TANEJA GROUP
TECHNOLOGY ANALYSTS

TECHNOLOGY VALIDATION

STORTRENDS 3500i

Solid State Performance for Everyday Business

DECEMBER 2013

Solid state storage technology – typically storage devices based on NAND flash – have opened up new horizons for storage systems over the past couple of years. The storage market has seemingly been flooded by new products incorporating solid-state storage somewhere within their product line while making promises of break through levels of storage performance. Vendors have found there are some challenges when it comes to putting solid-state technology in the storage system. Many storage products entering the market in turn have a few wrinkles beneath the surface.

Approximately a year ago, Taneja Group set out to test one of the quieter success stories in the storage market – the StorTrends 3400i array built by American Megatrends, Inc. (AMI) in Norcross, GA. At the time, we were surprised by an array that delivered uniquely comprehensive and well-architected mid-range storage at a highly competitive price point.

Just recently, AMI has introduced yet another SAN storage appliance. The StorTrends 3500i integrates a comprehensive solid-state storage layer into the StorTrends iTX architecture. The 3500i brings with it the ability to use Solid State Drives (SSDs) in multiple roles – as a full flash array or hybrid storage array. As a hybrid storage array, the SSDs can be utilized as cache, tier, or a combination of the two. Along with SSD caching and tiering, the StorTrends 3500i incorporates these performance features with the field proven storage architecture validated by more than 1,100 global installs. The net result is a high performance and cost effective storage array. In fact, the 3500i array looks poised to be one of the most comprehensively equipped storage system options for the mid-range enterprise storage customer looking for solid state acceleration for their workloads.

With this most recent storage system launch, StorTrends once again caught our attention, and we approached AMI with the idea of a hands-on lab exercise that we call a Taneja Group Technology Validation. Our goal with this testing? To see whether StorTrends truly preserved all of their storage functionality with the integration of SSD into the 3500i storage system, and whether the 3500i was up to the task of harnessing the blazing fast storage performance of SSD.

> ***Technology Validated:***
>
> Based on our hands-on assessment with the StorTrends 3500i array, AMI combines a complete and robust storage system architecture with solid-state storage technology to create one of the best combinations of storage management capabilities and well-behaved high performance for mid-range customers.

What did we find? Storage system architecture and maturity makes a world of difference. With SSD in tow, the AMI StorTrends 3500i preserved all of the storage capabilities we tested last year, including features like high availability, built-in WAN-optimized replication, thin provisioning, snapshots, auto-tiering, and VMware VAAI support. The 3500i also brings to market a comprehensive set of SSD integration tools that makes SSD in the 3500i add up to more efficient performance under a wider range of conditions than many competitive products. We validated the

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

use of each of these technologies and how well they delivered performance under a variety of conditions and workloads, and came away impressed – a few of the highlights are listed below in Table 1.

Validating StorTrends 3500i - A Sample of Our Findings		
Performance – tested number of mailboxes with Microsoft Exchange and Microsoft Exchange LoadGen	10,000	***SSD cache-accelerated performance of 12 SAS disks, with fewer than 50 outstanding tasks during test, and no outstanding work queue at end of test***
Performance – simultaneous boot of 300 Windows 7 Virtual Desktops using VMware View	29 seconds	***SSD auto-tiering accelerates the performance of rotational disk***
Performance – 8K random read IOPs	265,000	***From 16 dedicated SSD disks***
Integration – VAAI primitive support for Write Same, SCSI UNMAP, Block Copy, and Hardware Assisted Locking	12X	***VAAI integration created eager zeroed volumes – effectively thinning out repeated empty data blocks – 12X faster.***

Table 1: Highlights of our Technology Validation testing with StorTrend's 3500i storage array.

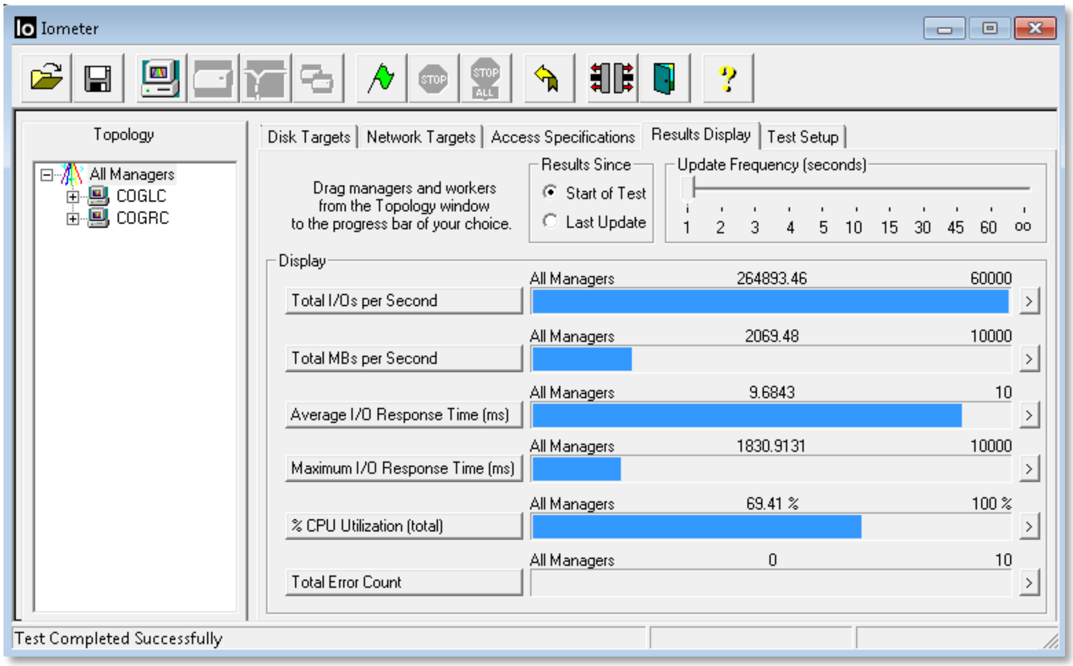

Figure 1: 265,000 8K random read IOPS measured during testing.

FOCUS ON STORTRENDS

Norcross, Georgia based AMI has a considerable pedigree in storage and compute technology. Over 28 years ago, AMI began building their substantial reputation in compute BIOS solutions, and rapidly expanded their focus with storage RAID technologies in 1989 where they defined the industry's leading RAID controller architectures. Since 1989, AMI has continued delivering deeply-engineered storage technologies, while they have evolved their technology portfolio to include a number of different solutions. Today, their history of storage innovation has equipped AMI with over 60 issued patents for storage technologies.

AMI's current StorTrends product line includes the small and branch office 1301i and 2400i arrays, the scalable 2401i and 3400i arrays, and their latest offering – the scalable SSD-equipped 3500i SAN. The StorTrends 3500i is a mid-range dual controller storage array that supports up to 16 SSDs (up to 32TB) in a pure SSD or SSD and disk hybrid system and uses those SSDs to deliver more than 250,000 sustained IOPs. Just as importantly, the StorTrends array maintains its cost-effective mid-range pricing.

MAKING STORAGE ARCHITECTURE MATTER

AMI's launch of the StorTrends 3500i array is particularly notable. Rarely have we seen tried and true storage architectures enable the deployment of SSDs without significant penalties on SSD performance or the need to develop entire new sets of functionality. AMI's promise is that they will leverage the field proven architecture of the StorTrends family (with over 1,100 deployments across a global collection of customers) to harness the full potential of solid-state storage technology.

This is a rarity in the blossoming performance storage market. We have seen complete abandonment of existing product versions as vendors introduce SSD, and more often than not, the SSD performance bragging rights fall to newcomers in the market who can start with completely fresh architectures that can be purpose built for the requirements of SSD.

