

## **STORTRENDS 3400I** Reinventing the value of storage

**NOVEMBER 2012**



Selecting a primary storage solution is undoubtedly one of the most critical decisions an IT department can make. As the foundational piece of the modern datacenter, it represents perhaps the single most important piece of IT infrastructure for businesses large, medium or small. Business critical applications will live and breathe on the performance of the selected storage system, and business data will be inevitably constrained by the capacity of that storage system.

In the mid-market, making a storage investment can be particularly daunting, as the stakes are higher, and the selection is harder. Compared to larger enterprises, in the mid-market, storage dollars are fewer and harder to come by. Precious and often limited IT staff time is spread across more systems and technologies, their core skills are often not rooted in storage, and technically vetting a storage system can be all but impossible. This makes storage a risky proposition for the small enterprise (SME) and SMB customer. We frequently hear tales of storage system purchases where I/O is not sufficient, features are missing (or require additional licenses and cost to acquire), or where architectural compromises create availability issues that regularly impact the entire business.

For several years, the developers of the StorTrends line of NAS/SAN solutions have been working hard to architect a storage system for the mid-market that puts an end to these risks and compromises. By harnessing the engineering expertise from their parent American Megatrends, Inc. (AMI) – an innovator in storage and BIOS technologies – StorTrends has been tackling the challenge of delivering abundant performance, robust reliability, and feature rich storage with the SMB and SME customer in mind. Their claim is that the StorTrends 3400i is both one of the most cost effective choices in the market, and one of the most well rounded.

In mid-2012, StorTrends caught our attention with these claims and a series of notable customer wins in a highly competitive market. To learn more, we approached StorTrends with the idea of a hands-on lab exercise, what we call a Technology Validation, to examine in more depth how StorTrends was delivering comprehensive value for customers in the mid-market space. Utilizing our proven validation methodology that included time spent at AMI headquarters in Norcross, GA, we put a set of StorTrends 3400i storage systems through the paces, with an eye toward examining several capabilities that StorTrends claims makes the 3400i one of the best value storage options in the mid-market. In a nutshell, these capabilities are:

1. **Availability and Resiliency** – StorTrends claims the 3400i is architected “top-to-bottom” for high availability, and that they offer best of class features that will meet any storage administrator’s demands.

### ***Technology Validated:***

At the end of our testing, we’re convinced – customers in the market for small and mid-range storage systems can rest confidently on a choice of StorTrend’s 3400i. Complete with easy expandability and a feature list that puts many other major vendor storage systems to shame, the 3400i is equipped to go the distance in terms of performance, reliability, and adaptability and be a lasting part of the SMB and SME customer’s IT infrastructure.

2. **Performant** – StorTrends, harnessing their expertise developing market-leading RAID solutions (AMI originally developed the MegaRAID product line, subsequently sold to LSI, Inc. in 2001) has now turned their attention to a high performance 3400i controller architecture and on-disk data layout that maximizes the performance of every spindle in the system, while still allowing customers to scale both performance and capacity as business needs change.
3. **Efficient** – Along with performance, StorTrends emphasizes a deep focus on multi-dimensional data optimization inside of their StorTrends arrays. A combination of auto-tiering, on-disk optimization (what StorTrends labels Zoned Bit Recording), and a data management policy engine helps leverage every ounce of performance and capacity to maximum effect. Specifically, StorTrends places data for maximum cost effectiveness or maximum performance depending on the business value of that data. It does this dynamically as IO is handled.
4. **Enterprise-class Storage Features** – StorTrends claims to deliver a full range of storage features within the 3400i (snaps, clones, provisioning tools, visibility tools) and claims to have architected the entire system so that these features are all available without impact on on-going IO.
5. **Fully Disaster Recovery Enabled, out of the box** – Finally, going one step beyond what most customers identify as a set of complete storage features, StorTrends also comes with a full set of replication technology with no additional licensing required. This replication technology can operate in many-to-many relationships, and is wrapped with WAN Optimization. StorTrend’s WAN Optimization is a powerful complement to replication, and significantly reduces bandwidth needed for replication. This stands to make disaster preparedness much more practical.

Validating StorTrends 3400i - A Sample of Our Findings		
<i>Availability and Resiliency</i> – detectable impact during controller reboots, disk removals, network interface outages, and power removal.	Total downtime avoidance	<b><i>StorTrends is architected for uptime and disruption-free operation</i></b>
<i>Performance</i> – tested number of mailboxes with Microsoft Exchange and Microsoft Exchange LoadGen	10,000	<b><i>Auto-tiering and ZBR optimize StorTrends performance</i></b>
<i>Performance</i> – streaming sequential read bandwidth on active/active controllers.	3.4GB/s	<b><i>Architected top to bottom for maximum performance</i></b>
<i>Optimization</i> – amount of random write performance increase as data was tiered up to fastest disk and short-stroked to optimize performance	2X	<b><i>StorTrends optimizes on-disk data placement, which helps stretch performance to new levels.</i></b>
<i>Features</i> – remote replication data efficiency improvement measured with WAN optimization Data Services deduplicating and compressing transmitted data	1.7X reduction in transmitted data	<b><i>StorTrends built-in WDS reduces data to 58% of the original volume size.</i></b>

**Table 1: Highlights of our Technology Validation testing with StorTrend’s 3400i storage array.**

What did we find as we undertook this lab exercise? During our time examining the technology, it became clear that the StorTrends 3400i storage system is extremely well rounded and mature. While at first glance StorTrends may seem a relative newcomer to the market, they’ve actually been developing and selling storage arrays for a number of years, with a number of notable, large scale wins. This tells the story of the subtle StorTrends difference – StorTrends is engineered to an extreme, and polished to a well-honed finish, with the needs of the SMB and SME in mind. At the conclusion of our test, we’re convinced that the small and mid-market storage customer alike will find exceptional value in StorTrends, and can buy StorTrends with confidence, knowing that this is one serious storage platform that can stand up to their expectations.

## FOCUS ON STORTRENDS

AMI has been delivering innovative storage technologies for years, beginning in the mid-nineties with the development of the MegaRAID controller card. This card eventually became the market leader and took the industry by storm. Since that day, AMI has maintained a serious footprint in componentry for x86 compute systems – from remote system management components to system BIOS components. Meanwhile, generations of innovation have been delivered under the banner of “StorTrends”, a division of AMI. Today, StorTrends sports a technology portfolio that consists of small and branch office 1300i and 2400i arrays, a scalable 2401i array, the SafeTrends continuous data protection solution, and the 3400i scalable mid-range iSCSI array.

### THE STORTRENDS 3400I ARRAY

As their flagship array, the 3400i iSCSI array starts with a 3U dual active-active controller that houses 16 disks and can scale up to 128 disks through the addition of expansion shelves. Fully active controllers, each built on industry-standard x86 8-core processors, support simultaneous connections to internal storage (single controller-to-volume connections, not dual controller-to-volume). These controllers come equipped with 2x1G Ethernet connections, and options to expand to multiple 1Gb and 10Gb interfaces (the controllers come equipped with 10Gb controller-controller connectivity, but external interfaces are added as an option). Each controller is equipped with 4GB 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.

The 3400i's capacity can be scaled with a modularly expandable system architecture that allows customers to grow over time as capacity needs change. Individual 16 drive SAS expansion shelves can be added to a 3400i controller shelf to bring additional nearline (NL-SAS) or high speed SAS drives into the array. Up to 7 shelves can be added for a total of 128 drives or as much as 256TB of capacity within a single 3400i. Expansion shelves are attached in daisy chain fashion starting with dual 6 Gb/s SAS 2.0 connectors on the 3400i and additional SAS expanders on each expansion shelf.

