

Intel® SSD DC S3500 Series Workload Characterization in RAID Configurations

White Paper



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2013, Intel Corporation. All Rights Reserved.



Contents

1.0	Revision History	5
2.0	Supporting Documentation	5
3.0	About This Guide	5
4.0	Overview	5
4.1	What Impacts SSD IO Performance.....	6
4.2	Queue Depth and Latency	7
4.3	Why Mixed Workload Is Important.....	7
4.4	Drive Endurance.....	8
4.5	Selection of RAID Controller.....	8
5.0	RAID 1	9
5.1	Test System Specifications	9
5.2	Intel SSD DC S3500 Series in RAID 1 Performance Characterization Data.....	10
5.3	RAID 1 Consistency	13
5.4	RAID 1 Performance Conclusions	13
6.0	RAID 5	14
6.1	Test System Specifications	14
6.2	Intel SSD DC S3500 Series in RAID 5 Performance Characterization Data.....	14
6.3	RAID 5 Performance Consistency.....	16
6.4	RAID 5 Performance Conclusions	17
7.0	Summary	18
8.0	Appendix	19
8.1	RAID Levels	19



Tables

Table 1 Typical Mixed Workloads in Data Center Applications 7

Figures

Figure 1 RAID 1 Random 100% Write @ 4KB Transfer Size with Average Latency 11
Figure 2 RAID 1 Random 70% Read @ 4KB Transfer Size with Average Latency 11
Figure 3 RAID 1 Random 90% Read @ 4KB Transfer Size with Average Latency 12
Figure 4 RAID 1 Random 100% Read @ 4KB Transfer Size with Average Latency 12
Figure 5 RAID 1 Maximum Latency for 2-drive and 8-drive Configurations 13
Figure 6 RAID 5 Random 100% Write @ 4KB Transfer Size with Average Latency 15
Figure 7 RAID 5 Random 70% Read @ 4KB Transfer Size with Average Latency 15
Figure 8 RAID 5 Random 90% Read @ 4KB Transfer Size with Average Latency 16
Figure 9 RAID 5 Random 100% Read @ 4KB Transfer Size with Average Latency 16
Figure 10 RAID 5 Maximum Latency for 3-drive and 8-drive Configurations 17



1.0 Revision History

Document Number	Revision Number	Description	Revision Date
329903	001	Initial release	December 2013

2.0 Supporting Documentation

For more information on Intel SSDs, see the corresponding documentation.

Document	Document No./Location
Intel® Solid-State Drive DC S3500 Series Product Specification	328860

3.0 About This Guide

This guide describes Intel® SSD DC S3500 Series performance characteristics in RAID configurations across multiple workloads, and provides analysis to help optimize performance.

The audience is technical IT professionals: Systems, Storage, Database, and Application Engineers.

4.0 Overview

The Intel SSD DC S3500 Series provides high random read and write storage Input/Output Operations per Second (IOPS) across mixed read and write workloads. This high random performance and the consistency of IOPS under workload deliver robust and scalable operation when used behind a RAID controller. Data centers can benefit in both performance and TCO by using the Intel SSD DC S3500 Series in the appropriate applications.

Compared to the approximately 200-300 random IOPS that a single 15K SAS hard disk drive (HDD) can provide, an Intel SSD DC S3500 Series operates at much higher IOPS; up to 75,000 IOPS for random 4KB reads and up to 11,500 IOPS for random 4KB writes, over the entire span of the SSD. The Intel SSD performance numbers are based on the Intel product specification sheet, as derived from internal Intel testing. With real-world workloads, the IOPS that any particular device can produce will vary depending on several factors: the application's ability to produce IOPS, the ratio of random to sequential access, the block transfer size, the queue depth, the read/write mix of the workload, and overall resource utilization in the server running the workload.

This guide presents data for RAID 1 and RAID 5 configurations due to their prominence in the datacenter. Additional RAID levels are currently being tested and will be presented in future revisions, or as separate papers.



A selection of workloads that represent both best-case performance and real-world performance are presented in this white paper. These scenarios give the IT professional a better understanding of the capabilities of the Intel® SSD DC S3500 Series drive when used in conjunction with a hardware RAID controller. More importantly, it helps the IT professional understand a variety of workloads and circumstances in which Intel SSD technologies will accelerate those workloads and provide business value for their organization.