The reason for this is simple – traditional storage controllers are often too easily pushed to their limits by the multi-faceted demands and capabilities of SSD storage. Not only is SSD capable of supporting phenomenal levels of small IOs that can crush storage controllers and buses, but SSD also has unique read, write, and wear characteristics that require special handling. Read performance is often vastly better than write performance, especially at high levels of capacity utilization, and this can break storage controller queues and caching algorithms. Meanwhile, SSD media is subject to significantly greater deterioration from wear, and controllers that can optimize the patterns in which SSD media are used can vastly extend media lifespan.

To tackle these challenges, vendors have turned to entirely new products and product architectures, and this means that many storage products designed for performance may not be field-proven, or may wield a partial storage feature set and interface that feels distinctly different than the business's traditional storage tools.

THE STORTRENDS 3500I ARRAY – ARCHITECTED FOR PERFORMANCE

AMI set out from day one to inject the 3500i with highly tunable performance alongside StorTrends traditional comprehensive feature set and easy manageability. While the 3500i remains equipped with a full set of wizardry that can let any customer breeze through a basic configuration, AMI has now opened the door to the use of SSD in a number of powerful and flexible ways, and has put granular controls at the fingertips of the customers who want to fully manage how SSD storage is used. The end product is a storage system that can deliver highly efficient and fully capable high performance.

AMI accomplished this with a multifaceted approach that included:

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

- Preserving and carrying forward traditional storage management capabilities like granular auto-tiering of data.
- Developing a number of SSD-specific features such as software kernel optimizations, read/write optimizations, and the ability to leverage part of the SSD capacity as non-tiering read and write cache.
- Leveraging their x86 processor controller architecture to maximize the performance of controller-shelf-installed SSD.
- Extending their caching architecture such that SSD for can be used with versatility and complete interoperability with AMI's existing and sophisticated auto-tiering and optimization technologies.

Prior generations of StorTrends' products included a built-in auto-tiering capability that granularly tiers data across multiple types of disk, and then optimizes the placement of individual data blocks on a single disk to ensure that blocks in higher demand benefits from higher access and throughput speeds based on their placement on disk. This data placement engine utilizes a workflow monitoring process in the background in order to make sure that performance was rapidly optimized but never interfered with application demand for storage performance. This functionality now carries forward to the 3500i where it works with SSD, along with advancements in software architecture, as well as a new caching layer that can leverage SSD devices as a read/write cache for accelerating any rotational disk storage volume in the system, alongside simultaneous use of other devices for tiering or as dedicated SSD volumes.

The 3500i continues to leverage a 3U dual active-active controller that houses 16 disks and can scale up to 256TB through the addition of expansion shelves. Fully active controllers, each built on industry-standard x86 8-core processors, support simultaneous connections to internal storage. These controllers come equipped with 2x1G Ethernet connections with options to expand to multiple 1Gb and 10Gb interfaces. Each controller is equipped with 16GB of system memory, and uses this memory to cache IO, relying on external UPS connectivity and top-to-bottom redundancy to protect the active cache footprint.

Figure 2: The StorTrends 3500i

What makes the 3500i a distinct product separate from the 3400i is the inclusion of SSD storage in the 16 bay storage controller. Customers at the time of acquisition can elect to purchase any even-number of drives up to 16, and can upgrade the SSDs in the future to increase the capacity of the 3500i. SSDs are offered in 200GB, 400GB, 800GB, and 2TB configurations, giving customers the option to deploy as much as 32TB of raw SSD capacity that can deliver over 250,000 IOPs of performance.

The 3500i remains highly scalable. While SSDs can concurrently only reside in the main controller chassis, individual 16 drive SAS expansion shelves can be added to a 3500i controller shelf to bring additional nearline (NL-SAS) or high speed SAS drives into the array. However, at the time of validation, AMI was engineering SSD support for their JBOD expansion shelves. Up to 7 shelves can be added for a total of 256TB within a single 3500i system. Expansion shelves are attached in daisy chain fashion starting with dual SAS 2.0 connectors on the 3500i and additional SAS expanders on each expansion shelf.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

When it comes to performance though, the 3500i uniquely scales with SSD. While up to 16 SSDs can be deployed in the controller shelf, those SSDs can be used in a variety of ways to deliver high performance for specific volumes or for all volumes in a system. SSDs can function as a dedicated tier of storage, integrated with rotational disk into an auto-tiering pool of storage that moves the most accessed blocks of data to SSD in 8MB chunks, or used as a protected read/write cache that can accelerate the performance of rotational disk on a LUN-by-LUN basis. Meanwhile, the full suite of StorTrends storage features continue to seamlessly function alongside SSD use – thin-provisioning, snapshots, replication, and others.

THE MANY USES OF STORTRENDS SSD

When used as a tier, AMI's proven data management policy engine can automatically migrate data between multiple tiers of rotational disk and SSD, and even optimize the placement of data blocks for maximum rotational disk performance. When used as cache, up to 4 dedicated SSDs (as much as 8TB of raw flash capacity) are automatically provisioned to temporarily hold reads and writes in separate spaces on the SSD. To maximize useable SSD capacity, write space is mirrored across devices to protect data, while read space is configured as a single stripe. Cached IO is optimally written to SSD in a log-structured type fashion in order to best match the characteristics of SSD media. As data access slows or data ages in cache, it is automatically flushed from cache in the background. This cache acceleration can be applied to any amount of storage in the 3500i, up to the maximum of 248TB of rotational disk and SSD tier capacity. Altogether, this architecture allows customers to flexibly and granularly configure how their storage performance is optimized, even electing one type of acceleration for some volumes, and another type of acceleration for others.

VALIDATING THE STORTRENDS 3500I

In 2013, AMI brought the impending launch of their StorTrends 3500i array to our attention. Clearly, AMI put significant thought into the integration of SSD into the StorTrends array family, and our thoughts immediately turned to revisiting prior testing Taneja Group completed in the summer of 2012. With advanced SSD technology, we wanted to see if AMI's field-proven storage architecture would bear the demands of solid state storage and still deliver on the full use of AMI's previously advertised comprehensive storage features.

In turn, Taneja Group turned to a test lab exercise and performed a variety of tests across several StorTrends 3500i configurations. Specifically, we set out to validate StorTrends 3500i performance claims, SSD resiliency, and how SSDs are integrated and used in the 3500i to create a complete storage solution that can be right-sized in performance and capacity. Secondarily, the StorTrends 3500i has also introduced a new software layer for the StorTrends storage array family, and with this new software, AMI has included support for a broad set of VMware VAAI primitives. In turn, as we conducted our storage performance testing, we leveraged VMware's vSphere product family heavily, with an eye toward validating StorTrend's virtual infrastructure capabilities.

TEST LAB ENVIRONMENT

Test Environment Equipment at Start of Test		
Equipment	**Description**	**Purpose**
(2) 3500i storage systems	Dual Active-Active HA Controllers 32GB of total memory 4x10G and 4x1G Ethernet 4x2TB SSDs 12x15k RPM 600GB SAS disks 16x7.2k RPM 2TB NL-SAS disks	Used for various IO and application testing, including IOMeter, VDI, and Microsoft Exchange

(1) 3500i storage system with 16 SSDs	Dual Active-Active HA Controllers 32GB of total memory 4x10G and 4x1G Ethernet 16x2TB SSDs	Used for IOMeter testing of the 16 SSD StorTrends 3500i configuration
(3) Windows 2008 R2 Servers, Standard	8 core, 32GB memory	Various tests
(1) Windows 2008 R2 Server, Enterprise	8 core, 48GB memory	10,000 user Exchange LoadGen test
(2) VMware ESXi Hosts	16 core, 48GB memory, hosting various Windows 2008R2 VMs	Used for replication testing and evaluation of vSphere plugins.
(2) Dell / Force10 10G Ethernet switches	24-port S2410 switches	Used for redundant iSCSI multi-pathing

Table 2: List of equipment used in our hands-on testing, with exception of general (non-iSCSI) network connectivity which is not listed.