### INNOVATING IN THE ARRAY CORE – iTX

While attention to hardware density and architecture is critically important, the storage industry has rapidly changed over the last few years. Now, innovative features and capabilities are increasingly defined by software. iTX is the all important software core within the 3400i. With a 64-bit architecture, iTX is responsible for handling all IO flowing through the system, effectively caching IO, and coordinating all storage operations, as well as serving up a web management GUI and interfacing with storage agents such as VMware vSphere plugins. But most importantly, iTX is about deeply optimizing IO transactions and the placement of data on disk to push IO performance to the physical limits of the 3400i's hardware componentry. iTX does this by leveraging much of the knowhow and expertise from AMI's long history in RAID and storage. Specifically:

***iTX delivers RAID-RT data protection.*** RAID-RT is StorTrend's special sauce for moving RAID from a dedicated controller onto industry standard processors. In the place of proprietary RAID controllers of the past, RAID-RT uses the MMX offload capabilities of modern processors to offload XOR operations and accelerate those operations with the latest multi-core architectures. Integrated into this RAID-RT technology is a layer of SMART predictive failure analysis that can pre-emptively notify users before a drive failure fully occurs.

***iTX delivers wide-striping.*** AMI uses the iTX software engine to wide-stripe data across the total pool of same-tier RAID disks, ensuring that performance increases as spindles are added to the array.

***iTX delivers dynamic tiering and optimization.*** Far different from most “time interval” auto-tiering technologies on the market today, iTX auto-tiers data dynamically to move it across different tiers of disk as access patterns change, thereby dynamically increasing performance for suddenly “hot” data.

But more than just auto-tiering across different types of disks, iTX also carves the disk drive into different zones, and places data on each individual disk so that the most important data receives optimal performance. As we'll see, this can have a significant impact on performance, and helps customers extract maximum value from the 3400i without being pressed to turn to the expensive SSD storage strategies that other vendors use to overcome array inefficiencies.

### ENGINEERED FOR STORAGE AND AVAILABILITY

The AMI engineering pedigree is self-evident throughout the architecture of the StorTrends 3400i, where each component is carefully matched to the next. RAID-RT volumes are distributed across controllers, so that either controller can serve up access to a volume in true active-active fashion, and one controller can take over all IO in the event of a controller failure. Controllers share an internal bus to access disks attached to the other controller and provide availability, thereby eliminating any dependence on external switches. That internal bus is designed to support full system performance, and reduces the impacts from any component failure. Meanwhile, nearly every other component in the system is redundant and/or dual-pathed, with redundant power supplies and SAS ports across the controller and all expansion shelves. Finally, attention to small details is tremendous – such as a StorTrends best practice that one power supply should be attached to an uninterruptable power supply (UPS) and the other must be attached to line power. This allows the 3400i to go beyond normal UPS monitoring, and monitor actual voltage input from the UPS. If the 3400i detects there's a failure or degradation (even one that may not be detected by the serial or USB monitoring of a UPS) then the 3400i automatically reverts to write-through caching that ensures all writes are immediately written to disk. This protects data in case of a full power failure.

### VALIDATING THE STORTRENDS 3400I

In mid-2012, a number of innovations from StorTrends caught our attention. Our assessment led to a conclusion that the 3400i had a number of features and capabilities that would make it a highly viable mid-range storage solution for SMB/SME customers. We in turn approached StorTrends with a validation exercise to evaluate whether or not the 3400i array would deliver a unique value proposition to the SMB/SME customer. Our validation methodology was to measure the 3400i against a core set of criteria – criteria we advise mid-market customers to use when considering a storage investment, in order to maximize the value and longevity of that investment.

### CRITERIA FOR MID-MARKET STORAGE SYSTEMS

What do these criteria come from? In essence, they are common sense. The SMB/SME customer is hard pressed to become a domain expert specifically for storage. Per employee, there are many more IT initiatives at work in the mid-sized business, leaving little room for high levels of expertise in one single infrastructure system like storage. Traditional, legacy storage systems can be a significant challenge for these typical SMB/SME storage customers, as they often impose a level of complexity and management that proves challenging for IT “generalists”. Worse yet, many storage solutions in the market fall short of delivering all of the capabilities an SMB/SME customer expects out of the box, and can become a licensing and add-on nightmare. Finally, new cutting edge solutions that seem more attuned to SMB/SME customer needs can turn out to be less well rounded than anticipated, and may deliver less sustainable or predictable performance and reliability when they're put into production use. Few SMB/SME customers have time to fully vet storage solutions in lab environments and ensure that they meet promises. These IT challenges in the SMB/SME business drive our criteria for mid-market storage. Storage systems must:

1. **Avoid complexity.** Mid-market customers shouldn't require storage expertise to deploy, provision, and utilize storage. Such events are infrequent, but can make or break IT agility in businesses without dedicated storage admins.

2. **Deliver heightened availability and resiliency.** Mid-market customers are just as dependent on IT system uptime as their large enterprise brethren, but they require a recipe for similar uptime protection that does not entail the complexity resident in enterprise-class FC storage arrays and fabrics.
3. **Efficiently provide both performance and capacity.** The value of storage is determined by a complex balancing act between performance and capacity. Extracting the maximum amounts of each, and providing the mid-range customer sufficient headroom to meet changing demands in performance and/or capacity can help stretch invested dollars further.
4. **Come with full features that do away with further integrations or on-going add-ons.** Finally, the enterprise is riddled with add-on licensing and pay-by-the-feature approaches to storage. Storage is a precious commodity, and seemingly innocuous features can turn storage for the mid-range customer into powerful tools that can keep data better protected, make storage easier to manage alongside other systems, or even enable levels of disaster planning that were never possible before. The breadth and depth of storage features are critical enablers for business critical computing in the mid-market customer.

## A TEST FOR MID-RANGE STORAGE

We used these criteria to define a test scenario and guide our hands-on testing of a StorTrends 3400i array. Working with several 3400i arrays provided for our evaluation at AMI headquarters in Norcross, GA, we walked through a number of real world scenarios to see how the 3400i measured up in deployment as well as on-going use. The equipment used in this test included several multi-core servers, including an Exchange Server, several Windows 2008 R2 Servers, and two VMware ESXi Hosts with several Windows 2008 R2 VMs. What follows is a recap of our exercise and our findings, as we progressed putting the StorTrends 3400i through a real world deployment and use exercise.

