FreeNAS 9.2 vs 9.3 v3

7th of JAN 2015, by J.Bronka [Kobetsu]

Table of Contents

Introduction

Reason to create this document and compare more detailed FN9.2 and FN9.3 was great mistake I did when without tests replace FN9.2.1.5 with new version FN9.2.1.8 on production server.

All works but number of issues about performance degradation forced me to go back to FN9.2.1.5.

With FN9.2.1.8 the same problems were with experimental kernel ISCSI target as well as traditional.

Benefits of FN9.3 as full VAAI implementation are interesting to consider migration… but… performance need to be checked.

I. Configurations

I.1 FreeNas 9.2 Test Configuration

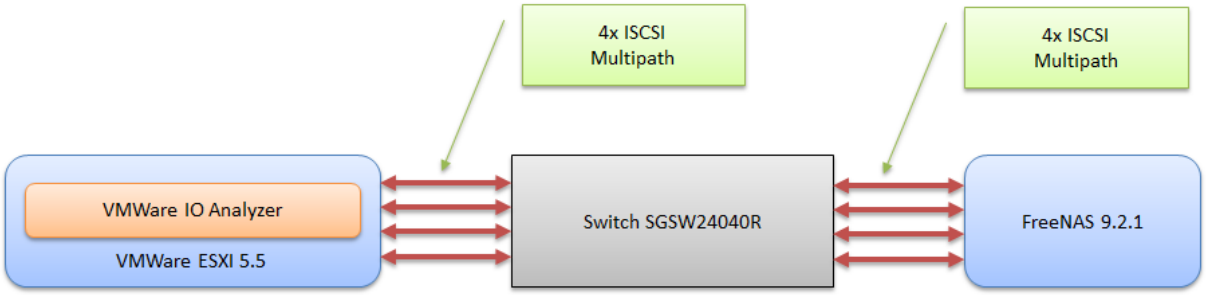

VMWare ESXI:

- Host: VMWare ESXI 5.5
- Test VM: VMWare Appliance IO Analyzer
- 1x XEON E5-2667 v3 / 3.2 GHz
- SuperMicro X10DRi
- 64GB RAM
- 2x NIC I350 / 1GBit (JUMBO 9000)

FreeNas:

- FreeNAS-9.2.1.5-RELEASE-x64
- I7 4820K / 3.7Ghz
- 64GB RAM
- 2x NIC Intel 82574L / 1GBit (JUMBO 9000)
- RAID LSI 9260-16i (There are no support for JBOD but each of disk is working as RAID0)
- RAID 10 (14x Western Digital Red 2TB [WD20EFRX])
- L2ARC (SSD)
- Autotune turned on

This production server I'm using – tested in this document as real working instance. My objective is to migrate this server to 9.3, just before I need to be sure it will work better after migration.

I.2 FreeNas 9.3 Test Configuration

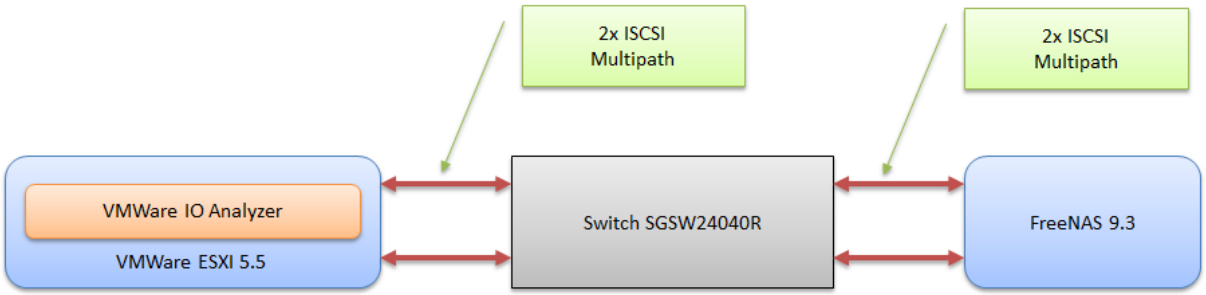

VMWare ESXI:

- Host: VMWare ESXI 5.5
- Test VM: VMWare Appliance IO Analyzer

- 1x XEON E5-2667 v3 / 3.2 GHz
- SuperMicro X10DRi
- 64GB RAM
- 2x NIC I350 / 1GBit (JUMBO 9000)

FreeNas:

- FreeNAS-9.3-STABLE-201412240734
- I7 3770K / 3.5Ghz
- 32GB RAM
- 2x NIC Intel 82574L / 1GBit (JUMBO 9000)
- SATA onboard controller
- Standard installation with autotune turned on

I.3 FreeNas 9.2 (#2) Test Configuration

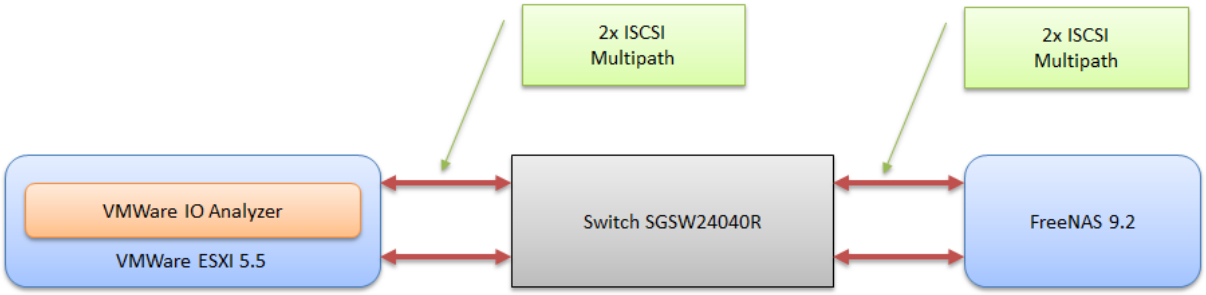

VMWare ESXI:

- Host: VMWare ESXI 5.5
- Test VM: VMWare Appliance IO Analyzer
- 1x XEON E5-2667 v3 / 3.2 GHz
- SuperMicro X10DRi
- 64GB RAM
- 2x NIC I350 / 1GBit (JUMBO 9000)

FreeNas:

- FreeNAS-9.2.1.5-RELEASE-x64
- I7 3770K / 3.5Ghz
- 32GB RAM
- 2x NIC Intel 82574L / 1GBit (JUMBO 9000)
- SATA onboard controller
- Default installation with autotune turned on

Real comparation can be done using test case I.2 and I.3 because both refer to the same hardware only FN is different.

I.4 FreeNas 9.3 (#2) Test Configuration

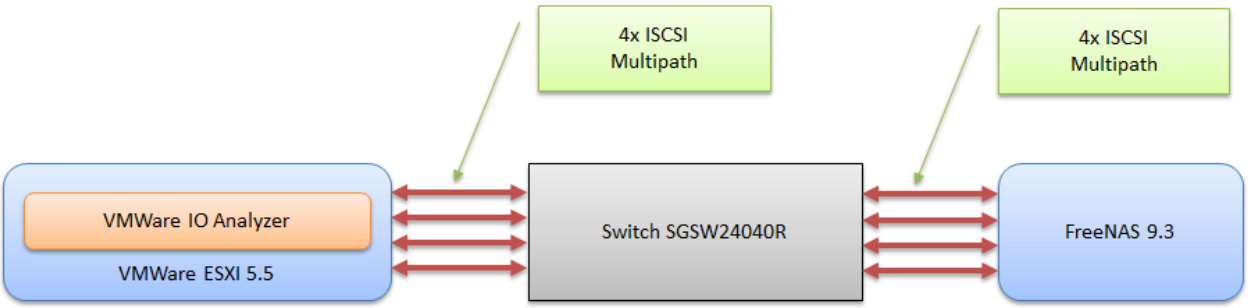

VMWare ESXI:

- Host: VMWare ESXI 5.5
- Test VM: VMWare Appliance IO Analyzer
- 1x XEON E5-2667 v3 / 3.2 GHz
- SuperMicro X10DRi
- 64GB RAM
- 2x NIC I350 / 1GBit (JUMBO 9000)

FreeNas:

- FreeNAS-9.3-STABLE-201412240734
- I7 4820K / 3.7Ghz
- 64GB RAM
- 2x NIC Intel 82574L / 1GBit (JUMBO 9000)
- RAID LSI 9260-16i (There are no support for JBOD but each of disk is working as RAID0)
- RAID 10 (14x Western Digital Red 2TB [WD20EFRX])
- L2ARC (SSD)
- Autotune turned on

This is my production (previously described in I.1) migrated to FN9.3

I.5 Disks

The following disks were tested in several configurations:

- Disk Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX)
- Disk Western Digital Red 2TB (WD20EFRX)

II. Test cases

The following cases were tested (regarding disks configuration):

- **FreeNAS 9.2**
  - **Case FN9.2#1**
    - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
    - Lun 500GB over RAID 10 (14x Western Digital Red 2TB [WD20EFRX])
    - ISCSI with the same settings in initiator/target configuration:

        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression
  - **Case FN9.2#2**
      - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
      - Lun 500GB over RAID 10 (14x Western Digital Red 2TB [WD20EFRX])
      - ISCSI with the same settings in initiator/target configuration:
        - FirstBurstLength = 1MB
        - MaxBurstLength = 1MB
        - MaxRecvDataSeqLength = 1MB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression
  - **Case FN9.2#3**
      - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
      - Lun 500GB over RAID 10 (14x Western Digital Red 2TB [WD20EFRX])
      - ISCSI with the same settings in initiator/target configuration:
        - FirstBurstLength = 64KB (default for FN9.2)
        - MaxBurstLength = 256KB (default for FN9.2)
        - MaxRecvDataSeqLength = 256KB (default for FN9.2)
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression

- **FreeNAS 9.3**
  - **Case FN9.3#1**
      - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
      - Lun 500GB over single HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX)
      - Default configuration of ISCSI in FreeNas
      - Compression LZ4
      - Initiator ISCSI settings:
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB

        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
- **Case FN9.3#2**
    - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
    - Lun 500GB over RAID1 (2x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
    - Default configuration of ISCSI in FreeNas
    - Compression LZ4
    - Initiator ISCSI settings:
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
- **Case FN9.3#3**
    - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
    - Lun 500GB over RAID10 (4x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
    - Default configuration of ISCSI in FreeNas
    - Compression LZ4
    - Initiator ISCSI settings:
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
- **Case FN9.3#4**
    - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
    - Lun 500GB over RAID1 (2x Disk Western Digital Red 2TB (WD20EFRX))
    - Default configuration of ISCSI in FreeNas
    - Compression LZ4
    - Initiator ISCSI settings:
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES

        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
    - **Case FN9.3#5**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over single Disk Western Digital Red 2TB (WD20EFRX)
        - Default configuration of ISCSI in FreeNas
        - Compression LZ4
        - Initiator ISCSI settings:
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
            - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
    - **Case FN9.3#6**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over RAID10 (6x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
        - Default configuration of ISCSI in FreeNas
        - Compression LZ4
        - Initiator ISCSI settings:
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
            - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1

- **FreeNAS 9.2#2**
    - **Case FN9.2#2#1**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over single Disk Western Digital Red 2TB (WD20EFRX)
        - ISCSI setting (rest, not listed = default)
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
        - No compression
    - **Case FN9.2#2#2**

        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over RAID1 (2x Disk Western Digital Red 2TB (WD20EFRX))
        - ISCSI setting (rest, not listed = default)
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
        - No compression
    - **Case FN9.2#2#3**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over single HDD Western Digital 1TB Caviar Black 7200 64MB ISCSI setting (rest, not listed = default)
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
        - No compression
    - **Case FN9.2#2#4**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over RAID10 (2x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
        - ISCSI setting (rest, not listed = default)
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB
            - MaxRecvDataSeqLength = 128KB
            - InitR2T = NO
            - ImmedaiteData = YES
            - DelayedACK = YES
        - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
        - No compression
    - **Case FN9.2#2#5**
        - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
        - Lun 500GB over RAID10 (4x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
        - ISCSI setting (rest, not listed = default)
            - FirstBurstLength = 256KB
            - MaxBurstLength = 256KB

        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression
  - **Case FN9.2#2#6**
      - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
      - Lun 500GB over RAID10 (6x HDD Western Digital 1TB Caviar Black 7200 64MB SATAIII (WD1002FAEX))
      - ISCSI setting (rest, not listed = default)
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression

- **FreeNAS 9.3#2**
  - **Case FN9.3#2#1**
      - Virtual Disk 128GB, eager zeroed, filled with random values before test (dd if=/dev/random ….)
      - Lun 500GB over RAID 10 (14x Western Digital Red 2TB [WD20EFRX]) + L2ARC 120GB SSD
      - ISCSI with the same settings in initiator/target configuration:
        - FirstBurstLength = 256KB
        - MaxBurstLength = 256KB
        - MaxRecvDataSeqLength = 128KB
        - InitR2T = NO
        - ImmedaiteData = YES
        - DelayedACK = YES
      - Additionally initiator has default VMware ESXi value of MaxOutstandingR2T=1
      - No compression

III. Tests methodology

1. Duration of every test = 120s.
2. Every test executed 5 times to calculate average value
3. Each test case was examined against 12 tests:
   - 512K_RW_RANDOM

     This test will generate 512k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 256K_RW_RANDOM

     This test will generate 256k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 128K_RW_RANDOM

     This test will generate 128k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 64K_RW_RANDOM

     This test will generate 64k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 32K_RW_RANDOM

     This test will generate 32k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 16K_RW_RANDOM

     This test will generate 16k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 8K_RW_RANDOM

     This test will generate 8k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - 4K_RW_RANDOM

     This test will generate 4k random IO with 50% probability of Read IO. Number of Outstanding IO is set to 16.
   - MAX_READ_THOUGHPUT

     This test will generate 512k sequential IO with 100% probability of Read IO. Number of Outstanding IO is set to 32. By sequentially reading blocks of large size we intend to mimic workloads similar to streaming video. This configuration reports the maximum throughput for your system.
   - MAX_WRITE_THOUGHPUT

     This test will generate 512k sequential IO with 0% probability of Read IO. Number of Outstanding IO is set to 32. By sequentially writing blocks of large size we intend to mimic workloads similar to video surveillance, video uploading and file backups. This configuration reports the maximum write throughput for your system.
   - MAX_READ_IOPS

     This test will generate 512b sequential IO with 100% probability of Read IO. Number of Outstanding IO is set to 32. By sequentially reading blocks of small size we intend to stress the storage controller cache. This configuration reports the maximum IOPS for your system.
   - MAX_WRITE_IOPS

