19th International Symposium INFOTEH-JAHORINA, 18-20 March 2020

The influence of workloads and depth queue on the performance of SSD disk RAID 0 level array

Nikola Davidović
University of East Sarajevo,
Faculty of Electrical Engineering,
East Sarajevo,RS, BiH
nikola.davidovic@etf.ues.rs.ba

Slobodan Obradović
Information Tehnology School
Belgrade, Serbia
slobodan.obradovic@its.edu.rs

Borislav Đorđević, Valentina Timčenko
Mihajlo Pupin Institute, University of Belgrade
Belgrade, Serbia
borislav.djordjevic@pupin.rs, valentina.timcenko@pupin.rs

Abstract—**The RAID 0 level represents a group of N independent disks. The data is written consecutively in blocks to various hardware units. This procedure increases the total storage space. Since the capacity of SSD devices is relatively small, connecting them into a N disks RAID 0 array results in a device N times larger than the smallest device. RAID 0 allows multiple devices to be accessed simultaneously, that is, to obtain write and read speed operations in proportion to the number of SSDs that are forming the array (if the bus speed allows). However, the performance of the RAID 0 array, besides the impact of the number of used units, is also affected by other parameters. The paper analyzes the effect of RAID 0 array size (number of SSD units), average block size, and queue depth to the RAID 0 array performance.**

RAID 0, capacity, SSD disk, write speed, read speed, average workload data block, queue depth.

I. INTRODUCTION

The growing need for the increase of the exchanged amount of data has implicated the necessity for providing larger storage space, as well as for the development of faster access and easier management of data. The performance of the secondary memory is significantly dependent on the available capacity, data access speed and reliability while storing and keeping safe the data. Improvements in the performance of secondary memory can be achieved either by using the parallel disks or/and by the application of some novel technology.

The development of the electronic nonvolatile memories began in the middle XX century. For decades, (Hard Disk Drive) HDDs dominate over other secondary computer memory technology. The realization of the ideas related to the electronic memory that could jeopardize the dominance of the HDD was partially achieved with the rapid development of flash memory, and their use for data storage and transfer. Semiconductor memory storage SSD (Solid State Disk) have one of the biggest advantages related to their use - they have better read and write data operation speed.

The RAID (Redundant Array of Inexpensive Disks) represents a solution where the data is stored on the disk architecture that represents the combination of several physical disks into one logical unit. The goal is to enhance performances and increase the storage capacity[1] [2].

When compared to all other RAID levels, RAID 0 offers the highest degree of storage space used for storing data. This RAID level provides the best read and write operation performance, but it does not offer redundancy. Despite this disadvantage, thanks to its unmatched performance, RAID 0 is used in systems where data access speeds and storage space size play the key role.

II. TEST CONFIGURATION

A. Hardware configuration

The hardware for the test environment is specified in Table 1. Tests were carried out on Microsoft Windows 10 Education suite, with the OS on SSD on ADATA SSD devices (characteristics are given in Table 2). We used an integrated RAID controller on the motherboard ASUS H97M-E to configure RAID 0.

TABLE I. HARDWARE CONFIGURATION DESKTOP PC

Hardware	**Specifications**
Motherboard	ASUS H97M-E
RAM	16 GB, 2*DIMM DDR3 8GB 1866MHz Kingston HyperX Fury
CPU Model	CPU Intel(R) Core(TM) i7-4790 CPU@3.6GHz, 3601 Mhz, 4 Core(s), 8 Logical Processor(s)
Secondary memory device with OS	SSD ADATA, Premier Pro SP900, 2.5" SATA 6Gb/s
Operational system	Microsoft Windows 10 Education, 10.0.17134 Build 17134

978-1-7281-4775-8/20/$31.00 ©2020 IEEE

TABLE II. SSD ADATA SPECIFICATION

SSD	Specifications
Model	SSD ADATA, Premier Pro SP900, 2.5" SATA 6Gb/s ASP900SS-128GM
Capacity	128 GB
Read	up to 555 MB/s
Write	up to 535 MB/s

B. *ATTO Disk benchmark*

ATTO Disk Benchmark is a freeware software for Windows which helps the measurement of the storage system performance [6]. ATTO identifies the performance levels of the hard drives, solid state drives, RAID arrays, as well as the host connection to the attached storage. One of the advantages of this benchmark is the ability to control the process of write and read operation, while the drawback is the inability to test the random data access speed.

