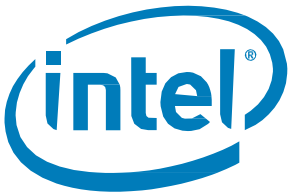

# Intel® Solid-State Drive DC S3700 Series – Quality of Service

***Technology Brief***

July 2013
Order Number: 329281-001US

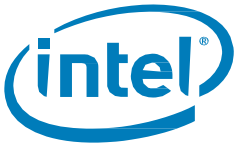

# Contents

# 1. Executive Summary

This paper covers the Quality of Service (QoS) and performance consistency of the Intel® SSD DC S3700, hereon referred to as S3700. It introduces the metrics on which Quality of Service and performance consistency are measured and characterized. It also focuses on the importance of QoS in multiple data center applications.

# 2. Quality of Service(QoS)

Quality of Service of an SSD consists of the predictability of low latency and consistency of high IOPS while servicing read/write I/O workload. This means the latency or the I/O command completion time needs to be within a specified range without having unexpected outliers. Throughput or I/O rate also needs to be within a tight regulation without causing sudden drop in performance level. QoS metric needs to be set for applications ranging from read intensive to write intensive workload.

Latency distribution of an SSD needs to hit certain service level, which can be measured by the latency of an application workload for 99.9% of the time or on a tighter scale 99.99999% of the time. SSD that optimizes around this metric will exhibit superior level of performance predictability.

Performance Consistency of an SSD is defined by the difference of average and the lowest IOPS during application run time. The ratio between the 99.9th slowest 1 second interval and the average input/output per second (IOPS) performance during the test can be used as a metric.

For example, performance consistency can be measured using the following data that was collected for 2000 seconds on an SSD.

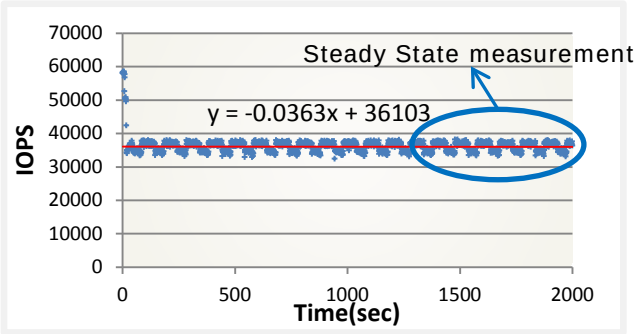


99.9th slowest 1-second interval = 34101 IOPS

Performance consistency = 99.9th slowest 1-second interval/Average IO per second

=[1 - (36103 – 34101)/36103] *100

= 94.5%

**Figure 1: Performance Consistency Illustration**

Once the steady state is reached in the last 600 seconds, we determine the 99.9th slowest 1-second interval which in this case returns 34101 IOPS and calculate deviation from the average. In the example above, there is 5.5% deviation from the average which translates to 94.5% consistency.

# 3. Quality of Service in HDD compared to SSD

The average Hard Disk Drive with rotating media will have an average latency of 5ms to 10ms. This latency consists of the rotational latency of the media and the actuators seek time. SSDs do not suffer from rotational and mechanical seek latencies which greatly improving application performance by reducing read/write latencies. Moving from an HDD based solution to an SSD based solution will provide magnitude of difference in latency time for applications. Typically SSDs are rated for less than 100usec of average latency. Similarly read/write IO throughput time for random access is more than 100x better for SSDs. For QoS, consistency of IOPS is important. Rotating Hard Disk Drives (HDD), such as 10K or 15K rpm drives, maintain a Logical Block Address (LBA) fixed to a specific position in the platter. This provides predictable access time for a given LBA during multiple rounds of operation. This helps to maintain consistency, however the level of IOPS is pretty low and latency very high making them unsuited for high performance applications.