### TEST LAB ENVIRONMENT

Test Environment Equipment at Start of Test		
Equipment	Description	Purpose
(2) 3400i storage systems with a single 3202j expansion shelf	Dual Active-Active HA Controllers 8GB of total controller memory 4x10G and 4x1G Ethernet 16 15k RPM 600GB SAS disks 16 7.2k RPM 2TB SAS disks	Used for our replication testing, and for various other tests.
(1) 3400i storage system with 7 3202j expansion shelves	Dual Active-Active HA Controllers 8GB of total controller memory 4x10G and 4x1G Ethernet 128 15k RPM 600GB SAS disks	Used for Active-Active streaming IO tests
NISTNET	Dedicated Linux server with NISTNET software	Used to emulate WAN connections between replication targets
(3) Windows 2008 R2 Servers	8 core, 32GB memory	Various tests
(2) VMware ESXi Hosts	16 core, 48GB memory, hosting various Windows 2008R2 VMs	Used for replication testing and evaluation of vSphere plugins.
(3) Windows 2008 R2 Servers	8 core, 32GB memory	Various tests
(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.**



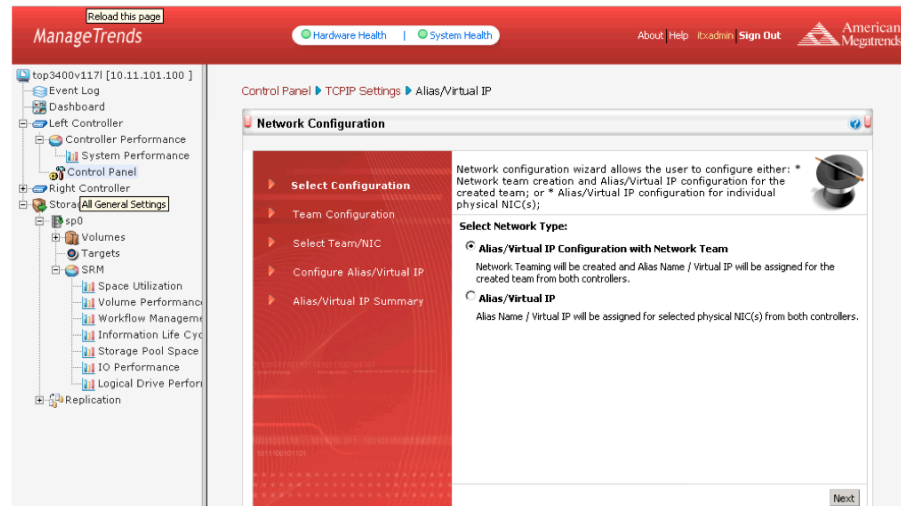
## OUT-OF-THE-BOX-EXPERIENCE – DEPLOYMENT EASE OF USE

We began testing of the StorTrends 3400i with the initial setup of a new 16 drive storage array. After the array was mounted in a rack, cabled to Ethernet switches, connected to line power and UPS power, and powered on, we began our test with initial deployment and provisioning of storage.

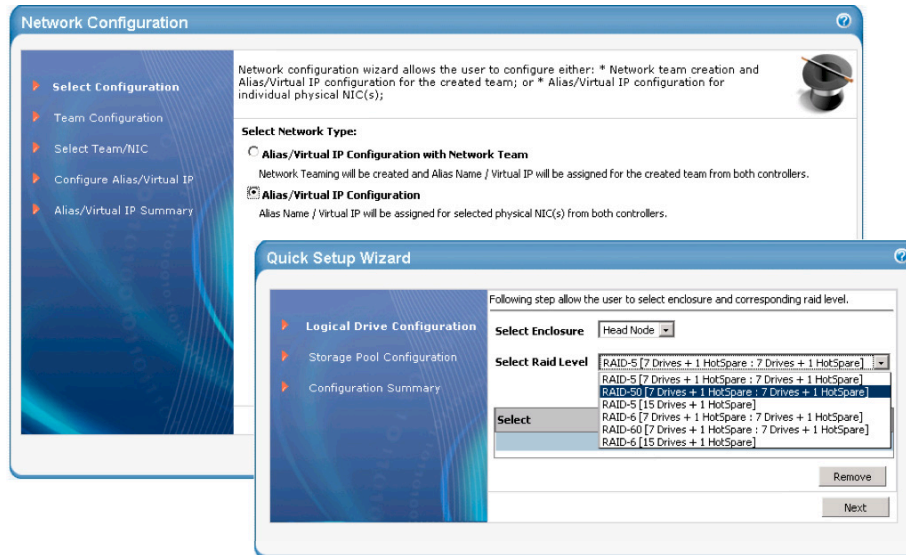
### DEPLOYMENT

Configuration was straightforward, and guided by a set of well-polished web GUI wizards that walked us through the

process of setting up initial IP addresses and storage volumes. On initial access, the management interface for the 3400i can either obtain an IP address through DHCP or have a static address assigned through a local console. After that address is configured, customers can access the management interface with a web browser, and will be greeted by initial configuration wizards. The first wizard walked through the network configuration information for storage interfaces, and the second wizard walked through the configuration of an initial storage volume. In both cases, wizards were full featured and presented us with all of the expected settings, including a range of RAID 50, 60, 5 and 6 options.

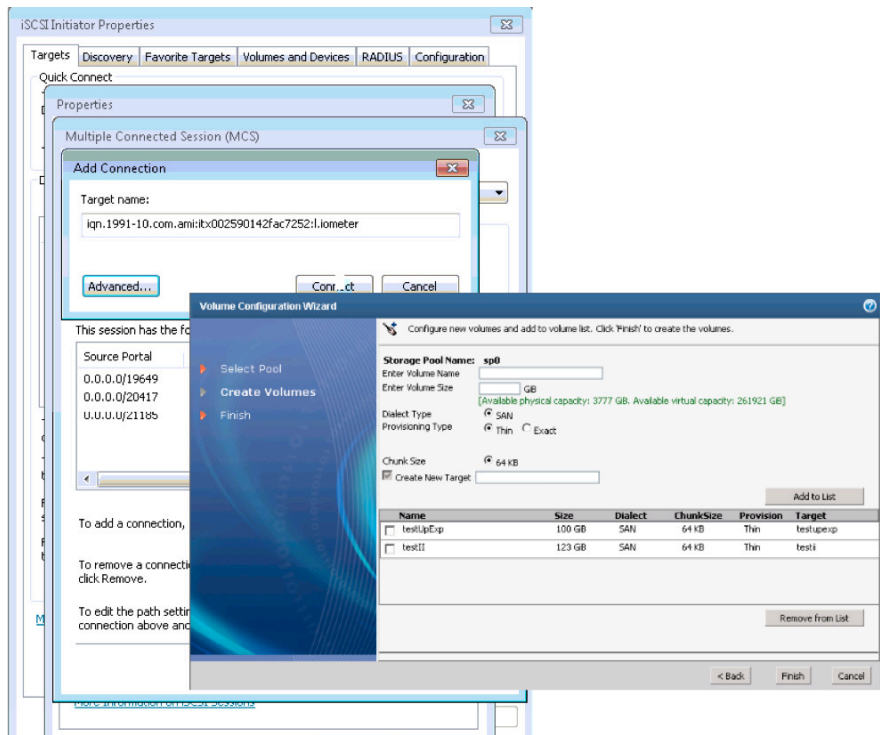


**Figure 1: The ManageTrends Web GUI is comprehensive and well-polished, and simplifies initial setup as well as on-going management.**



**Figure 2: From the outset, comprehensive ManageTrends wizards walked us through every storage setup and administration task.**

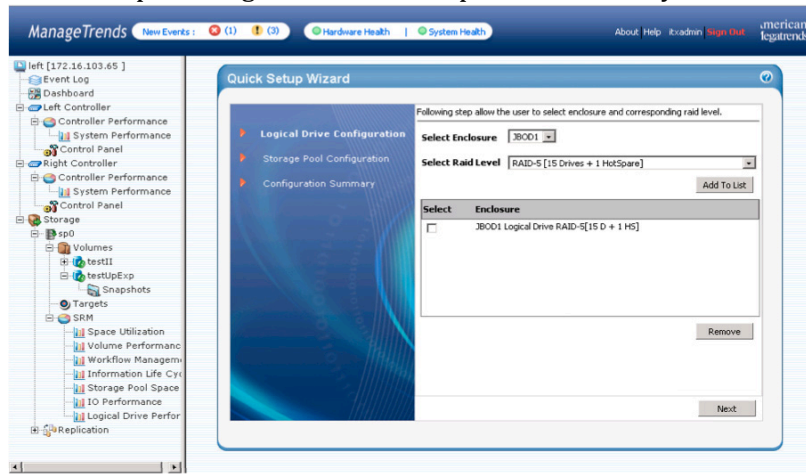
We next turned our attention to creating initial storage volumes and assigning them to targets. Once again, we continued with the easy to use wizard that walked us through all volume options, including features like thin provisioning. Using this wizard, we were able to automatically provision multiple volumes at once. After setting up volumes, we went to one of our Windows 2008 R2 servers and setup an iSCSI connection, as well as multipathed connections to our target. We then tested that target with IOMeter to verify connectivity. All told, this initial setup from system configuration to first storage volume connection took less than 5 minutes. Following this configuration wizardry, the volumes were ready to use immediately.