Some of the setting options over which ATTO Disk Benchmark can affect system performance or can isolate certain situations in practical work, are:

- File Size – specifies the test file length. Transfer lengths are from 64 KB to 32 GB
- Bypass write Cache – this option allows to bypass the drive write cache.
- Direct I/O – test drive is performed with no system buffering or caching.
- I/O Comparison – compares the input and output data to detect errors.
- Overlapped I/O – this option performs queued I/O testing. The factor that specifies the I/O overlapping is the queue depth.
- Neither – do not perform overlapped I/O or I/O comparisons.

III. TEST RESULTS

Certain restrictions were set in order to get as highest speed as possible for reading and writing during the test procedures.

The first configured limitation, is to force the use of the full capacity semiconductor disk for testing. The tests are using 128 GB of each semiconductor disks of the RAID 0. Thus, in the case for 2 and 3 semiconductor disks, an entry space is 256 and 384 GB, respectively.

Secondly, in all the tests the size of the stripe unit (SU) is configured to 16 KB, so the size of the full stripe for RAID 0 with 2 and 3 SSD are 32 KB and 48 KB, respectively.

Data caching feature could give wrong results so that tests results would not show the real performance of the SSD, but the performance of the cache. Because of this, when configuring, we have used the option to bypass the cache of the disks, as well as to generate the caching on the controller itself. In ATTO disk benchmark this is enabled by the Bypass Write Access and Direct I/O options.

The fourth limitation is the number of multiple transfer requests, which defines the maximum number of read/write commands that can be executed in one-time interval. The queue depth factor specifies the number of queue entries for overlapped I/O.

In the following test, we have used full capacity size NTFS partition, as the test file space was limited to 32 GB. The larger file was selected in order to get the better average values for large transfers, which can also be assumed as the sequential data access test.

Data given in the Table 3. contain values of read and write operation performances in MB/s for a single ADATA SSD and RAID 0 arrays with 2 and 3 SSDs, with and without I/O parallelism, and with queue depths of 1 or higher value (QD1, 2, 4, 8, 10, 16, 32, 64, 128 and 256). The testing is performed with workloads ranging from 512 B to 32 MB, with each consecutive workload being twice the size of the previous. First part of the table contains measurements for read (blue), while the second contains data for write (red).

The observed performances of a single SSD imply that read and write performances are increasing with the size of the workload, and the maximum values declared by the manufacturer are reached only with large workloads (2 MB and greater). For values ranging between 512 B and 4 KB, the obtained performances are far less than declared. With further increase of the workload, read and write performances increase as well. This is due to the fact that SSD is divided into 16 KB pages, causing every workload to be on a different page. Only workloads larger than 16 KB will contain two or more pages. The maximum performances are practically reached with workloads of 2 MB, while further increase of the block size has no significant effect on the performance.

Besides measured performances of read and write using a single ADATA SSD, the Table 3 values using RAID 0 with two drives (RAID 0 2) and RAID 0 with three drivers (RAID 0 3). Each row contains the Queue Depth (QD) parameter, with different workloads given in the table columns ranging from 0.5 KB to 32768 KB. Rows contain measured values for different queue depths 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 (QD). Measured reading and writing performances are given in absolute values, in MB/s, as illustrated in the Table 3. The diagrams given in Fig. 1. – Fig. 6. are based on these values.

Diagrams given in Fig. 1., Fig. 3. and Fig. 5. illustrate reading performance for a single SSD, RAID 0 array with two disks, and RAID 0 with three disks, for different workload sizes ranging from 0.5 KB to 64 MB, and for different values of QD (1, 2, 4, 8, 10, 16, 32, 64, 128 and 256). Likewise, diagrams given in Fig. 2., Fig. 4. and Fig. 6. illustrate measurements for writing performances.