     This test will generate 512b sequential IO with 0% probability of Read IO. Number of Outstanding IO is set to 32. By sequentially writing blocks of small size we intend to stress the backend disks or the write cache in large storage boxes. This configuration reports the maximum write IOPS for your system.
4. Two identical tests (listed in point 3 above) were never started one after the other
5. No other load on FreeNas as well as on ESXI

6. Tested virtual disk size = 128GB
7. Tested virtual disk type eager zeroed
8. Tested virtual disk fully filled with random data

IV. Results

IV.1 IOPS vs Block Size

Registered data:

	512K	256K	128K	64K	32K	16K	8K	4K
	IOPS							
FN9.2#1	41,208	59,348	81,796	95,454	109,434	122,83	142,414	157,824
FN9.2#2	22,942	31,662	41,61	51,87	56,616	59,59	70,686	68,16
FN9.2#3	23,726	34,212	49,014	59,504	66,364	71,48	82,864	81,396
FN9.3#1	24,452	38,402	58,322	93,318	118,182	125,69	152,54	165,964
FN9.3#2	41,608	67,122	108,996	175,216	218,724	263,236	287,642	294,306
FN9.3#3	73,964	125,406	211,136	334,732	408,312	474,752	531,572	560,538
FN9.3#4	36,072	57,684	92,366	155,692	198,624	231,604	278,174	289,062
FN9.3#5	21,304	32,152	47,116	79,754	101,022	116,186	136,068	146,792
FN9.3#6	103,15	182,954	315,956	443,428	577,668	652,654	749,448	806,594
FN9.2#2#1	19,968	33,984	37,728	52,1	62,674	66,752	70,358	68,072
FN9.2#2#2	28,88	38,418	43,126	51,006	59,622	62,176	70,044	69,716
FN9.2#2#3	33,746	46,34	51,818	77,525	92,426	97,12	103,088	99,466
FN9.2#2#4	28,122	39,034	52,38	58,028	66,046	71,994	82,534	82,264
FN9.2#2#5	34,288	48,43	51,842	67,43	72,782	74,754	84,388	81,48
FN9.2#2#6	46,564	63,272	74,136	93,902	100,126	104,778	117,708	117,602
FN9.3#2#1	137,49	254,078	446,802	592,686	780,694	931,912	1099,52	1235,474

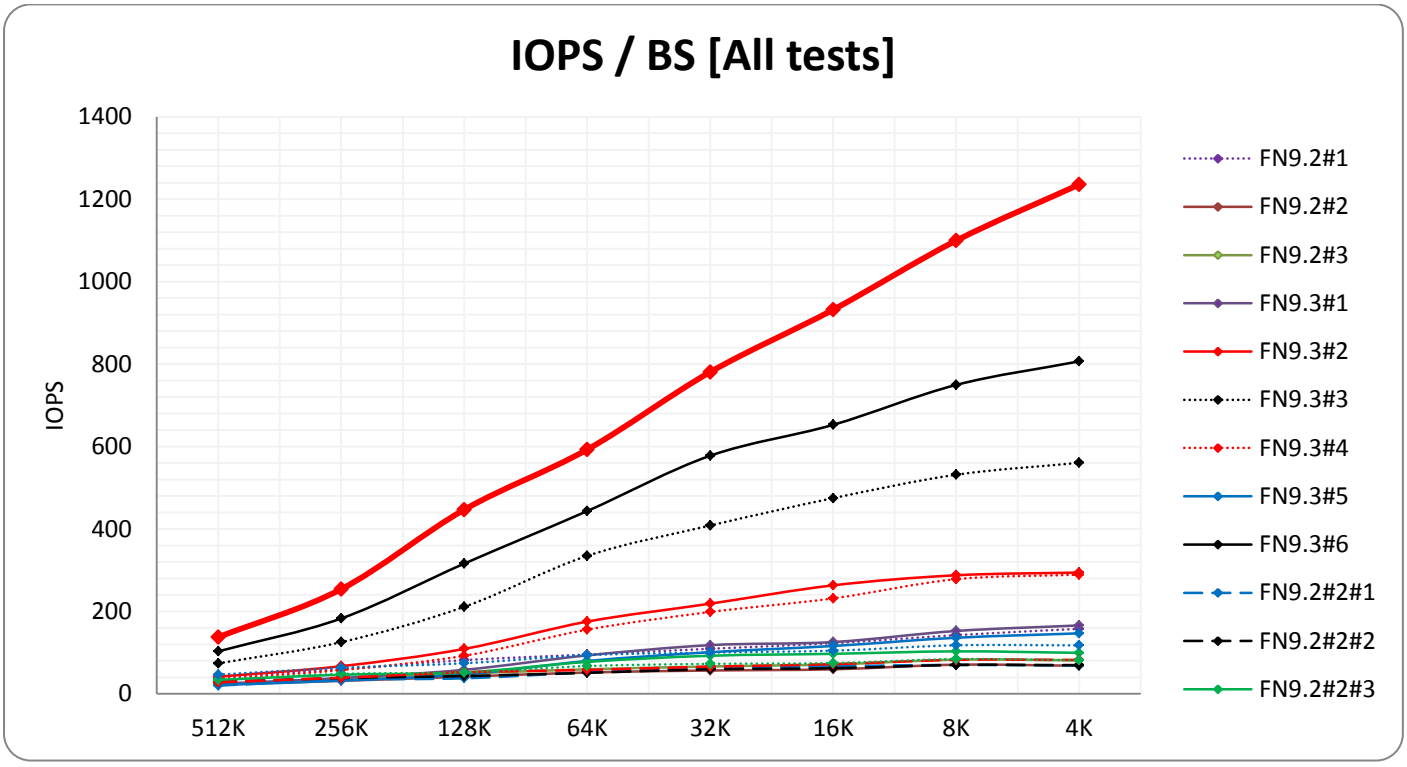

Figure 1: All tests

IV.1.1 General perspective

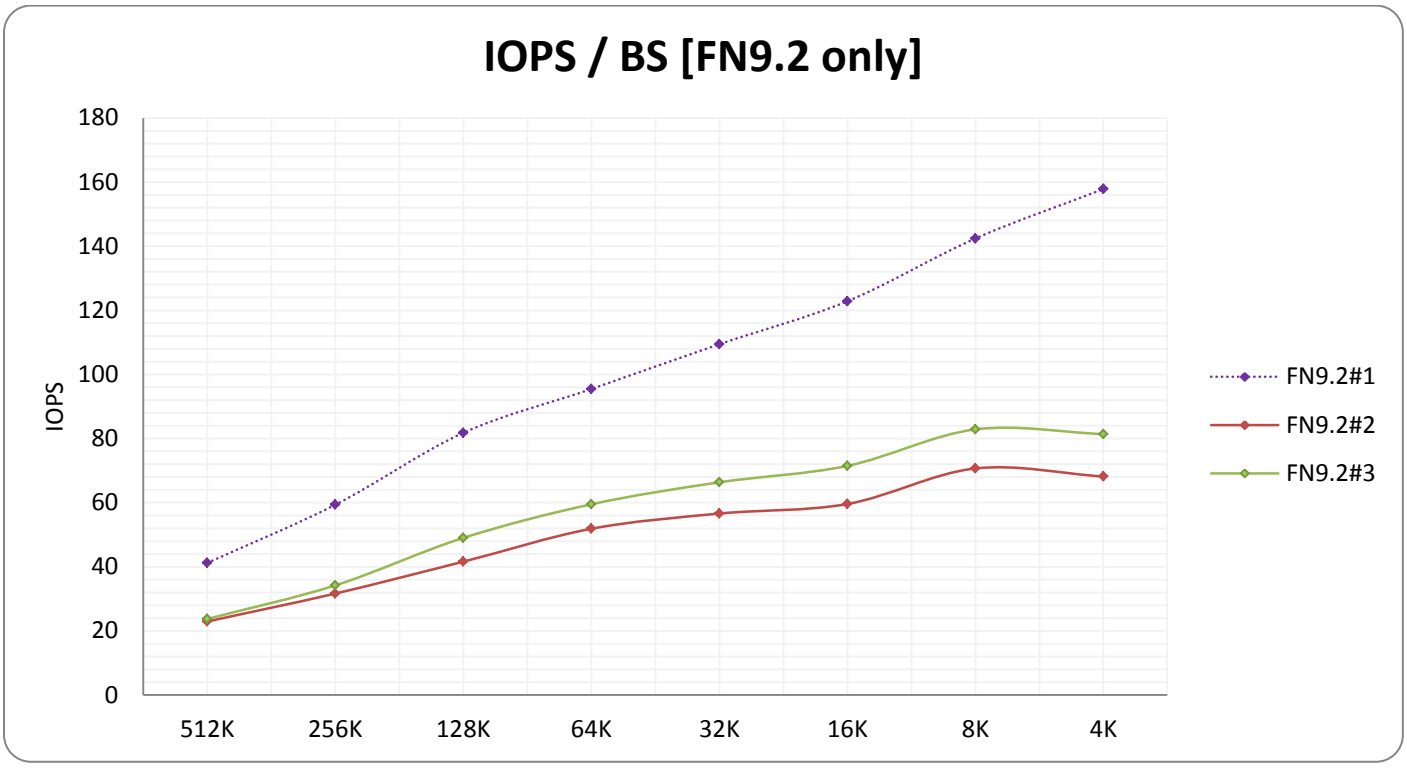

Figure 2: FN9.2 (production)

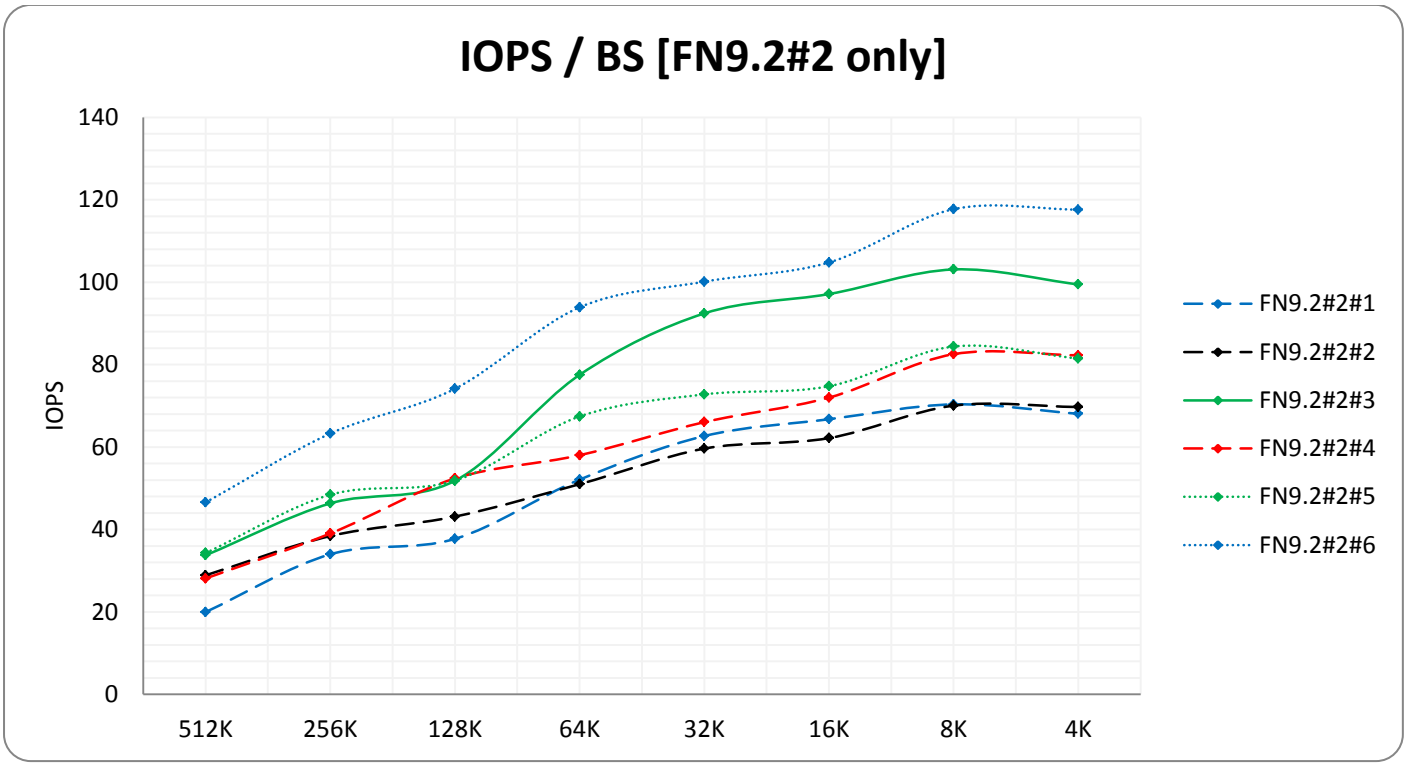

Figure 3: FN9.2#2 (test configuration)

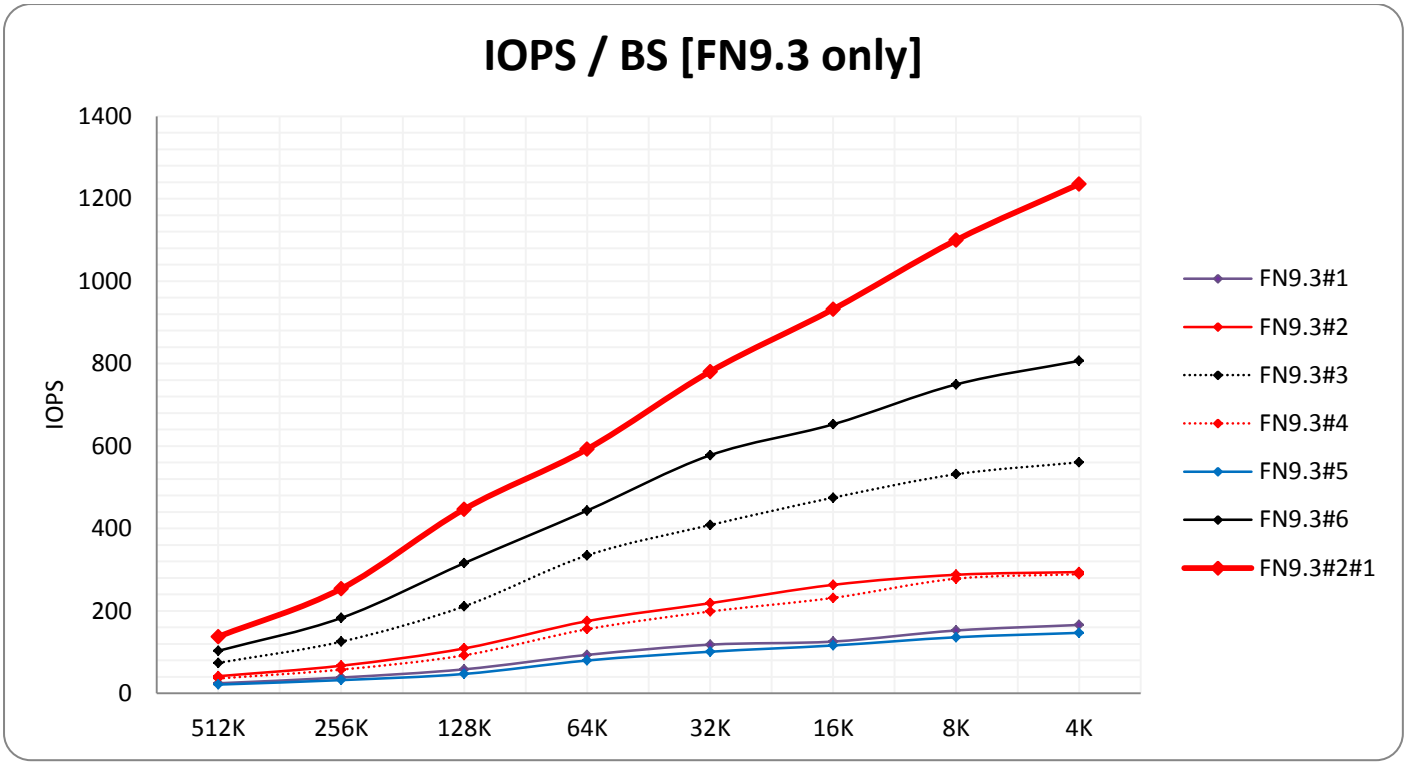

Figure 4: FN9.3 (test configuration)

