

Some Results of Measurement Server DISK I/O Performance.

For COMPASS Fron-end, Trigger and DAQ Workshop
2-3 March 2020

V.Frolov(JINR), K.Gritsaj(JINR)

Using Modern RAID Controllers for in High-performance systems.

- Modern RAID Controllers and HDDs allow us to build very fast disk systems with reasonable price, especially when we are using RAID0 disk configuration (up to 6Gb/s per port);
- Unfortunately the performance is depend on number of concurrent read/write processes very much. The tests which we have done on COMPASS last year showed us quite big performance during single write/read process (~800MB/s for write and 600 MB/s for read) and small R/W speed when 2 or more processes were running simultaneously (~ 100 MB/s);
- The R/W speed also depend on the filling level of the disks in array: at the beginning of the disk it is almost two times bigger than at the end of it.

The Test System

- For the tests we have used following hardware:
 1. SSG-6028R-X10DRHi Supermicro Superserver (2xCPU, 32GB RAM);
 2. AVAGO **MR936i-8i** MegaRaid controller (PCI-express 3.0 x8);
 3. 24 disks 2TB **Seagate BarraCuda ST2000DM008**.
- Operating System: **CentOS 7**
- We had tested RAID0 with following configurations 2x(6 disks), 3x24, 2x12 and single array with 24 disks;
- File system: **XFS**
- For measurement of the disk performances we have used following commands:
 1. For write tests:
dd if=/dev/zero of=/disk/fname bs=1M count=1M conv=fsync
 2. For read tests:
dd if=/disk/fname of=/dev/null bs=1M

6 disks RAID0 performance results:

2 RAID0 disk arrays with 6 HDDs in each array.

We have got following results:

Write speed:	1.0793 GB/s (<i>maximum</i>)
Read speed:	0.9024 GB/s (<i>maximum</i>)
Simultaneous write/write speed:	1.0754/1.0777 GB/s
Simultaneous write/read speed:	1.0796/1.0590 GB/s

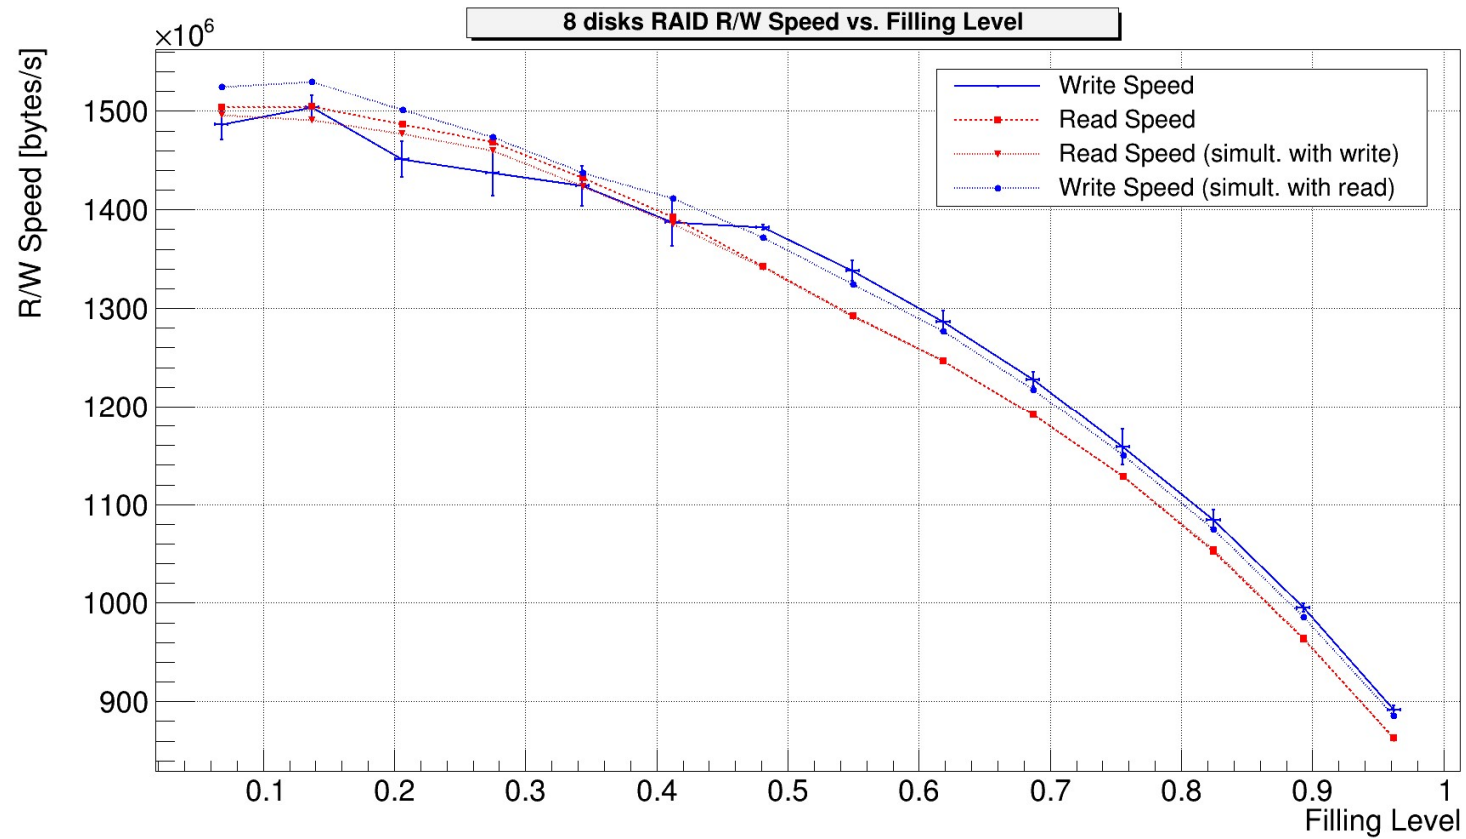
8 disks RAID0 performance results:

3 RAID0 disk arrays with 8 HDDs in each array.

We have got following results:

Write speed:	1.5 GB/s (<i>maximum</i>) 0.897 GB/s (<i>minimum</i>)
Read speed:	1.5 GB/s (<i>maximum</i>) 0.863 GB/s (<i>minimum</i>)
Simultaneous write/write speed:	(max= 1.5 /min= 0.886 GB/s)
Simultaneous write/read speed:	(max= 1.5 /min= 0.860 GB/s)

8 disks RAID0 performance results:



12 disks RAID0 performance results:

2 RAID0 disk arrays with 12 HDDs in each array.

We have got following results:

Write speed: **1.5 GB/s** (*maximum*)

1.3 GB/s (*minimum*)

Read speed: **1.9 GB/s** (*maximum*)

1.2 GB/s (*minimum*)

Simultaneous write/write speed: (max=**1.5**)/(min=**1.3** GB/s)

Simultaneous read/read speed: (max=**1.9**)/(min=**1.2** GB/s)

24 disks RAID0 performance results:

Single RAID0 disk arrays with 24 HDDs.

We have got following results:

Write speed:

1.6 GB/s (*maximum*)

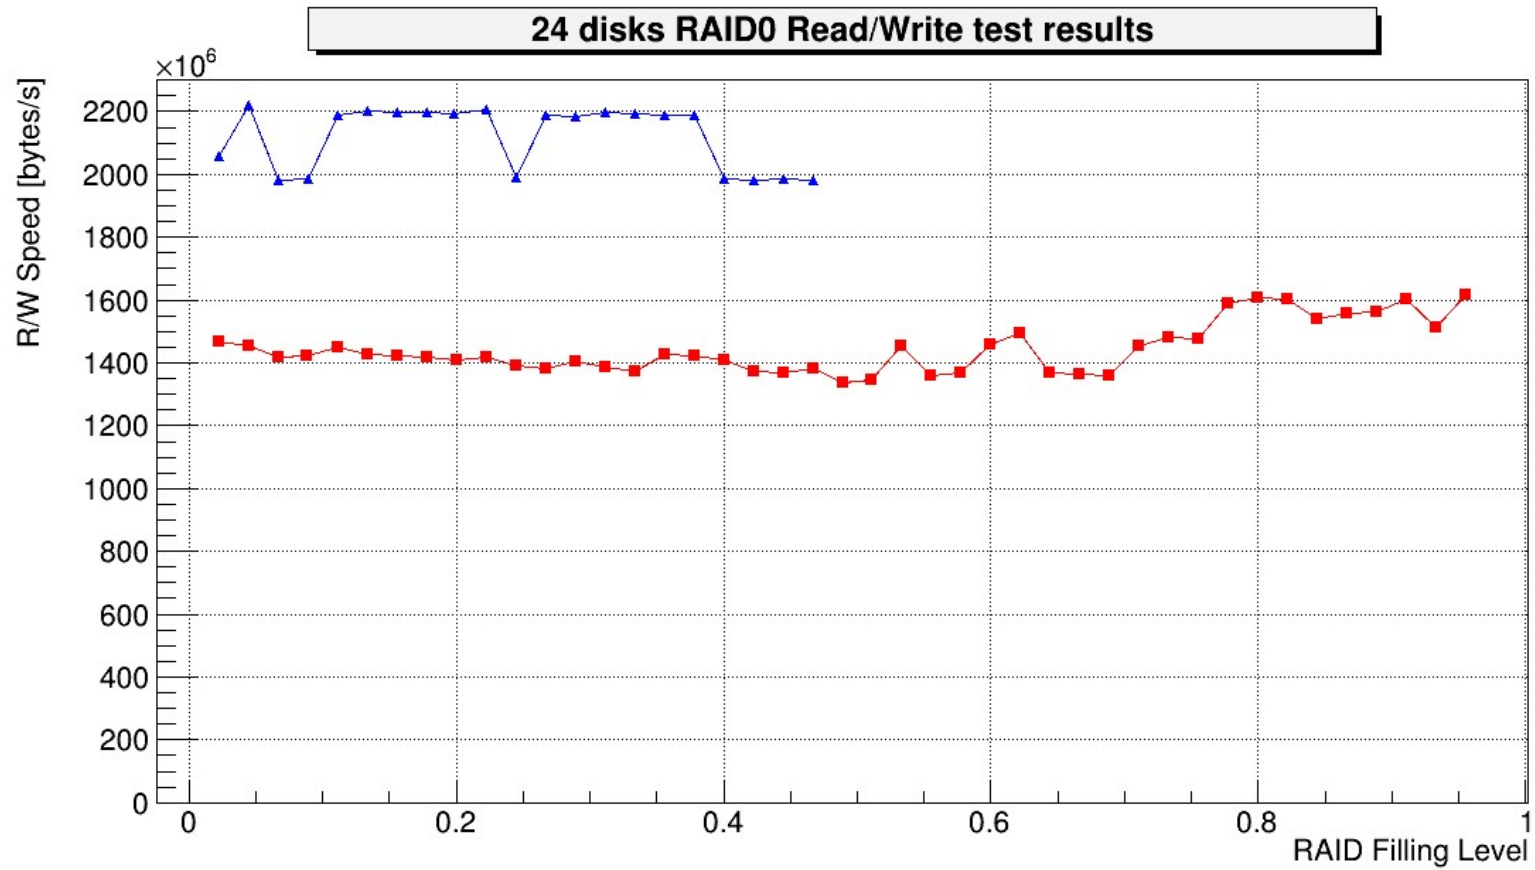
1.3 GB/s (*minimum*)

Read speed:

2.2 GB/s (*maximum*)

2.0 GB/s (*minimum*)


24 disks RAID0 performance results:



Conclusions:

1. The R/W speed of one RAID **does not depend on** activity of another RAID(s) of the same controller. So, for future high data rate experiments we can separate write and read processes;
2. The performance of the RAID with small number of the disks (<12HDDs) **depends on** its filling level (the R/W speed at the end of HDD is much smaller than at beginning of the disk);
3. Maximal write speed for the tested RAID controller **is saturating** on the level of **1.5 GB/s** (RAID5?);
4. Maximal read speed for the tested RAID controller is **2.2 GB/s**;
5. The R/W speed **almost does not depend on** filling level for the RAIDs with big number of the HDDs in array (>12 HDDs) and equal **1.5 GB/s** (mean write speed) and **2.0 GB/s** (mean read speed for 24 disks RAID);
6. We are suggesting to use **AOC-SAS3-9380-8E** RAID controller and **Supermicro SuperChassis 847E2C-R1K28JBOD** (supports up to **44** 3.5" SAS3/2 or SATA HDDs with 12Gbps throughput) to get at least 1GB/s R/W speed.

AOC-SAS3-9380-8E RAID controller



Supermicro SuperChassis 847E2C-R1K28JBOD