**Figure 3: After walking through a wizard to configure multiple targets, we simply copy and pasted target names into the Windows wizard to create 4 multi-pathed connections to the storage system.**

## SCALING

After initial deployment, we turned to scaling the capacity in the 3400i by adding a 3202j expansion shelf filled with 16 7200 RPM NL-SAS drives. We immediately turned our attention to this scaling exercise in order to bring on a mixture of SAS and NL-SAS disks for later testing. After cabling this shelf and powering it on, we were presented with yet another configuration wizard that walked us



through the process of configuring the disks into a RAID group and adding the capacity into the storage system. We chose to add this capacity to our existing total storage pool.

Additional storage capacity added to an array this way can be provisioned into an existing storage pool, or into a separate storage pool. Different types of disk provisioned into a single storage pool are automatically recognized as “tiers” of disk with each tier representing a different

**Figure 4: Adding an expansion shelf to our 3400i array.**

performance capability. Moreover, 15k RPM disks are automatically recognized as higher tiers – as discussed earlier in this report, StorTrends auto-tiering (or more specifically ZBR) places data on disk for optimal performance. ZBR works by segmenting the fastest tier into three different performance zones that are represented as three different tiers of storage performance. We’ll review this further later in this report.



**Figure 5: ManageTrends Web GUI displays everything from hardware indicators to capacity and performance trending.**

By adding this new storage shelf into our existing “storage pool” we were able to expand our total available capacity for existing and new volumes. Alternatively, a separate pool of storage could have been used to isolate this shelf for specific needs such as backup or archive, where data shouldn’t ever be tiered up onto faster disks. Once again, our entire exercise around adding an expansion shelf required no more than a few minutes.



In total, the collection of ManageTrends, StorTrends design, and various wizardry in the StorTrends interfaces make the 3400i array exceedingly easy to setup – clearly ranking with the best in the market. Furthermore, as we’ll review elsewhere, the same design philosophy that surfaces in setup makes the StorTrends 3400i equally easy to use over the long run. Key to this are a set of powerful, integrated storage features that make data protection, disaster preparedness, and virtualization simple and ready to use with the existing systems already operating in the business data center.


For the mid-range storage customer, the sum total of integrations and features create a significantly lower investment in specialized skills and reduces on-going time and effort spent on administration. Based on our informal assessment, our impression is that the StorTrends 3400i could easily require less than ½ the annual time and effort of similarly capable competitive arrays when the sum total of skillset development and on-going management is considered. For the SMB or SME customer, this can easily mean the difference between a storage strategy that supports the business, versus one that breaks the figurative back of the IT organization.

### ***Technology Validated:***

StorTrends 3400i setup, configuration, and management could easily require less than half the skill specialization and time and effort of competitive systems when considered over the course of a year. This is a byproduct of comprehensive design that streamlines and automates the most challenging parts of storage, and puts the right information right at administrator fingertips.

## AVAILABILITY AND RESILIENCY

While StorTrends has simplified setup, configuration, and management of the 3400i storage system, ease of use doesn’t go far without a similarly sharp focus on data integrity and performance. With this



**Figure 6: Conducting an IOMeter test with a failed drive (note the alert icon on drive bay).**

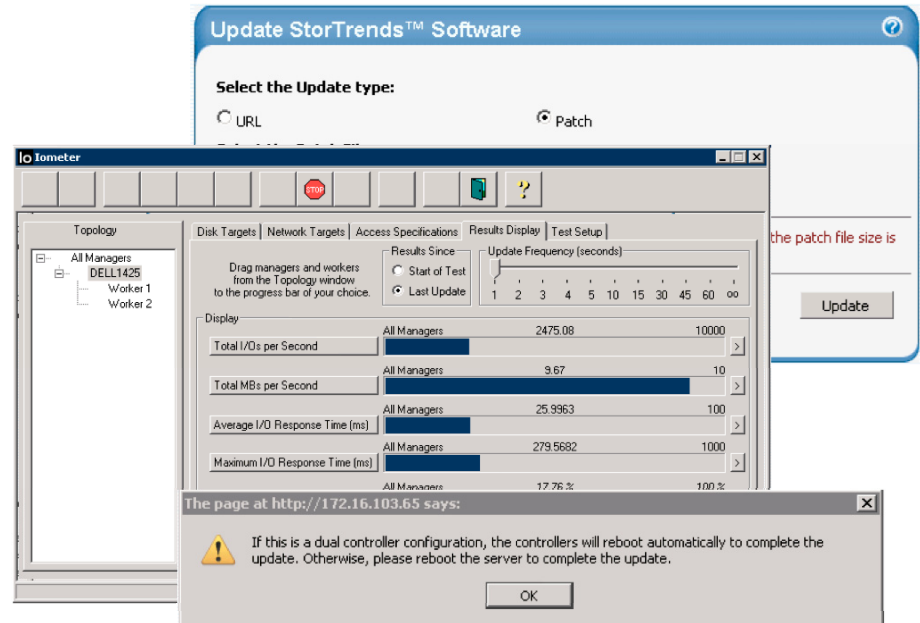
in mind, we next turned our focus to evaluating StorTrends availability and resiliency. In this regard, StorTrends claims the 3400i is architected “top-to-bottom” for high availability, and that they offer best of class features that will meet any storage administrator’s demands.

We evaluated storage performance during various failure scenarios by running IOMeter tests during failure events. We tested loss of a drive and loss of power with no substantial changes to performance.

But to carry our testing a little bit further, we set out to evaluate how controller failure might impact IO, and we evaluated this by running a firmware upgrade while also running an IOMeter test. Firmware upgrades execute on one controller at a time, and fail over all IO to the second controller while the upgrading controller reboots. Then the process is repeated on the second controller.

Our exercise was setup to simulate a scenario with no host multi-pathing and with all IO connecting to a single controller of the array. We wanted to verify the failover of the complete identity of one controller, as some customers may configure a system this way, and avoiding worst-case disruption from such a configuration would represent an above par availability feature. During our failover, we experienced a period of approximately 20 seconds during which IO was halted while the controller identity and active sessions were assumed by the secondary controller, and this time was largely a matter of network switch tables returning to a consistent state.