IV.1.2 Disk type perspective – WD RED

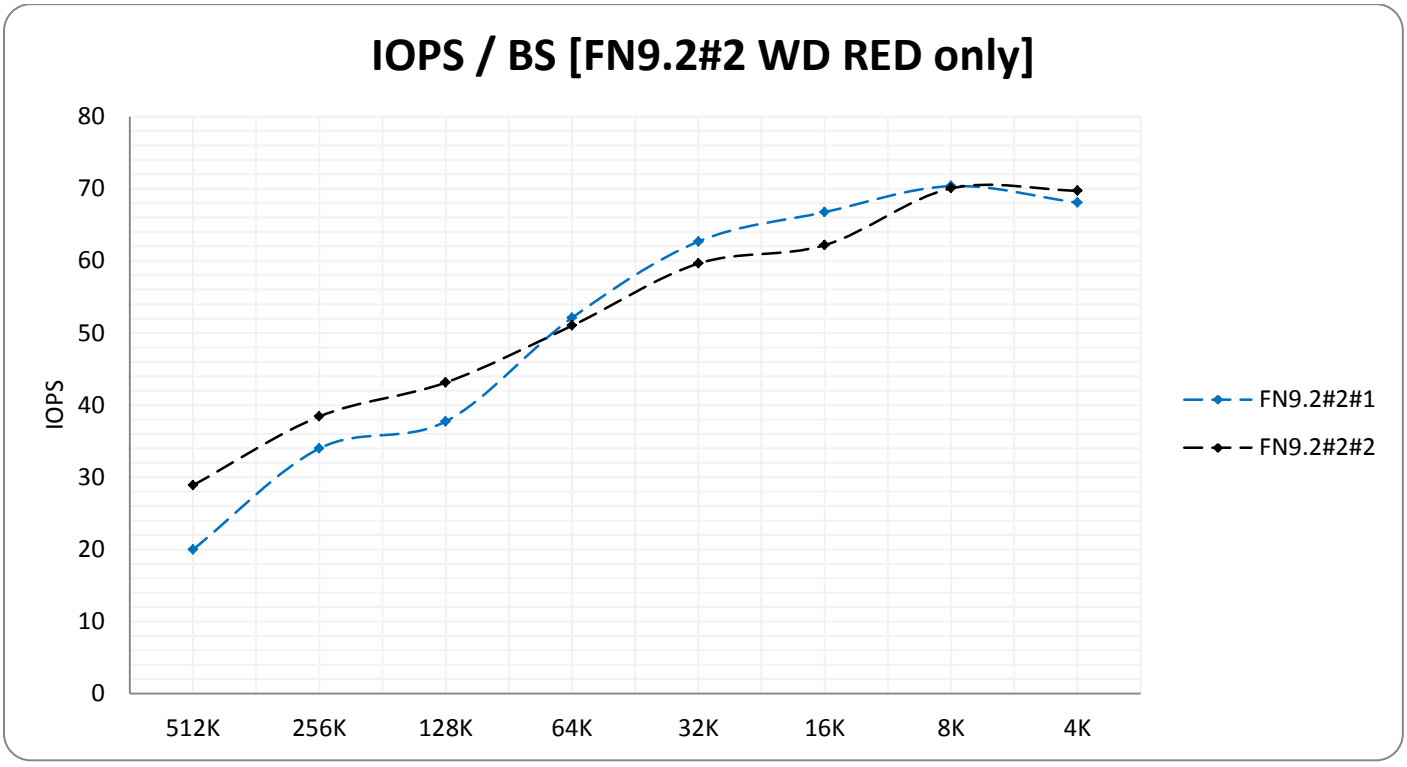

Figure 5:FN9.2 (test configuration)

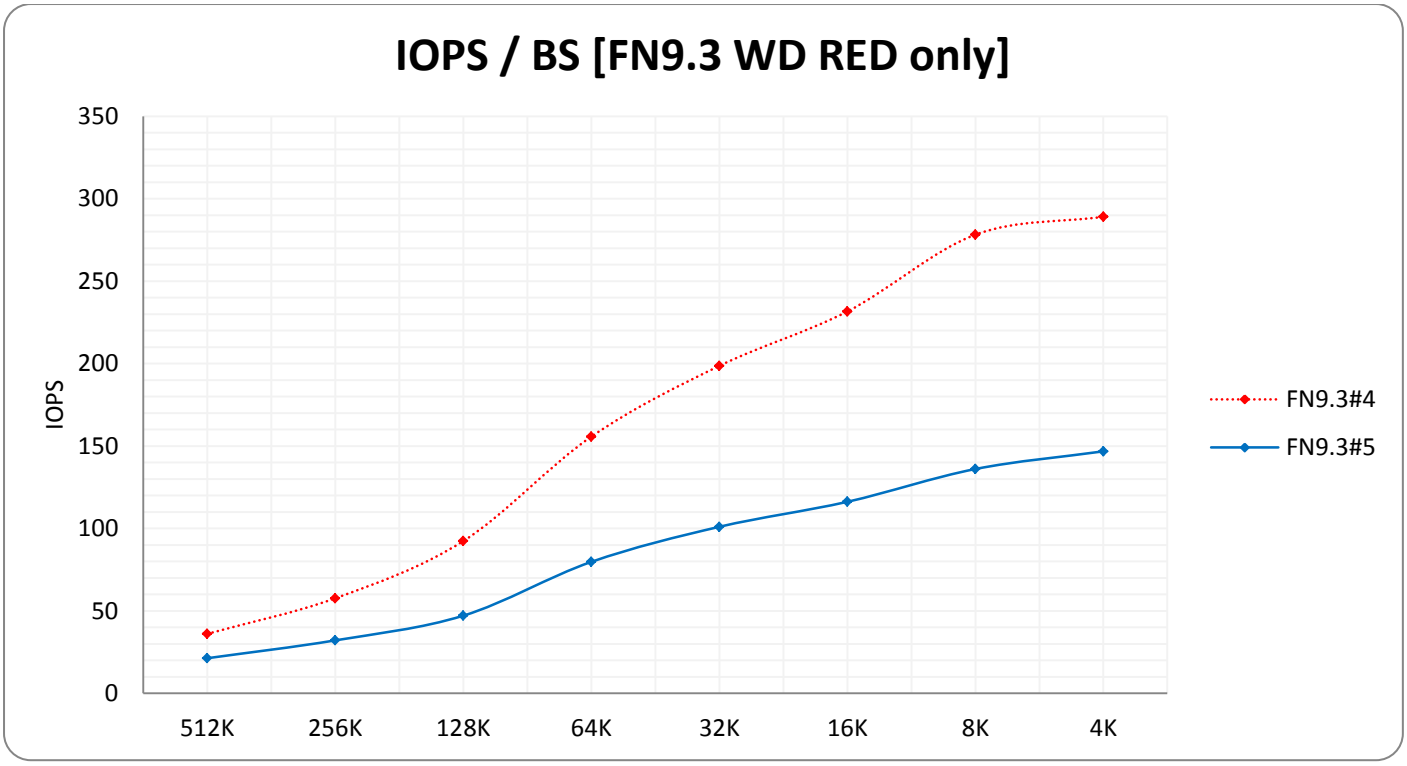

Figure 6:FN9.3 (test configuration)

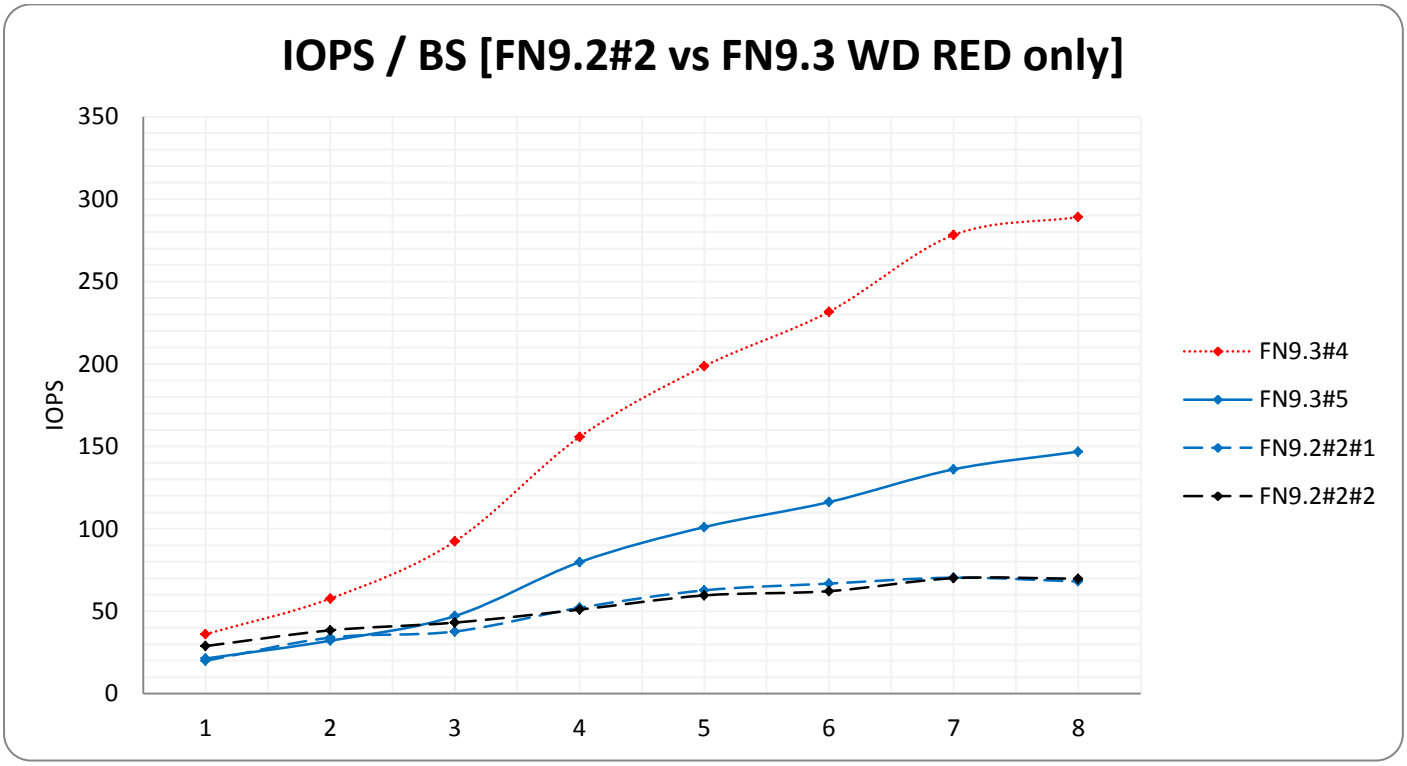

Figure 7:Comparation FN9.2 vs 9.2 (test configuration)

IV.1.3 Disk type perspective – WD BLACK

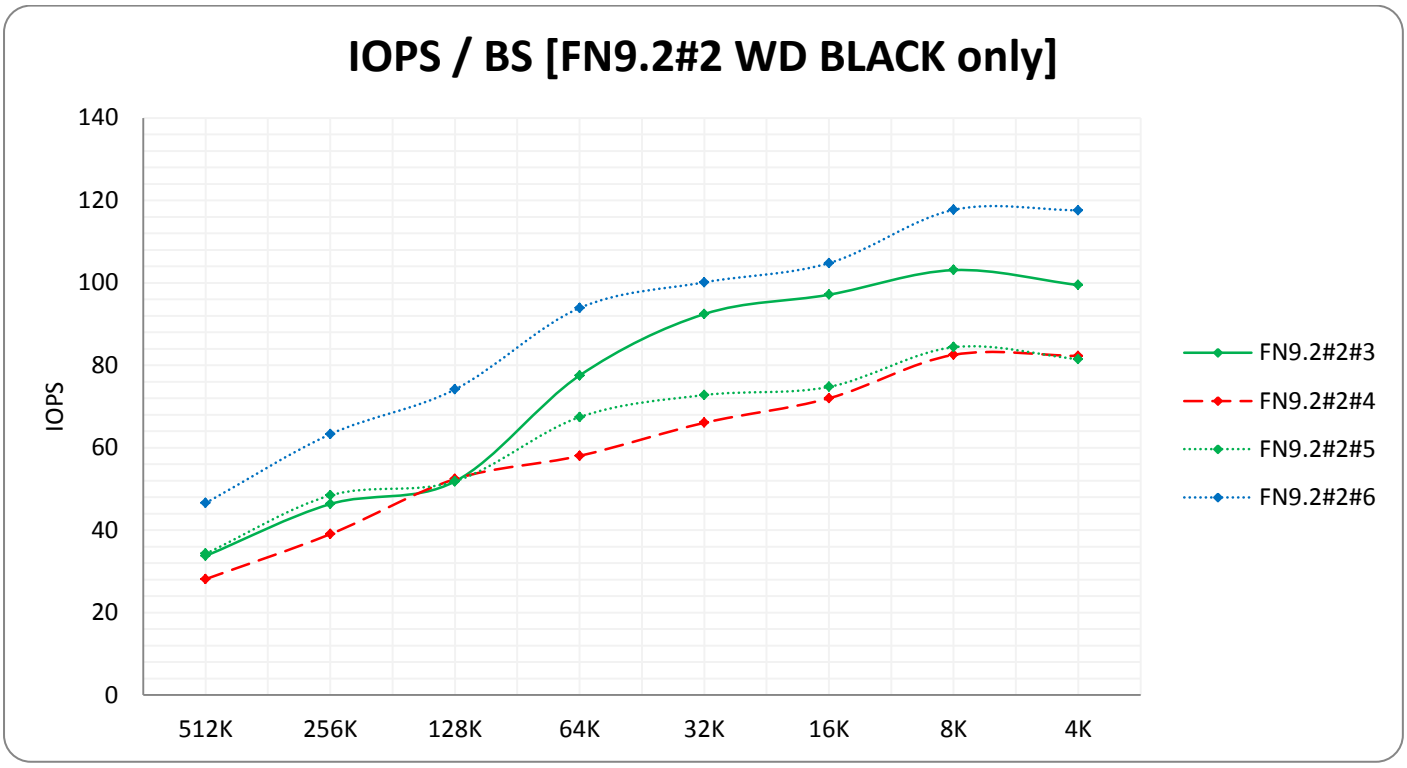

Figure 8:FN9.2 (test configuration)