STORTRENDS 3500I PERFORMANCE

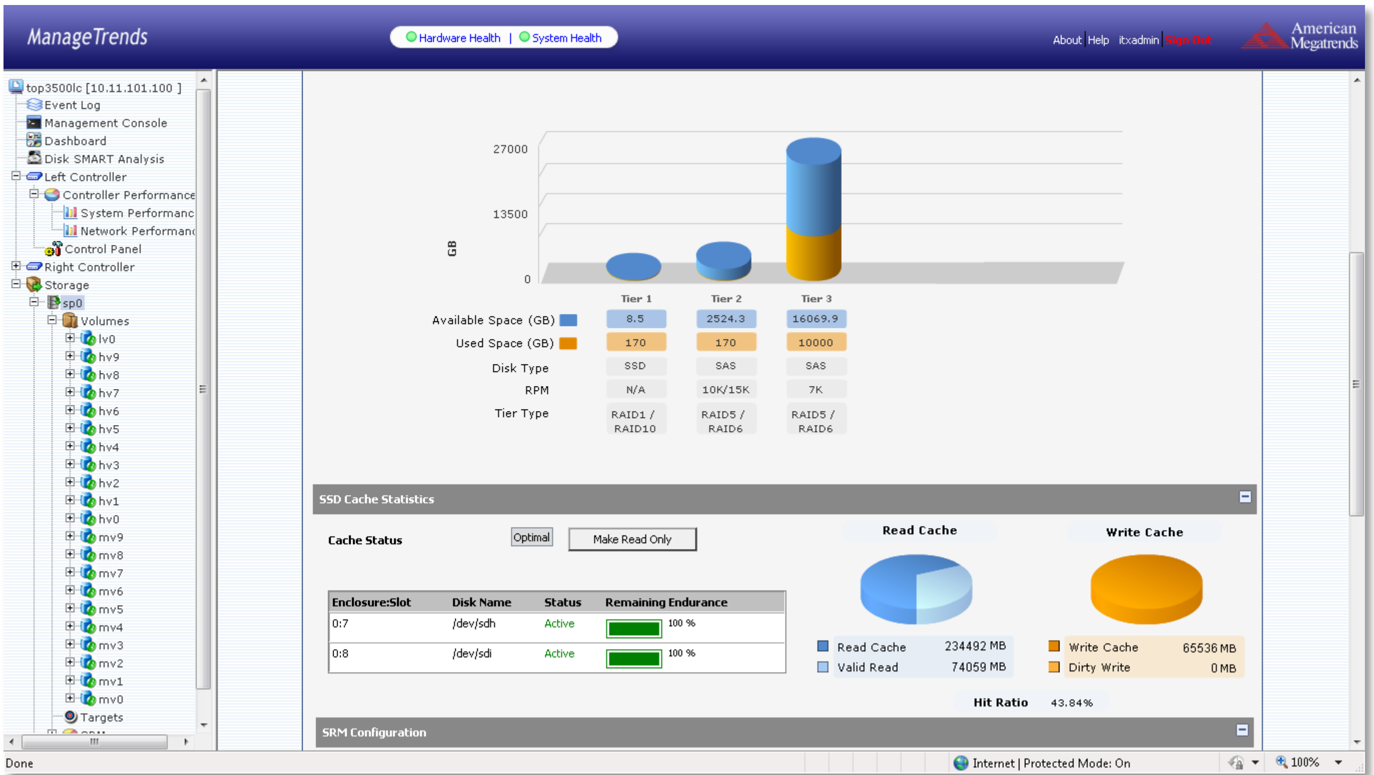

Figure 3: Our StorTrends 3500i was using 2 200GB SSDs configured into 64GB of write cache and 234GB of read cache, shown on the bottom of the image above.

We began testing of the StorTrends 3500i with a variety of IOMeter workloads to evaluate read and write performance capabilities in isolation. First, we tested a hybrid 3500i with 4 SSDs. In this system, 2 200GB SSDs were allocated as cache, with the storage system automatically configuring these two disks into 64GB of mirrored write cache space, and 234GB of read cache space. We conducted our testing with IOMeter, and ran it against 170GB of storage on 15K SAS disks to evaluate the performance of StorTrends 3500i caching. This 170GB of storage was provisioned to our host as ten 17GB storage volumes, in order for us to queue enough IO to enough disk targets to saturate performance. As configured – using only two SSDs for cache in front of rotational disk – this system represented AMI's smallest SSD configuration, but we decided that this would set our expectations on whether AMI can deliver the 250K+ IOPS they claim. We began testing this system by running through a series of random IO tests – 100% read, then 100% write, then 70/30 read write. The results are listed in Table 3, below.

IO Profile	Description	Performance
100% 8K Random Read	At start of test (44% cache hits)	28,000 IOPS
100% 8K Random Read	After 1 hour of warm up (51% cache hits)	42,000 IOPS
100% 8K Random Write	Note, write cache is mirrored across two devices, so performs at the speed of one device.	11,500 IOPS

Table 3: IOMeter testing included a number of different profiles, but the pure read and pure write most clearly outlined the performance capabilities of a StorTrend 3500i with a single set of 2 SSDs configured as read and write cache.

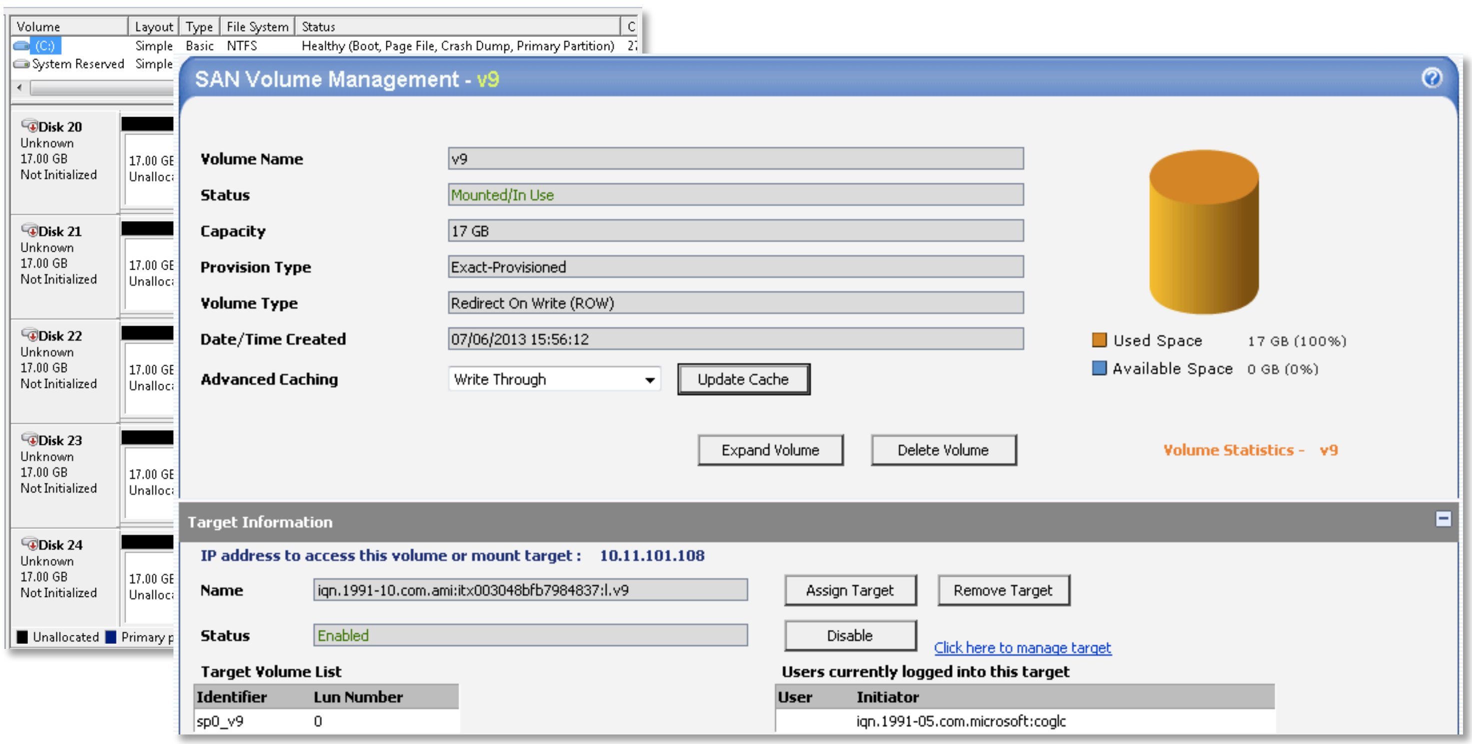

Figure 4: We used 10 17GB iSCSI storage volumes attached to a single Windows Server 2008 R2 host for IOMeter testing. These 10 connections allowed us to queue enough IO to drive the storage system to saturation.