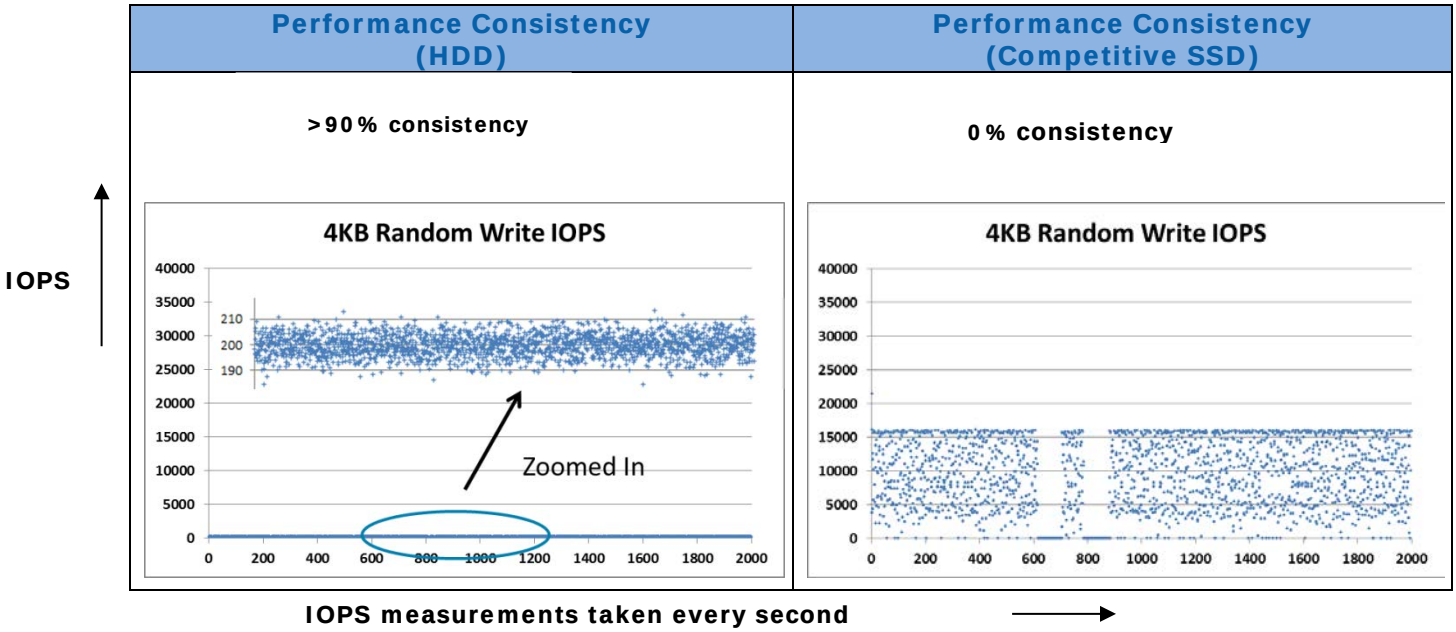


**Figure 2: Performance Consistency of an HDD compared to a competitive SSD**

On the other hand, performance consistency of many SSDs available in the market today is not as predictable as shown in Fig 2. Typically SSDs do not have a fixed LBA to physical location mapping which may cause difference in performance run over run. The background tasks such as wear leveling and reclaiming spare space can often consume much of the SSD's internal bandwidth, temporarily starving host IOs and creating significant performance and latency variation. NAND based SSDs also require additional background tasks to deal with technology limitations such as retention and disturbs which can also impact QoS.

Many datacenter applications benefit from a fast and consistent performing storage sub-system. Intel SSD DC S3700 was engineered not only for better

performance but also for superior QoS and performance consistency to meet the needs of the datacenter market segment.