Beyond these lab test measures, we also reviewed 3400i system architecture in depth, and noted that the 3400i is both thoroughly engineered for availability and equipped with a number of features that can further enhance availability and resiliency beyond the capabilities of a single array. While we will touch on some features later in this report – specifically asynchronous remote replication – the StorTrends iTX software is also equipped with synchronous volume mirroring (available in certain StorTrends products) for the local LAN that can allow customers to set HA protection across two storage arrays, and further enhance the availability for clusters and mission critical applications. Altogether, the 3400i made a suitable showing of high availability and resiliency characteristics, and convinced us that it represents best-in-class capabilities for mid-range customers.



**Figure 7: Applying a patch, rebooting, and failing over IO to the other controller during an IOMeter test. The IOMeter test shows IO approximately 10 seconds after failover completed.**

### **Technology Validated:**

Our testing and review of the StorTrends 3400i architecture convinced us that the 3400i represents best-in-class availability and resiliency for dual controller mid-range storage systems, and this architecture passed our tests with flying colors.

## STORTRENDS: SOLVING A PERFORMANCE AND CAPACITY EQUATION


Storage efficiency is a complex equation, and a constant thorn in the side of most storage administrators. Storage is simultaneously about sufficient performance and sufficient, but not excessive, capacity. Managing both together add up to a nearly impossible to manage balancing act. Performance typically comes from adding extra rotational disks that are each responsible for churning out a limited amount of IOs per second. But such disks are expensive, and adding disks can easily result in more capacity than a business could ever use. Meanwhile, the constant creation of more data by today's productivity worker – from Microsoft Office files to video content to design data from products like AutoCAD – means that capacity is under constant pressure, especially in shared storage systems. Fighting capacity constraints may lead a business to add ever larger hard drives without a realization that capacity from big disks may come without sufficient performance. Efficiency comes from constantly managing performance and capacity in balance with each other, such that there are never excess power consuming amounts of either, but always sufficient amounts to support bursty business demands.

### CAPACITY EFFICIENCY WITH THIN PROVISIONING AND CAPACITY MANAGEMENT

StorTrends capabilities aim to create a complete capacity-efficiency recipe for the mid-range customer. One ingredient is a thin-provisioning volume-virtualizing architecture that aims to fully maximize the usable capacity from a storage system. First StorTrends 3400i can virtualize all attached disk into a single pool of capacity, from which any sized volume can be provisioned. Then, these volumes can be thin-provisioned, making it appear to the client that all of the volume's capacity is available, while in fact only enough capacity is consumed from the storage array to hold the real data within that storage volume. Then finally, StorTrends makes it possible to easily see how much storage is actually consumed and avoid the risk of running out of capacity behind these thin volumes with a ManageTrends report that clearly displays capacity utilization.

### OPTIMIZATION DRIVES PERFORMANCE

With capacity highly optimized, StorTrends then applies a rather unique set of technologies to deliver comprehensive performance efficiency. First, StorTrends' multi-threaded, multi-processor RAID-RT and wide-striping can harness the multiple controllers, processors and spindles that make up a 3400i



**Figure 8: The Capacity Utilization report and planner – ManageTrends displays historical utilization, including volume space as well as snapshot space, and provides a forward-looking forecast alongside a tool that allows administrators to see how long a capacity addition might last given current growth rates.**

array. RAID-RT is then paired with auto-tiering technology – what StorTrends calls Information Lifecycle Management (ILM) – as well as on-the-spindle data placement technology that StorTrends calls Zoned Bit Recording (ZBR). Both of these technologies are designed to dynamically adjust data placement, and thereby storage performance, in response to changing data access patterns. This allows the 3400i to extract maximum performance out of every storage spindle.

To put StorTrends’ claims about ILM and ZBR effectiveness to the test, we constructed a test approach that would intentionally place a test volume on slower SATA disks attached to our 3400i in an expansion shelf. We did this by provisioning a large “filler” volume that occupied all of the space on our fast SAS disks. Then we provisioned the test volume. Our theory was that as we started generating IO against the test volume, the blocks for that volume would be automatically migrated to the fastest disk, displacing some of our unused “filler” volume blocks.


Fortunately, StorTrends provides a detailed reporting interface that shows just what is happening with the blocks behind disk activity, and this allowed us to see this migration taking place (Figure 11).

To trigger the migration, we ran IOMeter against our test volume, allowing it to run for approximately 10 minutes. After 10 minutes, we monitored the progression of data tiering as ILM ran in the background and rearranged the placement of our data. Once all blocks were re-tiered, we ran IOMeter again to evaluate the improvement in performance.


In the background, even during our testing, data blocks began to rapidly move because we were not fully utilizing system resources. The 3400i monitors disk and processor activity and makes sure that system processes or background activities do not interfere with IO performance – StorTrends calls this feature “Workflow Management”.

Shown in Figure 10, Workflow Management is equal parts a load monitor and a predictive engine – by routinely profiling system utilization, Workflow Management can identify peak times of use and make sure background processes are scheduled around those typical peaks. Since our test system had no regular use patterns and only a light load, Workflow Management allowed tiering of data to rapidly

execute during our test.



**Figure 10: StorTrends Workflow Management profiles load on both controllers to intelligently schedule maintenance and background tasks like data tiering.**



**Figure 9: ZBR divides the disk platter into multiple zones, and uses the outside tracks – the rose colored tracks above – to hold the hottest data blocks in order to improve latency by reducing head movement, while also optimizing for better streaming throughput.**





**Figure 11: StorTrends 3400i ILM interface shows the auto-tiering of data across storage types and on the individual spindles themselves (each SAS disk is divided into three zones within Tier 0 above, to help optimize the access time and throughput for critical data). At the beginning of our test, 1.2K 1MB blocks underlying our test volume were stored on SATA disk (Tier1). By the end of our test all 1.2K 1MB blocks were on Tier 0.**

Of course, behind auto-tiering is a full range of settings for ILM that can restrict this auto-tiering functionality to the most important workloads, restrict it to certain times of day, or allow auto-tiering to constantly optimize everything.

ILM and ZBR proved to be quite effective during our testing, irrespective of the access pattern we threw at it. Before and after migration we sampled 4 different access patterns, using the averaged IOMeter results from 4 minute runs for our final measurement. The four patterns are below.

Auto-Tiering: StorTrends ILM and ZBR Performance Impact				
	Read – Random 8KB	Write – Random 8KB	Read – Seq. 1MB	Write – Seq. 1MB
Before (data resident on SATA)	3,766 IOPS	640 IOPS	497 MB/s	131 MB/s
After (data tiered to SAS and optimized on SAS by ZBR)	7,083 IOPS	1,286 IOPS	599 MB/s	157 MB/s
Improvement	+88%	+101%	+21%	+20%

**Table 3: Improvement in StorTrends 3400i performance as ILM and ZBR is used to re-optimize data from 7,200rpm SATA disk to 15,000rpm SAS disk.**

Clearly ILM and ZBR had significant impact, delivering as much as a 101% performance improvement in random write with a 8KB access pattern. We consider performance optimization the reason why our StorTrends 3400i environment was also able to deliver on one other notable accomplishment – a successful performance test of 10,000 Exchange users using Microsoft Exchange 2010 and Microsoft’s Exchange Loadgen test tool, as illustrated in Figure 12.