4.1 What Impacts SSD IO Performance

Although Intel SSDs excel in delivering random read and write IOPS, it is important to remember that more IO activity at the application level results in higher CPU utilization in the applications' host. In addition to the abilities of the SSD, IO performance in any particular situation is dictated by how the particular application scales, and the IO profile of the workload produced by the application.

The following workload-specific characteristics have a direct impact on the ability of the SSD to produce IO:

- Read/Write Mix – NAND programming (writes) and read timing (reads) differ significantly at the hardware level. Because of the higher controller overhead required for processing writes, the number of read IOPS are often higher than write IOPS. Real world workloads are most often a mix of read and write.
- Random/Sequential Mix – IOPS can vary depending on the ratio of sequential versus random accesses. With higher random write workloads, more data movement and greater data management activity occurs in the drive. As random write activity increases, the IOPS serviceable to the host typically decreases.
- Queue Depth - Higher queue depths typically allow the SSD to generate higher IOPS through concurrent processing of commands. However, as the queue size increases, latency will be negatively impacted.
- Random Transfer/Block Size - With a smaller transfer size, the SSD controller has to work harder to maintain the logical-to-physical address mappings. In addition, the smaller the transfer size, the larger the logical space needed for its mapping. Once logical space constraint is reached, background re-mapping will take place. These frequent events slow IOPS.
- Available Spare Area – A larger spare area directly impacts random write and mixed read/write performance by minimizing the frequency of reclaim activities and freeing up processor cycles to support more host read/write requests. You can increase the spare area by over-provisioning the SSD. See the *Intel® High Performance SATA Solid-State Drive Over-Provisioning an Intel® SSD White Paper* for more information.

In summary, the following principles of storage are often true concerning queue depth, block size, randomness, and per-IO transactional latency:

- As queue depth increases, IOPS increase, and latency increases.
- As block size increases, throughput increases, and latency increases.
- As randomness increases, IOPS decreases, and latency increases.



4.2 Queue Depth and Latency

Latency – The amount of time needed to service one outstanding IO to the drive, measured in milliseconds (ms) or, with SSDs, microseconds (μ s).

The Intel® SSD DC S3500 Series supports a maximum queue depth of 32 per drive. In a RAID array, the queue depth is multiplied by the number of drives in the RAID set. Example: In a RAID 5 set of 8 drives, the maximum total queue depth would be 256 (8 X 32). As more commands are queued in the SSD, average latency is impacted. Our internal testing indicates that average latency increases sharply with queue depths beyond 8. However, these high queue depths can increase IOPS with read intensive workloads.

Obtaining the best performance for a particular application requires balance. The challenge is to achieve high speed or IOPS at an acceptable latency level. This white paper presents lower queue depths of 1, 2, 4 and 8 per drive. The results shown demonstrate favorable speed and IOPS generation without pushing latency to extreme levels.

4.3 Why Mixed Workload Is Important

Mixed random workloads are predominant in data center and enterprise applications. Intel SSDs have been deployed in a variety of these applications ranging from content delivery and video on demand networks, to Internet datacenter portals and database management servers. Although these applications see unique IO traffic across the storage drive, there are commonalities in their usage of random and read/write mixed workloads.

Table 1 shows an overview of transfer sizes, read/write mixes and randomness in commonly used workloads in data center applications. These are based on commonly available industry information and information available through such benchmarks as, TPC-C, TPC-E, TPC-H, and TPOX, which attempt to mimic these real world applications.

Table 1. Typical Mixed Workloads in Data Center Applications

Application	Transfer Size	%Random	%Read
Web-Servers	4KB/8KB/16KB+	~75%	~95%
Exchange Email	4KB	~95%	~70%
Database OLTP	4KB/8KB	~95%	~70%
Decision Support	16KB+	~95%	~95%
Video On Demand	16KB+	~95%	~95%
Search Engine	4KB/8KB/16KB	~95%	~95%
Cache	16KB+	~95%	~95%
Content Delivery Network	16KB+	~95%	~70%-95%

Based on these usage trends, small transfer sizes—ranging from 4KB to 16KB and above—are common in enterprise and data center applications. Also, much emphasis is placed on random accesses, and although there are varied levels of read and write



mixes, read-intensive workloads are more prominent. In summary, it is important to select the proper SSD for a particular workload.