THE SSD AS IN-LINE CACHE

Caching is obviously not instantaneous, and this test illustrated that caching improves an active workload's performance as time passes. As a cache becomes "warmed up" with data that will be accessed again, the cache can accelerate more IOs than when it is empty and data must be accessed directly from rotational disk. The most obvious illustration of this happened with our read testing, when we allowed the test to run for an hour – during this time our performance increased from 28,000 IOPS to 42,000 IOPS as our cache hits rose from 44% to 51%.

More importantly, our testing illustrated one of the challenges of SSD performance, and our additional testing went on to systematically investigate how AMI addresses this challenge; specifically, the nature of NAND flash means SSD write performance is always significantly lower than read performance. This is because a write requires multiple NAND operations to clear and re-write data. Storage systems are often challenged to optimize the efficiency of SSDs for the best write performance possible, especially as the number of SSDs scale and vendors attempt to leverage historical data protection technologies like RAID.

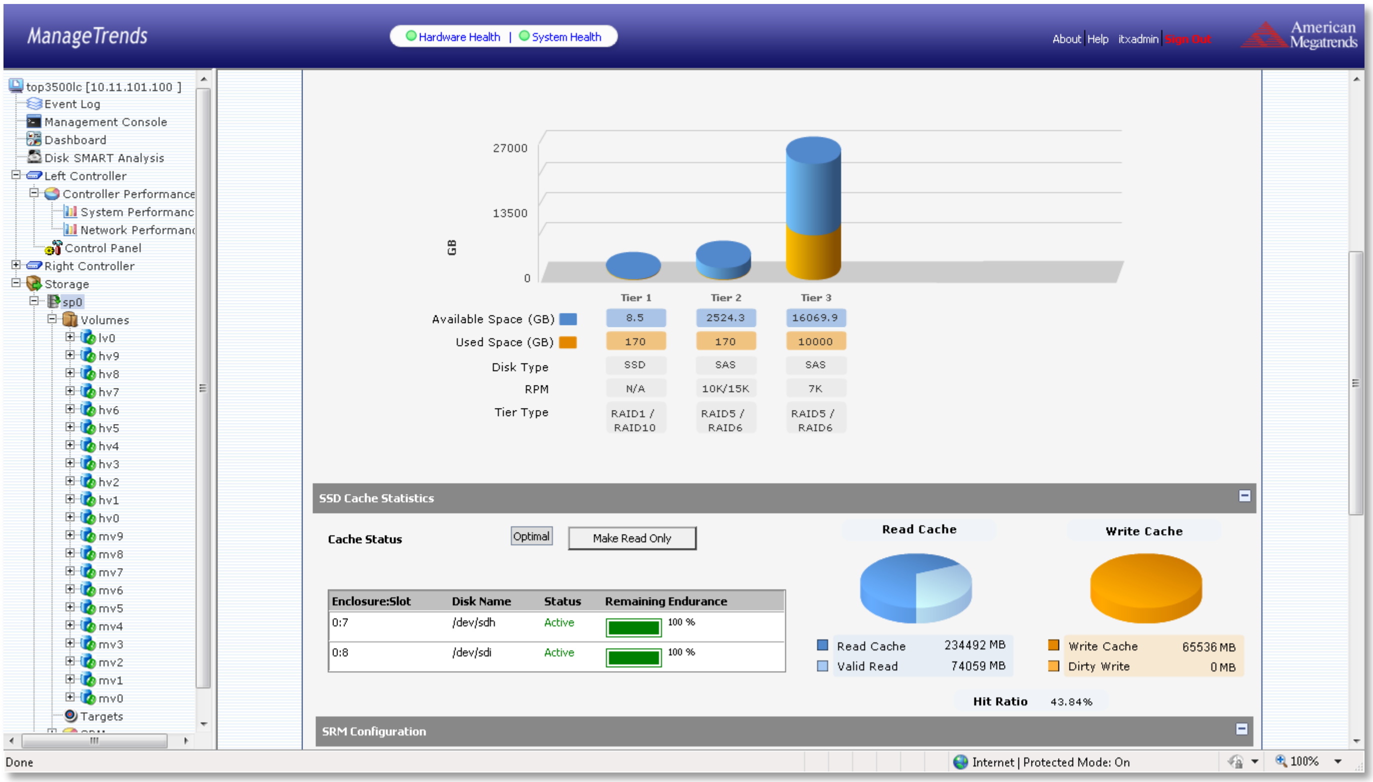

Figure 5: Our read cache hit ratio at the beginning of our IOMeter test of 2 SSDs configured as cache. The Hit Ratio is identified as 43%, meaning that 43% of the IOs on our 100% read test were serviced from cache.

THE SSD AS A TIER

We next examined how the StorTrends 3500i could make use of SSD as a storage tier, and performed similar tests with a small pool of 4 SSDs configured as a single RAID10 tier of storage. Those results are again listed in the table below, and demonstrated the type of smooth scalability we hoped to see – the write performance reflected a fairly linear increase in speed over the single drive performance we witnessed with our cache testing, and delivered performance that equaled three SSDs (with some likely minor trade-off for duplicate write overhead in this RAID10 configuration).

IO Profile	Performance
100% 8K Random Read	125,000 IOPS
100% 8K Random Write	31,000 IOPS
100% Sequential 8K Read	277,000 IOPS

Table 4: We turned our IOMeter testing to examining the performance of 4 SSDs configured as a single tier of storage. In this configuration, we used 10 storage volumes provisioned from a RAID10 SSD disk group containing 4 disks.

Finally, following our tests against 4 SSDs, we turned to a 16 SSD StorTrends 3500i array and repeated our test again. As we began this final test, we added a second Windows 2008 R2 host to our test bed and mounted another 10 volumes in order to distribute IO across controllers and fully saturate the SSD within the system.

IO Profile	Performance
100% 8K Random Read	265,000 IOPS
100% 8K Random Write	153,000 IOPS

Table 5: Our final IOMeter test focused on a 16 SSD StorTrends 3500i system. We did not repeat the sequential read test as our testing with 4 SSDs already saturated the available bandwidth (2.2GB/s).

Once again, the devices scaled linearly, with 14 active drives (2 drives were used for parity) delivering approximately 14X the random write performance we witnessed in our cache testing which effectively profiled the write performance of a single SSD. Read IO only approximately doubled versus 4 SSD random read, but this was because the Read IO reached the bandwidth limitations of our 10Gb Ethernet connections.

StorTrends' SSD capabilities can clearly deliver outstanding performance, and the StorTrends' ability to mix and match SSD as a dedicated SSD storage tier, SSD read/write cache, or both. This provides customers with an outstanding ability to harness high performance for their most important workloads, no matter how they want to manage the storage of that workload's data (e.g. even data stored on cost optimal nearline SAS can reap the benefits of high performance SSD cache).

Just as importantly, our preliminary testing demonstrates that StorTrends 3500i has the architectural underpinnings to make the most challenging aspect of SSD performance – write performance – sing. Our testing confirmed linear scalability in random write performance across all of these drives, at an 8K block size that is typical of demanding applications like databases.

