade to 10 Gbps,
- ✓ Few files could not complete,
- ✓ Beginner in EOS
- ✓ Rebuilding was not working.

Transfer requests (add new request)								
10122			- any -					Filter
ID	Path	Target SE	Status	Progress	Files	Total size	Started	Ended
10122	MC_output_with_single_copy	ALICE::KOLKATA::EOS	Error		89585	24.58 TB	13 Apr 2017 14:04	03 Oct 2017 12:36
1 requests					89585	24.58 TB		

Requests per page: 500

Expanding of Kolkata Tier2 EOS

Hardware for Disk Server (FST)

- ❖ 7 nos. of Dell PowerEdge R730XD servers. Each server contains:-
 - 2 * Intel Xeon E5-2630 v4 @ 2.20GHz 10Core Processors.
 - 16 * 10TB NLSAS HDD and 2*480GB SSD.
 - 4 *10GB Ethernet (Fibre).
 - 8 * 16 = 128GB DDR4 RAM.
 - RAID: Server are configure with 2 types of RAID:-
 - RAID-0 in 16 nos 10TB NLSAS HDD.
 - RAID-1 in 2 nos 480 GB SSD.

Hardware for Management Server (Manager)

- ❖ 2 nos. of Dell PowerEdge R640 1U Rack Mount Servers. Each server contains:-
 - 2* Intel Xeon Processor Silver 4112 @ 4C 2.60 Ghz Processors
 - 64GB RAM
 - 1 * 480GB SSD and 4*1.2TB SAS HDD.
 - 2 * 10GB Ethernet (Fibre)
 - RAID: Server are configure with 2 types of RAID:-
 - RAID-5 in 3 nos of 1.2TB SAS + one hot spare.
 - RAID-0 in 480 GB SSD.

- ❖ Name of Instance: ALICE::Kolkata::EOS2
- ❖ Procurement : During 2018
- ❖ Space: 1.1PB.
- ❖ No. of Management Server: 2 (Master/Slave)
- ❖ No. of Disk Server: 7Nos.
- ❖ Each Disk Capacity : 10TB
- ❖ Hard Disk Type: NLSAS
- ❖ Nos. of Disk in each disk server: 16
- ❖ H/W RAID type: 0 (Zero)
- ❖ OS: CentOS 7.x
- ❖ EOS version: EOS CITRINE (4.4.23)
- ❖ RAIN type (Redundant Array of Independent nodes) : RAIN6
- ❖ In-Memory Namespace


Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Kolkata::EOS	OK	128.9 TB	26.94%	94.17 TB	34.73 TB	475.3 K	FILE	FAIL
ALICE::Kolkata::EOS2	OK	1.1 PB	6.223%	1.032 PB	70.09 TB	1.324 M	FILE	OK
ALICE::Kolkata::SE	OK	76.39 TB	32.14%	51.84 TB	24.55 TB	458.7 K	FILE	OK

