INTEGRATED DISASTER RECOVERY:

Technologies for Comprehensive Storage Array Protection

SEPTEMBER 2012

DR has long been particularly challenging for the midmarket customer. It usually requires multiple layers and components, host-based software, replication gateways or appliances, and often array-based functionality that is licensed and managed separately. Add to this complexity the need for robust bandwidth or an expensive WAN optimization approach and it's no surprise that DR can have a significant impact on both OPEX and CAPEX budgets.

The cost to manage all of these different elements can dwarf the cost of a primary storage system itself. The enterprise faces many of the same challenges, but they also have bigger budgets and more specialists to manage the complexity. Midmarket businesses and organizations may not have the same level of budget and specialists, but they certainly face the same challenges.

Recently, Taneja Group Labs put the StorTrends 3400i array through a Technology Validation exercise to evaluate how StorTrends measured up as a SMB/SME storage solution in terms of ease of use, performance, availability, adaptability, and innovative features. Over the course of our Technology Validation exercise, it was clear that one particular StorTrends capability rose above all others: StorTrends built-in multi-site WAN optimized data replication. Specifically, StorTrends suite of replication functionality looks poised to equip SMB and SME customers with the tools that for the first time makes robust DR really achievable. In this report, we'll highlight what we found, and why it stood out.

FOCUS ON STORTRENDS

StorTrends, a division of Norcross, GA-based American Megatrends Inc. (AMI), has been delivering innovative storage technologies since 1989. The current StorTrends product line was developed by the same team at AMI that developed the AMI MegaRAID card. That card eventually helped AMI become the largest RAID controller provider in the market. Leveraging their foundational background in RAID, AMI has maintained a serious footprint in storage as well as a range of componentry for x86 compute systems – from remote system management components to system BIOS components. The current storage portfolio under the name of StorTrends 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 array.

THE STORTRENDS 3400I ARRAY

As their flagship array, the 3400i starts with a 3U dual active-active controller array that houses 16 disks and scales up to 128 disks. 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, 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 8GB of system memory, and uses this

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

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 IO connectors on the 3400i and further connecting through SAS expanders on each expansion shelf.

StorTrends: DR for the Mainstream, No Strings Attached

One of the most notable features of StorTrends' array is 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.

In addition to establishing an initial replication point, SAR uses a delta-snapshot approach for on-going replication. After the first snapshot is taken and replicated, all future snapshots only capture the data that is different from the original snapshot. The replication snapshot occurs as part of the normal snapshot process. It creates a read/write point in time, with all data changes being directed to new blocks. These new blocks are easily identified and tracked. Then the next snapshot freezes these blocks, and makes it easy to transmit one entire set of changes. Since new data is then being redirected (redirect-on-write) to new blocks, there is no risk of data corruption or overwrite. This helps to eliminate performance overhead sometimes associated with other snapshot techniques. All of this in turn makes SAR a particularly powerful data protection tool.

Challenging Issues for Mid-Range DR	The StorTrends Solution
Replication requires add-on cost and complexity	Built-in replication technologies. StorTrends offers replication options including HA mirroring of volumes across systems or controllers, synchronous replication, and asynchronous replication.
Inconsistent Application Consistency	Snapshot-Assisted Replication (SAR). A point-in-time snapshot identifies and replicates a single replication point. Theses snapshots define application consistent points in time for good data integrity and consistency. Consistency groups can deliver consistency across multiple volumes by synchronizing snapshots.
Snapshot Volume Sizes	SAR's delta-snapshot approach. After the first snapshot is taken and replicated, all future snapshots only capture the data that is different from the original snapshot. The replication snapshot occurs as part of the normal snapshot process.
Limited WAN Bandwidth	WAN optimization with WDS. StorTrend's WAN-optimization Data Services (WDS) reduces transmit data through deduplication and compression, optimizes network characteristics for fast transport, and encrypts replicated data for security.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Limited Multi-Site Options	Multi-target replication. SAR supports a multi-site configuration with many-to-one, and bi-directional replication. StorTrends also, supports a mixture of platforms and disk types within the StorTrends product family, enabling customers to configure a cost- and capability-optimized solution.