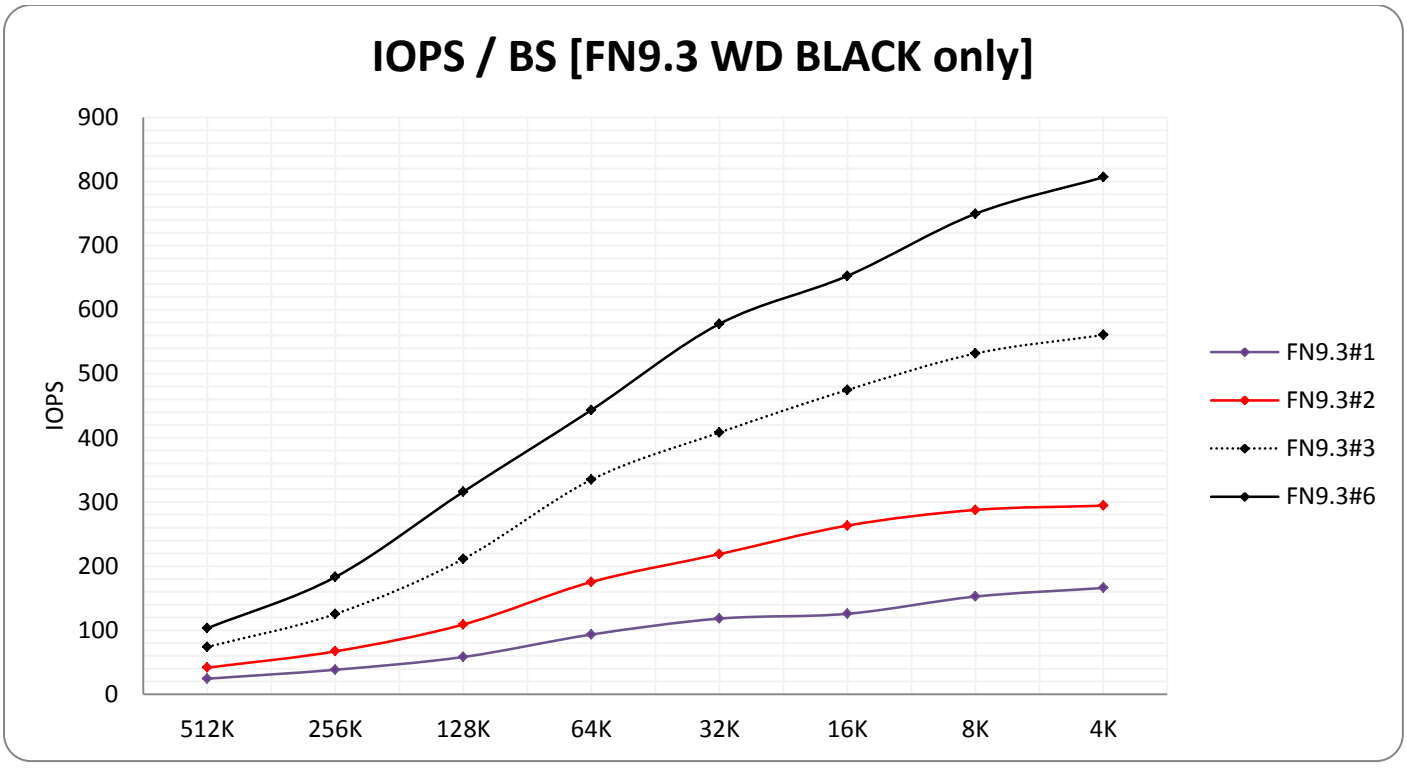

Figure 9:FN9.3 (test configuration)

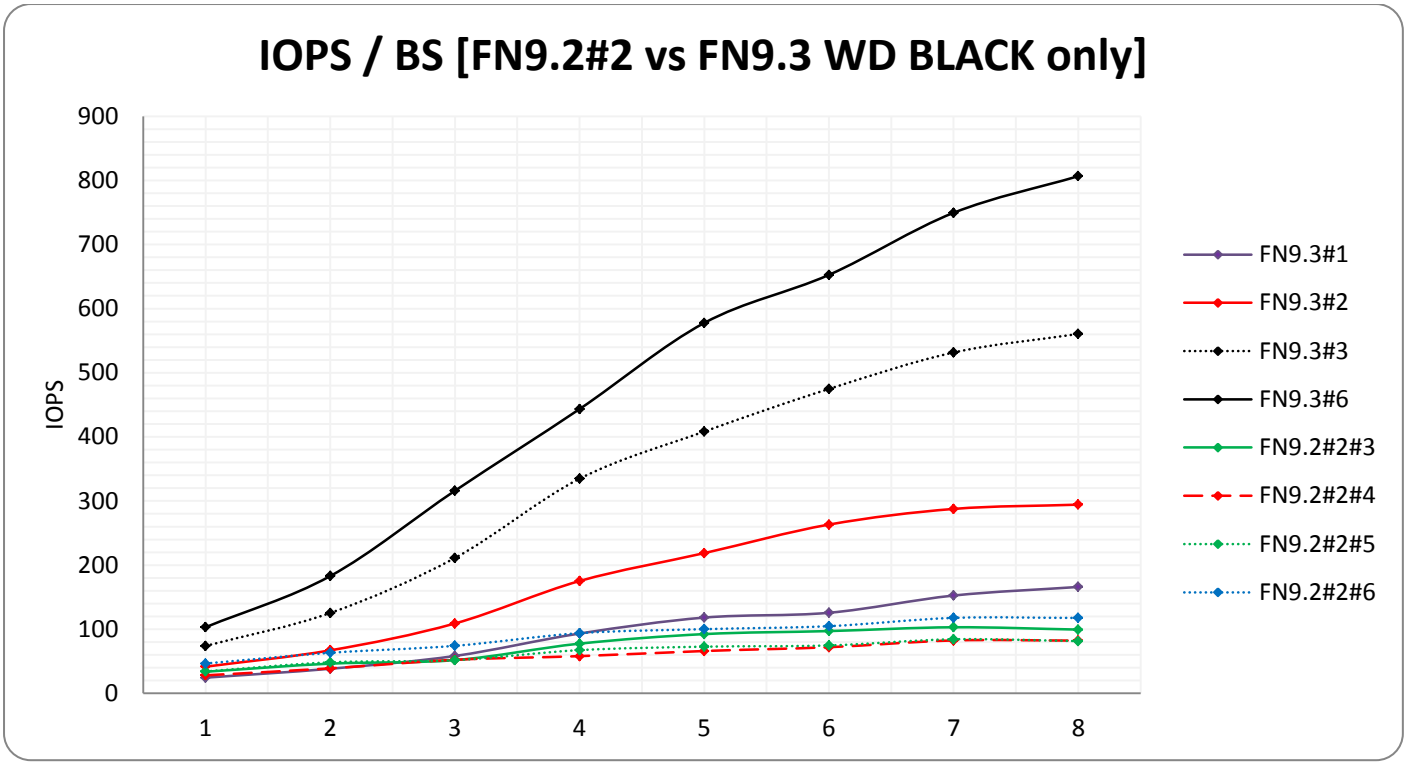

Figure 10:FN9.2 vs FN9.3 (test configuration)

IV.2 Max IOPS

Registered data:

	MAX READ IOPS	MAX WRITE IOPS
	IOPS	
FN9.2#1	54450,086	36843,198
FN9.2#2	43030,788	33616,902
FN9.2#3	53900,968	34887,692
FN9.3#1	103323,062	57570,278
FN9.3#2	103585,838	53121,27
FN9.3#3	103711	47126,676
FN9.3#4	100404,416	47878,604
FN9.3#5	97592,23	49471,17
FN9.3#6	108762,282	48757,162
FN9.2#2#1	68115,05	12297,774
FN9.2#2#2	65334,048	41887,136
FN9.2#2#3	62735,592	10987,976
FN9.2#2#4	62724,324	4486,93
FN9.2#2#5	64641,95	38665,604
FN9.2#2#6	65193,85	44737,484
FN9.3#2#1	70938,198	65629,512

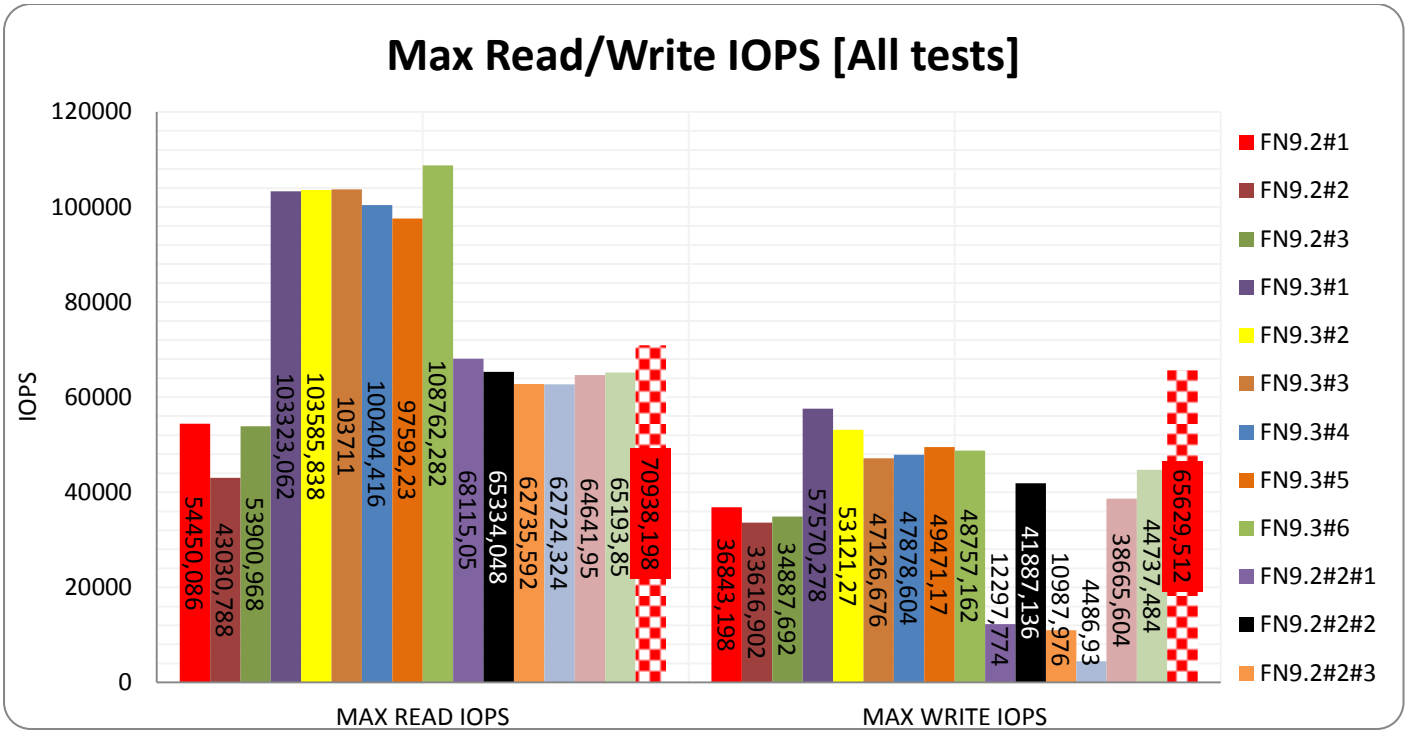

Figure 11:All tests

IV.2.1 General perspective

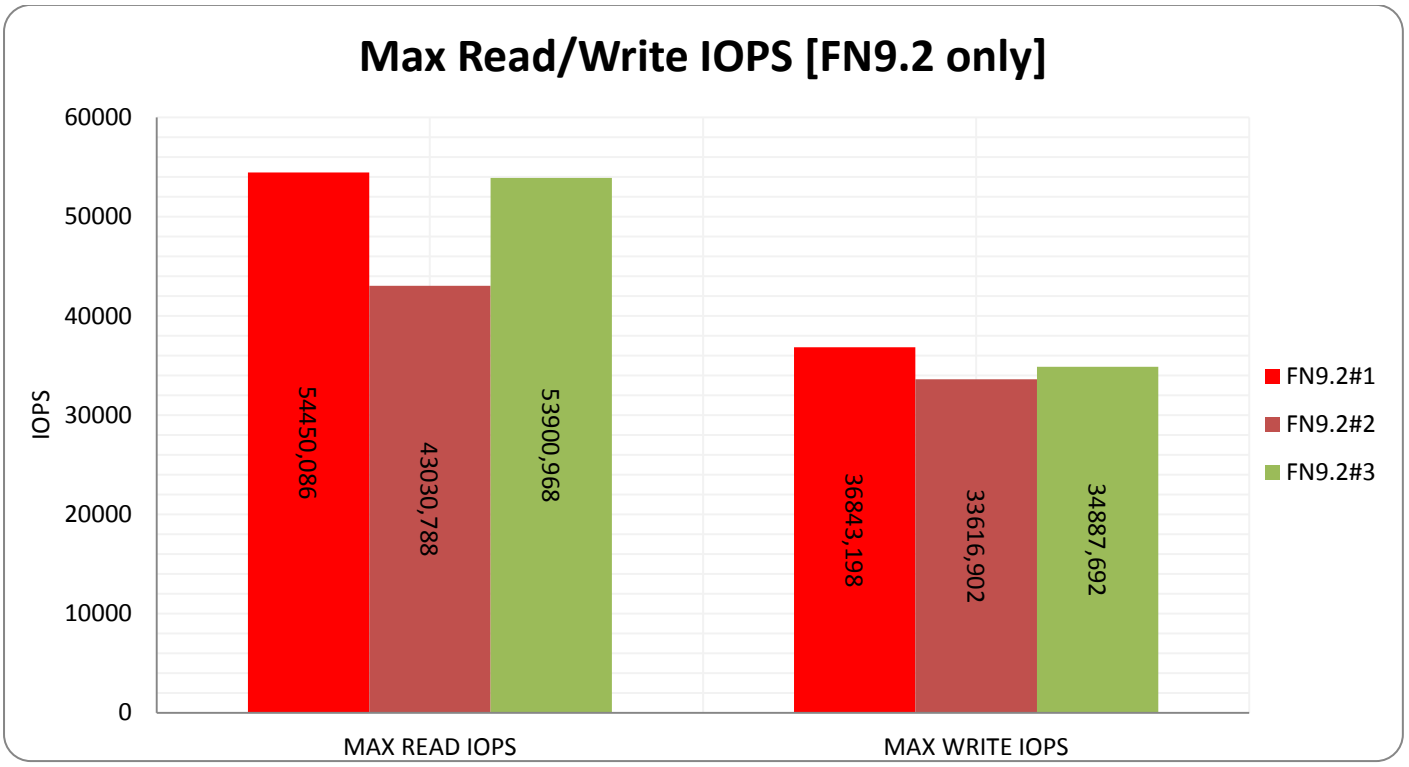

Figure 12:FN9.2 (production)