RAID and RAIN

RAID6

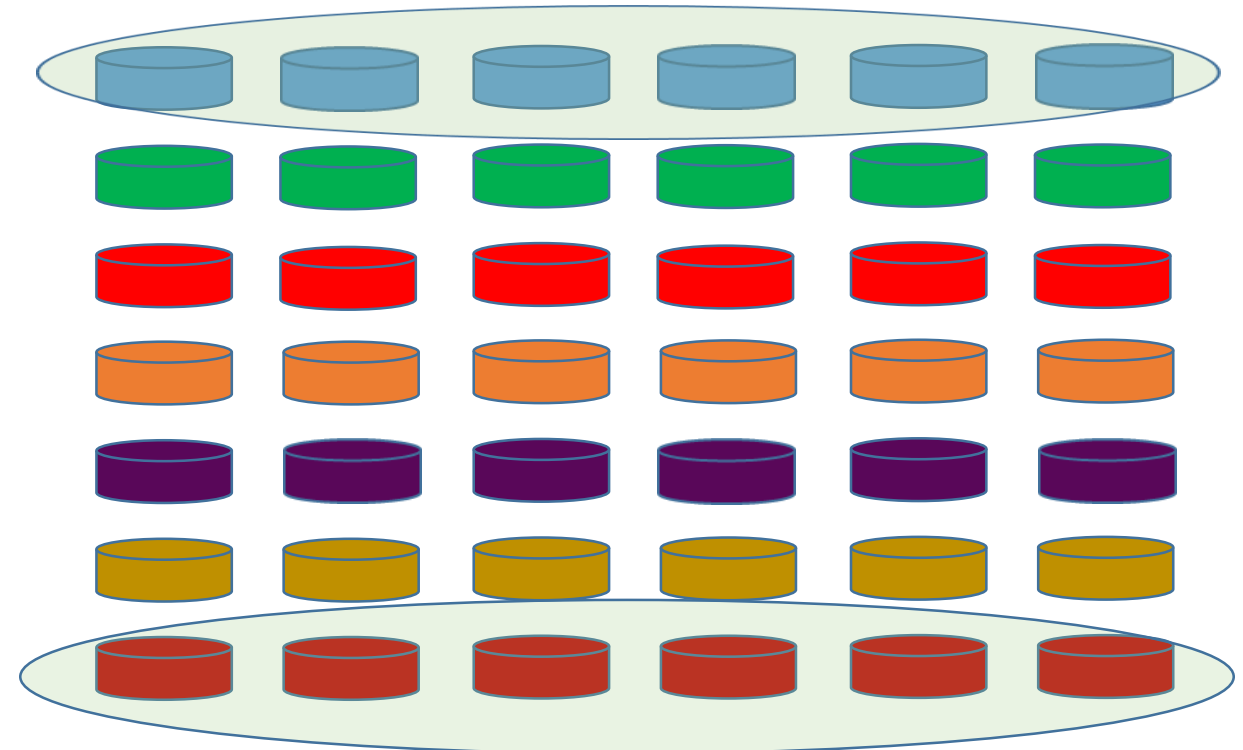
Each Node1



2 Parity on same node within RAID6

RAIN6

Node1 Node2 Node3 Node4 Node5 Node6



2 Parity on different nodes within RAIN6

Disk Grouping of 7 disk across 7 Disk Server to make RAIN6



Disk Grouping of Hard Disk with 7 disk within 7 Disk Server to make RAIN6 – (2)

```
[root@eos-mgm ~]# eos -b group ls
```

type	name	status	N(fs)	dev(filled)	avg(filled)	sig(filled)	balancing	bal-shd	drain-shd
groupview	default.0	on	7	1.67	19.03	0.79	idle	0	0
groupview	default.1	on	7	1.66	14.59	0.76	idle	0	0
groupview	default.10	on	7	1.03	13.19	0.48	idle	0	0
groupview	default.11	on	7	1.14	16.02	0.55	idle	0	0
groupview	default.12	on	7	1.05	16.82	0.49	idle	0	0
groupview	default.13	on	7	1.72	14.51	0.79	idle	0	0
groupview	default.14	on	7	1.15	13.36	0.52	idle	0	0
groupview	default.15	on	7	1.11	12.82	0.51	idle	0	0
groupview	default.2	on	7	1.18	18.30	0.56	idle	0	0
groupview	default.3	on	7	1.06	13.56	0.53	idle	0	0
groupview	default.4	on	7	0.94	13.46	0.44	idle	0	0
groupview	default.5	on	7	1.20	12.93	0.53	idle	0	0
groupview	default.6	on	7	1.24	17.84	0.57	idle	0	0
groupview	default.7	on	7	1.32	13.11	0.58	idle	0	0
groupview	default.8	on	7	1.07	13.06	0.50	idle	0	0
groupview	default.9	on	7	1.12	13.00	0.53	idle	0	0

Screenshot of 16 disk group in ALICE::Kolkata::EOS2

Screenshot of 7 nodes group in ALICE::Kolkata::EOS2

```
+++++[root@eos-mgm ~]# eos -b node ls
```

type	hostport	geotag	status	activated	txgw	gw-queued	gw-ntx	gw-rate	heartbeatdelta	nofs
nodesview	eos04.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos05.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos06.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos07.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos08.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos09.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16
nodesview	eos10.tier2-kol.res.in:1095	Kolkata::EOS2	online	on	off	0	10	120	0	16

Namespace synchronisation between Master and Slave MGM

After successful installation of eos, we check the namespace sync in both mgm i.e.:-

#eos -b ns

➤ In eos-mgm:-

```
# -----  
ALL   Replication           mode=master-rw state=master-rw  
master=eos-mgm.tier2-kol.res.in configdir=/var/eos/config/eos-mgm.tier2-  
kol.res.in/ config=default mgm:eos-slave.tier2-kol.res.in=ok  
mgm:mode=slave-ro mq:eos-slave.tier2-kol.res.in:1097=ok  
# -----  
ALL   File Changelog Size    1.03 GB  
ALL   Dir  Changelog Size    515.61 MB  
# -----
```