The examples presented here use 100% write and 100% read workloads to show the maximum performance in these areas. Also, 70%/30% read/write and 90%/10% read/write are used in order to simulate typical workloads in the datacenter.

4.4 Drive Endurance

Drive endurance, or wear, is an important consideration when selecting an SSD for a particular application. The Intel® SSD DC S3500 Series drive is a standard endurance, enterprise class drive, designed for read-heavy workloads. It is important to understand how drive wear is affected by the RAID level.

A RAID level that uses dedicated parity, such as RAID 4, will write all parity to a single drive. This can potentially cause the parity drive to wear faster than the other drives in the set. Distributed parity RAID levels (RAID 5 and RAID 6) reduce this issue.

RAID 1 and RAID 5, as tested in this example, shows very consistent wear across all drives in the RAID sets. This is due to the tests using the full LBA space of the RAID set, thereby not creating any hotspot activity.

4.5 Selection of RAID Controller

There are many quality RAID controllers on the market today, with varying levels of performance, features and price points. Below are the important features considered in selecting the RAID controller for this sample test:

- RAID levels available
- Controller chipset
- PCIe* version
- SAS/SATA speed
- Internal/External ports
- Compatibility with SSDs

Within these categories, the LSI* MegaRAID 9265-8i was chosen for the following reasons:

- RAID 0, 1, 5, 6, 10, 50, 60 are supported
- LSISAS2208 Dual-Core RAID on Chip (ROC)
- 1 GB 1333MHz DDR3 SDRAM Cache
- x8 PCIe 2.0
- 6Gb/s per port
- 8 internal SAS ports
- SSD support
- Relative availability/popularity in the industry



5.0 RAID 1

5.1 Test System Specifications¹

The system used for RAID 1 testing include the following:

- Intel® R2208GZ4GC-IDD 2U rack mount server
- Intel® S2600GZ server board
- 2x Intel® Xeon® E5-2690 8-core CPUs (2.9 GHz)
- Intel® C602 chipset
- 192GB DDR3-1333 memory
- Microsoft Windows Server 2008 R2*, 64-bit
- LSI MegaRAID 9265-8i* controller card
- 2x up to 8x Intel® SSD DC S3500 Series 480GB drives

BIOS configuration changes:

- Hyper-Threading disabled²

RAID controller configuration:

- 256KB Striping (default)
- No Read-Ahead³
- Write-Through³
- Direct I/O

Windows Drive Configuration:

- Basic disk
- GUID partition table
- Simple volume
- Use full available space
- NTFS format

Test Software configuration:

- IOMeter 2009.10.22
- 1x worker per drive in RAID set⁴

Notes:

1. The system was selected to make sure the performance of the RAID card and the SSDs would not be inhibited by the server.
2. Hyper-Threading is disabled in this test system specifically due to additional latency introduced during benchmark testing. In any practical application, Hyper-Threading would NOT be disabled.
3. In the configuration of the RAID set, No Read Ahead and Write-through are used due to the speed of the SSDs. Read and write caching was designed for use with HDDs. Caching with SSDs introduces additional overhead thus interfering with the SSDs performance.
4. One thread, or worker, per drive was used in order to simulate the manner in which many applications utilize storage and also to attempt to saturate the communication channels to the SSDs.



5.2 Intel® SSD DC S3500 Series in RAID 1 Performance Characterization Data

This section provides performance characterization data for the Intel® SSD DC S3500 Series in RAID 1 configurations.

To establish baseline expectations for IOPS, the Intel SSD DC S3500 Series 480GB drives were evaluated in RAID 1 sets of 2, 4, 6 & 8 drives. The data collected was based on a different mix of read and write random and sequential workloads. Since higher queue depths can sometimes yield higher IOPS, queue depths of 1, 2, 4, & 8 per drive were chosen in the test setup. Multiple transfer sizes were tested, however, only selected data is presented here. All tests were done using the entire LBA range of the virtual drive. Tests were repeated at least twice to validate results.


