ROBERT FRANCES GROUP
Business Advisors to IT Executives
46 Kent Hills Lane, Wilton, CT 06897 Phone: 203-429-8951

The SSD Revolution

RFG Perspective: The solid-state drive (SSD) Revolution has begun. The economic value of SSDs exceeds high performance hard disk drives (HDDs) in cost and price/performance by reducing total costs by more than 70 percent while magnifying performance by up to 37x. Hence, over the next few years the high performance HDD install base will spiral downward as enterprises move to server side and cache controllers for high performance data access. IT executives need to understand the SSD value proposition – the return on investment (ROI), total acquisition cost (TCA), and total cost of ownership (TCO) – for their various target architectures and develop roadmaps for transforming these architectures to be SSD driven. A no-brainer first step IT executives should consider is employing SSDs for server side caching to improve performance without disrupting current HDD shared-storage SAN environments.

The low latency requirements of applications such as real-time analytics, database queries, online transaction processing (OLTP) and infrastructure plays such as virtual desktop infrastructure (VDI) are fueling the demand for higher performance from storage subsystems. SSD or solid state drive (SSD) alternatives can provide users with up to a 50x latency performance improvement that can enhance the way businesses interact with customers, thereby increasing revenues, improving loyalty and minimizing risk exposures. One major provider leading the way in this market is Intel with its SATA-attached Data Center Solid-State Drive S3700 storage or PCIe-attached P3700 storage offerings combined with its Cache Acceleration Software (CAS) and Intel Raid Storage Technology enterprise (RSTe) software products. The value of caching software is that it allows applications to benefit from the performance of non-volatile flash memory (NVM) SSDs without any code changes. The value of RSTe is that is allows customers the benefits of RAID technology on non-volatile flash (NVM) SSDs without the additional latency or costs associated with traditional RAID host-bus adapters. These solutions enable IT executives to slash acquisition and operational costs and resources while providing the business with a potential competitive advantage.

The Storage Market - HDD

The traditional HDD market consists of two different subsets with dissimilar characteristics: high performance disk drives and capacity HDDs. High performance HDDs are designed for more and quicker responses to data requests than the capacity architected drives, which means limited data storage and higher RPM devices. On the other hand, capacity HDDs are suited for holding large volumes of data, which translates to fewer and slower responses to requests. This is summarized in the following chart:

Drive type	RPMs	Capacity	IOPS (I/Os per second)	Drive Latency	Space utilization	SAS/SATA	Price
High performance HDDs	15k	146 – 300 GB	175 – 300	2 – 2.5 ms	30 - 50%	SAS	~$2.05/GB
	10k	300 – 900 GB	100 - 150	3 – 3.5 ms	30 - 50%	SAS/SATA	~$0.72/GB
Capacity HDDs	7.2k	4 - 6 TB	75 - 100	4 – 5 ms	70 – 80 %	SATA	~$0.23/GB

As can be implied from the chart, there are multiple costs associated with the need for high performance drives. To achieve greater performance disk makers had to trade off capacity for performance and users had to extend that tradeoff by short-stroking the disk drives. This tradeoff resulted in sprawling storage farms that drove up administration, power, software, and space costs. To solve the problems caused by storage sprawl and to satisfy the multiple, disparate storage requirements vendors constructed a storage tiering architecture that addressed the various needs.

- **Tier 1 storage** consists of the highest performance drives and is used for critical business transactional applications. These are the most expensive 15k RPM drives that cost approximately $1/GB.
- **Tier 2 storage** refers to drives used for general business applications such as emails. Subsecond responses here are not required but speed is still important. The 10k RPM drives, which cost much less than the tier 1 drives, fill this need.
- **Tier 3 storage** is for non-critical business applications where response time is less important. The low cost capacity HDDs fit this scenario.
- **Tier 4 storage** holds large volumes of data that is usually kept for historical purposes. High-density SATA drives or tape are normally used for these use cases.

In all cases the cost metric based on $/GB or $/TB of storage became the *de facto* standard for evaluating the value proposition.