➤ In eos-slave:-

```
# -----  
ALL   Replication           mode=slave-ro state=slave-ro master=eos-  
mgm.tier2-kol.res.in configdir=/var/eos/config/eos-mgm.tier2-kol.res.in/  
config=default mgm:eos-mgm.tier2-kol.res.in=ok mgm:mode=master-rw mq:eos-  
mgm.tier2-kol.res.in:1097=ok  
# -----  
ALL   File Changelog Size    1.04 GB  
ALL   Dir  Changelog Size    525.66 MB  
# -----
```

Few common trouble-shootings

Problem-1 > Sometimes FSTs could not contact MGM. FST's xrdfst.log shows error like:-

```
=====
180910 20:55:37 time=1536593137.299353 func=Cleaner          level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598b2fe700 source=Cleaner:43      tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

180910 20:55:40 time=1536593140.302278 func=Publish        level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598bfff700 source=Publish:90      tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

180910 20:55:42 time=1536593142.285195 func=Remover        level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598ddf700 source=Remover:45       tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

180910 20:55:42 time=1536593142.291149 func=WaitConfigQueue level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598b3ff700 source=Drainer:48      tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

180910 20:55:42 time=1536593142.291251 func=WaitConfigQueue level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598bef700 source=Drainer:48      tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

180910 20:55:42 time=1536593142.299520 func=Cleaner          level=INFO logid=static..... unit=fst@eos01.tier2-
kol.res.in:1095 tid=00007f598b2fe700 source=Cleaner:43      tident= sec=(null) uid=0 gid=0 name=- geo="" Snoozing ...

@@@@@@ 00:00:00 op=shutdown msg="shutdown timedout after 15 seconds"

@@@@@@ 00:00:00 op=shutdown status=forced-complete
=====
```

Solution : eos-log-repair

```
eos-log-repair files.eos.tier2-kol.res.in.mdlog.300617 files.eos.tier2-kol.res.in.mdlog
Header status: OK (version: 0x1, content: 0x1)
Elapsed time: 0 m. 44 s. Progress: 1.235 GB / 1.246 GB Scanned:          9500282
Healthy:          9500281
Bytes total:      1338847232
Bytes accepted:   1338847096
Bytes discarded:  136
Not fixed:        1
Fixed (wrong magic): 0
Fixed (wrong checksum): 0
Fixed (wrong size): 1
Elapsed time:     0 m. 44 s.
=====
```

After repairing of log, again restart eos daemon.

Few common trouble-shootings cont...

Problem-2> Due to large namespace (nos. of files to store is large) sometime eos daemon crashed and then took too long to boot or hanged.

EOS down due to corruption of /var/eos/md/directories.eos-mgm.tier2-kol.res.in.mdlog and/or /var/eos/md/files.eos-mgm.tier2-kol.res.in.mdlog files.

Solution: eos-log-compact

Repair: eos-log-compact to repair namespace files and it deletes bad records. But there reference files were sometimes not deleted.

=====

```
[root@eos-mgm ~] eos-log-compact /var/eos/md/directories.eos-mgm.tier2-kol.res.in.mdlog /var/eos/md/directories.eos-mgm.tier2-kol.res.in.mdlog.compact
```

And

```
[root@eos-mgm ~] eos-log-compact /var/eos/md/files.eos-mgm.tier2-kol.res.in.mdlog /var/eos/md/files.eos-mgm.tier2-kol.res.in.mdlog.compact
```

.....

.....

Records kept: 6854362 out of 23597126

Elapsed time: 0 m. 53 s. Records written: 6854362 out of 6854362

Records updated 22855024

Records deleted: 742099

Records total: 23597126

Records kept: 6854362

Records written: 6854362

Elapsed time: 0 m. 53 s.