Drives were prepared using IOMeter to fill the entire user area of the drive with data. Then, the first workload of each type (Sequential or Random) was 100% write performed for 120 minutes. Each subsequent workload was run for 12 minutes, with average IOPS collected over the last 10 minutes of the run.

The following figures show the different levels of performance for a selection of configurations.

Note: The scale of the IOPS charts is variable to clarify the changes that occur as drives are added.



Figure 1. RAID 1 Random 100% Write @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013


Notes:

Figure 1 - The write performance of the two drive RAID 1 set matches the write performance of a single Intel® DC S3500 drive. This indicates very low latency introduced by the RAID controller.

As more drives are added, the write performance scales linearly. At four drives, the performance is 2x that of a single drive, at six drives, it is 3x higher and at eight drives it is 4x. This is true at all queue depths tested and at transfer sizes from 4KB to 128KB. In this case, queue depth does not affect performance significantly.

At a queue depth of 1, the average latency for 100% write at 4KB transfer size is less than 200 μ s. Latency increases as the queue deepens, ending at 1.4ms for a queue of 8. It is interesting to note that latency is not affected by the number of drives.


Figure 2. RAID 1 Random 70% Read @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013




Figure 3. RAID 1 Random 90% Read @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013

Figure 4. RAID 1 Random 100% Read @ 4KB transfer size with Average Latency



Intel internal testing, October 2013

Notes:

Figures 2, 3 - In mixed workloads, 70% read and 90% read, IOPS increase with additional drives and show slightly exponential growth with deeper queues.

Figures 2, 3, 4 - Average latency for 70% read starts out similar to 100% write, but the progression is not as steep through deeper queues, ending between 500-600 μ S. Average latency for 90% read and 100% read continue to improve due to the higher speed of reads over writes.

Figures 2, 3, 4 - The latency of the two drive set is lower than other drive counts as read percentage increases due to the manner in which the LSI controller deals with the additional drives. This is expected behavior. For more information, please contact LSI for details.

Figure 4 - At 100% read the IOPS performance scales linearly. In other words, the IOPS for four drives is double that of two drives, six drives is triple that of two drives, and eight drives is four times that of two drives.




5.3 RAID 1 Consistency

Consistency behind a RAID controller is very important because the performance of any RAID set is limited by the lowest performing drive. As a RAID set increases in number of drives, the probability of any given drive performing poorly increases. Therefore, if the model of drive used is inconsistent in its performance, the inconsistency increases with the size of the RAID set.

The Intel® SSD DC S3500 Series drive has shown excellent consistency when used as a single drive. Figure 5 illustrates the consistent of DC S3500 in RAID sets. Notice that the maximum latency is grouped very tightly for both the two-drive and eight-drive RAID sets, indicating there is very little change in consistency as more SSDs are added.

Figure 5. RAID 1 Maximum Latency for 2-drive and 8-drive Configurations



Intel internal testing, October 2013

5.4 RAID 1 Performance Conclusions

RAID 1 is a very good choice for data needing robust replication. The RAID controller used shows good bandwidth with low latency causing little to no effect on read and write speeds of the SSDs. The linear scaling of read and write performance with additional drives shows that adding more drives would provide good ROI in most applications. The highest throughput seen in this test was 2300 MB/s during 100% read using eight drives with transfer size of 128KB and a queue of 8 per drive. This means the theoretical bandwidth limit of the PCIe lanes was not reached (4000 MB/s for x8 PCIe 2.0). It is theoretically possible that more than eight drives could be used and obtain an increase in performance, depending on the latency introduced by the necessary SAS expander.

The consistency of the drives is well demonstrated in these tests and shows that Intel SSD DC S3500 Series drives provide high performance with excellent stability, even behind a RAID controller.



6.0 RAID 5

6.1 Test System Specifications

The system used for RAID 5 testing was identical to the system used for RAID 1 testing except the following changes:

- 2x Intel Xeon E5-2680 8-core CPUs (2.7 GHz)
- 3x up to 8x Intel SSD DC S3500 Series 800GB drives

Note: For this test, 800GB drives were used. The rated performance of the Intel SSD DC S3500 Series drive in 800GB, 600GB, and 480GB capacities are nearly identical, per internal Intel testing.