87 Elm Street, Suite 900 • Hopkinton, MA 01748 • T: 508.435.2556 • F: 508.435.2557
www.tanejagroup.com

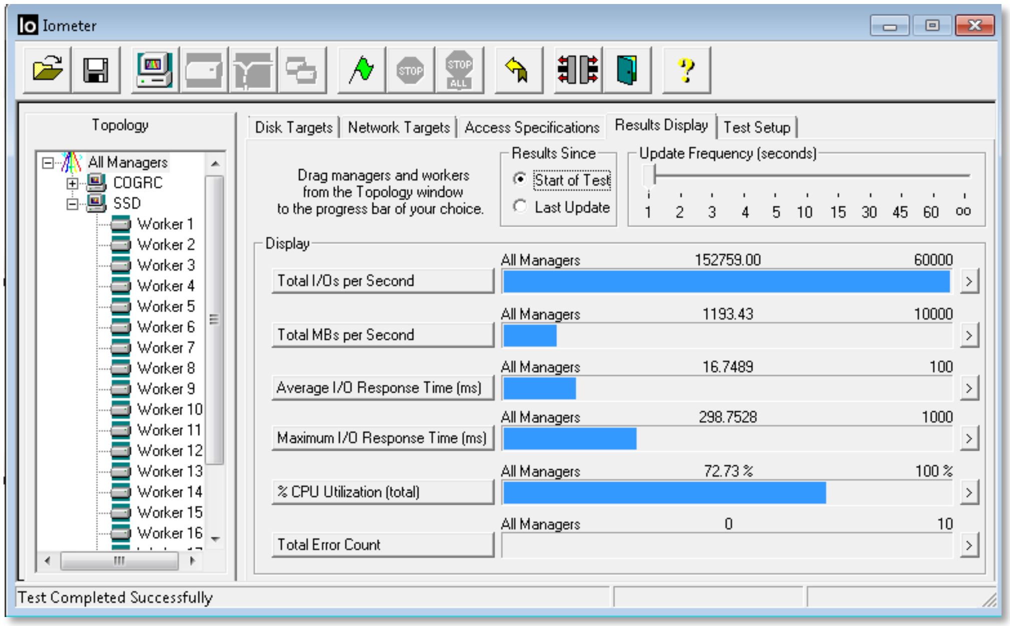

Figure 6: Testing 8K 100% random writes on a 16 SSD LUN generated 153,000 IOPS.

ADAPTIVE TIERING – THE SPEED OF ADAPTATION

During our testing with the AMI StorTrends 3500i, we also wanted to shine a light upon how fast storage systems are able to adapt when using auto-tiering technologies that promote data blocks to faster tiers of disk when IO increases. To shed light on the StorTrends 3500i's adaptation, we placed a 100GB SAS volume under a 100% 8KB random read workload, saturating the performance of 12 15K rpm SAS disks, and then turned on the 3500i's ILM feature. Arguably, this is a worst-case workload, as the SAS disks were under maximum and constant load, providing little IO headroom for block movement. Moreover, the AMI StorTrends 3500i has a workload management engine that is highly intelligent about system load, and avoids any impact to active IO. This throttled back the data movement versus a less consistent IO pattern that would allow more headroom for data movement.

What we observed was complete volume migration in 6 hours (100GB of data) and partial data movement (enough to create acceleration) in approximately 30 minutes. Is this enough to adjust to data center workloads in a timely manner? We think it probably is. First, constant IO is likely to be lower and less regular, potentially further accelerating how fast data can be moved. Second, under real world conditions, 100GB is a pretty big active data set. Even large SQL Servers with several 100GB storage footprints often only have an active data set closer to 5 to 10GB. Tiering those blocks up to SSD would happen much faster. Moreover, as we see from this exercise, even partial tiering can create significant acceleration and address changing business demands while still reaping the benefits of sharing SSD across larger pools of data.

The graph below illustrates the data points we sampled during this exercise – note, we only sampled a few data points and are interpreting the data points in between, on the assumption that data movement was fairly consistent throughout. Given the data below, a couple of conclusions can be drawn. First, for a large data set like this 100GB, the array can completely adapt to extreme performance within a few hours. Second, partial acceleration that makes a difference (10% performance improvement) happens within 30 minutes. Third, for smaller data sets that require less data movement, the optimization will happen much faster. In this case, if the data set had been 10GB, the entire data set would have moved into the SSD tier and would be delivering the 1GB Ethernet line maximum of approximately 13,000 IOPs in only 37 minutes.

87 Elm Street, Suite 900 • Hopkinton, MA 01748 • T: 508.435.2556 • F: 508.435.2557 www.tanejagroup.com

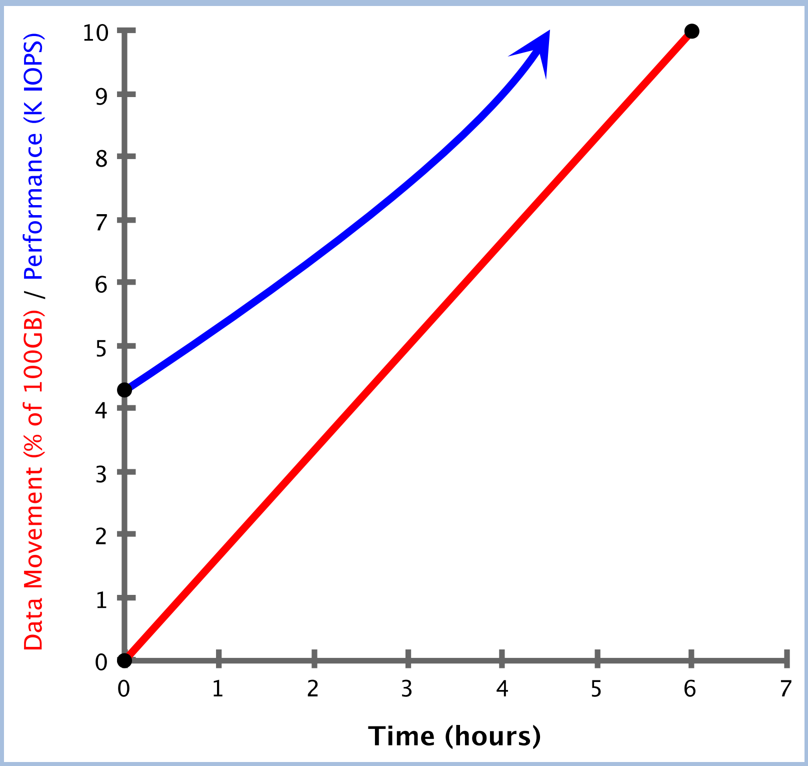

Figure 7: Data tiering progression and performance acceleration over time. Note, extrapolated from a few data points, and the assumption data movement remains linear.

PERFORMANCE IN THE REAL WORLD

With the drastic alteration in disk behavior that is made possible with solid state technology, it is increasingly difficult to say with any assurance that a vendor's demonstration of IO numbers with synthetic tools will hold up in the real world. We consequently took the opportunity to examine the StorTrends 3500i array's performance with a couple of workloads more indicative of real world behavior.

One such application was Microsoft's Exchange Loadgen tool. Loadgen is a Microsoft tool that simulates real world user activity and interactions with a Microsoft Exchange Server including many simultaneous email clients connecting to the server and interacting with each other. On typical rotational disk storage systems, an aggressive test often simulates 1,000 to 5,000 users. For our testing of the StorTrends 3500i, we simulated 10,000 users connecting to Microsoft Exchange Server 2010. We used our Windows 2008 R2 server as the Exchange Server, with all data stored on a tiered SSD configuration of 4 SSDs and 12x15K 600GB SAS disks. We simulated 10,000 very heavy Outlook 2007 users over a 10 hour test.

The Exchange Server passed, with plenty of breathing room. As we monitored the execution queue, it rarely surpassed 50 tasks, and the test finished at exactly 10 hours, with no cache on the Exchange Server requiring time to purge. These observations are notable, as in our experience Exchange Server 2010 tends to generate fairly heavy write IO and at an unusual average IO size of approximately 12KB. This IO pattern is often challenging for storage arrays.