SSD – A Storage Game Changer

While SSD storage has been around for quite a while, it did not gain general acceptance in enterprise systems until the last half of the last decade. Initially SSD appeared as Tier 0 storage and was used as SSDs in storage arrays, replacing a set of the HDDs. But over the past few years SSD storage forked into multiple tiers:

- **Storage array side storage** in either hybrid HDD/SSD environments or all SSD storage arrays
- **SSD controllers** that enabled storage to be shared amongst servers though network or clustered systems
- **Server-side storage** in the form of memory extensions or PCIe cards

The addition of SSD to the storage ecosystem revolutionized storage systems and how to evaluate them. These solutions provide a quantum leap in performance capabilities (ranging from 14x to 50x improvement) and are doubling in size almost annually while the pricing plummets. Moreover, since all of the SSD is usable without a performance hit and its usage is for performance purposes, a total capacity to capacity comparison is not a meaningful metric. When looking at SSD and comparing it to HDD arrays, $/IOPS is a more appropriate measurement. The following charts summarize SSD characteristics and offer an SSD vs HDD comparison:

Drive type	RPMs	Capacity	IOPS R/W	Drive Latency	Space utilization	Maximum performance	Price	Price/IOPS
High performance HDDs	15k	146 – 300 GB	175 – 300	2 – 2.5 ms	30 - 50%	600 MB/s	~$2.05/GB	$2.05
	10k	300 – 900 GB	100 - 150	3 – 3.5 ms	30 - 50%	600 MB/s	~$0.72/GB	$4.32
Capacity HDDs	7.2k	4 - 6 TB	75 - 100	4 – 5 ms	70 – 80 %	600 MB/s	~$0.10/GB	$5.78
SATA DC S3700	N/A	100, 200, 400, 800GB	75k/19k (100GB), 75K/36K (800GB)	50us R 65us W	80%	500MB/s read, 460MB/s write	~$1.75/GB	$0.02
PCIe – P3700	N/A	400, 800, 1600, 2000GB	450k/75k (400GB), 450K/175K (2000GB)	20us	>80%	2800MB/s read, 2000MBps write	~$3/GB	$0.01

Today enterprise-class storage-side SSDs can be acquired for as little as $1.75/GB (although $2/GB is a more realistic price) while server-side SSDs can be bought for less than $3/GB. Interestingly, users can expect the capacity to double every year over the next three years while the price is projected to drop 25 percent per year. On the other hand, high-performance HDDs costs a little less than half the price of SATA SSDs on a $/GB basis. But this is a false comparison – it is like comparing typewriters to word processors. While data entry on either device consumes about the same time and effort, word processors are far more efficient due to its editing capabilities, spell and grammar checking, reuse and archiving functions. Thus, while typewriters were cheaper, they have virtually disappeared and replaced by personal computers with document productivity tools.

Most recent to market is the directly connected SSD PCIe card. It outshines HDDs in multiple dimensions. It is able to execute data reads of up to 450,000 IOPS while being totally independent of the operating system driver stack and any requirement for a host bus adapter (HBA). The Intel P3700 can hold up to 2 TB of data and has a bandwidth transfer rate of 2.8 GB/second using four PCIe GEN3 lanes. Furthermore, the SSD PCIe cards consume less than 10 percent of the power of HDDs and have a significantly longer mean time to failure rate. For most databases this translates to being able to put all the data on the server and slashing administrator and operating costs, while for applications with a number of large databases, it means putting all the hot data on the PCIe and drastically cutting storage administrator and operating costs. For example, At 200 IOPS per HDD, 2,300 15K 300 GB HDDs would be needed to match the performance of one Intel DC P3700. Or more realistically, 15 TBs for data on 100 short-stroked 15K 300 GB HDDs could be replaced by as few as four PCIe SSD cards. Additionally, in both cases it means more workload handled per server and per administrator.

Thus, the first price/performance metric to consider when evaluating high performance storage is $/IOPS, or $/relative IOPS. As shown in the previous table, SSDs offer a 100x improvement in cost per IOPS. Moreover, since SSD storage revolutionizes storage topography, one must also evaluate server workload and administrator impacts and floor space and power savings. The below TCO examples demonstrate the quantum impact SSD is having on where data is stored and its impacts upon the associated workloads.

Storage Hierarchy Management

A fallout from the creation of all the storage tiers was the need for an efficient data management methodology that placed highly used data in the fastest SSD storage tier and allocated the low priority data to appropriate tiers based on use and need. The storage management tool had to also swap data from one tier to the next so as to accommodate the ever-changing access demands on individual records. Intel's CAS is the software that performs this function for any of the Intel SSD options.

For those organizations that want to take advantage of SSD solutions, one of the easiest, non-disruptive approaches is to add a CAS SSD combination to an existing HDD environment. This can be installed transparently to the applications and the performance gains will be felt immediately. For those applications that require the least latency server side PCIe SSD provides the most economical, best performing solution – with a 50:1 performance ratio over HDD.