6.2 Intel SSD DC S3500 Series in RAID 5 Performance Characterization Data

This section provides performance characterization data for the Intel SSD DC S3500 Series in RAID 5 configurations.

To establish baseline expectations for IOPS, the Intel SSD DC S3500 Series 800 GB drives were evaluated in RAID 5 sets of 3, 4, 5, 6, 7 and 8 drives. The data collected was based on a different mix of read and write random and sequential workloads. Since higher queue depths can sometimes yield higher IOPS, queue depths of 1, 2, 4, and 8 per drive were used in the test setup. Multiple transfer sizes were tested, however, only selected data is presented here. All tests were done using the entire Logical Block Address (LBA) range of the virtual drive. Tests were repeated at least twice to validate results.


Drives were prepared using IOMeter to fill the entire user area of the drive with data. Then, the first workload of each type (Sequential or Random) was 100% write performed for 120 minutes. Each subsequent workload was run for 12 minutes, with average IOPS collected over the final 10 minutes of the run.

The following figures show different levels of performance for a selection of configurations.

Note: The scale of the IOPS charts is variable in order to clearly show the change as drives are added.



Figure 6. RAID 5 Random 100% Write @ 4KB Transfer Size with Average Latency




Intel internal testing, October 2013

NOTES: There are gains in write performance as drives are added to the RAID 5 set. The change at queue depth 1 from three drives to six drives is approximately 58% increase in IOPS. For eight drives, the change is 97% increase in IOPS over the three drive set.

At a queue depth of 1, latency increases as more drives are added to the RAID set, most likely caused by the additional overhead of calculating parity and striping across more drives. This increases as the queue deepens.


Figure 7. RAID 5 Random 70% Read @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013




Figure 8. RAID 5 Random 90% Read @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013

Figure 9. RAID 5 Random 100% Read @ 4KB Transfer Size with Average Latency



Intel internal testing, October 2013

Notes:

Figures 7, 8, 9 - As the workloads become more read intensive, there is a steady increase in performance both as drives are added and as the queue deepens.

Figures 7, 8, 9 - As read percentage increases, the exponential increase in latency is not as prominent with deeper queues. This is due to the speed at which reads are performed.


6.3 RAID 5 Consistency

Consistency behind a RAID controller is very important because the performance of any RAID set is limited by the lowest performing drive. As the number of different drives in a RAID set increases, so does the likelihood of any given drive performing poorly. Therefore, if the model of drive used is inconsistent in its performance, the inconsistency increases with the size of the RAID set.

The Intel® SSD DC S3500 Series drive has shown excellent consistency when used as a single drive. Figure 10 illustrates that the DC S3500 is also very consistent in RAID sets. The three-drive and eight-drive data shows the maximum latency is grouped very tightly; indicating that adding more drives would have little impact on consistency.



Figure 10. RAID 5 Maximum Latency for 3-drive and 8-drive Configurations



Intel internal testing, October 2013

6.4 RAID 5 Performance Conclusions

The RAID 5 write performance data illustrates the additional processing power required of the RAID controller to calculate parity and stripe data across multiple drives. There is diminished performance gain by adding drives when compared to RAID 1. Intel's data also shows that in mixed workloads and in pure reads, RAID 5 performs well, reaching over 300K IOPS in 100% read at a queue of 8 per drive on eight drives. As the workloads become more read heavy, latency drops from a high of 2.2 ms (100% write) to a low of 140 μ s (100% read). The highest throughput achieved was 2400 MB/s with eight drives, 100% read, 128KB transfer size, and queue depth of 8. This leaves room for possible improvement by adding more drives to the RAID set.

In configurations where RAID 5 would traditionally be used, SSDs would provide significant performance gain over HDDs. Additionally, RAID 5 with SSDs could be used in situations where RAID 5 with HDDs would not perform well.

The consistency of the drives is well demonstrated in these tests and shows that Intel® SSD DC S3500 Series drives consistently offer higher performance with excellent stability, even behind a RAID controller.