TABLE III. MEASURED DATA – SINGLE SSD AND RAID 0 WITH 2 AND 3 ADATA SSD

VB SSD	MB/s	0.5 KB	1 KB	2 KB	4 KB	8 KB	16 KB	32 KB	64 KB	128 KB	256 KB	512 KB	1024 KB	2048 KB	4096 KB	8192 KB	16384 KB	32768 KB
ADATA single disk, read	QD 1	9	19	37	73	126	193	198	257	248	314	383	439	464	481	490	495	497
	QD 2	9	19	36	71	103	168	177	224	249	312	380	431	463	481	490	494	496
	QD 4	18	35	68	131	199	276	313	361	435	480	503	517	522	522	522	522	522
	QD 8	26	52	103	203	312	387	419	453	488	509	521	524	523	524	523	523	522
	QD 10	33	65	127	231	321	405	435	466	488	494	524	523	523	523	523	523	522
	QD 16	34	68	135	267	358	419	462	488	504	517	523	524	523	523	521	522	522
	QD 32	34	68	137	272	364	439	483	503	517	523	515	524	523	523	523	522	522
	QD 64	34	69	131	270	364	438	483	503	522	523	522	523	523	523	523	522	522
	QD 128	35	69	139	276	363	439	484	503	517	523	518	524	522	519	518	518	519
	QD 256	35	69	138	276	365	439	481	499	517	516	522	524	523	523	523	522	522
ADATA RAID 0 2 SSD read	QD 1	32	63	122	221	403	738	878	875	876	876	875	879	879	881	881	874	875
	QD 2	26	54	95	182	371	691	875	877	878	878	877	881	881	881	881	876	871
	QD 4	63	125	245	489	875	879	878	880	878	873	879	879	879	881	879	879	877
	QD 8	61	127	248	511	877	879	878	880	878	875	879	879	879	881	878	878	876
	QD 10	62	125	242	471	877	881	878	877	878	876	881	881	881	881	878	877	876
	QD 16	67	129	250	463	879	881	878	880	878	873	879	879	879	879	878	878	877
	QD 32	63	126	257	478	879	881	881	880	880	871	879	881	879	877	878	878	877
	QD 64	67	125	244	526	881	881	881	880	878	874	879	879	877	875	878	878	877
	QD 128	64	127	246	484	881	881	881	880	878	876	879	875	877	878	878	878	877
	QD 256	64	127	250	503	881	881	878	877	878	870	875	875	880	877	879	878	876
ADATA RAID 0 3 SSD read	QD 1	31	65	124	222	413	660	1146	1172	1179	1170	1179	1185	1193	1198	1140	1190	1176
	QD 2	23	49	104	193	366	637	1133	1169	1170	1173	1179	1168	1193	1194	1185	1191	1175
	QD 4	66	130	254	481	992	1176	1169	1182	1198	1196	1202	1202	1200	1203	1190	1192	1196
	QD 8	67	134	266	525	936	1180	1180	1194	1195	1208	1205	1202	1210	1201	1153	1199	1196
	QD 10	68	134	259	491	984	1176	1188	1197	1204	1199	1199	1208	1172	1198	1203	1198	1194
	QD 16	68	135	270	518	1055	1188	1188	1200	1207	1199	1199	1199	1158	1199	1205	1198	1195
	QD 32	68	135	261	482	984	1195	1200	1200	1210	1196	1190	1199	1170	1201	1207	1200	1203
	QD 64	67	134	267	508	842	1203	1205	1197	1207	1205	1196	1205	1212	1199	1203	1199	1202
	QD 128	67	135	248	531	963	1200	1122	1204	1198	1205	1190	1188	1200	1201	1205	1202	1202
	QD 256	69	135	262	496	1035	1208	1200	1201	1201	1205	1199	1202	1174	1203	1203	1201	1202
ADATA single disk, write	QD 1	11	21	38	98	164	249	297	320	324	385	426	441	458	468	468	476	478
	QD 2	11	19	36	93	143	207	256	272	309	369	408	433	452	459	468	472	472
	QD 4	15	30	57	202	327	412	452	478	487	500	500	498	499	499	498	497	497
	QD 8	25	48	88	285	390	451	472	483	487	498	496	499	496	493	492	491	491
	QD 10	26	56	108	316	416	449	482	487	489	495	495	495	496	494	494	495	492
	QD 16	34	68	118	271	429	464	486	491	493	500	498	498	498	498	498	498	497
	QD 32	37	70	127	310	431	464	485	491	494	495	499	500	499	499	498	498	498
	QD 64	37	71	127	283	430	463	482	487	489	496	496	496	497	498	496	496	495
	QD 128	36	66	127	256	429	461	483	493	493	490	494	496	496	483	485	495	491
	QD 256	37	71	126	349	426	465	482	489	494	485	493	495	496	497	495	496	495
ADATA RAID 0 2 SSD write	QD 1	12	22	38	101	168	251	459	614	730	770	862	897	913	930	931	937	939
	QD 2	12	22	40	100	168	251	455	613	737	768	866	905	915	926	931	936	938
	QD 4	17	33	63	226	424	729	808	882	926	943	955	962	964	971	973	970	972
	QD 8	36	58	111	390	727	758	898	935	950	950	955	964	971	978	970	968	971
	QD 10	38	58	130	376	748	717	902	939	954	955	966	973	976	978	971	971	969
	QD 16	57	102	182	517	816	883	904	937	959	955	971	971	976	976	973	967	971
	QD 32	60	118	244	477	850	891	904	938	959	964	978	978	976	971	973	968	969
	QD 64	64	122	241	520	902	891	904	946	973	973	978	976	972	972	972	973	971
	QD 128	64	130	256	484	887	891	906	951	964	964	969	967	968	974	972	966	969
	QD 256	67	134	266	523	861	891	915	952	964	973	968	965	968	967	972	969	969
ADATA RAID 0 3 SSD write	QD 1	11	20	39	98	167	251	451	622	850	935	1032	1087	1134	1148	1151	1166	1173
	QD 2	11	21	39	100	166	250	451	619	852	919	1030	1094	1115	1159	1146	1170	1146
	QD 4	17	33	60	223	425	660	1039	1122	1192	1185	1199	1199	1200	1208	1214	1239	1239
	QD 8	38	61	120	381	689	1012	1013	1179	1198	1196	1190	1217	1215	1128	1244	1240	1238
	QD 10	42	67	128	365	822	1055	1096	1191	1192	1199	1193	1211	1215	1238	1221	1234	1240
	QD 16	54	100	166	520	914	1035	1166	1166	1195	1199	1188	1208	1235	1236	1227	1231	1230
	QD 32	60	115	222	469	793	1145	1156	1194	1201	1193	1188	1214	1227	1240	1223	1232	1228
	QD 64	66	124	228	506	908	1161	1161	1198	1198	1199	1244	1235	1240	1238	1240	1243	1199
	QD 128	65	128	252	512	914	1153	1153	1191	1235	1214	1229	1224	1230	1231	1223	1237	1192
	QD 256	65	131	258	480	818	1135	1174	1191	1225	1225	1214	1232	1232	1238	1238	1225	1236