The Economics of SSD - Example

There is an above reference to a comparative PCIe vs HDD environment where 15 TBs of data resides on 100 short-stroked 15K 300 GB HDDs that could be replaced by four or six PCIe cards, using RAID 5 and RAID 10 respectively. The known variable costs[1] associated with these configurations are shown on the chart on the following page. There are three key takeaways that can be gleaned from the TCO analysis. First is that for less than a four percent increase in costs an existing system can improve performance 9x by adding one PCIe card to the mix. Secondly, a migration from traditional HDDs to six PCIe cards using the same RAID 10 storage recording technology cuts the costs by almost 60 percent while increasing performance by 37x. Lastly, by using the more efficient RAID 5 method one can obtain the same performance level (37x better) while reducing the TCO by 3.5x (72 percent savings). Or another way to look at it is that for little more than the cost of upgrading 20% of the HDD infrastructure to the current technology, one can replace the entire HDD farm with SSDs. Thus, it is easy to see why, with the economics so favorable to SSDs, the high performance HDD market will fade away over the next few years.

[1] For the purposes of this paper the server, deployment, migration, networking personnel, development/test systems and procurement costs are ***not included***. This is a steady-state comparison of the impacted storage ecosystem from acquisition through total cost of operations.

	PERFORMANCE			
Intel CAS-DC P3700 + SATA vs all HDD	*good*	*better*	*best - alternative 1*	*best - alternative 2*
Configuration Cost Breakdown	*All HDD*	*HDD+ (1) Intel P3700 + Intel CAS*	*ALL Intel P3700 + Intel RSTe*	*ALL Intel P3700 + Intel RSTe*
Storage SubSystem Configuration	100 x 300 GB 15K SAS HDDs	100 x 300 GB 15K SAS HDDs + 1 x Intel 3700 PCIe SSD 800GB	6 x Intel 3700 PCIe SSD 2.0 TB	4 x Intel 3700 PCIe SSD 2.0 TB
RAID Level for data protection (RAID10-50% for mirroring/stripping, RAID5- 1 drive of space used for parity)	RAID10	RAID10	RAID10	RAID5
Total Protected Storage (HDD-overprovisioned)	15TB	15TB	6TB	6TB
Amount of active capacity allocated (short stroked HDD, 40% of total) (100% SSD utilized)	6TB	6TB	6TB	6TB
% of Hotdata that could benefit from a caching solution?		10%		
Size of caching device based on capacity used and % of of hot data		600GB		
Drive Costs	$32,900	$35,399	$38,388	$25,592
Enclosure Costs	$11,600	$11,600	$0	$0
HBAs	$814	$814	$0	$0
Additional Software (Intel CAS or Intel RSTe)		$300	$150	$150
Storage SubSystem Costs (drives,enclosures,hba,software)	$45,314	$48,113	$38,538	$25,742
Maintenance and support (est 10%/yr)	$4,531	$4,811	$3,854	$2,574
Annual Facilities Cost (space for enclosures)	$16,000	$16,000	0	0
Storage Power Costs(24hr/365dy, $.107/KW)	$4,124	$4,148	$141	$94
Total Costs (Storage HW+ Power+Space) over 3 yrs	$119,281	$122,990	$50,521	$33,746
Costs (Savings) after 3 years due to lower power			($68,760)	($85,535)
Max 4K Read IOPs (Theoretical)**	21,000	368,000	1,350,000	1,350,000
Theoretical Performance Increase	1	18	64	64
Performance Impact based on DB Workloads: Single P3700 used in test	1	9X	37X	>37X

Summary

SSD alternatives should cause IT executives to reevaluate their storage cost and price/performance metrics just like virtualization forced re-evaluation of server investment returns. Once business and IT executives experience the performance gains SSD delivers and understand the new pricing paradigm, there is no doubt that SSD will become the primary high performance storage platform for business applications.

RFG POV: SSD, in all its forms, will be the high performance storage device of choice and supplant high performance HDDs within the next three to five years. The high performance HDD share of market will slowly spiral downwards and eventually disappear. Server-side SSD will dominate systems that are not clustered or operate independently while storage-side SSD will be predominant in clustered or networked systems or in shared storage environments. To remain competitive IT executives need to adapt to the new technologies and begin now to pilot and incorporate SSD solutions into their storage architectures. Furthermore, IT executives should develop self-funding roadmaps for adopting SSD without disrupting current ecosystems.

Intel Corp. sponsored this report. This document exclusively reflects the analysis and opinions of Robert Frances Group (RFG), who has final control of its content.