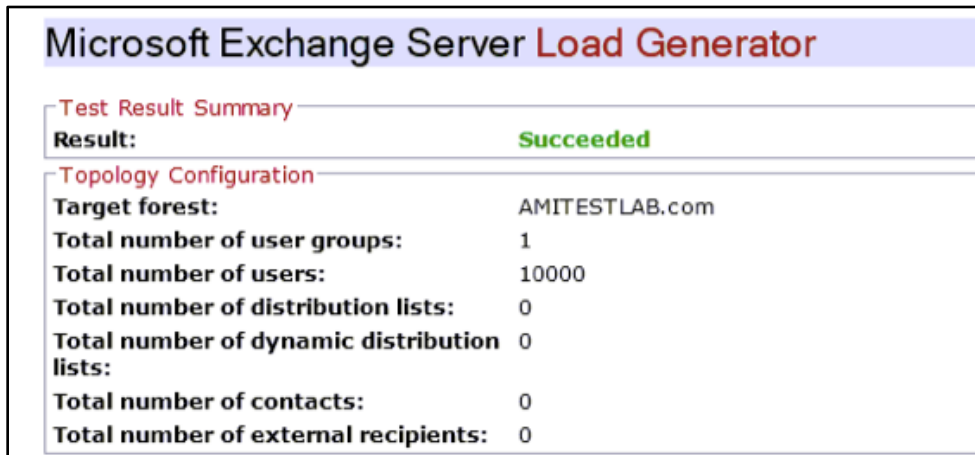


Figure 12: During our testing with the 3400i, we ran an Exchange Loadgen test using 10,000 users and an Outlook 2007 average user profile. The 3400i was able to successfully complete this test.

In our assessment, the collection of technologies within the StorTrends 3400i – from thin-provisioning to ILM and ZBR – deliver a powerful and comprehensive recipe for performance and efficiency, and will help StorTrends customers realize exceptional value from their storage systems that last longer and stretch further than typical mid-range arrays.

**Technology Validated:**

StorTrends 3400i is clearly designed to deeply optimize underlying disk while intelligently managing tiering activities. The 3400i surpassed our expectations in responding to bursty demands to deliver excellent performance.

**STORAGE FEATURES: ENTERPRISE-CLASS IN A MID-RANGE ARRAY**

While our testing to this point uncovered a number of notable StorTrends features – including easy provisioning, thin-provisioning, robust multi-pathing with active-active controller paths, and ILM+ZBR auto-tiering – a few features remained uninvestigated. We turned our attention next to storage management features such as snapshots, management plugins, and replication. As we began this exercise, it rapidly became clear that replication deserved a special focus, so we’ll turn next to first reviewing our testing of StorTrends snapshots and management plugins.

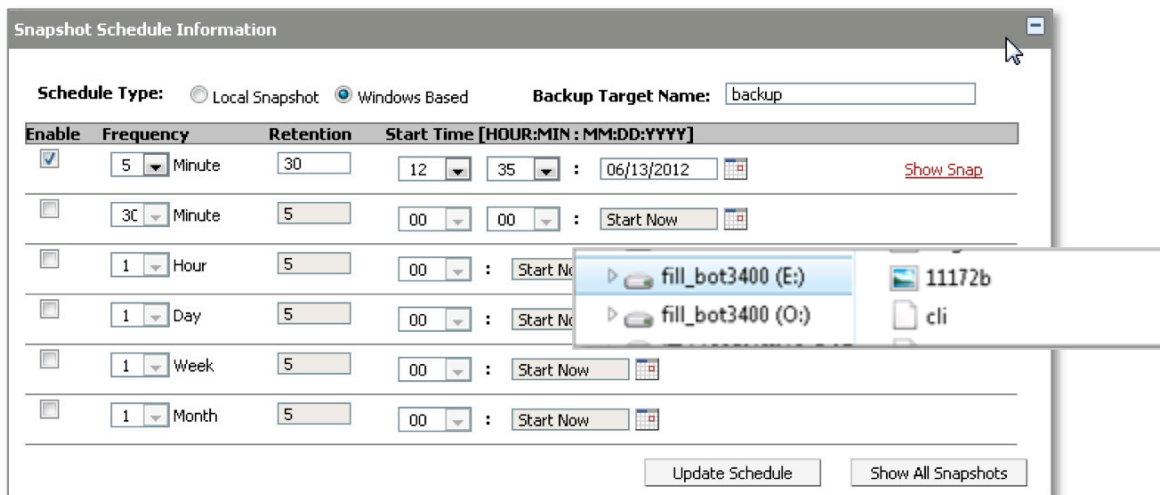


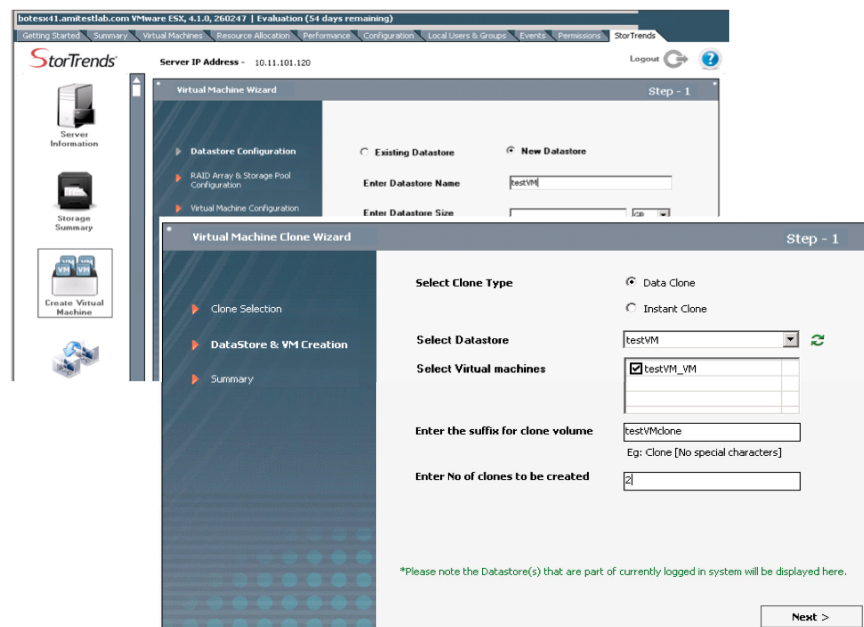
Figure 13 (above): A screenshot of setting up a 3400i snapshot for our Windows-based file server. The overlay box shows a file browse of the data being served up by this file server via a snapshot on a remote 3400i array (where the snapshot was remotely replicated).

StorTrends comes equipped – at no additional cost – with a robust set of snapshot and cloning technology and plugins, including Windows Volume Snapshot Service (VSS) integration for the protection of a wide-range of applications. StorTrends snapshot technology is built on a redirect-on-write approach that freezes existing blocks for reads, and redirects all new writes to new data blocks. This helps maintain space efficiency with snapshots by eliminating unnecessary duplicate data and minimizing the data that changes and grows inside of each snapshot over time. This redirect-on-write approach also allows administrators to create many snapshots, makes it possible to use snapshots as read-only or read-write, and helps maintain maximum performance from each snapshot.

StorTrends VSS integration can quiesce all popular Microsoft applications – such as Exchange, SQL Server, and Hyper-V guests – any time a 3400i snapshot is executed. This is particularly powerful as it guarantees a consistent protection point but still allows administrators to use the 3400i interface and scheduling engine as their single point of management.

But StorTrends capabilities don't stop there, as StorTrends is also deeply integrated with VMware's vSphere. Specifically, StorTrends comes with a vSphere management plug-in that enables vSphere-based execution of nearly all storage provisioning and management tasks, including snapshots and replication, as well as direct access to StorTrends management interfaces.