=====

After compactification, we again restored compacted mdlog to it's actual location.

```
cp -rv /var/eos/md/files.eos-mgm.tier2-kol.res.in.mdlog.compact /var/eos/md/files.eos-mgm.tier2-kol.res.in.mdlog
```

And

```
cp -rv /var/eos/md/directories.eos-mgm.tier2-kol.res.in.mdlog.compact /var/eos/md/directories.eos-mgm.tier2-kol.res.in.mdlog
```

In-Memory Namespace to Quarkdb (Disk based Namespace)

- ❖ Name of EOS Instance: ALICE::Kolkata::EOS2
- ❖ During Last week of April 2021
- ❖ Space: 1.1PB.
- ❖ No. of Quark DB-Member Servers: 3
- ❖ No. of EOS servers: 2 Nos.
- ❖ No. of Disk Server: 7Nos.
- ❖ OS: CentOS 7.x
- ❖ EOS version: EOS CITRINE (4.8.46)
- ❖ QuarkDB version: 0.4.2.1
- ❖ Added 960 GB SSD (RAID-0) for storing QuarkDB files.

QuarkDB Nodes Details:-

```
[root@eos-mgm ~]# redis-cli -p 7777 raft-info
```

```
1) TERM 101244
2) LOG-START 433001111
3) LOG-SIZE 483298647
4) LEADER eos-slave.tier2-kol.res.in:7777
5) CLUSTER-ID 8370db63-fdf7-4df9-807a-36a7c5692232
6) COMMIT-INDEX 483298646
7) LAST-APPLIED 483298646
8) BLOCKED-WRITES 0
9) LAST-STATE-CHANGE 1817550 (21 days, 52 minutes, 30 seconds)
10) -----
11) MYSELF eos-mgm.tier2-kol.res.in:7777
12) VERSION 0.4.2
13) STATUS FOLLOWER
14) NODE-HEALTH GREEN
15) JOURNAL-FSYNC-POLICY sync-important-updates
16) -----
17) MEMBERSHIP-EPOCH 1111
18) NODES eos-mgm.tier2-kol.res.in:7777,eos-slave.tier2-kol.res.in:7777,eos-qdb.tier2-kol.res.in:7777
19) OBSERVERS
```

```
20) QUORUM-SIZE 2
```

```
[root@eos-mgm ~]#
```

Improvement after migrating from In-Memory namespace to Quarkdb

```
root@eos-mgm ~]# eos -b ns
# -----
# Namespace Statistics
# -----
ALL Files          38506903 [booted] (0s)
ALL Directories    59201
ALL Total boot time 3 s
# -----
ALL Replication    is_master=true master_id=eos-mgm.tier2-kol.res.in:1094
# -----
ALL files created since boot 337874
ALL container created since boot 26
# -----
ALL current file id 78647024
ALL current container id 60663
# -----
ALL eosxd caps      0 c: 0 cc: 0 cic: 0 ic: 0
ALL eosxd clients   0
ALL eosxd active clients 0
ALL eosxd locked clients 0
# -----
ALL File cache max num 40000000
ALL File cache occupancy 643624
ALL In-flight FileMD 0
ALL Container cache max num 5000000
ALL Container cache occupancy 36196
ALL In-flight ContainerMD 0
# -----
ALL eosViewRWMutex peak-latency 0ms (last) 0ms (1 min) 0ms (2 min) 0ms (5 min)
# -----
ALL QClient performance (RTT) 0ms (min) 1ms (avg) 272ms (max)
# -----
ALL memory virtual 16.13 GB
ALL memory resident 12.23 GB
ALL memory share 35.57 MB
ALL memory growths 13.74 GB
ALL threads 404
ALL fds 506
ALL uptime 1902543
# -----
ALL drain info thread_pool=central_drain min=10 max=100 size=10 queue_size=0
ALL fsck info thread_pool=fsck min=2 max=20 size=2 queue_size=0
ALL converter info thread_pool=converter min=16 max=100 size=16 queue_size=0
# -----
ALL tracker info tracker=balance size=161
# -----
[root@eos-mgm ~]#
```

- ❖ Boot time reduced from few 10s of minutes to few seconds.
- ❖ Availability of namespace increased,
- ❖ 3 Quark DB Servers provide higher reliability,
- ❖ Which comes last become the Leader,
- ❖ Stable and smooth operations.

Problems due to no spare

Problem-1: if any disk was failed, then eos became READ Only → ADD Test failed till faulty HDD for the scheduling group was not replaced with right HDD and it's correctly added to same scheduling group with proper way and permission.

Problem-2: Cannot update firmware or Kernel. For the same need to take downtime.

Problem-3: Also we observed that after replacing of faulty disk, balancer and converter were not working and correcting the disk. Draining of disk was not done. As there is no extra file systems which can be used to recover any broken or missing stripes.

