

## **Synchronous replication in a heterogeneous storage server environment**



© Copyright 1998-2008 American Megatrends, Inc.  
All rights reserved.  
American Megatrends, Inc.  
5555 Oakbrook Parkway, Building 200  
Norcross, GA 30093

## **TRADEMARK AND COPYRIGHT ACKNOWLEDGMENTS**

This publication contains proprietary information that is protected by copyright. No part of this publication can be reproduced, transcribed, stored in a retrieval system, translated into any language or computer language, or transmitted in any form whatsoever without the prior written consent of the publisher, American Megatrends, Inc. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. American Megatrends, Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

## **FOR ADDITIONAL INFORMATION**

Call American Megatrends at 1-800-246-8600 for additional information. You can also visit us online at [ami.com](http://ami.com).

## **LIMITATIONS OF LIABILITY**

In no event shall American Megatrends be held liable for any loss, expenses, or damages of any kind whatsoever, whether direct, indirect, incidental, or consequential, arising from the design or use of this product or the support materials provided with the product.

## **LIMITED WARRANTY**

No warranties are made, either express or implied, with regard to the contents of this work, its merchantability, or fitness for a particular use. American Megatrends assumes no responsibility for errors and omissions or for the uses made of the material contained herein or reader decisions based on such use.


## **DISCLAIMER:**

Although efforts have been made to assure the accuracy of the information contained here, American Megatrends expressly disclaims liability for any error in this information, and for damages, whether direct, indirect, special, exemplary, consequential or otherwise, that may result from such error, including but not limited to the loss of profits resulting from the use or misuse of the information contained herein (even if American Megatrends has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to American Megatrends at the address shown on the back cover of this document. American Megatrends provides this publication "as is" without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a specific purpose. Some states do not allow disclaimer of express or implied warranties or the limitation or exclusion of liability for indirect, special, exemplary, incidental or consequential damages in certain transactions; therefore, this statement may not apply to you. Also, you may have other rights that vary from jurisdiction to jurisdiction. This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. American Megatrends may make improvements and/or revisions in the product(s) and/or the program(s) described in this publication at any time.

## Introduction

Today, Disaster Recovery (DR) and Business Continuity (BC) are inarguably the most important and talked about topics when it comes to Storage Servers. A DR plan is no longer a luxury but is mandated by various business needs and regulatory and compliance requirements. Availability and safeguarding of Data which is the lifeblood of storage servers is very important. There are various ways in which a DR plan can be deployed. The two most important parameters that govern this are the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO). These two attributes essentially are a measure of data loss and down time respectively. Some critical Business segments like Banking and Airlines Industries cannot tolerate any loss of data i.e. an immediate RPO is mandated. Time is money, and thus it is always desirable to keep the RTO very low to contain the 'down time' to a minimum. For such critical business applications, Synchronous replication is the only solution for a viable DR strategy.

As is obvious, due to the stringent requirements imposed by such a DR solution, Synchronous Replication is one of the most expensive to implement and deploy. So to stay within the conditions of requirements, various replication solutions are erected to drive the cost low. The most common practice followed to keep the TCO under control is to cut the cost down at the recovery site. After all, synchronous replication is a peace-of-mind insurance policy and the recovery site is used only during disasters. So in the event of total outage at the primary site, the recovery side can be failed-over to and relatively less expensive and under-performing resources can be deployed to keep the business going. Once the situation is corrected, operation can be Failed-back to the primary storage site. Another common practice to keep the cost low is to have a central recovery site, where many primary sites can be replicated to. So in essence, this central recovery site works as a consolidated recovery storage server. StorTrends, with its thin provisioning and provision tracking features along with efficient ROW snapshots is an ideal candidate for such a solution.