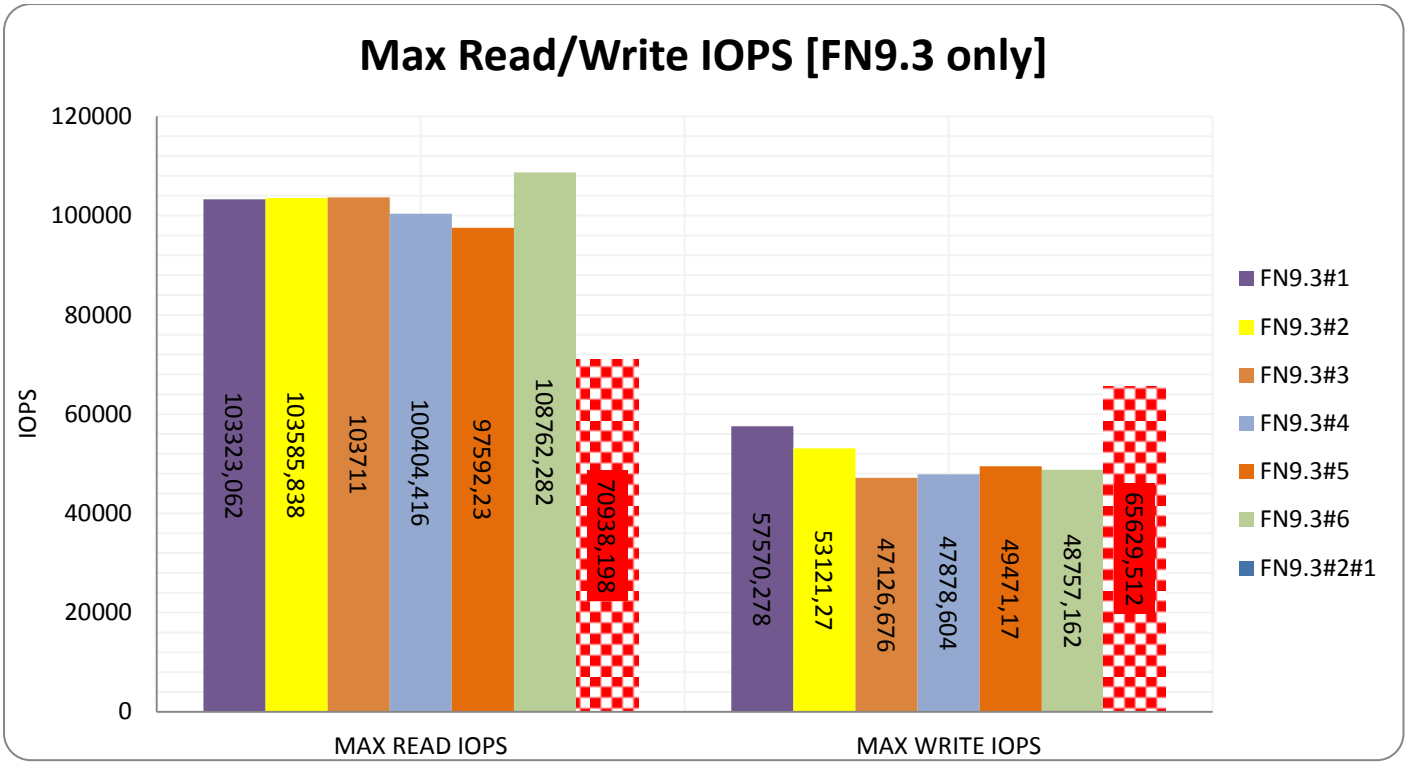

Figure 13: FN9.3 (test configuration)

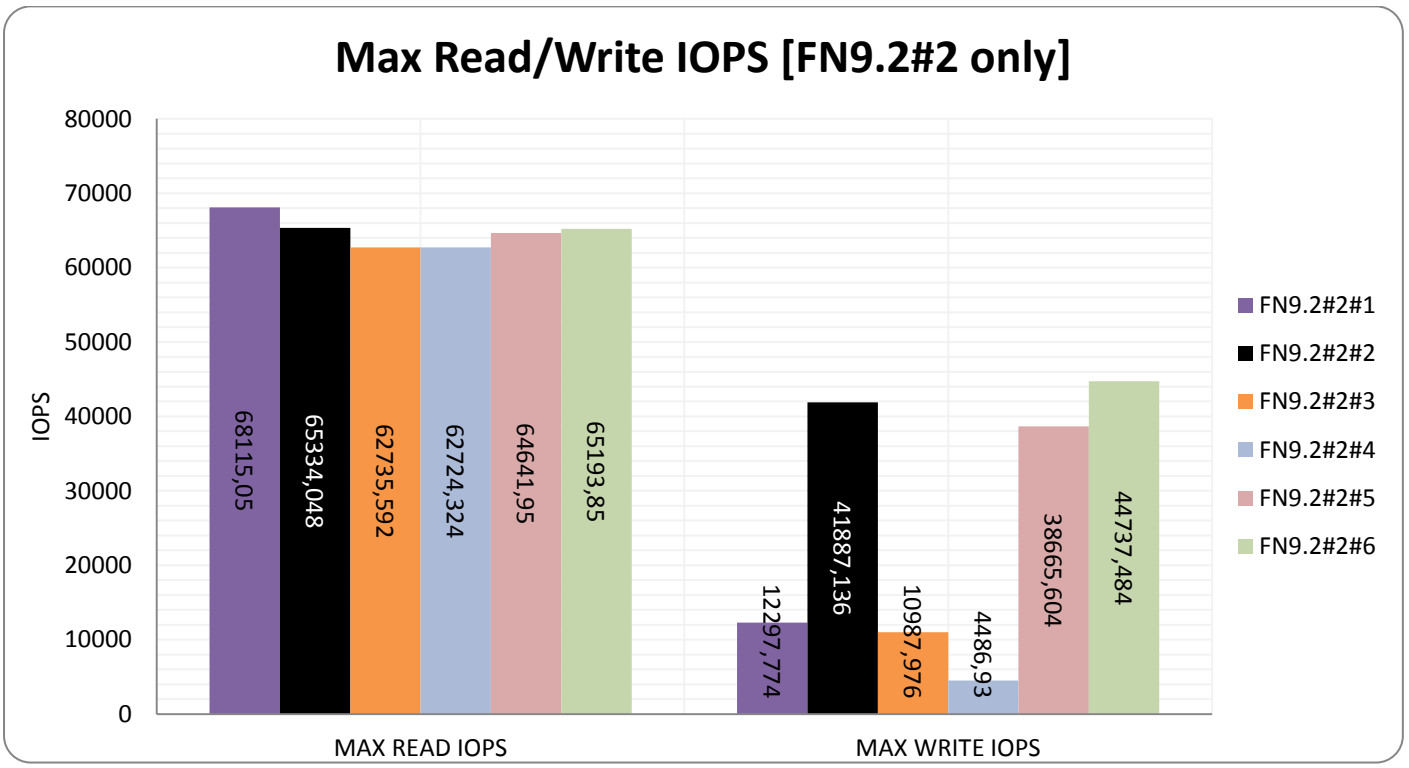

Figure 14:FN9.2 (test configuration)

IV.2.2 Disk type perspective – WD RED

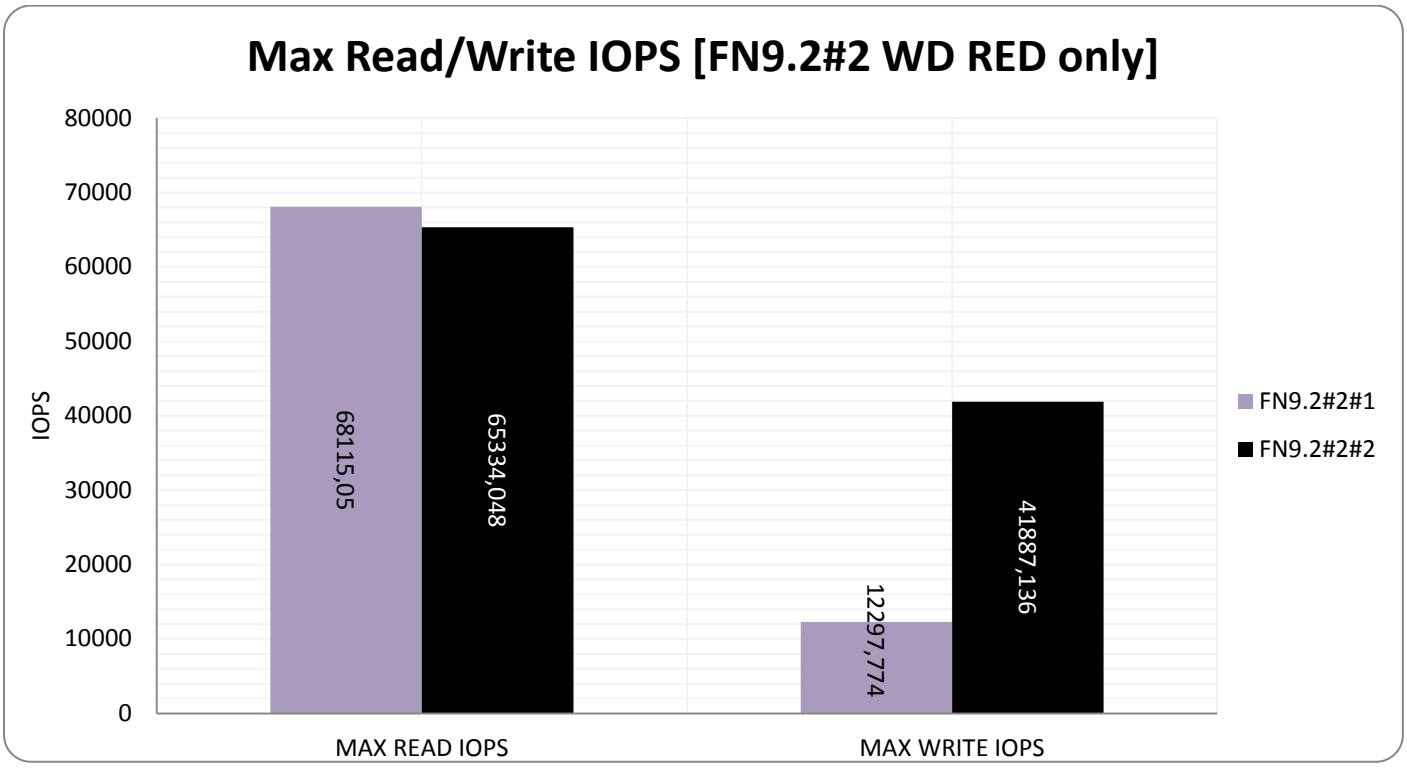

Figure 15:FN9.2 (test configuation)

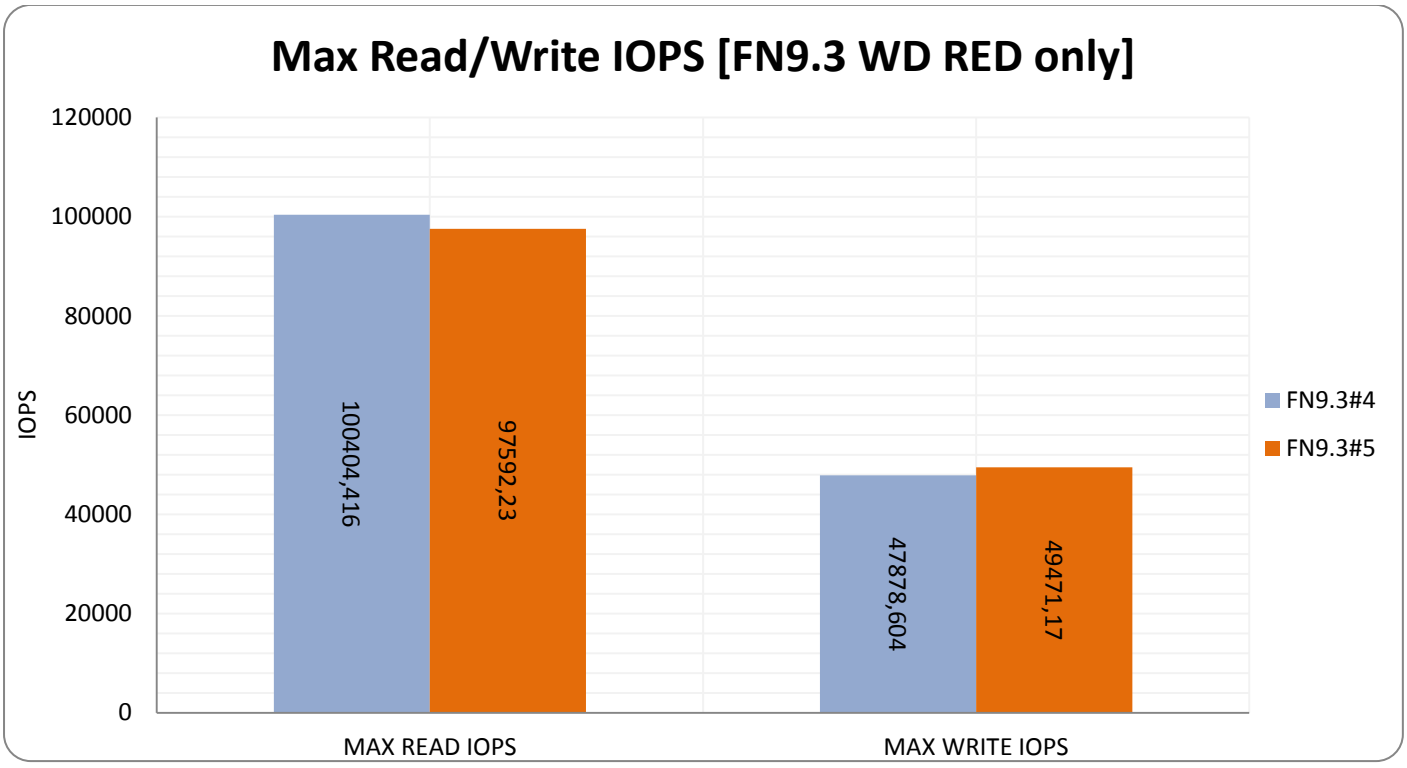

Figure 16:FN9.3 (test configuration)