# 4. Intel SSD DC S3700 – QoS and Consistency Metrics

## 4.1 QoS

The S3700 QoS is best-in-class with 4K Random reads/writes taking less than 0.5 ms (500uS) 99.9% of the time for Queue Depth (QD) of one as listed in Table 1. Product is also specified to meet 99.9999% latency targets.

| Specification | Unit | Intel SSD DC S3700 | | | |
|---|---|---|---|---|---|
| | | Queue Depth=1 | | Queue Depth=32 | |
| | | 100 GB | 200/400/800 GB | 100 GB | 200/400/800 GB |
| **Quality of Service[1,2] (99.9%)** | | | | | |
| Reads | ms | 0.5 | 0.5 | 1 | 1 |
| Writes | ms | 0.5 | 0.5 | 15 | 10 |
| **Quality of Service[1,2] (99.9999%)** | | | | | |
| Reads | ms | 10 | 5 | 10 | 5 |
| Writes | ms | 10 | 5 | 20 | 20 |

**Table 1: Quality of Service Specification on a 4KB workload**

## 4.2 Performance Consistency

Table 2 lists the performance consistency of the S3700. The S3700 performance consistency is greater than 90% for 4KB reads/writes on 200GB, 400GB and 800GB capacities and greater than 85% on 100GB capacity. Quality of Service combined with best-in-class performance consistency makes S3700 desirable for datacenter applications.

| Specification[4] | Unit | Intel SSD DC S3700 | | | |
|---|---|---|---|---|---|
| | | 100 GB | 200 GB (2.5"/1.8") | 400 GB (2.5"/1.8") | 800 GB |
| Random 4 KB Read (up to)[3] | % | 90 | 90 | 90 | 90 |
| Random 4 KB Write (up to) | % | 85 | 90 | 90 | 90 |
| Random 8 KB Read (up to)[3] | % | 90 | 90 | 90 | 90 |
| Random 8 KB Write (up to) | % | 85 | 90 | 90 | 90 |

**Table 2: Measured performance consistency on a 4KB workload**

1. Device measured using Iometer. Quality of Service measured using 4 KB (4,096 bytes) transfer size on a random workload on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability.
2. Based on Random 4KB QD=1, 32 workloads, measured as the time taken for 99.9(or 99.9999) percentile of commands to finish the round-trip from host to drive and back to host.
3. Performance consistency measured using Iometer* based on Random 4KB QD=32 workload, measured as the (IOPS in the 99.9$^{th}$ percentile slowest 1-second interval)/(average IOPS during the test). Measurements are performed on a full Logical Block Address (LBA) span of the drive once the workload has reached steady state but including all background activities required for normal operation and data reliability

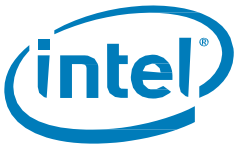

# 5. QoS in Transactional Processing

QoS is an important metric for transactional databases. With better QoS and consistency, it is possible to scale up user base without impacting service level agreements. An example case can be found in an Oracle* Timesten* database case study published here - http://www.intel.com/content/www/us/en/mission-critical/mission-critical-scalability-oracle-intel-brief.html.

To demonstrate the importance of QoS, Intel devised a test using a proprietary transaction processing tool to inject different amount of latencies at varying frequencies. The resulting change in transactions were measured and listed in Table 3. A delay of 90ms for every 10,000 writes (1/10000 = 99.99% of the time) translates to a reduction in the transactions per second by 28% and an increase in average latency by 155%. However, if the frequency is reduced to every 1,000,000 writes (1/1000000 = 99.9999% of the time), the impact is significantly less.These measurements demonstrate the overall performance impact that can result due to size and frequency of SSD latency delays and helps to validate the importance of SSD QoS metrics.

| Size of delay | Number of writes between delay | Impact to transactions per second | Impact to average latency |
|---|---|---|---|
| 0 ms | N/A | 0% | 0% |
| 90ms | 10,000 | -28% | 155% |
| 90ms | 100,000 | -5% | 22% |
| 90ms | 1,000,000 | -1% | 3% |
| 50ms | 10,000 | -15% | 96% |
| 20ms | 10,000 | -3% | 43% |
| 10ms | 1,000 | -7% | 120% |
| 10ms | 10,000 | -1% | 12% |

**Table 3: Impact of QoS in Database Transactional Processing Applications**

# 6. Intel SSD DC S3700 QoS and Consistency Plot

This section highlights the variation of QoS and performance among the S3700. Figure 3 & 4 shows the QoS and performance consistency of S3700.