But SAR becomes even more powerful because of one additional packaged AMI technology: WAN optimization in the form of StorTrend's WAN-optimization Data Services (WDS). Every StorTrends 3400i comes packaged with WDS functionality under the covers, which deduplicates, compresses, and encrypts data that is replicated to another StorTrend's node, and can intelligently profile network characteristics and optimize the level of deduplication and compression to make the most of network bandwidth, no matter the link size or the latency characteristics.

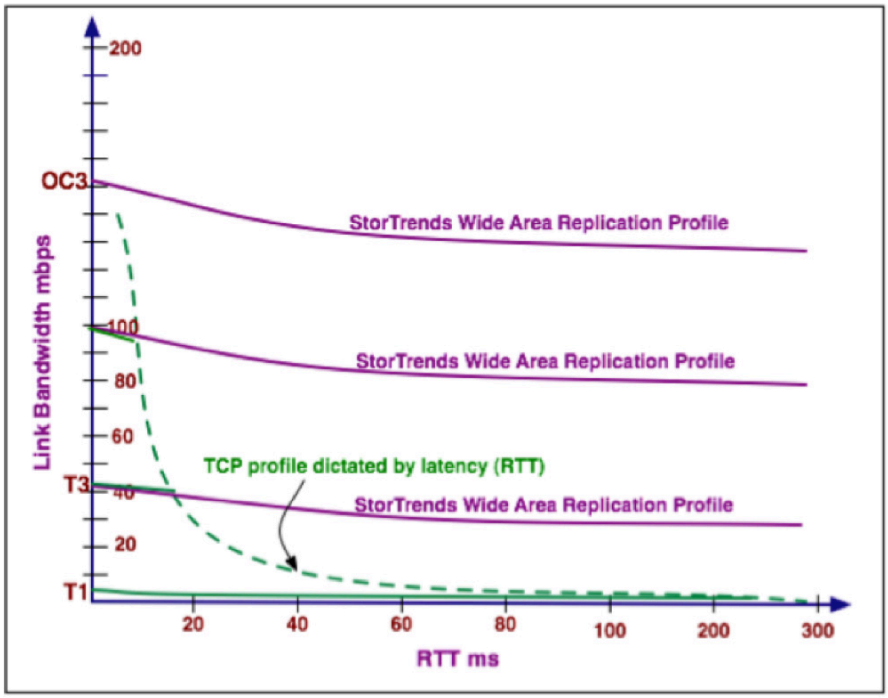

Figure 1: AMI StorTrend's illustration of the impact of latency (RTT ms) on effective bandwidth. At 50ms of latency, the rate of transmission can be greatly reduced. By compensating for round trip acknowledgements and pushing more data over the wire at a time, a technology like StorTrends WAN optimization Data Services can deliver higher effective throughput.

Latency characteristics actually have a far greater impact on long distance data transmission than available bandwidth. As is illustrated in Figure 1 above, because of in-flight buffer sizes and round trip acknowledgement impact, a marginal increase in latency can greatly reduce effective data throughput given typical packet sizes and network systems. Deduplicating and repackaging this data, as well as tuning the transmission characteristics can optimize how much data is successfully transmitted.

Of course any 3400i replication task, whether optimized or not, can be throttled down to a limited amount of available bandwidth to avoid disruption to other applications, and/or can be scheduled during specific parts of the day.

Finally, since snapshots are per-volume constructs, SAR can also be used in a multi-site configuration, with multiple remote systems replicating individual volumes onto a centralized storage system. Utilizing this volume based model, StorTrends' SAR allows for customers to host live volumes at each of their locations providing bi-directional as well as multi-site replication. Also note, while snapshots happen on a per-volume basis, volumes can be configured in consistency groups that can be snapshotted simultaneously to protect multiple volumes at the same time (Figure 2).