Each of these diagrams illustrates performances for a single SSD, RAID 0 with two SSDs, and RAID 0 with three SSDs, with a queue depth of 1 (QD 1, marked with triangle symbols: ▲), and with queue depth of 2 (QD 2, marked with rhombus symbols: ♦).

The diagrams given in Fig. 1. and Fig. 2. illustrate measured reading and writing performances for a single SSD for different QD values and different workload sizes. The performances are increasing with increasing the workload size. The trend is noticeable for smaller workloads (less than 16 KB) for all values of QD. Maximum performances are practically reached when working with large workloads (2MB and larger). Further increase of the workload has no significant effect on the performances.

By analyzing the performances for QD from 2 to 256, a conclusion can be made that for queue depth of 2 (marked with ♦) there is practically no change in performances using a single SSD. On the contrary, measured performances are lower for all workloads.

For queue depths greater than 2, increasing the queue depth consequently increases the performances. Maximum values are reached with smaller workloads than without parallelism (QD 1). This is a consequence of the fact that the performances are increasing up to the maximum declared values.

Measured reading performances indicate that there is practically no point in increasing the queue depth over QD 16. The diagrams for QD 16, 32, 64, 128 and 256 are overlapping. When writing, even with QD 4 diagrams are almost matched with larger queue depths. With QD 16 in both cases maximum performances are reached. Maximum reading performance of over 520 MB/s is reached with block size of 128 KB. Maximum writing performance of 500 MB/s is reached with workloads of 32 KB. Maximum writing and reading performances are higher by 4% and 5% is QD greater than 2 is used, as compared to QD 1 and QD 2.