IV.2.3 Disk type perspective – WD BLACK

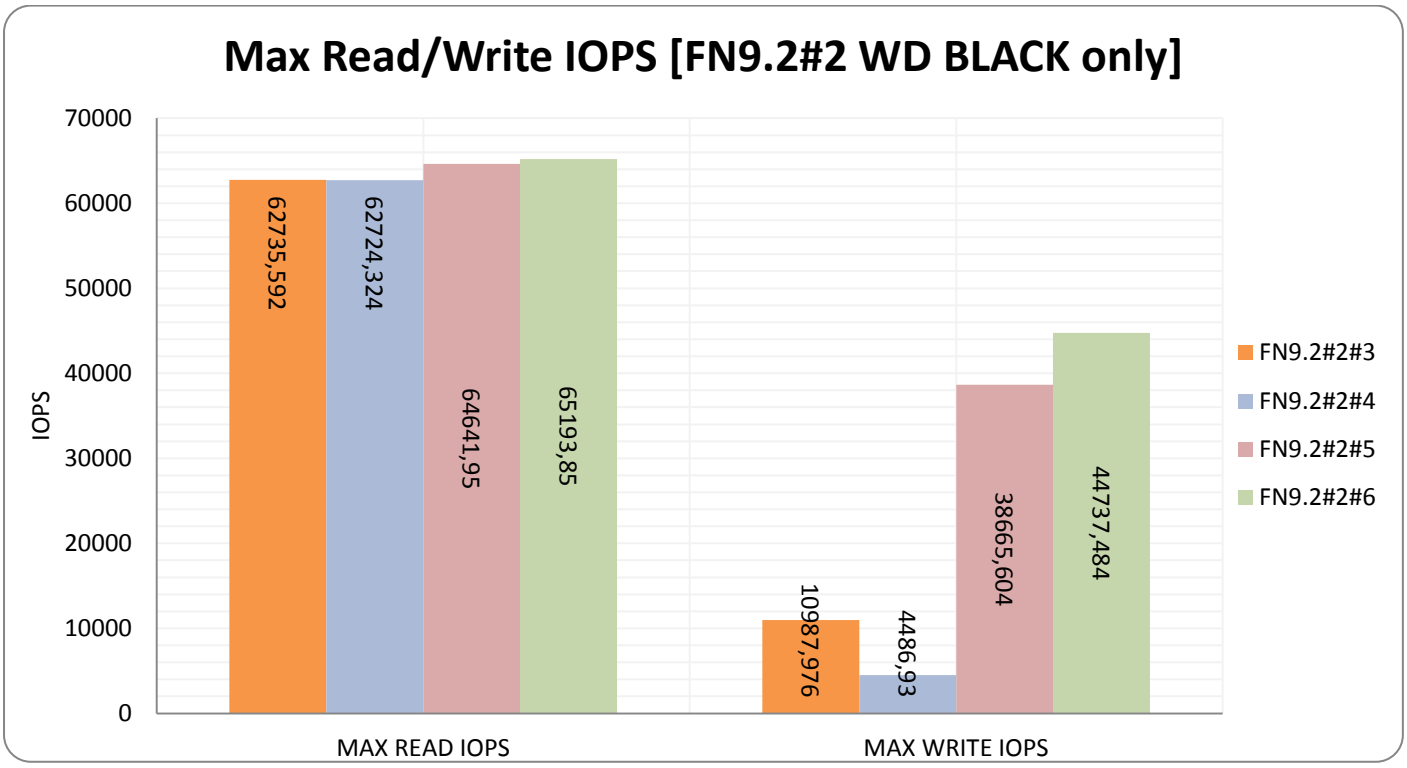

Figure 17:FN9.2 (test configuration)

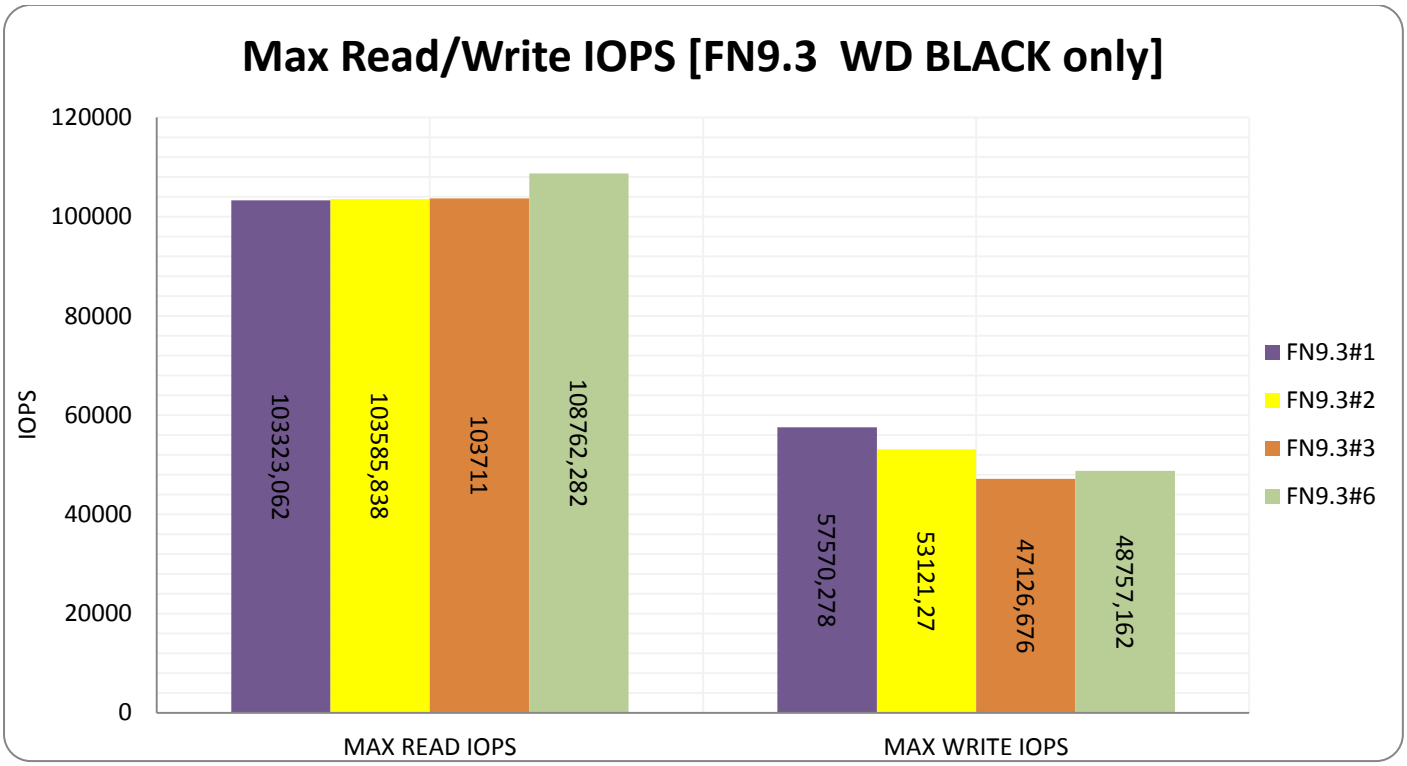

Figure 18:FN9.2 (test configuration)

IV.3 Max Thoughput

Registered data:

	MAX READ THROUGHPUT	MAX WRITE THROUGHPUT
	IOPS	
FN9.2#1	482,646	458,564
FN9.2#2	243,512	563,144
FN9.2#3	335,096	433,956
FN9.3#1	228,78	227,072
FN9.3#2	333,556	262,702
FN9.3#3	412,22	371,262
FN9.3#4	265,512	236,732
FN9.3#5	313,844	264,092
FN9.3#6	421,84	440,604
FN9.2#2#1	107,59	185,572
FN9.2#2#2	178,566	211,882
FN9.2#2#3	119,648	202,434
FN9.2#2#4	117,78	183,108
FN9.2#2#5	157,068	274,846
FN9.2#2#6	195,016	351,496
FN9.3#2#1	504,012	538,238

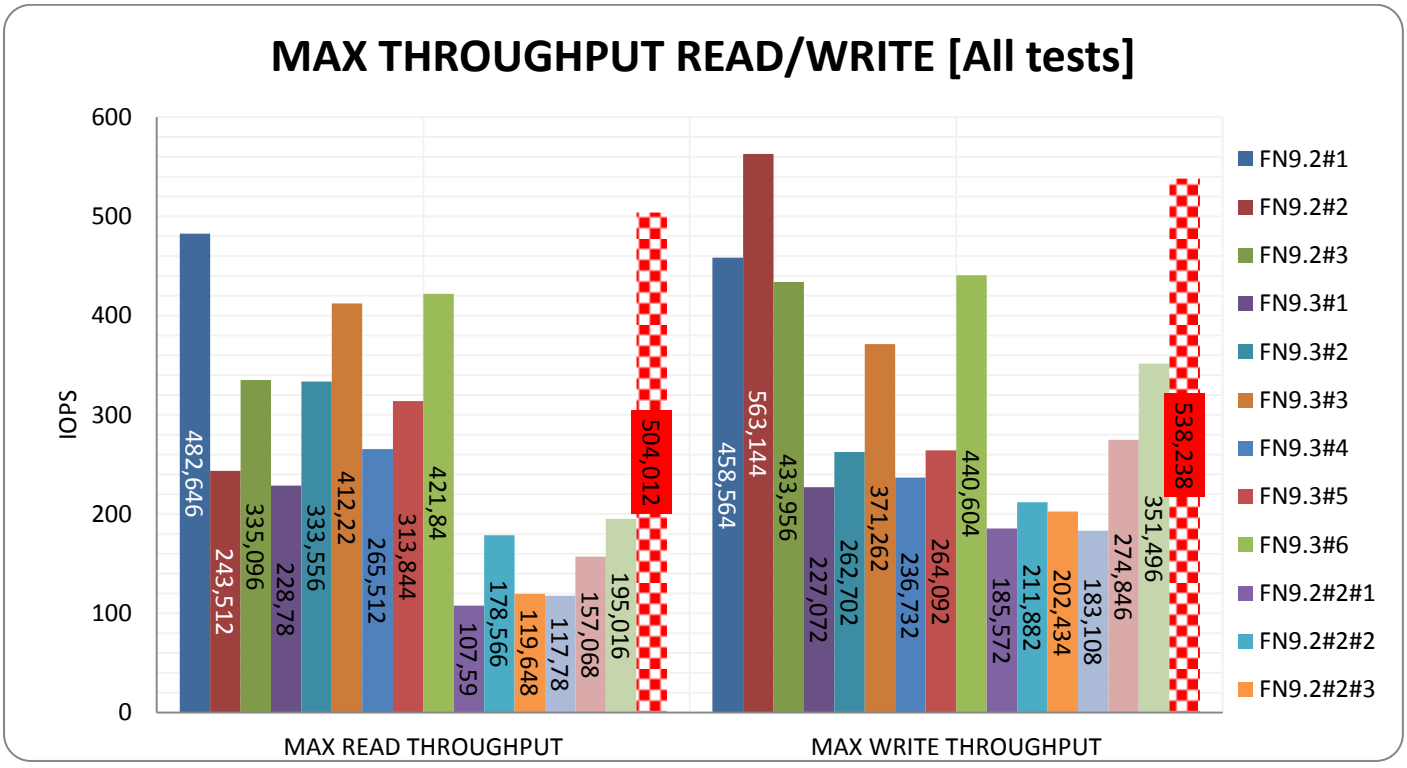

Figure 19:All tests

IV.3.1 General perspective

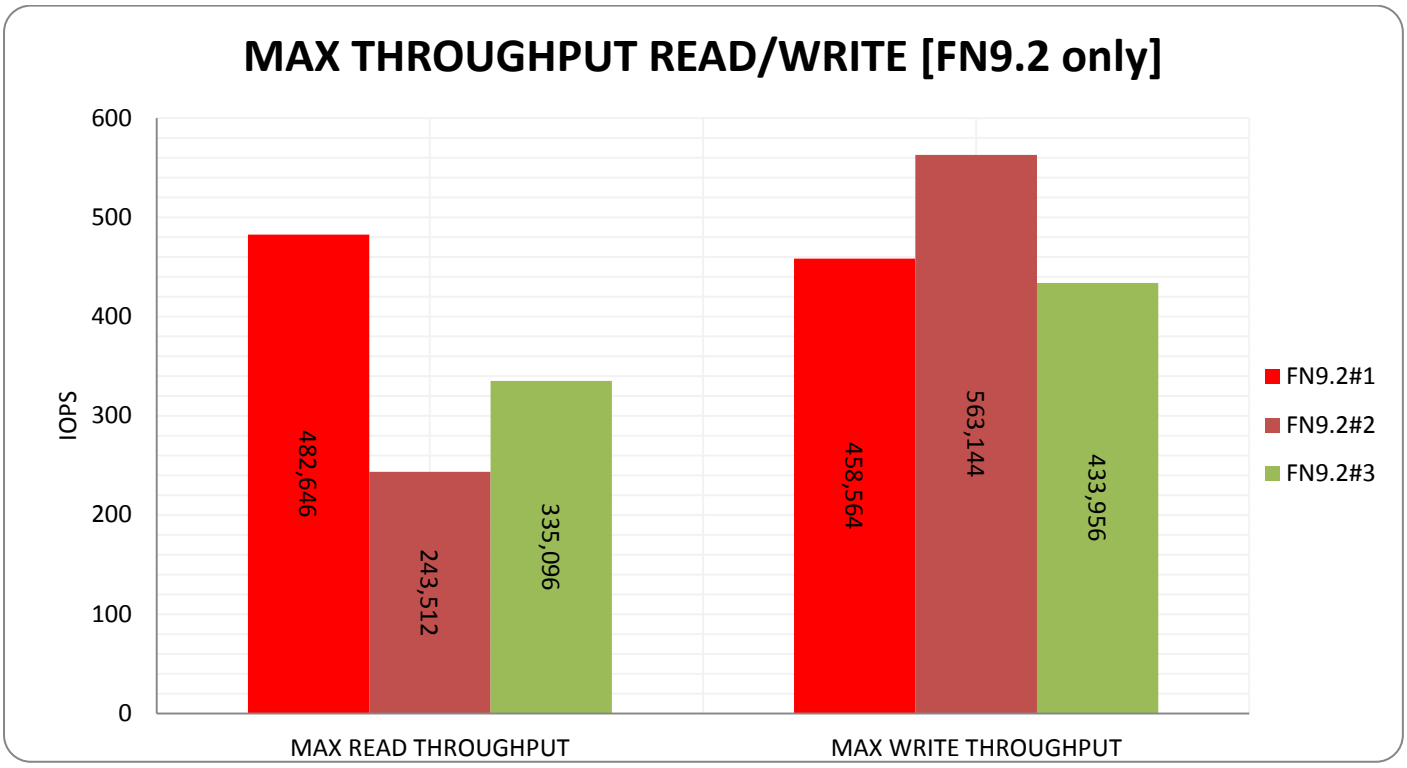

Figure 20:FN9.2 (production)