87 Elm Street, Suite 900 • Hopkinton, MA 01748 • T: 508.435.2556 • F: 508.435.2557
www.tanejagroup.com

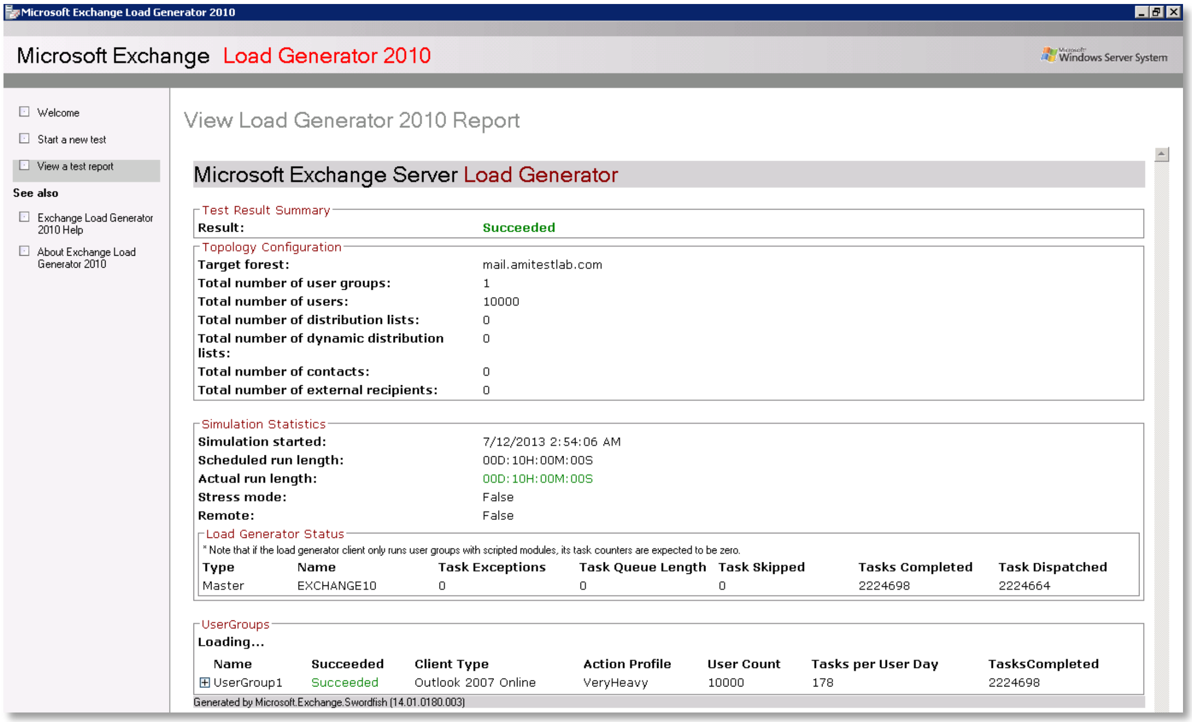

Figure 8: Results from a 10,000 very-heavy user Exchange test.

We also tested virtual desktops using a VMware Horizon View environment, and specifically evaluated boot time across many simultaneously booting VMs. We'll look at these tests in further depth as we examine virtual infrastructure integrations.

PERFORMANCE-ENHANCING INTEGRATIONS

With the StorTrends 3500i, AMI has also introduced and improved a number of features that will help SSD performance go even further. Two specific sets of features stood out in our opinion. One was an increase in the breadth of support for VAAI, now including Atomic Test and Set, Block Copy, Write Same, and UNMAP. These primitives allow the hypervisor to better interact with the array when using shared VMFS storage volumes, and allow the hypervisor to offload key operations to the array for high speed execution by the controller with minimal back-and-forth communication. Another feature that caught our attention were on-going improvements in StorTrends storage control. Specifically, this includes performance thresholds and tiering policies that control how data is automatically migrated across tiers. The changes in this area provide the administrator with more settings to differentiate between low and high importance workloads, as well as the ability to more granularly control the thresholds at which IO causes blocks of data to migrate to more or less performant tiers of disk.

We examined each of these sets of capability in turn.

VAAI support

We turned to our examination of VAAI support by looking at the four core VAAI primitives that the StorTrends 3500i supports – Write Same, SCSI UNMAP, Block Copy, and Hardware Assisted Locking.

To test a number of these features, we used vSphere system settings to turn on and off VAAI support, and compared time and operation efficiency with and without VAAI support.

WRITE SAME AND UNMAP

To test the effect of the Write Same VAAI primitives, we created a new 70GB eager-zeroed volume with VAAI support turned on, and then with VAAI support turned off (the vSphere DataMover.HardwareAcceleratedInit tag set to zero). During a volume creation, multiple empty volume blocks are laid out in a sequence that can be accelerated by VAAI – with VAAI support for Write Same, the writing of these empty blocks are offloaded to the array controller, and this can significantly accelerate performance. The results of our testing with the StorTrends 3500i are listed in Table 6 below.

VAAI Write Same On/Off	Time to Complete 70GB VMDK Creation
VAAI Write Same On	57 seconds
VAAI Writes Same Off	12 minutes, 35 seconds

Table 6: We tested the time to complete the creation of a new volume with VAAI Write Same turned on and off. Support for VAAI Write Same improved the performance of this operation by more than 12X.

During this same process, we were able to draw some conclusions about SCSI UNMAP as well. Because an eager-zeroed volume writes out an entire volume of empty data as fully provisioned space, that space is immediately consumed on the array irrespective of storage system policies. SCSI UNMAP allows an array to release empty blocks when the hypervisor has acknowledged they are empty. When we provisioned this volume with VAAI turned on, it was still provisioned in accordance to our storage system policy, as a thin volume, and only occupied 30MB of space. Moreover, SCSI UNMAP also works as VM volumes free up blocks in the future – such blocks will be automatically returned to free space.

BLOCK COPY

We next turned to testing VAAI's Block Copy primitive. The Block Copy primitive offloads the creation of duplicate blocks when data is being copied. Turning VAAI back on, we executed a clone operation against a 26GB VM, and then we repeated the test with VAAI turned off. For this testing, we used an array connected over 1GB Ethernet. Because of this 1GB Ethernet connection, Block Copy had an even more significant impact. With Block Copy turned off, we could see read/write operations approaching the theoretical maximum throughput of our network link (we observed approximately 110MB/s). Consequently, Block Copy offloaded this task and accelerated execution of the clone significantly, as shown in Table 7 below.

VAAI Block Copy On/Off	Time to Complete 26GB VM clone
VAAI Write Same On	1 minute, 59 seconds
VAAI Writes Same Off	8 minutes, 15 seconds

Table 7: We tested a VM clone operation with VAAI support turned on and off, and observed more than a 4X improvement in execution time.

HARDWARE ASSISTED LOCKING

Hardware Assisted Locking proved more difficult to test in a meaningful way in a relatively small-sized test environment. Hardware Assisted Locking has the greatest impact at scale, where many

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

hypervisors are attempting to simultaneously access a single shared VMFS datastore. But StorTrends logs allowed us to take a deeper look at what was happening, and prove AMI's support for VAAI's Hardware Assist primitive. Specifically, by turning VAAI support off, we observed an extremely large number of SCSI reservation releases as the hypervisor repeatedly requested control of the volume; these interactions disappeared with VAAI turned on. We observed this output while booting 5 VDI desktops on each of 2 hosts connected to the same shared volume. During this boot interaction, we estimate we captured more than 500 SCSI reservation interactions.