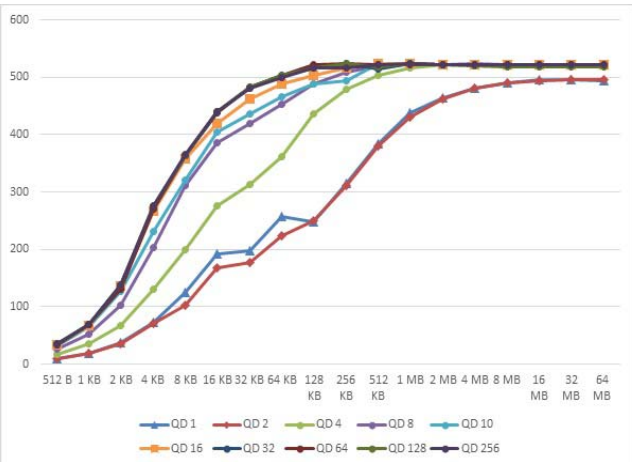

Fig. 1 *Writing performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for single SSD ADATA.*

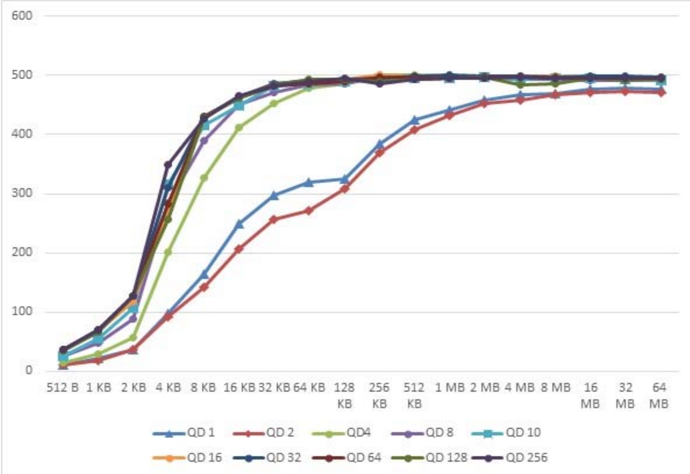

Fig. 2 *Reading performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for single SSD ADATA*

The diagrams given in Fig. 3. and Fig. 4. illustrate measure values of reading and writing performance for RAID 0 with two SSDs for different values of QD and workload sizes. Performances increase with increasing the workload size.

Performance with queue depth of 2 (marked with ♦) are practically the same as with queue depth 1 (marked with ▲). Maximum reading performance is reached for workloads of 32 KB, which is 64 times less than with a single drive. Writing performances reach maximum at workloads of 1MB and larger.

By analyzing the performances using queue depths from 4 to 256, we can conclude that for QD greater than 2, diagrams for both read and write practically overlap. The performance gain is very noticeable for smaller workloads, until the maximum values are reached. The maximum values for read are reached with very small workloads, starting with 8 KB. There is practically no gain in reading performance when increasing the queue depth over 4, because the diagrams for QD from 4 to 256 overlap.

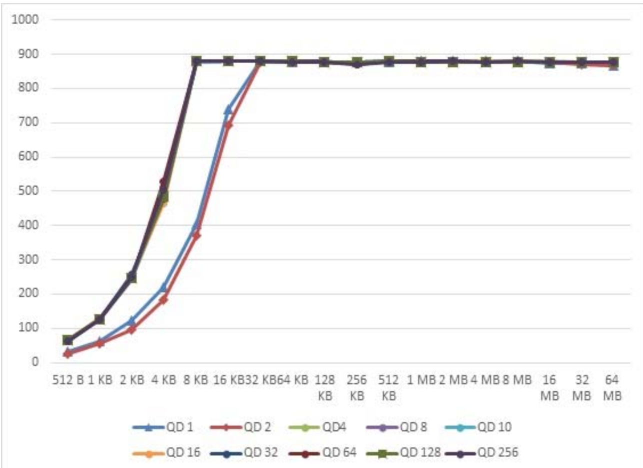

Fig. 3 Reading performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for RAID 0 with 2 SSDs