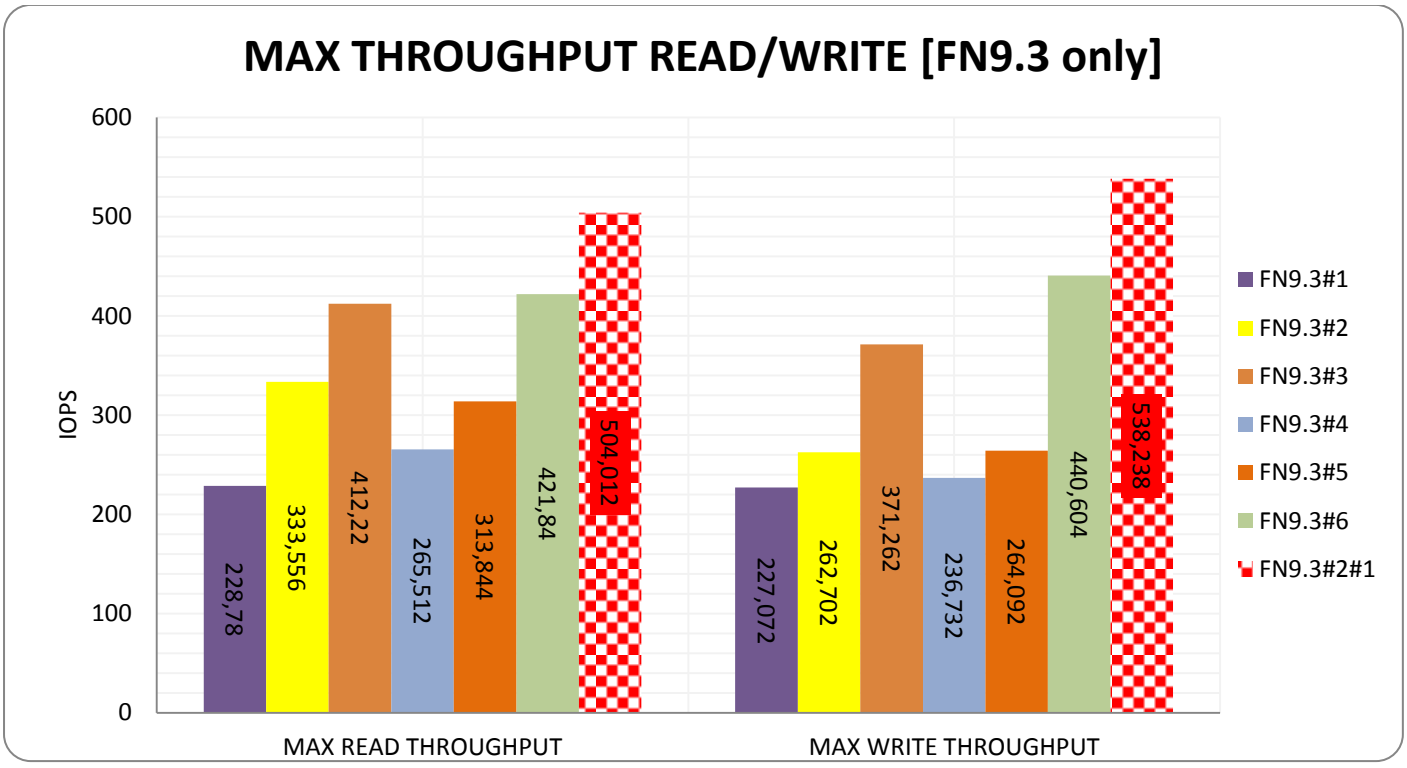

Figure 21:FN9.3 (test configuration)

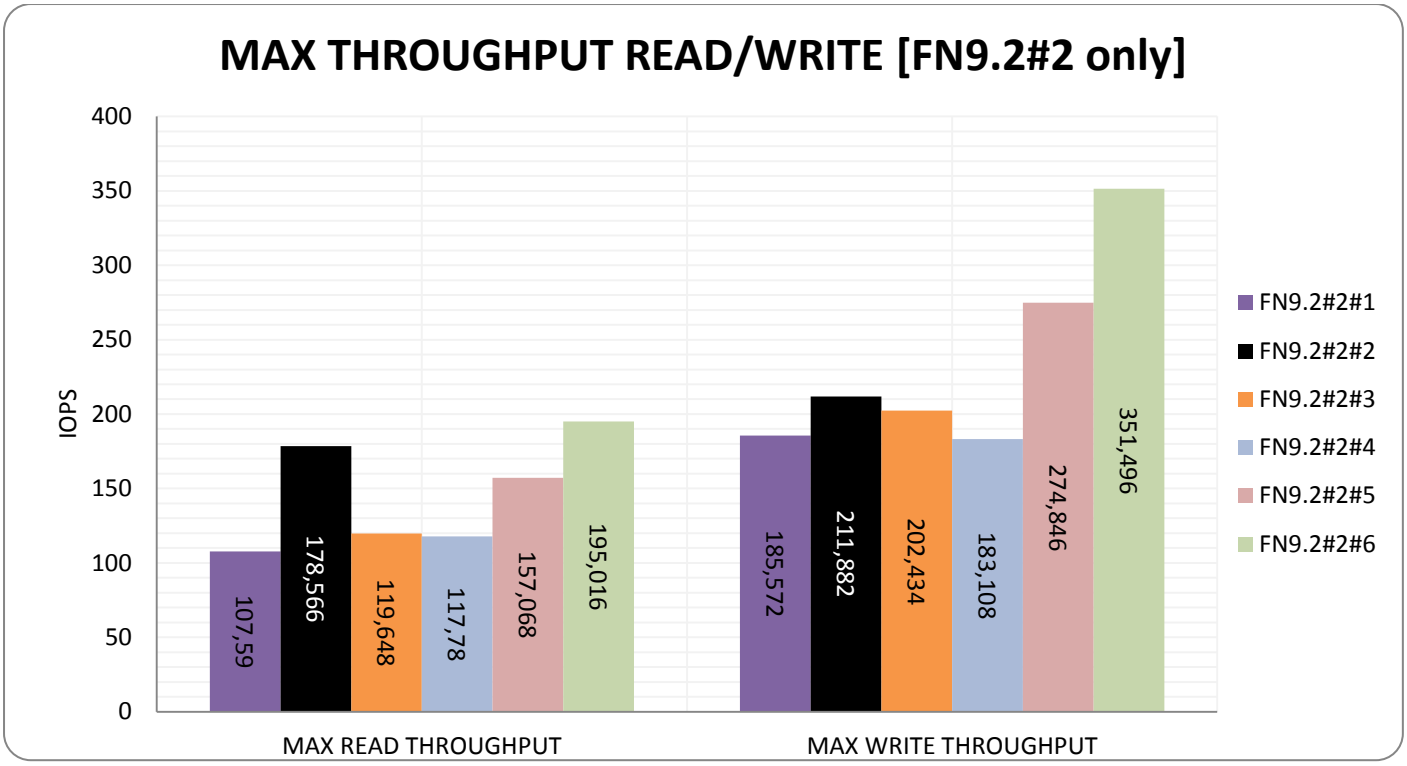

Figure 22:FN9.2 (test configuration)

IV.3.2 Disk type perspective – WD RED

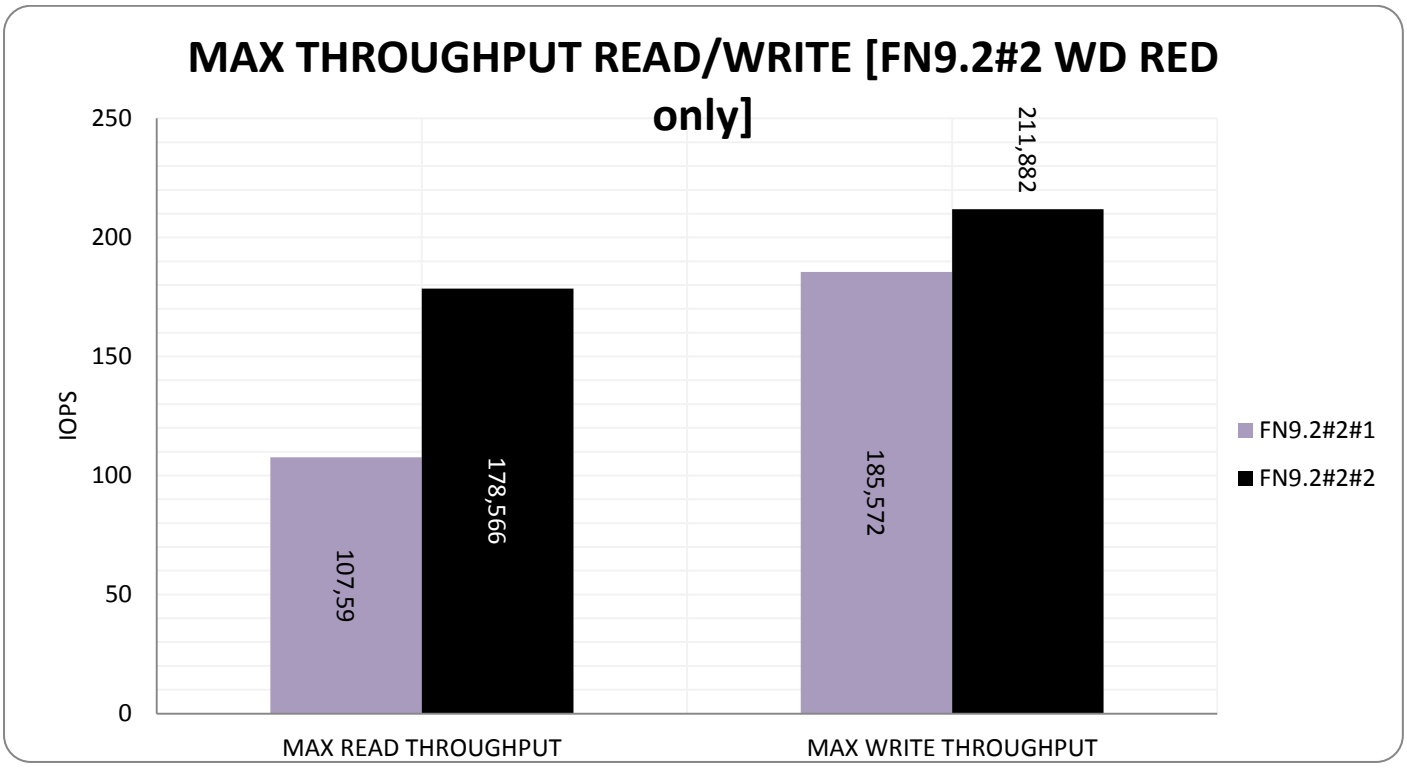

Figure 23:FN9.2 (test configuration)

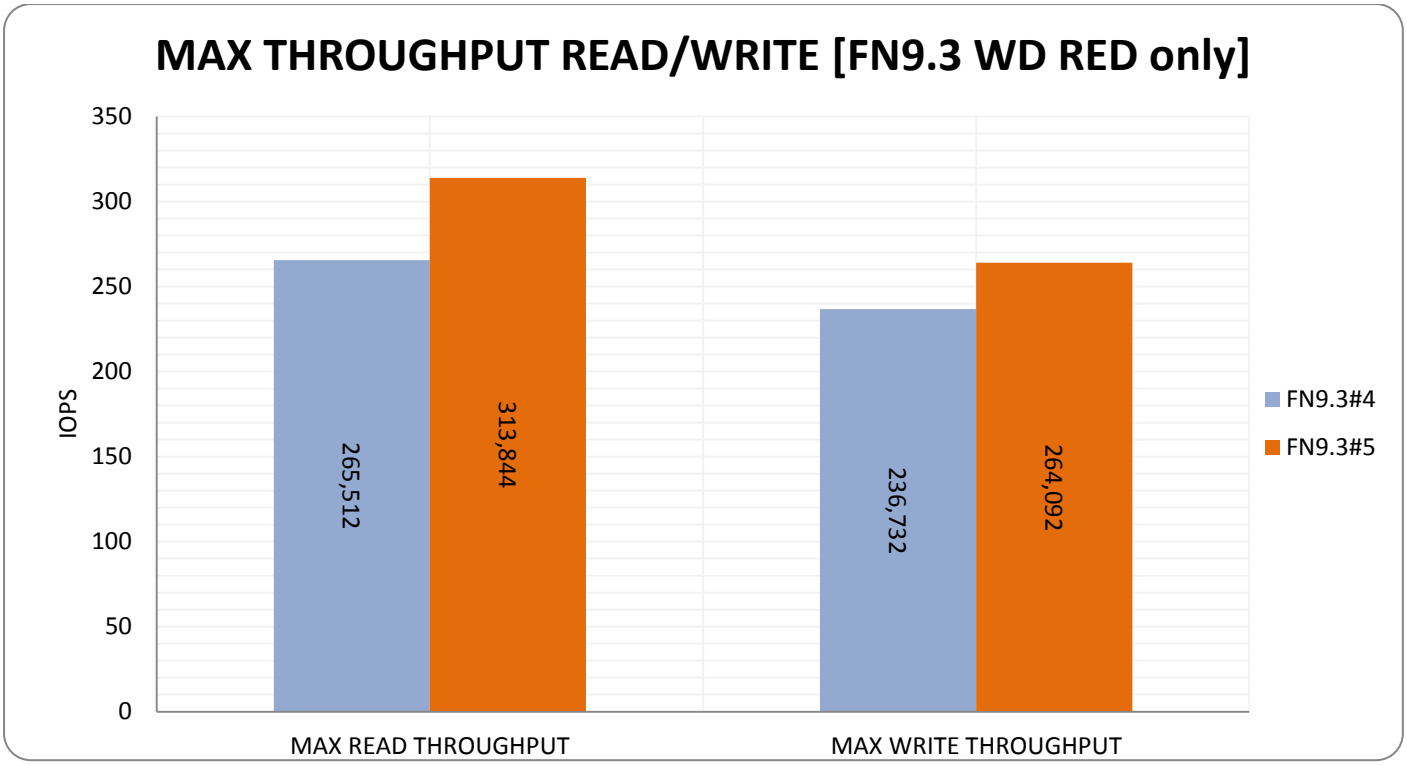

Figure 24:FN9.3 (test configuration)