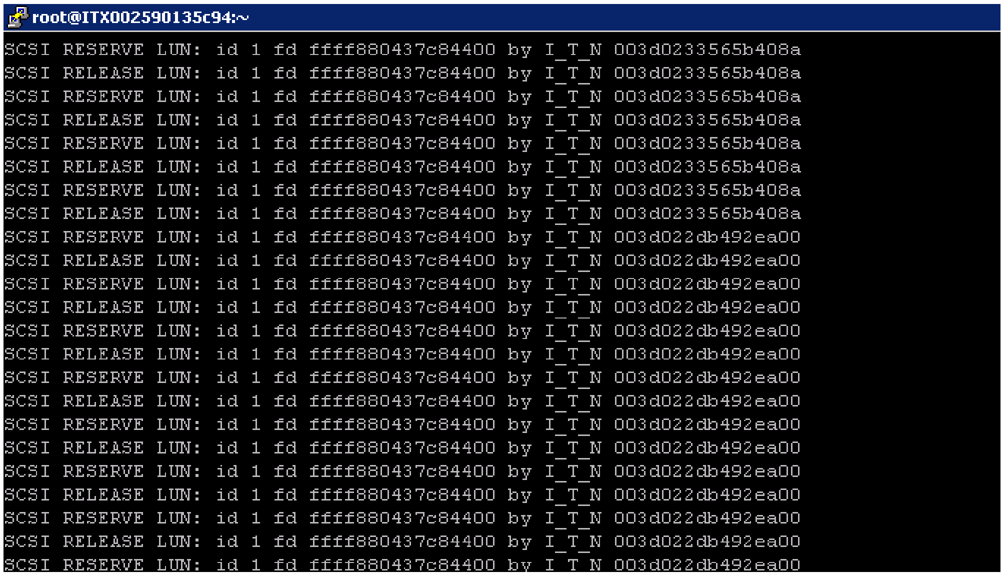

Figure 9: Turning VAAI support off demonstrated the chattiness of the SCSI reservation process that Hardware Assisted Locking is designed to decrease.

REALWORLD VIRTUAL PERFORMANCE

To cap off our testing of the StorTrends 3500i performance in a more real-world virtual environment, we performed one more virtualization-specific test: the simultaneous boot of 300 VDI desktops. We used a dedicated storage array and server infrastructure for this testing, consisting of 8x48GB 16 core servers, and an array with 4 SSDs configured as cache, and 12 15K SAS drives. We provisioned two separate datastores, consisting of 13 underlying LUNs each to maximize throughput. We accessed these two storage pools through two separate controllers, and used each datastore as a separate linked clone pool. We cloned 300 VMs as persistent desktops, using a single 40GB Windows 7 VM in each linked clone pool as a golden image. We then simultaneously booted these VMs with a script and monitored boot time.

From this test, we observed a 28.8 second boot time. Many arrays would crumble under this load, or minimally take much longer to boot this number of VMs – we typically observe SSD-equipped FC arrays averaging boot times of multiple minutes for similar loads.

Continued Improvements in Auto-tiering

Finally, not to be overlooked, StorTrends has continued to mature their leading auto-tiering technology. As readers of our prior test report will recall, AMI's auto-tiering technology (what AMI calls Information Lifecycle Management and Zoned-Bit Recording) is comprehensive. Not only does Information Lifecycle Management (ILM) migrate data across different types of disk, but Zoned -Bit Recording (ZBR) works in tandem to also tier data on rotating disks and optimally place the most important data on the fastest part of the rotating spindle. In our previous testing, ILM and ZBR together were able to nearly double random IO performance.

With the StorTrends 3500i, AMI has continued to mature their ILM and ZBR technologies, wrapping them in a more granular and easier to configure policy engine, as well as make them SSD-aware. The interface now includes a drag and drop interface for setting policies against individual tiers, and policies can be set globally (once) as well as on a LUN-by-LUN basis. In our view, we anticipate the enhancements here will be key enablers for customers running mixed SSD/HDD configurations. The StorTrends interface surpasses or matches other competitors in the marketplace, and the level of granularity gives customers a sophisticated amount of control. When using SSD in a tiering configuration, customers could easily use this interface to restrict the promotion of less important data, ensure that important data was rapidly tiered to SSD, and keep data there for an ideal period of time.

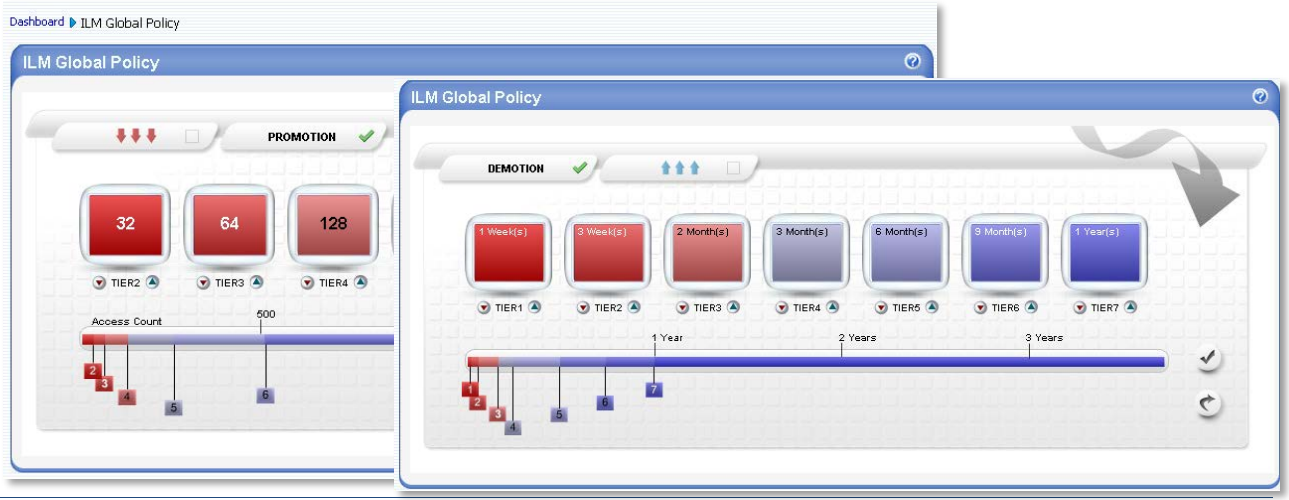

Figure 10: In the AMI StorTrends 3500i, the tiering interface becomes considerably easier to use and more granular. This will let SSD go further for many customers, as well as allow customers to more aggressively accelerate their most critical workloads.

ARCHITECTED FOR RESILIENCY

SSDs introduced into traditional storage systems also bring on challenges around media endurance and failure patterns. Many first generation attempts have in turn frustrated customers. AMI has an unusual pedigree in this regard, as they've long focused on engineering highly resilient storage components, as well as the multitude of low level details involved in interfacing with failure-prone rotational disk devices. Consequently it is little surprise that AMI StorTrends brings a full complement of SSD-supporting availability technologies to the table, in the form of their already established highly available system architecture, and a new StorTrends 3500i Wear Level Tracker.

DEVICE FAILURE

To revisit one of our proof points from our prior StorTrends 3400i failure testing, we took a look at what would happen to the StorTrends 3500i performance during a disk failure. Specifically, we unplugged a device being used in a 4-drive SSD cache configuration to note the impact on performance. When a device fails in this configuration, volumes will switch to non-cached and data will be automatically flushed and redistributed to fully protect the write data. Administrators can then re-enable caching before or after the disk is replaced (in a 4 disk configuration experiencing a single failure).

In this case, we removed one of the devices while running an 8K 100% random write IOMeter test on a cached volume. We didn't expect failure to be transparent, but we were specifically monitoring to see what the impact was, and whether a data loss event would occur.

When the drive failed, our IOMeter test was creating 19,500 IOPS – about the maximum random write performance of the two active SSDs (since two drives are mirrors for write data). When we removed an SSD, all caching was turned off and IO became initially erratic as response time greatly increased, but then recovered to 3,000 IOPS, which was the expected performance of our 12 15k SAS disks. We were then able to turn caching back on in the StorTrends ManageTrends interface. Write performance then returned to 9,000 IOPs, about the maximum performance of a single 2TB SSD. This behavior matched our expectations, and this process took place with no loss of data.

We'll note that we previously examined controller/fabric failure in the 3400i and came away satisfied, and the 3500i continues to use this same system architecture.