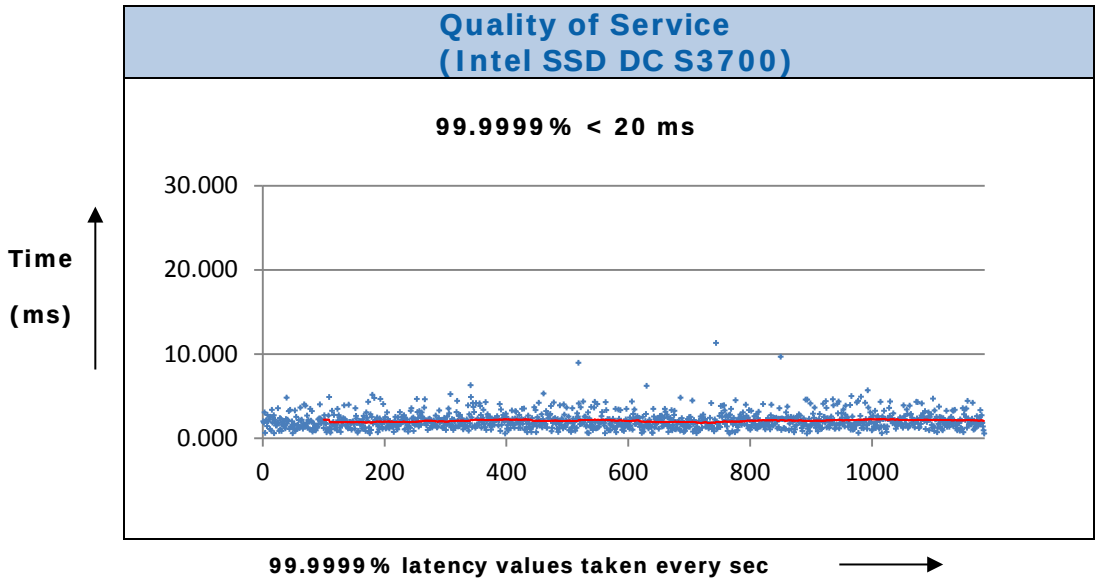


**Figure 3: QoS of Intel S3700 on a 4K write workload, QD = 1**

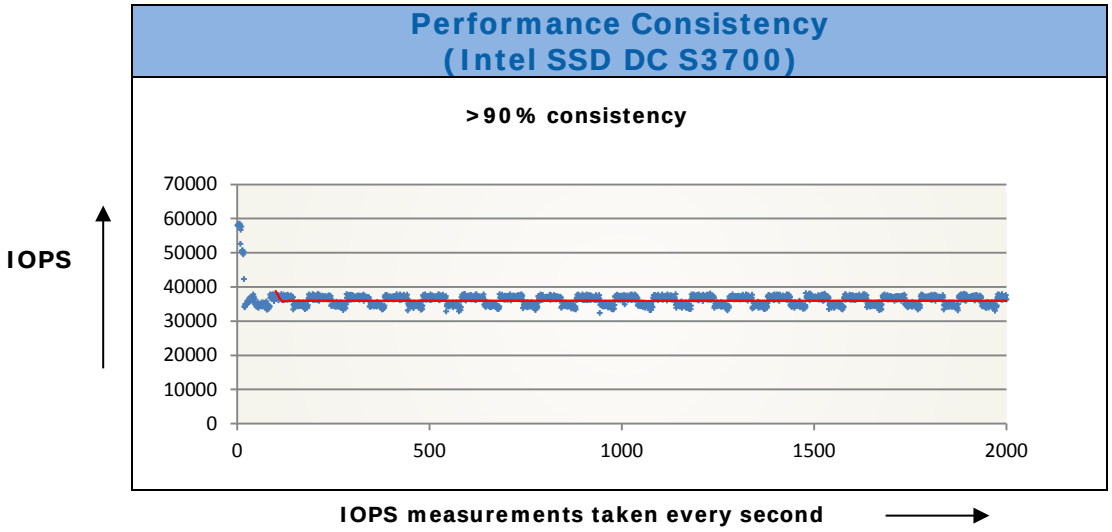


**Figure 4: Performance Consistency of Intel S3700, QD=32**

The S3700 finished 99.9999% of all the commands within 20ms to complete 36,000 write commands on a 4KB write workload in a given second. The consistency measured is greater than 90% on this workload. Intel SSD DC S3700 is designed for best-in-class Quality of Service and performance consistency as specified in table 1 and 2.

## 7. Applications of QoS and Performance Consistency

QoS and performance consistency is an important metric in Data Center, Cloud and virtual server storage environments. Many data center applications, such as databases, High Performance Computing (HPC), Big Data, On-Line

Transactional Processing (OLTP), web services, mail exchange, and data warehouse benefit from improved QoS and performance consistency.

## 7.1 Cloud Storage

Two decades back, the different segments of datacenters had their own hardware resources provisioned to meet the peak demand of their work load. In the last decade, cloud based services have grown at a phenomenal rate, as more and more workloads share common data center resources. This growth has created a burden on the hardware budget, datacenter space and power capability. Cloud applications are random in nature and can easily create I/O bottleneck problems with traditional storage system. Once the I/O bottleneck is removed, the full potential of the servers can be realized. Doing so enables customers to create different tiers of service to their end-users with different Service-Level-Agreements (SLAs). For example, a premium tier with DC S3700 and a standard tier can be created with a standard SSD. Different solutions such as High Performance Computing (HPC), Database (DB), Internet Portal Data Centers (IPDC) and Financial Service Institute (FSI) will be direct beneficiaries of this performance tier enabled through DC S3700 as they are IO bound today. The S3700 helps in removing the storage bottleneck when customers create tiered solution for their end-users.