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

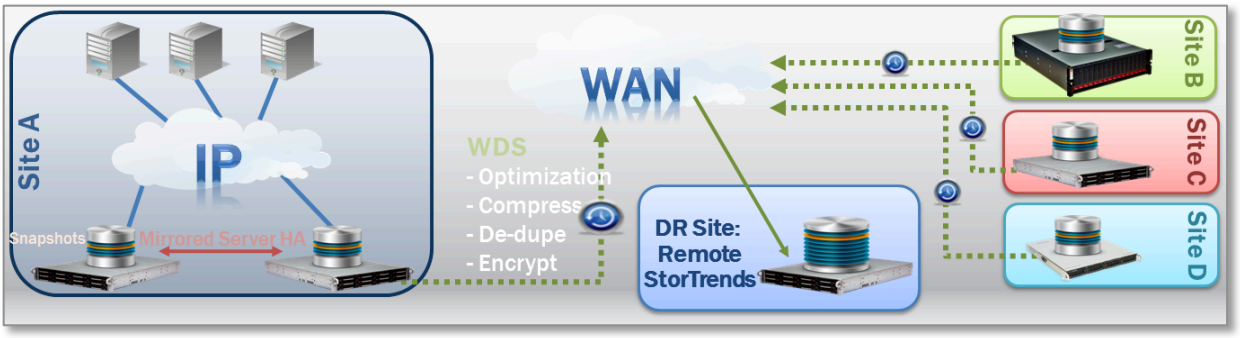

Figure 2: AMI StorTrends synchronous replication can be used to support HA applications across separated but same-site storage systems, while Snapshot-Assisted Replication can be used in a multi-site configuration to consolidate branch office data back to a centralized location.

For the VMware virtual infrastructure, StorTrends provides a full-featured set of plug-ins that perform block-level snapshots, datastore-SAN management and internal cloning operations to ensure that data is recoverable on a remote storage system. StorTrends also provides Windows VSS writers that can perform the same quiescing for Windows applications - SAR works just as well for Windows apps as it does for VMware.

Testing StorTrends 3400i WDS and DR

In a recent Taneja Group Technology Validation exercise, we turned a lens toward StorTrends replication capabilities, with our focus on evaluating how replication delivers value to real world StorTrends customers. We set out to compare non-optimized replication to StorTrends optimized replication, using WAN emulation software to create real world conditions, and a test data set consisting of a VMware vSphere 5 datastore housing several Windows 2008 R2 virtual machines.

TEST ENVIRONMENT

Test Environment Equipment at Start of Test	
System	**Specifics**
2 StorTrends 3400i's	16 disks –15k 146GB SAS disks configured in two 7 disk RAID 5 plus 1 hot spare disk groups configured in Active/Active mode
WAN Emulation	NISTNET software to emulate a 20Mbps link between these two systems with 50ms of delay, a 1% packet drop, and a 1% packet duplication rate.
Replicated Data	vSphere VMFS datastore volume that was 17.35GB in size, and contained several VMs.

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 4, next page). Notably, 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.

87 Elm Street, Suite 900 • Hopkinton, MA 01748 • T: 508.435.2556 • F: 508.435.2557

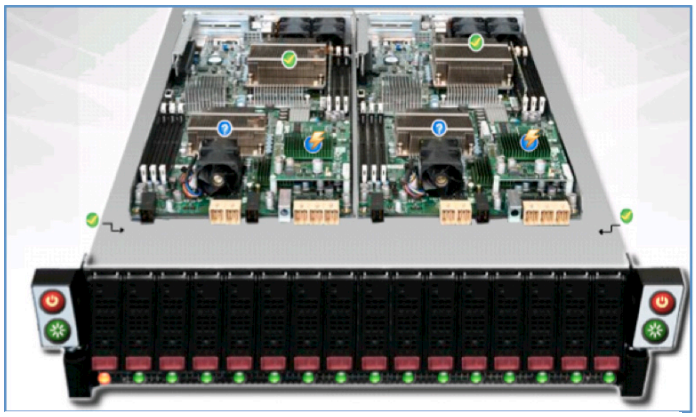

Figure 2: The StorTrends 3400i storage system used in our testing. We used two of these systems, connected through network links to a Linux server running NISTNET software that emulated a real world-like WAN connection.

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.