**Figure 1: Many-to-one Replication**

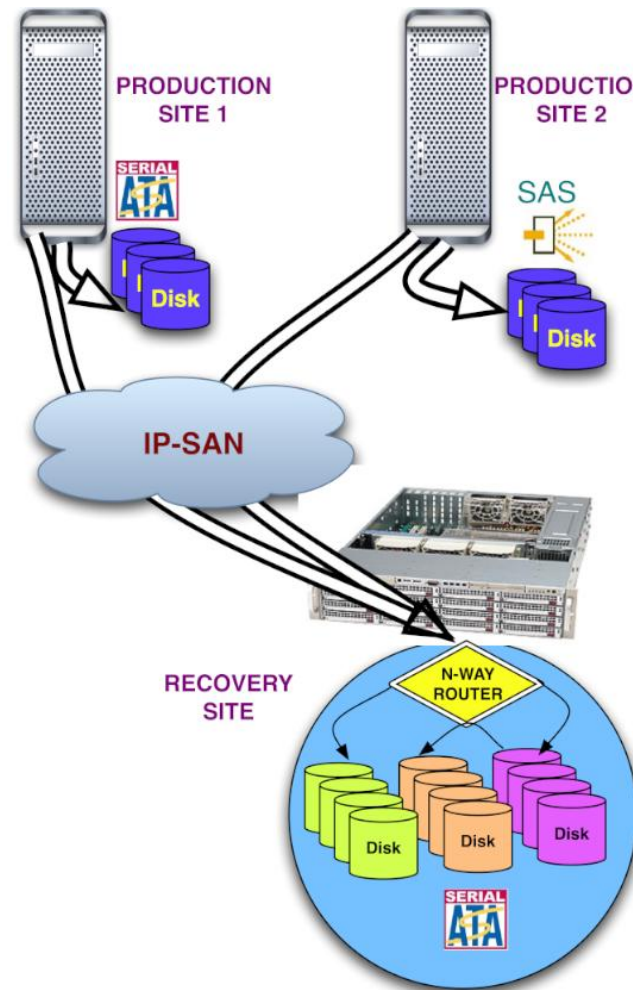
## Synchronous Replication and Performance issues

The biggest challenge here is the I/O latency. In synchronous replication every write I/O has to be replicated to and acknowledged from the recovery server before the I/O can be completed to the application host. Typically an average latency in excess of 25-30 ms stretches the limit of tolerance and may render the system unusable. Also, most of the applications that require synchronous replication are 'spindle hungry' and transactional in nature. So it is best for the primary site to have multiple high performance disk drives. The 15K or 10K RPM SAS drives with 4-5ms latencies can constitute a truly enterprise class storage server at the production site. To cut a balance between capacity and performance StorTrends with 'tiered' storage architecture becomes an ideal solution for the primary site. Now to effectively and synchronously replicate this, the recovery site also requires an expensive installation with SAS disks. This drives the cost of installation through the roof.

## StorTrends, the affordable solution for Synchronous Replication

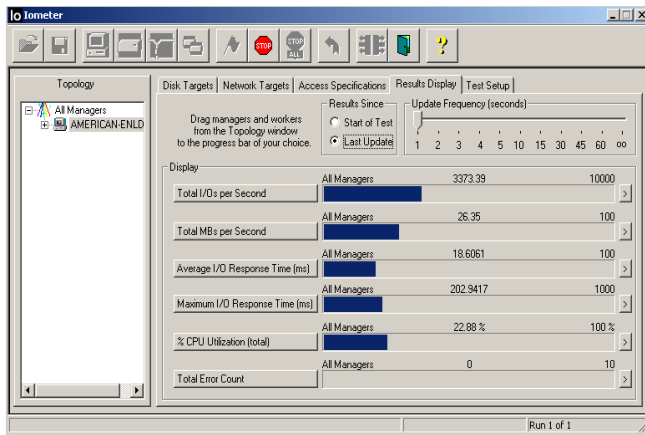
StorTrends, with its patent pending storage stack with run time configurable N-Way router provides a simple and cost effective solution to the problem. This allows a relatively low-cost recovery server with large and inexpensive SATA drives even with fewer spindles to effectively replicate enterprise class production servers with multiple SAS disks. This also allows a central and affordable recovery site to host SR for multiple primary sites. Another obvious advantage of such a configuration where large and inexpensive SATA drives are used is to have many more snapshots on the recovery site. Here all the incoming replica I/Os are fed through an efficient N-WAY path router that analyzes the incoming I/Os and schedules them to appropriate devices for rendering high performance and low latency. This routing is transparent to the outside world apart from the fact that the replicated IOS complete quickly. During this mode of operation, the recovery storage node cannot field host I/Os for the replicated volumes. Essentially the replicated storage auto-configures itself to a fast 'replication only' mode which precisely matches the need of a SR deployment. Added to this an intuitively designed I/O gating algorithm allows the recovery server even to function in write-back mode. With this combination the latency is further reduced. StorTrends utilizes yet another performance enhancing feature where multiple outstanding small writes can be packaged as a big I/O and sent across. This, together with support for Jumbo Frames (wherever applicable), improves the performance and CPU utilization by reducing the number of interrupts. In the event of a primary site outage, the recovery server stack with its workflow analysis module quickly readies itself to a state where it can field host I/Os after a failover. This allows the recovery stack to deliver low latency replication and low RTO for failover.