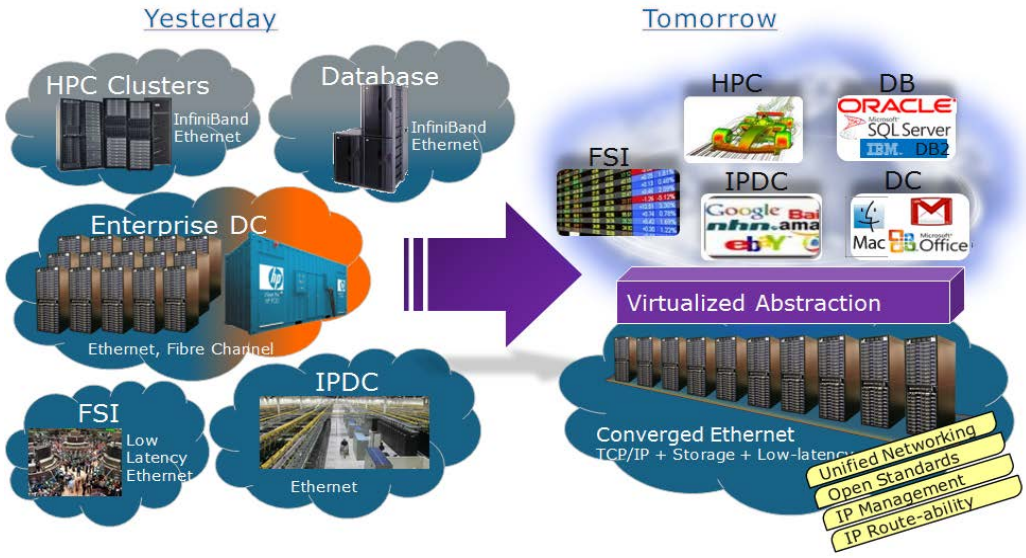


**Figure 5: Convergence to Cloud increases the need for better QoS and consistency**

## 7.2 Virtualized Environment

In today's datacenters, for many applications the performance is not bounded by CPU processing, but by the I/O latency and throughput of the storage subsystem. In virtualized environment, the need for random fast access is high and traditional storage is incapable of meeting the need.

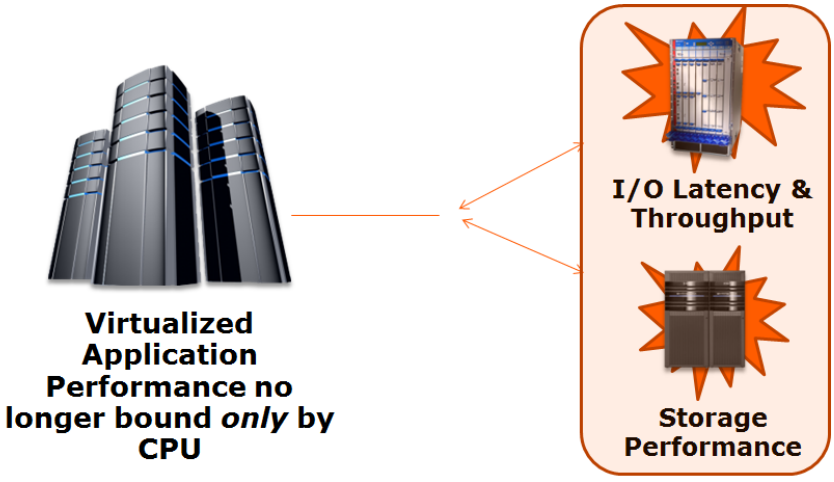


**Figure 6: QoS and consistency for Storage—an important factor for Virtualized applications**

The superior QoS and consistency, of the S3700, not only breaks the storage I/O latency and throughput bottleneck, but also increases the number of Virtual Machines(VM) that we can run in a given datacenter footprint. When compared with higher capacity HDDs that have been "short stroked" to achieve improved performance, the S3700 offer an alternative at a lower price point thereby giving a better overall TCO for the system. The S3700 is best suited for Virtualized Environment as it opens up the full I/O bandwidth with consistent performance and QoS.

## 7.3 Scale-Out Storage

Data centers can respond to the dramatic growth in data by combining traditional storage, with new types of storage specifically designed to respond to high performance needs by a Scale-Out Storage Architecture. Proximity based storage is well suited for applications where workloads performed on large amounts of data will perform better when co-located with the storage of that data. These include business processes, decision support analysis, and High-Performance Computing on dedicated servers, together with collaborative processes, applications and web infrastructure on virtualized servers.