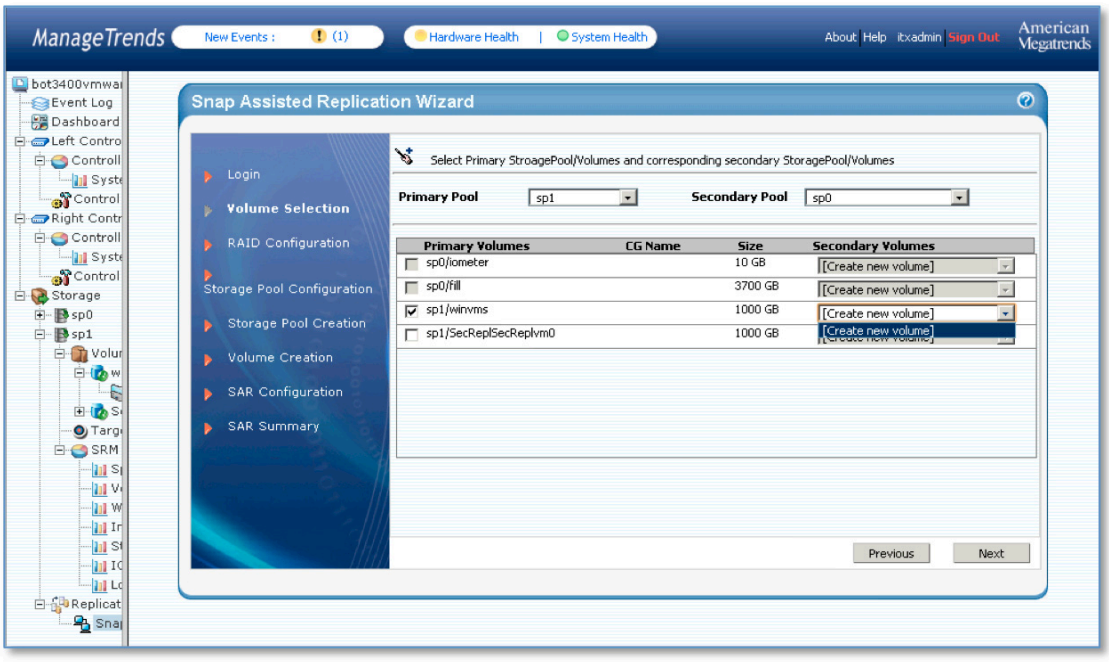

Figure 4: The ManageTrends web GUI makes configuration of replication easy, allowing an administrator to select a replication source and automatically create a replication target on a remote array from a single screen (of course replication can be configured to use an existing volume as well).

TESTING WAN-OPTIMIZED DATA SERVICES

We then turned on StorTrend's 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 hr, 3 minutes, and 52 seconds. But more importantly, we performed this replication with less aggregate bandwidth, even though we elected full utilization of the connection. All told, we transmitted slightly less than 8GB of data yielding a higher effective throughput of 4.64 MB/s, while consuming 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 although WDS is likely to still 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 5).

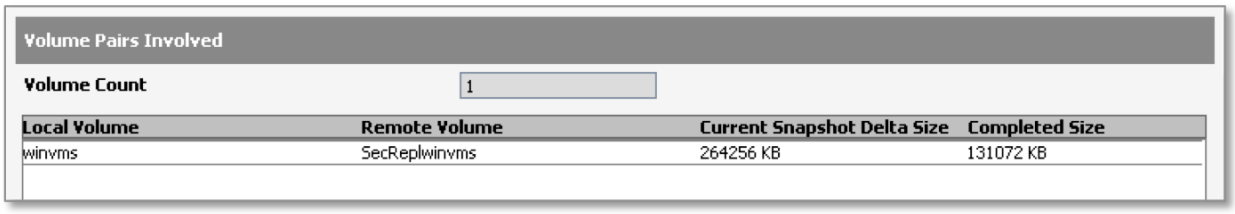

Figure 5: Replicating a second copy of the original volume resulted in a transmission of only 264MB of data, as the primary data was already there from the first replication task. This demonstrates the benefits of delta-snapshots and the SAR approach to replication. Data can easily be kept updated, keeping 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. This can allow an organization to recover from extreme tragedies by recovering to a data 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.