Added one more FST Disk Server as a spare

for RAIN6 nstripes=7
there are 7 FST's and each group has 7
disks.

```
[root@eos-mgm ~]# eos attr ls
```

```
/eos/alicekolkata/grid
```

```
.....
```

```
sys.forced.layout="raid6"
```


```
sys.forced.nstripes="7"
```

```
.....
```

```
[root@eos-mgm ~]#
```

- **During May 2021, added one FST**
- **eos11.tier2-kol.res.in**
- **RAIN6 (5 data + 2 parity + 1 spare)**
- **Balancer and converter are working properly.**
- **Draining of Disk also started**
- **Online firmware update is possible**
- **Less downtime.**
- **Less add test error**

Smooth running since May 2021



Size

	Series	Last value	Min	Avg	Max
1.	Kolkata - EOS2	1.115 PB	997.7 TB	1.029 PB	1.115 PB
Total		1.115 PB		1.029 PB	

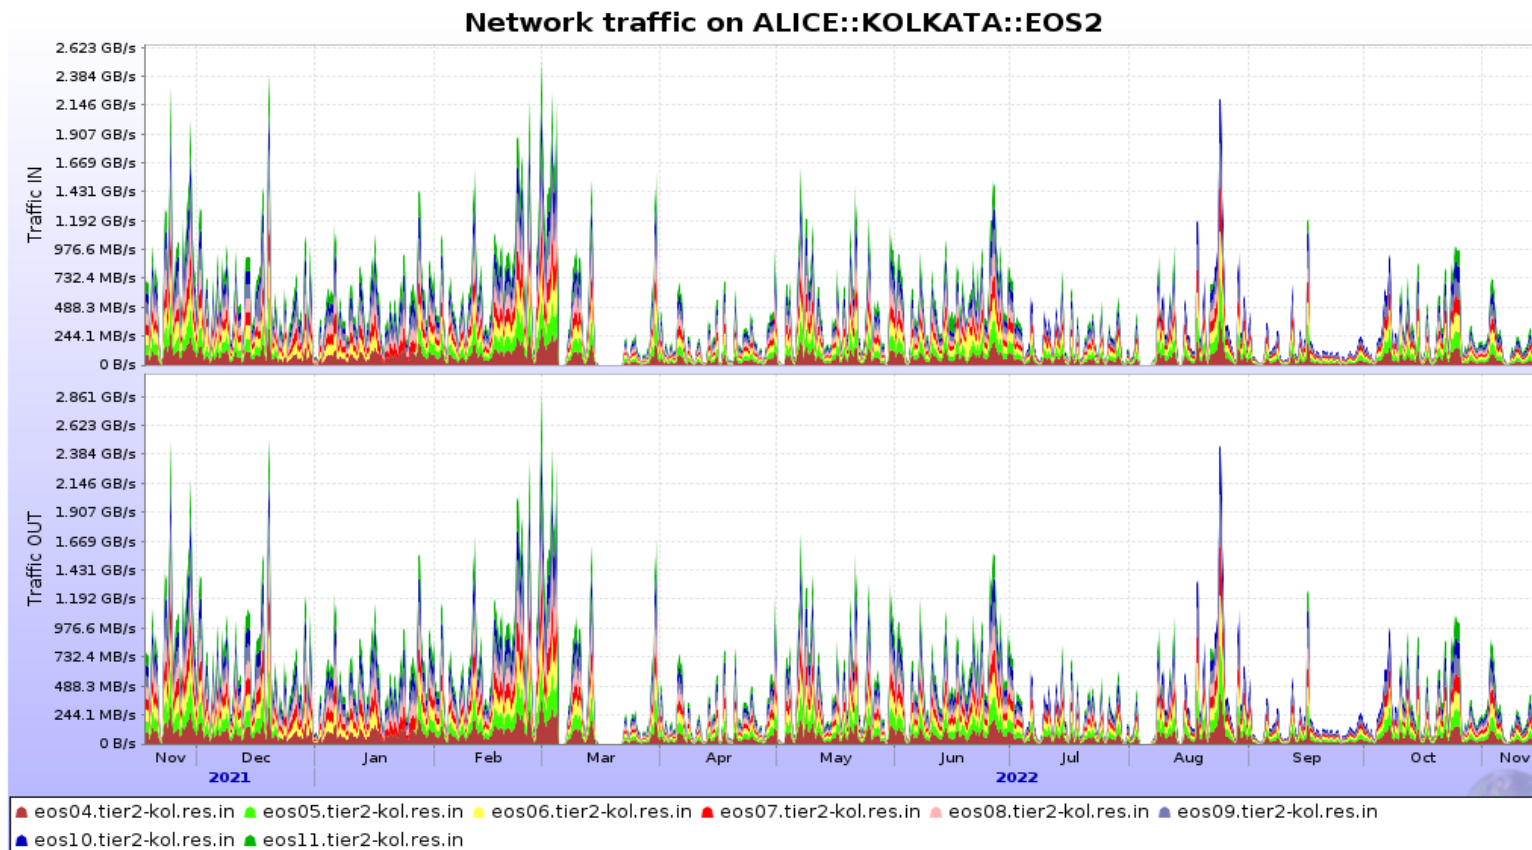
Used

	Series	Last value	Min	Avg	Max
1.	Kolkata - EOS2	223.3 TB	96.53 TB	149.4 TB	223.4 TB
Total		223.3 TB		149.4 TB	

Number of files

	Series	Last value	Min	Avg	Max
1.	Kolkata - EOS2	7013122	2744137	5896256	10711953
Total		7013122		5896256	


Traffic on Kolkata::EOS FSTs during the last one year



Traffic IN					
Series	Last value	Min	Avg	Max	Total
1. eos04.tier2-kol.res.in	4.391 MB/s	12.8 B/s	61.99 MB/s	686 MB/s	1.817 PB
2. eos05.tier2-kol.res.in	3.251 MB/s	15.14 B/s	54.39 MB/s	644 MB/s	1.594 PB
3. eos06.tier2-kol.res.in	4.081 MB/s	11.22 B/s	62.08 MB/s	698.2 MB/s	1.819 PB