The following figures illustrate the performance of a SAS storage unit under random I/O load that is typical of a database server. With 8 Drive RAID 0 volume of size 30 GB the SAS storage unit was giving about 3300 IOPS with a latency of about 18 ms. However, when a contemporary synchronous replication was setup between this SAS storage unit and a SATA secondary unit the performance dropped dramatically to about 576 IOPS with an average latency of about 800 ms. These figures when compared to StorTrends synchronous replication shows the benefits derived from StorTrends replication. In this it can be seen that the StorTrends synchronous replication from SAS to SATA nodes give about 3100 IOPS with a latency of about 18 ms. Thus StorTrends replication clearly scores over contemporary replication technologies delivering near-SAS performance benefits.

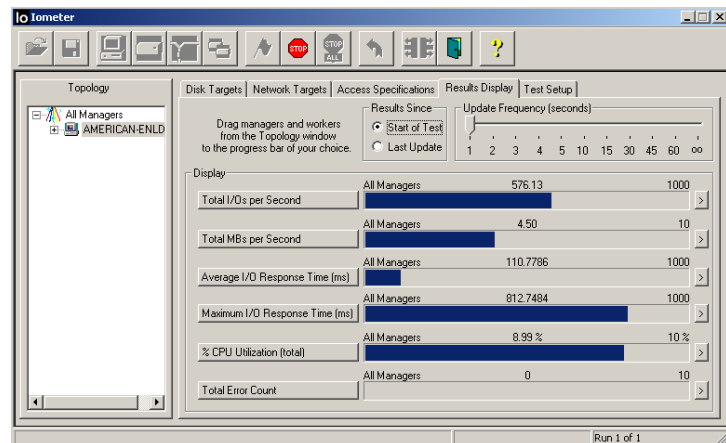


**Figure 2: Heterogeneous SR configuration using N-Way Router**

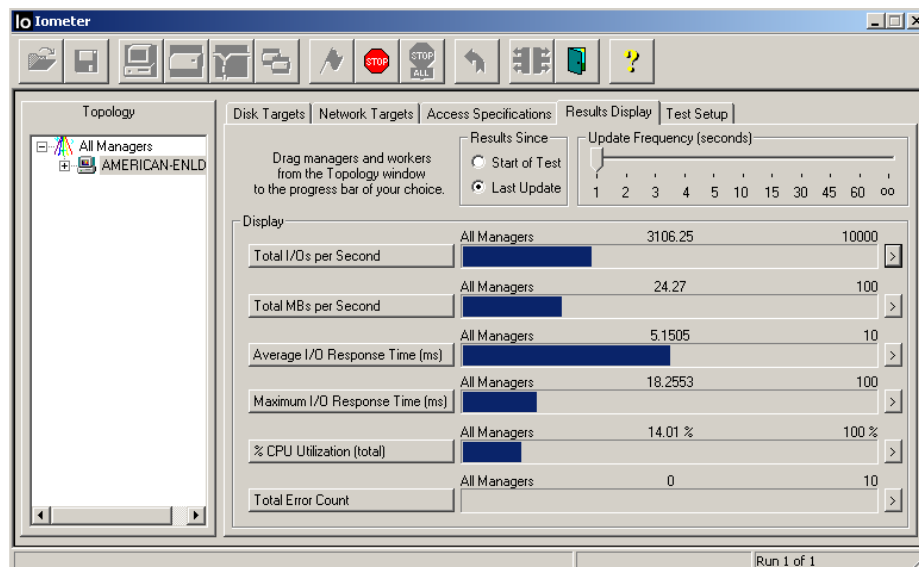
The following figures illustrate the performance of a SAS storage unit under random I/O load that is typical of a database server. With 8 Drive RAID 0 volume of size 30 GB the SAS storage unit was giving about 3300 IOPS with a latency of about 18 ms. However, when a contemporary synchronous replication was setup between this SAS storage unit and a SATA secondary unit the performance dropped dramatically to about 576 IOPS with an average latency of about 800 ms. These figures when compared to StorTrends synchronous replication shows the benefits derived from StorTrends replication. In this it can be seen that the StorTrends synchronous replication from SAS to SATA nodes give about 3100 IOPS with a latency of about 18 ms. Thus StorTrends replication clearly scores over contemporary replication technologies delivering near-SAS performance benefits.