7.0 Summary

The Intel® SSD DC S3500 Series drive has proven itself in many applications where speed and reliability are essential. The data in this paper shows that this drive is very robust behind a RAID controller.

In RAID 1 configurations, write performance is as expected for this setup; that is, a two drive set matches the single drive specification, and increases linearly as drives are added to the set. With eight drives at queue depth 1 per drive, this configuration processes over 50K write IOPS with 4KB blocks. The RAID controller adds very little latency, but latency does increase as the queue depth grows and, as more drives are added to the array. Read performance increases with additional drives and with queue depth, reaching over 200K read IOPS with eight drives, a queue of 8 per drive and 4KB blocks. More importantly, the performance on mixed workloads was excellent and increased as more drives were added. The 70% read workload topped out at close to 110K IOPS with 4KB blocks, and the 90% went to nearly 150K IOPS with 4KB blocks (both at queue of 8 per drive). The highest throughput seen in this test was 2300 MB/s during 100% read using eight drives with transfer size of 128KB and a queue of 8 per drive. This means the bandwidth limit of the PCIe lanes was not reached (4000 MB/s for x8 PCIe 2.0). It is theoretically possible that more than eight drives could be used and an increase in performance obtained.

In RAID 5 configurations, write operations increase both as drives are added and as the queue depth increases. The increase in write performance with queue depth is somewhat surprising and is most likely attributed to the RAID controller and the scaling effect of the cache in each drive. This may be due to the way the controller writes the stripes to the drive set, possibly consolidating the 4KB blocks into the 256KB stripes. Read performance is also very good, with eight drives reaching 300K IOPS with 4KB blocks and queue depth 8 per drive. Mixed workloads show very good performance with 70% read hitting 75K IOPS at queue of 8 per drive and 90% read coming in at almost 140K, both with 4KB blocks. Latency on all workloads is very manageable, although, as the queue depth increases, so does the latency. The graphs show that as queues grow, latency increases at an increasingly higher rate.

To summarize:

- In both RAID 1 and RAID 5, the Intel SSD DC S3500 Series drive shows excellent scalability, performance, and consistency.
- Very little latency was introduced by the RAID controller in RAID 1. In RAID 5, the overhead and latency are slightly higher.
- In random, mixed read/write workloads, SSDs perform significantly (as much as 100 times) better than HDDs in a similar situation.
- With this RAID controller, there is the possibility of greater performance by adding more than eight drives in both RAID 1 and RAID 5 configurations.



8.0 Appendix

8.1 RAID Levels

RAID (Redundant Array of Independent Disks), developed in 1988 to improve performance, reliability and scalability of hard disk storage systems has become a standard in datacenters because of these qualities. There are many types, or levels, of RAID.

RAID 0 uses block level striping to span one or more drives. This does improve performance, and increases capacity when more than one drive is used. However, there is no fault tolerance, so failure of any one drive will cause full data loss.

RAID 1, also called mirroring, writes data identically to two drives, producing a mirrored set. Reads can be serviced by either drive, and writes occur in unison on both drives. If one drive has a hardware failure, the data is protected in the mirrored copy. RAID 1 requires two drives. Many modern RAID controllers support RAID 1 sets of more than two drives, however, the original specification was for only two. Because of the 50% overhead, RAID 1 is the most expensive RAID type.

RAID 2 uses bit-level striping with dedicated Hamming-code parity. This is a theoretical model and not used in practice.

RAID3 uses byte-level striping with dedicated parity. This level is not commonly used.

RAID 4 uses block-level striping with dedicated parity. All parity data is on a single drive. I/O requests are handled in parallel, increasing performance.

RAID 5 uses block-level striping with distributed parity. Data and parity are distributed among all drives and requires all but one drive to be present. RAID 5 requires at least three drives and can survive a single drive failure.

RAID 6 uses block-level striping with double distributed parity. Identical to RAID 5 in the way it writes data, however, parity is written twice in different locations. RAID 6 can survive 2 drive failures; therefore it is often used for larger sets of drives.

RAID levels can also be nested for improved performance or fault tolerance. RAID 10, 0+1, 50 and 60 are common combinations.