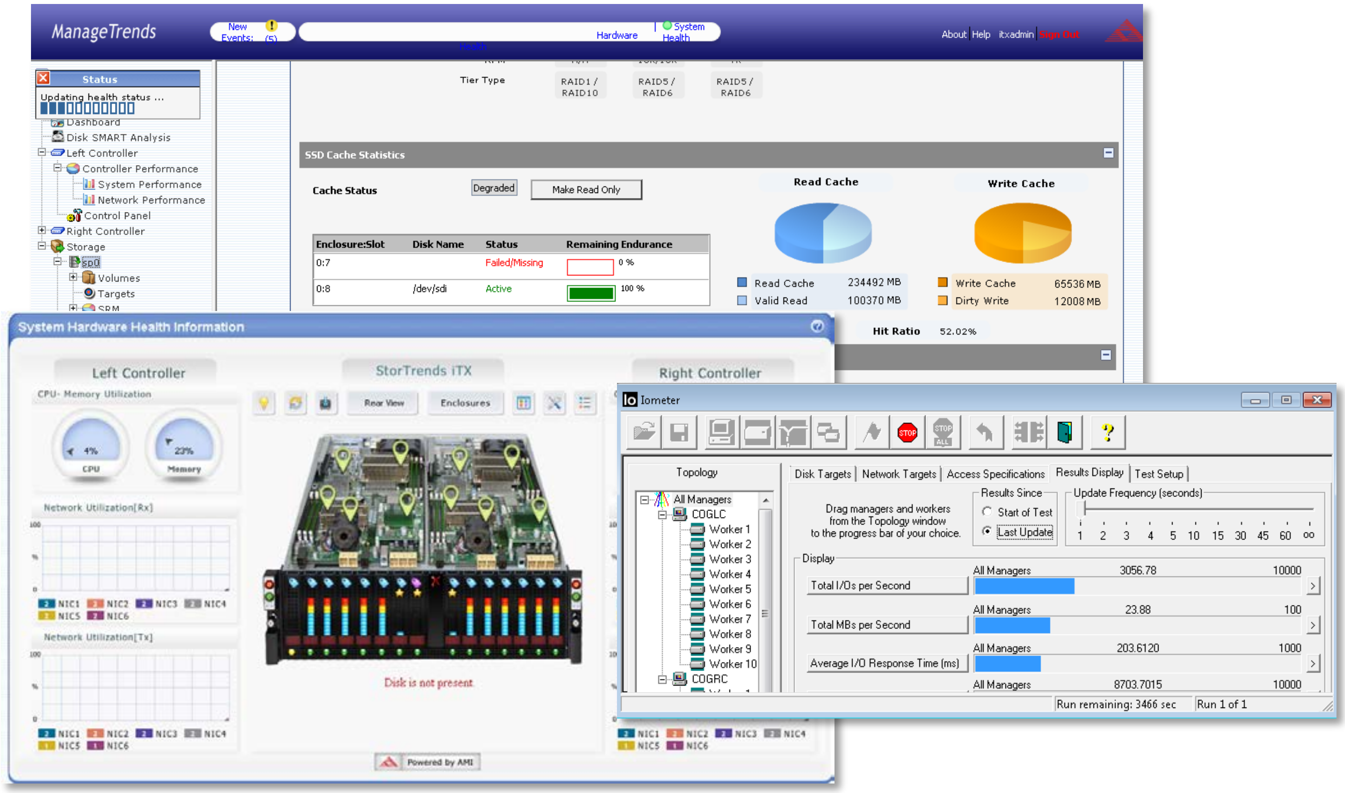

Figure 11: In the StorTrends 3500i, the tiering interface becomes considerably easier to use and more granular, allowing customers to deeply control and optimize which storage volumes receive the benefit of SSD. This will let SSD go further for many customers, as well as allow customers to more aggressively accelerate their most critical workloads.

WEAR LEVELING AND REPORTING

Finally, and not to be overlooked, AMI has also been highly attentive to wear leveling and understanding and reporting on the wear of the underlying devices. SSDs have limited lifecycles based on the number of times specific NAND cells are written. Within the iTX 2.8 OS, AMI made RAID handling optimizations to perform wear leveling across underlying SSDs and make sure wear was distributed evenly across all devices. AMI also developed a reporting interface that would track the wear on individual devices, and use a profile of that specific device (since AMI uses multiple types of SSDs) to report on the expected remaining life of an SSD. In turn, AMI can proactively alert customers to the need for SSD replacement from its StorAid support team.

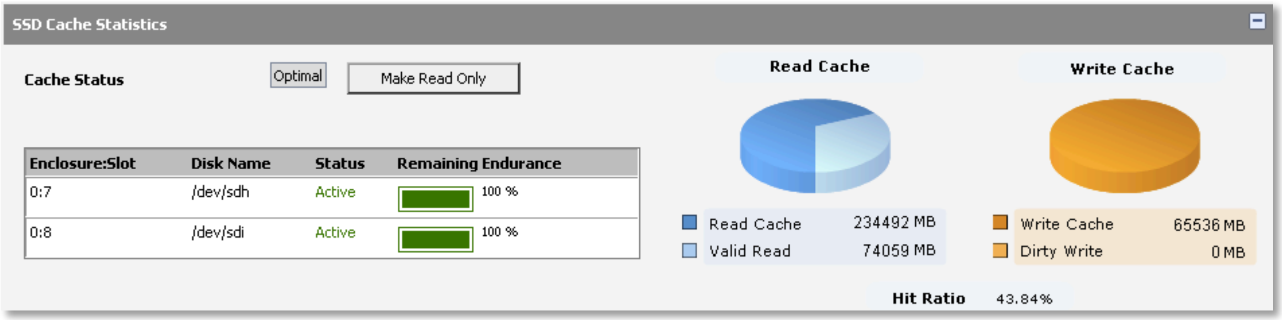

Figure 12: Wear level reporting in the StorTrends 3500i ManageTrends interface.

TANEJA GROUP OPINION

Accelerated storage has suddenly leapt from being a luxury to being the baseline that all IT practitioners should be considering for any storage acquisition. But the reality is not all storage solutions are the same, and customers are steadily learning that the mileage from SSD, as well as the storage capabilities from systems that contain SSD, vary. Given the relative youth of SSD-accelerated storage arrays, it is challenging to build systems that contain SSD, fully integrate that SSD to seamlessly and flexibly accelerate any storage within the array, and are fully mature in storage in all of their storage features.

AMI is an exception, and frankly it is little surprise – AMI has after all honed their storage chops for longer than the vast majority of vendors in the storage industry today. If anyone is equipped to build full-featured storage, and do all the engineering it takes to integrate SSD into a well-polished hybrid SSD/HDD solution, it is probably AMI.

Our assessment is that AMI's StorTrends 3500i array brings to market one of the most comprehensive, versatile, and cost effective solid state systems on the market today, and it looks to shake-up the average business's ability to access SSD-accelerated storage. The 3500i equipped with 4 SSDs gives customers the incredible flexibility of provisioning any storage volume as either SSD cached, SSD-tiered, or a combination of the two using their caching and tiering technologies. For customers previously trying to reach SSD, AMI looks like they'll be making this technology truly accessible for the first time. Meanwhile, the StorTrends 3500i also provides the option for all 16 SSDs delivering blistering performance - up to the 265,000 IOPs we saw in our hands-on Taneja Group Technology Validation. Even with these 16 SSDs, the StorTrends 3500i still looks to be priced more cost effectively than similar performance and capacity SSD systems. This may be a key enabler of StorTrends 3500i success in this market, as many of the all-SSD solutions available today build for extremes and consequently exceed the price point of many customers who could use a little less performance, but still benefit from SSD capabilities; a bulk of the market demands for high performance fit easily within this 265,000 IOPs footprint of the 16 SSD StorTrends 3500i.

The NetNet is there's no need to look for alternatives to make solid state more affordable and practical for the mainstream - with the AMI StorTrends 3500i, it is clearly here. Moreover, there's no compromise here – the AMI StorTrends 3500i is ready for prime time.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com