IV.3.3 Disk type perspective – WD BLACK

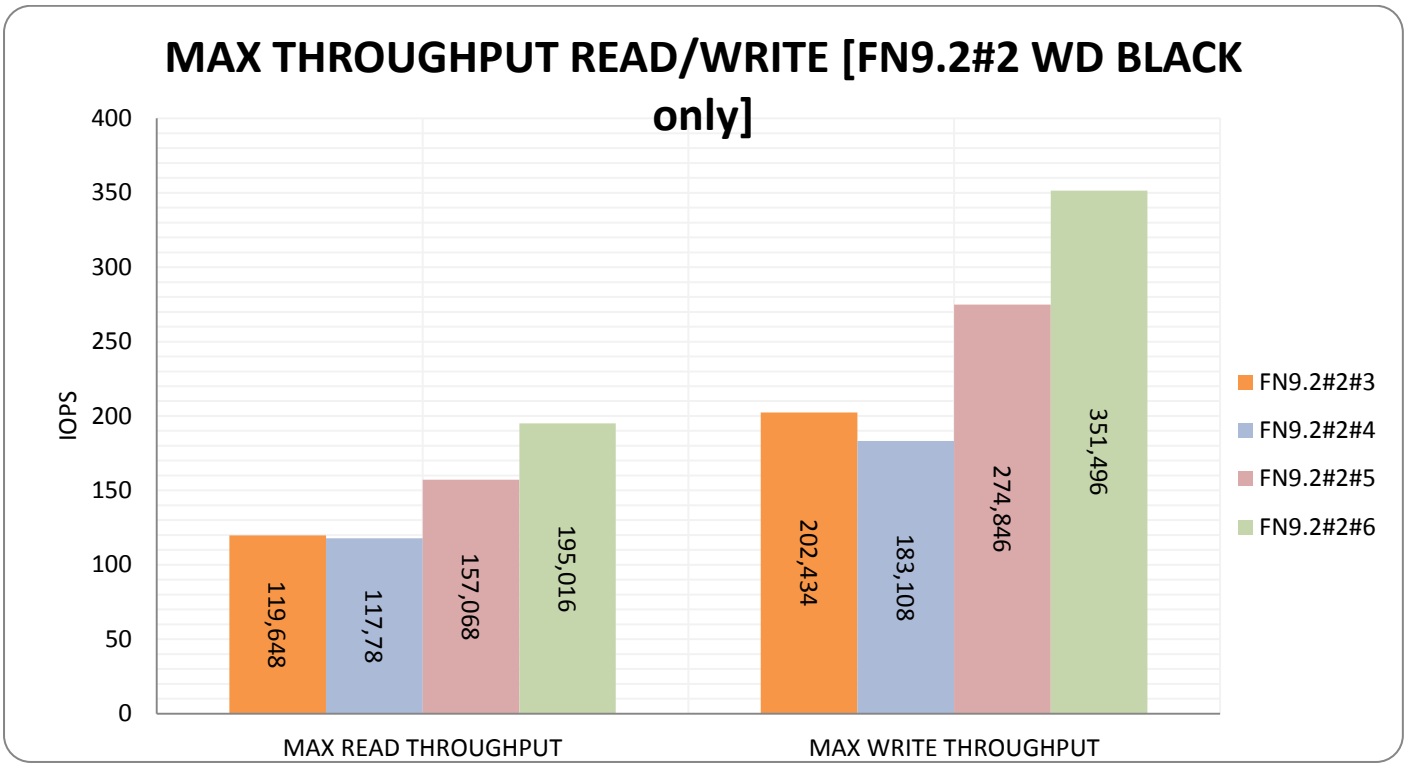

Figure 25:FN9.2 (test configuration)

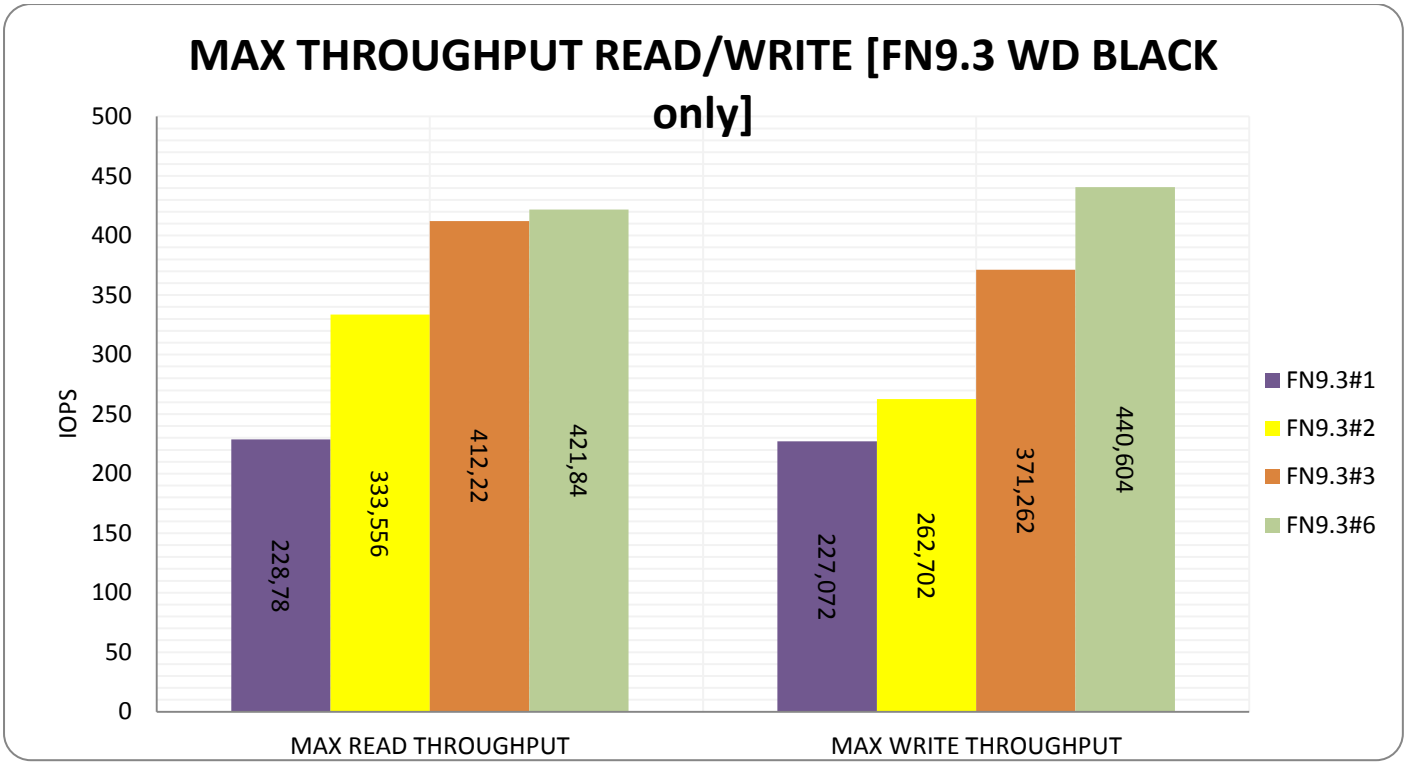

Figure 26:FN9.2 (test configuration)

V. Conclusions & problems

1. Performance of FN9.3 is significantly better than FN9.2 (Figures 4,7 and 10) according IOPS per BS. It is important to remind we are talking about random access workload (not sequential workload). Random workload is typical for ESXI storage access.

2. On figure 5 we can find interesting case – increasing number of disks inside RAID10 does not cause significant increase IOPS. We can even say increase is minimal. If we add to this observation fact that on production server (having 14 WD RED as RAID10) results are also not impressive (figure 2) – specially that hardware of production server is better than test configuration (100% more ram, better processor, LSI controller, L2ARC SSD) situation is more strange. If we compare these results with results of FN9.3 (figure 6) conclusion is again FN9.3 is significantly better.

3. Similar to conclusion 2, if we compare WD Black of FN9.2 (figure 8) and FN9.3 (figure 9) conclusion is the same.

4. Interesting observation we have when gstat is started during FN9.2 testing. We can see that %busy of disks is below 40% (practically during all test cases) while zvol %busy is 100%:

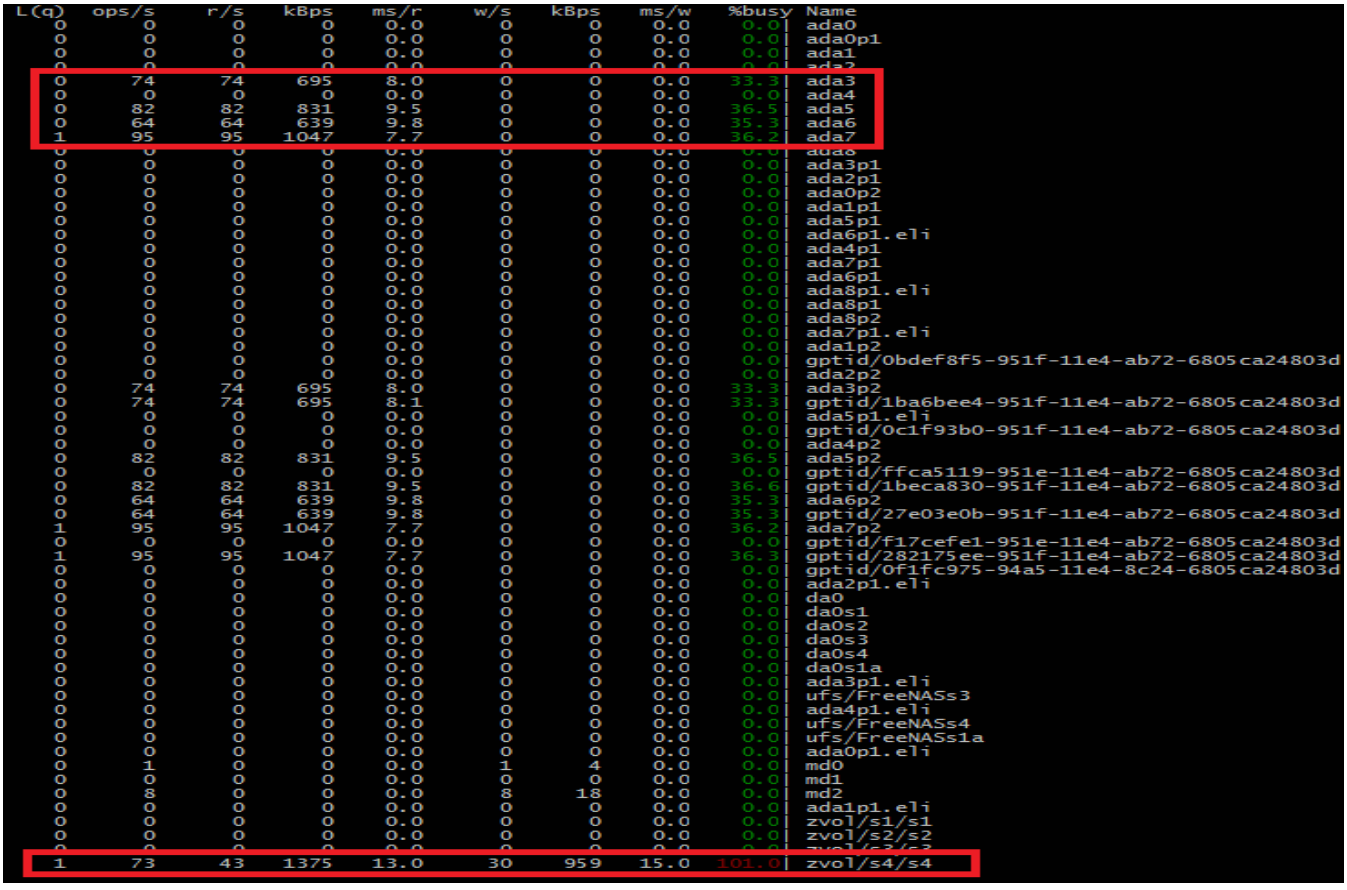

5. Taking into account conclusion 2, 3 and 4 it looks like performance of FN9.2 depends on general hardware performance not size of RAID.

6. Max IOPS (read/write) visible on figure 13 (for FN9.3) and 14 (for FN9.2) illustrate performance with very small packages (512 bytes). FN9.3 in this tests again is faster than 9.2.

7. Interesting issue we can find on figure 15 – as we tested 1 disk and then 2 disks (as mirror) is was expected that reading should be better in case of mirror and write should be the same as on 1 disk. What we see is something opposite – evolution of reading is minimal but evolution of writing is very huge. Looks like caching and FN implementation cause this.

   Similar analyze in case of FN9.3 (figure 16) shows that FN9.3 compensate differences between 1 disk and mirror pair.

8. Similar to conclusion 7, if we take a look for BLACK disks – figure 17 (for FN9.2) and 18 (for FN9.3) brings the same conclusion (as conclusion 7) but –because it was possible to test more disks—we can find not nice issue – FN9.3 (as we can see on figure 18) reports decrease of max write IOPS when we increase number of disks. This is really not nice and need to be study more deeply.

9. FN9.2 and FN9.3 regarding max throughput reports similar conclusions as 7 and 8 listed above. Important however is to say that having in test configuration only 2 paths (2x NIC) was not enough to fully test this cases because always FN network report shows maximal saturation of both NICs.

10. Tests made on production server (previously on FN9.2) after migration (to FN9.3) are visible on figure 4. As we can see performance of the same hardware in the same workload just with FN9.3 kick results significantly up (test FN9.3#2#1 on figure 4).