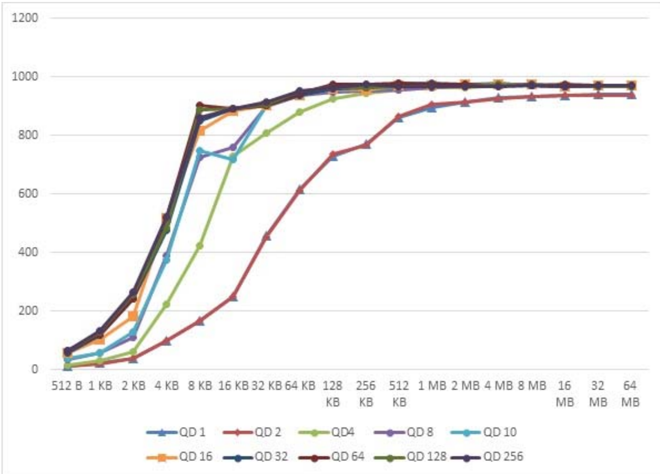

Fig. 4 *Writing performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for RAID 0 with 2 SSDs*

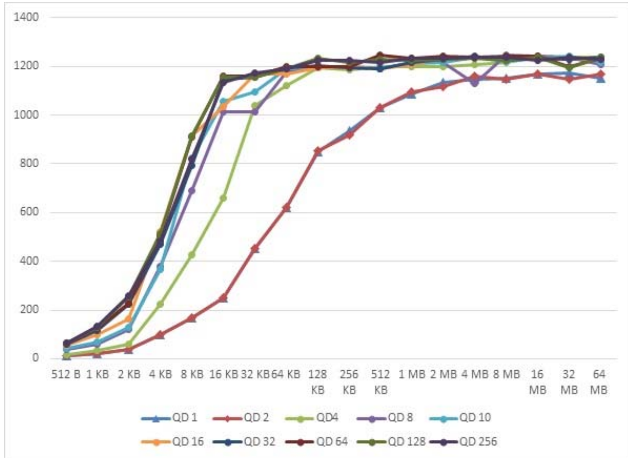

Fig. 6 Writing performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for RAID 0 with 3 SSDs

For writing, maximum performance is reached with QD 16 as the diagrams overlap with queue depths of 16 and greater. Maximum reading and writing performances are reached at workloads of 8 KB. Maximum reading performance is around 880MB/s, which is 1.7 times greater than with a single SSD, which is significantly lower than twice the performance of a single SSD (in theory, 1040 MB/s was expected for RAID 0 2). Maximum writing performance is practically reached with workloads of 32 KB, and are around 970 MB/s. This is 1.95 greater than with a single SSD (expected theoretical value).

The diagrams in Fig. 5. and Fig. 6. illustrate measured values of reading and writing performances for RAID 0 with three SSDs for different values of queue depths and workloads. The performances increase by increasing the workload. Performances using a queue depth of 2 (marked with ♦) are practically the same as for queue depth of 1 (marked with ▲).

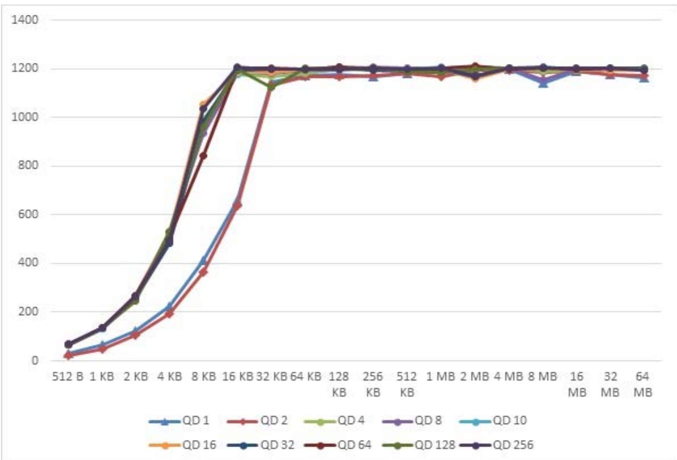

Fig. 5 Reading performances for QD 1, 2, 4, 8, 10, 16, 32, 64, 128 and 256 for RAID 0 with 3 SSDs

The maximum reading performance is practically reached at workload sizes of 32 KB. The maximum reached reading performances with QD 1 and 2 are insignificantly lower than with QD 4 or higher, which is 64 times smaller workload that when using a single SSD, and the same as for RAID 0 with two SSDs. Maximum writing performances for QD 1 and QD 2 are reached at a very large workloads, an order of magnitude of MB and greater. The maximum writing performance for QD 1 and QD 2 are 7% lesser than for higher values of QD (4, 8, etc.).