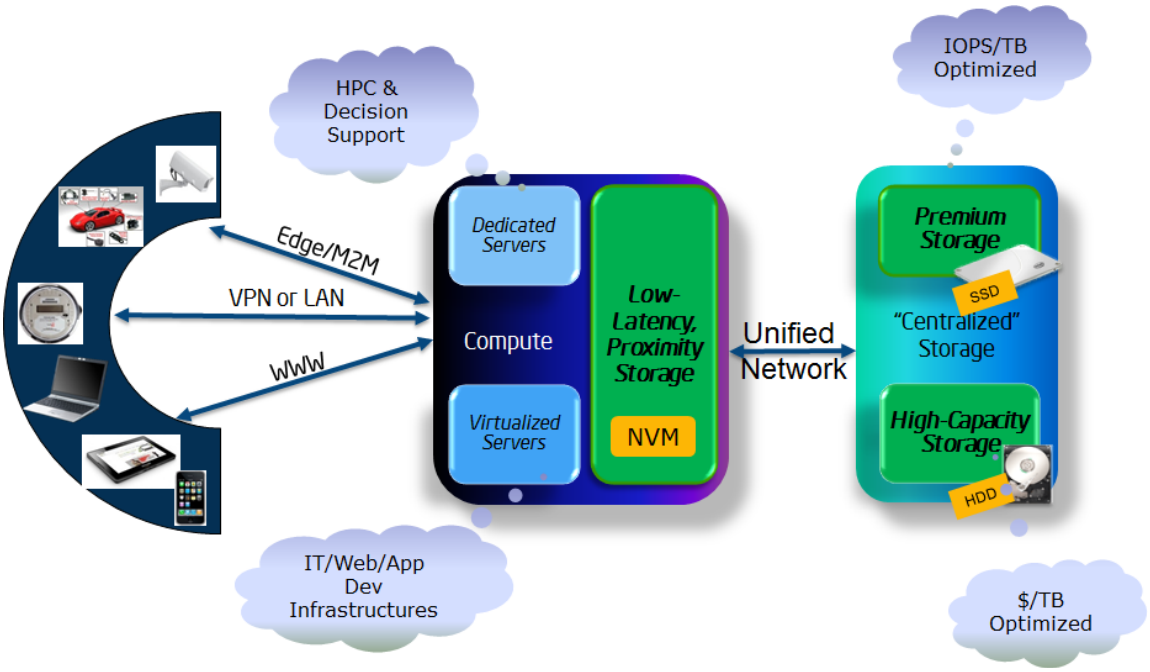


**Figure 7: QoS and consistency preferred in proximity storage**

The S3700 provides consistent high performance QoS which is highly desirable for proximity or direct attached storage. For example, employing S3700 in a Big Data solution such as Hadoop in either temp/shuffle/sort storage or in a dedicated development or real-time query cluster as primary storage can greatly accelerate map-reduce jobs. This change, combined with higher frequency CPUs and 10 Gigabit Ethernet can further accelerate Hadoop jobs, or make possible much smaller datacenter footprints and power/cooling expense.

## 7.4 Other Applications

The S3700 QoS and performance consistency also benefits other application such as the local VM cache for VM hosts, scratch pads for databases and HPC workloads for quick computations, journaling or metadata drives for distributed storage solutions. The metadata and the journals change frequently and if they are isolated to a high performing, low latency cache drive such as the S3700, the overall performance of the storage solution improves.

IT system engineers can look to the S3700 to meet a variety of high performance storage system needs. The application of the S3700 will help deliver higher SLAs for modern datacenter applications.

# 8. Conclusion

The Intel® SSD DC S3700 has superior Quality of Service as measured by consistent low latency IO operations and superior IOPS. The S3700 enables the cloud service providers to define Service Level Agreement (SLA) with their customers based on performance tiering of their storage. Solution providers for High Performance

Computing, Databases, and the Financial Services segment are among many segments that can take advantage of consistent performance. The S3700 breaks the I/O bottleneck thus enabling the Virtualized Environments to run more Virtual Machines for the same data center footprint. The S3700 provides the best proximity solution for Scale-Out Storage enabling complex computational tasks such as Hadoop clusters crunching of large data sets. The Intel SSD DC S3700 provides a superior solution for today's complex datacenter needs, unleashing the full potential of every datacenter.

# 9. Reference Documents

1) Product Specification

   http://download.intel.com/newsroom/kits/ssd/pdfs/Intel_SSD_DC_S3700_Product_Specification.pdf

2) New Levels of Scalability for Real-Time Analytics and Transactions

   http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/ssd-dc-s3700-oracle-intel-paper.pdf

# 10. Additional Information

For more information on Intel SSDs, go to: intel.com/go/ssd.

# 11. Revision History

| Date | Revision | Description |
|---|---|---|
| July 2013 | 001 | Initial release |