4. eos07.tier2-kol.res.in	4.284 MB/s	11.22 B/s	60.82 MB/s	704.3 MB/s	1.782 PB
5. eos08.tier2-kol.res.in	2.034 MB/s	7.377 B/s	50.82 MB/s	616.1 MB/s	1.489 PB
6. eos09.tier2-kol.res.in	4.103 MB/s	42.44 B/s	65.88 MB/s	700.6 MB/s	1.931 PB
7. eos10.tier2-kol.res.in	4.309 MB/s	11.71 B/s	61.7 MB/s	676.9 MB/s	1.808 PB
8. eos11.tier2-kol.res.in	5.119 MB/s	43.9 B/s	58.14 MB/s	605.3 MB/s	1.704 PB
Total	31.57 MB/s		475.8 MB/s		13.94 PB


Traffic OUT					
Series	Last value	Min	Avg	Max	Total
1. eos04.tier2-kol.res.in	3.969 MB/s	20.1 B/s	70.01 MB/s	779 MB/s	2.052 PB
2. eos05.tier2-kol.res.in	3.238 MB/s	20.67 B/s	60.71 MB/s	724.7 MB/s	1.779 PB
3. eos06.tier2-kol.res.in	4.876 MB/s	7.049 B/s	66.82 MB/s	754 MB/s	1.958 PB
4. eos07.tier2-kol.res.in	4.315 MB/s	20.34 B/s	68.1 MB/s	740.4 MB/s	1.996 PB
5. eos08.tier2-kol.res.in	3.666 MB/s	20.03 B/s	56.34 MB/s	631.6 MB/s	1.651 PB
6. eos09.tier2-kol.res.in	5.052 MB/s	20.13 B/s	69.88 MB/s	809.3 MB/s	2.048 PB
7. eos10.tier2-kol.res.in	3.837 MB/s	20.03 B/s	68.88 MB/s	744 MB/s	2.019 PB
8. eos11.tier2-kol.res.in	4.991 MB/s	20.03 B/s	61.9 MB/s	641.6 MB/s	1.814 PB
Total	33.94 MB/s		522.6 MB/s		15.32 PB

Present Status for the last 18 months (>95% Writable)



Statistics						
Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
Kolkata::EOS2	31 Jul 2021 13:27	17 Nov 2022 12:27	10930	455	96%	95.99%

Past 2019-20 Status (Many failures due to multiple reasons.)