By analyzing the diagrams given in Fig. 5. and Fig. 6. we can conclude that for queue depths from 4 to 256, diagrams for both read and write practically overlap. The reading performances reach maximum at very small workloads, starting with 16 KB. There is practically no gain in reading performances when increasing the queue depth over 4. For writing, maximum performance is reached with QD 16 as the diagrams overlap with queue depths of 16 and greater. Therefore, the maximum values can be reached by selecting the QD of 16, for both read and write. The maximum reading performance of 1200MB/s is reached with workloads of 8 KB. The value of 1200MB/s is 2.3 greater than when using a single drive, which is significantly lower than triple the performance of a single SSD (in theory, for RAID 0 3 expected reading performance is 1560 MB/s). Maximum writing performance of 1240MB/s is practically reached with workloads of 16 KB, which is almost 2.5 times the value for a single SSD (in theory, for RAID 0 3 expected reading performance is 1480 MB/s).

IV. CONCLUSION

The performance diagrams with QD 1 and QD 2 practically overlap. In some cases, performances with QD 2 are lower than with QD 1. Therefore, there is no point in setting QD to value 2 neither for a single SSD, neither for RAID 0 array of 2 nor 3 SSDs.

Reading and writing performances indicate significant dependency on the workload size. By increasing the workload size, the performances increase as well, and are asymptotically approaching the maximum values. For single SSD, the gain in performance by increasing the workload size is greater for write than read. The maximum performances for write are reached at workloads of two or more times smaller than for read.

On the contrary, with RAID 0 with 2 or 3 SSDs, reading performances are increasing faster than writing performances, and are reached at smaller workloads compared to write.

Reading and writing performances indicate significant dependency on the queue depth for smaller workloads, between 512 B and 16 KB.

For workloads larger than 16 KB reading and writing performances for QD 16-256 practically overlap, they are greater compared to QD 1 when there is no I/O operations parallelism and are asymptotically approaching the maximum values.

By increasing the queue depth, maximum performances are reached at smaller workloads. Optimal choice of QD is not the same for read and write, however for the analyzed devices there were no effects in increasing the performances with QD higher than 16.

With RAID 02, maximum reading performance is 1.7 compared to a single SSD. Increase in writing performances with RAID 02 is almost 2 times higher compared to a single SSD, which is significantly better result than increase in writing performance for READ 0 3 of less than 2.5 times compared to a single SSD.

ACKNOWLEDGMENT

The work presented in this paper has partially been funded by the Ministry of Education, Science and Technological Development of the Republic of Serbia.

REFERENCES

[1] Patterson, David; Gibson, Garth A.; Katz, Randy, “A Case for Redundant Arrays of Inexpensive Disks (RAID)”, www.eecs.berkeley.edu/Pubs/TechRpts/1987/CSD-87-391.pdf, december 2018.

[2] S. Obradović, B. Milosevic, “Comparative measurements of some performances of hard and ssd disks connected on the raid 0 level”, International scientific conference Unitech’2014, 20-21 nov. Gabrovo, ISSN 1313-230Xvol.2, pp.II-333-II-338

[3] V. Timčenko, B. Đorđević, S. Obradović, N. Čorni, “Impact of disk cache buffers on the performance of solid-state drives”, INFOTEH-JAHORINA, Vol. 12, pp. 1002-1005, 20-23 March 2013.

[4] V. Timčenko, B. Đorđević, S. Obradović, N. Čorni, “Impact of Internal Controler on SSD Disk Performance”, XIX YUINFO, Kopaonik, Mart 2013, pp. 433-436, 3-6 Mart 2013.

[5] https://www.adata.com/upload/downloadfile/Datasheet_SP900-EN-20150213.pdf

[6] https://www.atto.com/disk-benchmark/