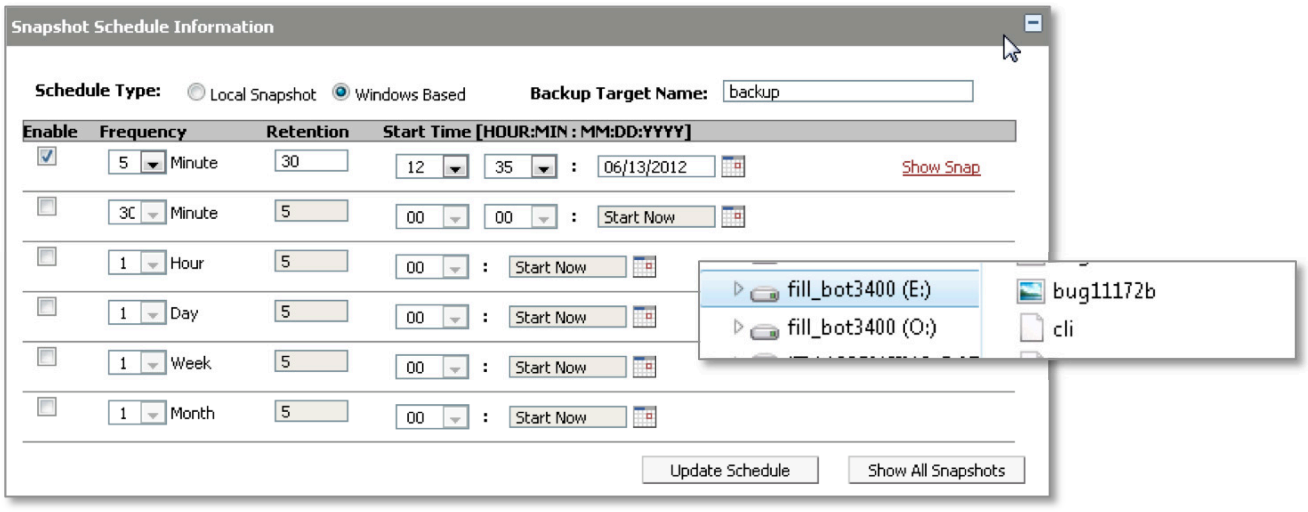

Figure 6: A screenshot of setting up a Windows-based snapshot for our file server, and simultaneous browsing of file data served up by this file server running from a snapshot on the remote replication target as well as the original file server (in small overlay box).

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

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 WDS (Figure 7).

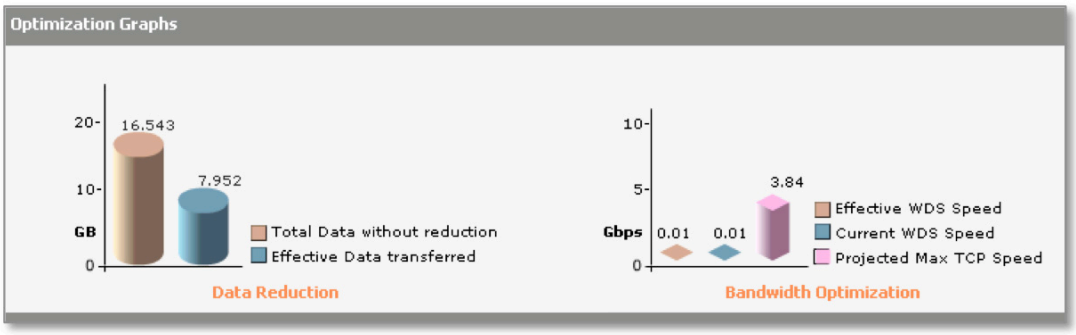

Figure 7: The reporting tools inside the ManageTrends web GUI show summary information on the effectiveness of WDS for our replication test. The original replication data was reduced by 52%, utilizing minimal bandwidth.

Taneja Group Opinion

StorTrends 3400i is a full-featured solution for the midmarket storage customer. 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.

As we've discussed elsewhere, the 3400i array comes with considerable versatility. Multiple types of storage can be easily mixed to address different storage demands (capacity or performance) and the data across these disks will be optimally placed with very little effort by the administrator, as the 3400i constantly tunes storage pool performance with a collection of wide-striping, auto-tiering, and disk tuning technology (called Zoned Bit Recording). This adaptability and optimization can make the 3400i a versatile system for multi-purpose storage, and it is easy to imagine how this versatility coupled with replication can extend the value of the 3400i even further.

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. And 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:

StorTrends 3400i's built in replication and WAN optimization technologies can be an indispensable boon to businesses trying to increase their resiliency and disaster preparedness - we are hard pressed to identify any other storage system with similar levels of single system scalability and performance that also come fully enabled for such a broad range of HA and multi-site DR capability with no additional cost, nor with similar ease of use and simplicity.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557