But particularly powerful is StorTrends integration with major vSphere tasks – StorTrends has fundamentally extended the integration and coordination between VM management and storage management. StorTrends, with a belief that the virtual infrastructure is better off from 1-1 connections between a storage volume and a virtual machine, has automated the task of provisioning storage volumes and made it possible to deploy a VM on a dedicated storage volume in a single wizard, *with no further interaction*. This makes it possible to conduct per-VM operations such as creating snapshot-based instant clones or full copy-based data clones while harnessing all of the performance and efficiency of the storage system, as well as ensuring that tools like ILM+ZBR are used against the most important workloads, without leaving the vSphere interface. Provisioning storage is just as fast as using a shared datastore, but made significantly more powerful because it allows the storage system to accelerate provisioning tasks, while keeping storage efficient.

















**Figure 14 (above): The StorTrends vSphere management plugin, showing a Virtual Machine creation wizard (background) and a Virtual Machine cloning wizard. Note, the cloning wizard can duplicate the volume within StorTrends or create a volume snapshot and add an instantaneous VM clone.**

During our testing, we evaluated the full range of functionality in StorTrends' vSphere plugin, using it to create virtual machines and datastores, view various storage statistics, and clone multiple copies of a virtual machine as both a full clone and a snapshot-based instant clone. Instant clones executed in less than a minute, and boot times for virtual desktops were consistently under two minutes, even when simultaneously booting more than ten instant clones.

At the end of our testing, it was clear that the StorTrends vSphere integrations compete with best-of-breed features from the biggest enterprise-class storage systems, where we see these technologies deliver significant returns by reducing the complexity and time and effort involved in virtual storage management.

### FULLY DR ENABLED, OUT-OF-THE-BOX

 testVM_VMtestVMclone1	Powered Off	47.00 GB	45.00 GB
 testVM_VMtestVMclone2	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant1	Powered Off	47.00 GB	45.00 GB
 suse11x64	Powered Off	512.00 MB	512.00 MB
 testVM_VMInstant2	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant3	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant4	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant5	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant6	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant7	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant8	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant9	Powered Off	47.00 GB	45.00 GB
 testVM_VMInstant10	Powered Off	47.00 GB	45.00 GB
 win2K8x64	Powered Off	44.00 GB	40.00 GB

**Figure 15: We created 9 StorTrends snapshot-based instant clones of an original Win2K8 R2 Server. Simultaneous booting of these VMs remained very fast.**

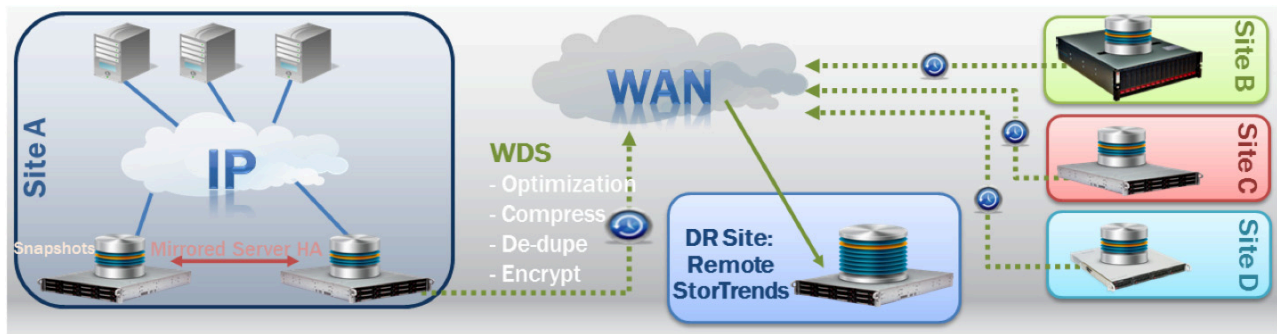
### Technology Validated:

StorTrends 3400i is equipped with a diverse set of storage features that rival or surpass bigger competitors, and will help StorTrends customers drive up the efficiency of their infrastructure through ease of use as well as the acceleration of key storage operations like VM cloning.

Moving beyond these core integrated features, we were notably surprised when we next turned our attention to replication. The StorTrends software comes with an unusually rich suite of replication technologies. These technologies include HA mirroring of volumes across systems or controllers, synchronous replication, and asynchronous replication. The asynchronous replication is particularly interesting, as StorTrends uses a snapshot-based approach that they call Snapshot-Assisted Replication (SAR). SAR uses a point-in-time snapshot to identify a replication point, and then only that data point is replicated. Since StorTrends snapshots can quiesce host applications and Hyper-V VMs through VSS enablement, snapshots are more likely to define an application consistent point in time where data is known to have good integrity and consistency. Moreover, SAR asynchronously keeps data up to data through delta-based snapshots that only transmit the data that has changed.

But that isn't all of the 3400i replication functionality – StorTrends also goes a step further in packaging data deduplicating and compressing WAN optimization technology into the 3400i. The StorTrends WAN-optimization Data Services (WDS) can be layered on top of SAR, and significantly reduce the amount of bandwidth utilized over WAN connections. This makes DR protection possible for nearly any office or business, with any type of connection.





**Figure 16: AMI StorTrends synchronous replication can be used to support HA applications across separated but same-site storage systems, while Snapshot-Assisted Replication can be used in a multi-site configuration to consolidate branch office data back to a centralized location.**

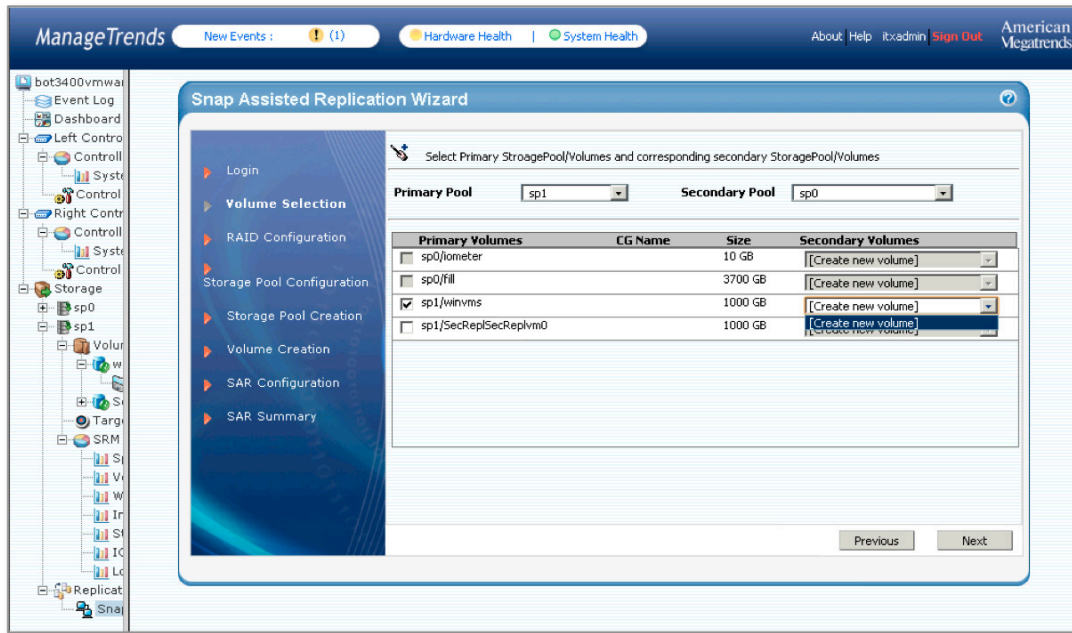
### TESTING A STANDARD NETWORK

SAR works with or without WDS turned on, so we first tested straight data replication without optimization to establish a baseline for un-optimized data replication.