Statistics						
Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
Kolkata::EOS2	18 Nov 2019 13:00	31 Jul 2020 13:59	2448	3696	39.84%	39.85%

A step towards Performance Evaluation of Disk based server

- Started finding out performance of such disk server. Needed for new procurement.
- Such server from any OEM comes with RAID + Host Bus Expander Card.
- No OEM have more than 8 channels of max 12 Gbps speed per RAID Cards. Will this sufficient?
- Is provided “RAID + Expander card” limits on the performance? (Bottleneck may be PCIe3.0)
- Try to gather throughput or IOPS. Simple Available tools:-
 - xdd
 - dd
 - rsync (Running multiple rsync in parallel on all the disks)
 - IOZone
 - IOR
- Per Disk 7.2K RPM NLSAS speed is approx 200-300MBps
https://www.cloudbyte.com/wp-content/uploads/Whitepaper_CloudByte_Measuring-Storage-Performance.pdf

Playing with different tools

```
dd if=/dev/urandom
of=/root/vol9/100g
bs=2M count=50000
```

volume	write rate
vol9	31.3MB/s
vol10	31.1MB/s
vol11	31.2MB/s
vol12	31.1MB/s
vol13	32.2MB/s
vol14	31.2MB/s
vol15	32.1MB/s
vol16	31.3MB/s





rsync in parallel
or
simultaneously

```
xdd.linux -op write -targets 1 /dev/sdb -mbytes
5000 -blocksize 1048576 -dio
```

	OUTPUT RATE(write)	OUTPUT RATE(read)
RAID0 (1*10 TB)	326.146 mb	250 mb
RAID5(7*10 TB)	1787.418 mb	1310 mb
RAID6(8*10 TB)	1784.587 mb	1321 mb

volume	file size	copy rate
vol1	104.88gb	107.96 Mbyte/sec
vol2	104.88gb	108.07 Mbyte/sec
vol3	104.88gb	108.29 Mbyte/sec
vol4	104.88gb	108.41Mbyte/sec
vol5	104.88gb	108.52Mbyte/sec
vol6	104.88gb	108.63Mbyte/sec
vol7	104.88gb	108.63Mbyte/sec
vol8	104.88gb	108.74Mbyte/sec
vol9	104.88gb	108.86Mbyte/sec
vol10	104.88gb	108.97Mbyte/sec
vol11	104.88gb	109.08Mbyte/sec
vol12	104.88gb	109.08Mbyte/sec
vol13	104.88gb	109.20Mbyte/sec
vol14	104.88gb	109.31Mbyte/sec
vol15	104.88gb	109.42Mbyte/sec
vol16	104.88gb	109.54Mbyte/sec

Write speed on Disks in a Disk Server via IOZone



Preliminary findings:-
For different RAID Configurations,
By varying number of disks,
Performance is limited by RAID Card (12Gbps)

Thank You

Queries and Questions



- At present Kolkata Tier2 has no unused IPv4 IPs.
We want to expand EOS in future without IPv4, only IPv6.
- How to optimize EOS to utilize parallel read, write on different Groups?
- How to add new disk servers under existing EOS? (Can expand the group.)
- How to scale EOS? We have few more similar disk servers.
- Graphical Monitoring of EOS will help for optimizing and debugging?

```
[root@eos-mgm ~]# eos space status default
```

```
# -----  
# Space Variables  
# .....  
autorepair      := off  
balancer        := on  
balancer.node.ntx      := 20  
balancer.node.rate    := 500  
balancer.threshold   := 1  
converter       := on  
converter.ntx       := 80  
drainer.node.nfs     := 20  
drainer.node.ntx     := 50  
drainer.node.rate    := 200  
drainperiod        := 86400  
fusex.hbi         := 10  
fusex.qti         := 10  
geobalancer       := off  
geobalancer.ntx     := 10  
geobalancer.threshold := 5  
graceperiod       := 86400  
groupbalancer     := on  
groupbalancer.engine := std  
groupbalancer.max_threshold := 5  
groupbalancer.min_threshold := 3  
groupbalancer.ntx   := 40  
groupbalancer.threshold := 0.5  
groupdrainer       := on  
groupmod           := 24  
groupsize          := 8  
headroom           := 5.10 GB  
inspector          := on  
inspector.interval := 86400  
lru                := on  
lru.interval       := 604800  
quota              := off  
scaninterval       := 1209600  
scanrate           := 100  
tapeawaregc.minfreebytes := 0  
tracker            := on  
wfe                := off  
wfe.interval       := 10  
wfe.ntx
```