We set up SAR with a GUI Wizard in the ManageTrends interface. This wizard comprehensively covered all steps (Figure 18, next page). In particular, this wizard fully automates the setup of snapshot replication (SAR) without any manual management of the storage system at the remote location. From one system, an administrator can choose an existing remote volume for replication or can select to automatically create a new volume.

We connected our source and destination volumes over a local LAN to a Linux server running NISTNET software. We used NISTNET to emulate a WAN link, constraining bandwidth to 20mbps and introducing a 50ms delay, a packet drop rate of 1%, and a packet duplication rate of 1%. This in turn simulated a WAN distance similar to what is commonly found between cross-country branch offices with a few slight packet errors to show a worst-case scenario.

Over this link StorTrends delivered a replication time of 1 hour, 39 minutes, and 41 seconds, with an effective transmission rate of 2.97 MB/s for our 17.35GB datastore volume. This is above the theoretical maximum of this link because of dedupe and compression, but would have had the undesirable effect of depriving a business of sending any additional traffic. This could quite possibly disrupt branch to branch or Internet communication, or could trade off completion time and leave a business unprotected for longer than desired.



**Figure 17: The ManageTrends web GUI makes configuration of replication easy, allowing an administrator to select a replication source and automatically create a replication target on a remote array from a single screen (of course replication can be configured to use an existing volume as well).**

**TESTING WAN-OPTIMIZED DATA SERVICES**

We then turned on StorTrends’ WDS, electing the full suite of encryption, deduplication, compression, and most importantly acceleration. Once again we turned to the SAR wizard and created a brand new replication relationship and remote data volume from the same snapshot point in time (we needed to capture the full volume over again, and using our existing relationship would only capture data that changed after the initial run).

This time we transmitted the same data volume in 1 hour, 3 minutes, and 52 seconds. But more importantly, we used less aggregate bandwidth, even though we configured the array to use as much bandwidth as needed. All told, we transmitted less than 8GB of data with a higher effective throughput of 4.64 MB/s, while using bandwidth of 2.5 MB/s with StorTrends’ WDS. Consequently this would have been significantly less disruptive to other network traffic.

Of course, as time goes on these replication jobs will shrink in size, and time to completion will speed up, but WDS is still likely to deliver much faster time to completion. We took our next snapshot after we ran a variety of IOMeter benchmarks and routinely used the VMs over a course of a day. This new snapshot added up to 264MB of data (Figure 19).

Volume Pairs Involved			
Volume Count			
1			
Local Volume	Remote Volume	Current Snapshot Delta Size	Completed Size
winvms	SecReplwinvms	264256 KB	131072 KB


**Figure 18: Replicating a second copy of the original volume resulted in a transmission of only 264MB of data, as the primary data was already there from the first replication task. This demonstrates the benefits of delta-snapshots and the SAR approach to replication. Data can easily be kept updated, keeping the business protected against disaster or data loss, with delta-snapshots.**

## DATA RECOVERY

Throughout the course of our testing, we verified that snapshots could be used for both application and file-level data recovery. For application testing, we used Microsoft Exchange as well as Windows 2008 R2 virtual machines running on ESXi. Snapshots worked as expected, but we also noted that snapshot history was retained on the remote storage array, and thereby preserved multiple data protection points at the remote site as well as on the source system. Customers would in turn be able to recover from extreme tragedies by restoring data to a point further back in time. One such scenario that might make this necessary would be data corruption followed by the complete loss of a storage system at a primary site.

## VISIBILITY AND MANAGEMENT – REPLICATION UTILIZATION

ManageTrends also provides detailed analytics on the data replication and WDS processes. The below figure is the ManageTrends dashboard analysis of WDS bandwidth utilization. Here, the WDS analysis provides an estimate of bandwidth that would be used with and without WDS. While the maximum bandwidth number reflects a point-in-time measurement that is higher than we actually achieved, the data reduction reflects the effectiveness of WDS (Figure 20).



**Figure 19: The reporting tools inside the ManageTrends web GUI show summary information on the effectiveness of WDS for our replication test. The original replication data was reduced by 52%, utilizing minimal bandwidth.**

The inclusion of data availability technologies in the form of HA synchronous mirroring, asynchronous remote SAR, and WDS also make the 3400i a key enabler for business protection, without additional acquisition costs, licensing costs, or management complexity. The 3400i can enable a level of protection and data movement for consolidation or other purposes that has seldom been available to mid-range storage customers. For typical SME customers, matching traditional technologies added to other storage systems can easily approach, and sometimes surpass, the \$100,000 mark.

The 3400i could easily be an attractive disk-to-disk backup target for many customers, and abundant capacity could make the 3400i a deep, long-term archiving platform. With replication, all of that consolidated data can easily be protected against disaster, or simply consolidated from the branch back to a primary data center. As we've highlighted, that replication can happen with extreme efficiency that is likely to make DR a better fit for the bandwidth constrained branch office, as well as control the costs of using that bandwidth for DR.

### **Technology Validated:**

The StorTrends 3400i's built in features deliver comprehensive storage functionality at no additional cost. Moreover, some of these features, like replication and WDS replace entire stacks of hardware and software that previously were out of the realm of affordability for most SMB and SME customers. These packaged features promise to enable a new level of data center capability like the SMB and SME have seldom been able to put together before. In our assessment, this gives the 3400i a class-leading ranking in delivering comprehensive value.

## TANEJA GROUP OPINION

During the course of this exercise, we set out to assess StorTrends 3400i storage against four criteria: reduction and avoidance of complexity, availability and resiliency, performance, and the extent of features that better enable data management and protection. As we went through our exercises it became rapidly clear that the StorTrends 3400i delivers against these criteria and is a highly integrated and extremely well-rounded system for mid-range storage customers – a solution that altogether would be incredibly difficult for competitive solutions to outpace.

Not only are 3400i core features tightly woven together to reduce administrative complexity, but the system is simultaneously balanced out with a level of deeper control and visibility that can let administrators who are interested peer further into the storage system. Simultaneously, integration goes far beyond a polished and highly usable surface set of management tools, but runs throughout the 3400i storage system architecture and across all of the underlying technologies like RAID-RT, thin-provisioning, ILM, ZBR, and SAR that are all dependent parts of a comprehensive set of storage technology. That set of technology altogether delivers outstanding availability, a uniquely efficient recipe for performance, and complete storage management without additional costs. Finally, to top it all off, StorTrends has packaged in a comprehensive set of replication functionality that looks like a separate solution unto itself – but with StorTrends, it is all part of the 3400i package.

That total package leaves very few spots unturned, and looks set to equip the mid-range storage customer with the most comprehensive set of storage and data management that we've seen from any vendor selling a traditional hardware array. For the agile business customer looking to dispense with complexity, but still obtain all the storage firepower they need to run real business applications, the StorTrends 3400i looks a serious contender that would be hard, if not impossible to beat. The value is certainly tremendous – the StorTrends 3400i simply goes further than typical storage arrays we see in the mid-market today.

### ***Technology Validated:***

At the completion of our testing, we're convinced – customers in the market for small and mid-range storage systems can rest confidently on a choice of StorTrend's 3400i. Complete with easy expandability and a feature list that puts many other major vendor storage systems to shame, the 3400i is equipped to go the distance in terms of performance, reliability, and adaptability and be a lasting part of the SMB and SME customer's IT infrastructure.

---

*NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.*