**Figure 3: SAS Storage Unit under Random OLTP Load**



**Figure 4: SAS Storage Server in Contemporary SR with SATA Storage Server under Random OLTP Load**

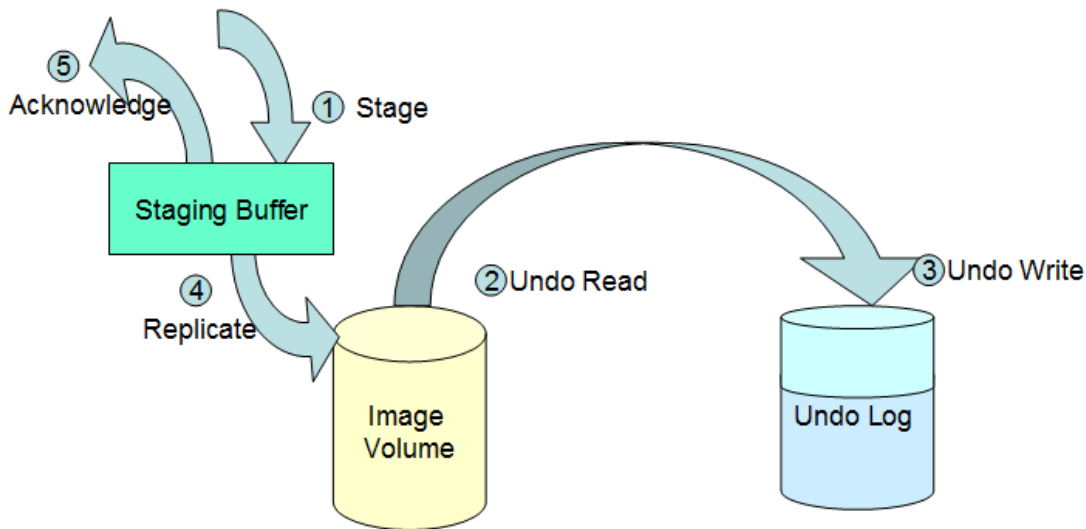


**Figure 5: SAS Storage Server in StorTrends SR with SATA Storage Server under Random OLTP Load**

## Reducing RPO in the event of data loss during logical disasters

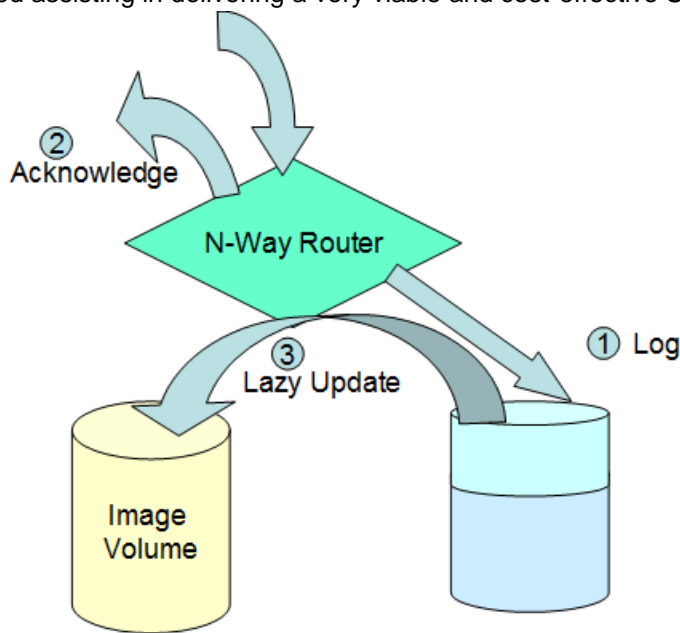
Synchronous Replication, by its very nature, keeps the production and recovery sites totally synchronized essentially offering immediate RPO. The only plausible crack this leaves open is that 'logical data-errors' can not be corrected. This is because acts of accidental deletions, erasures or virus attacks are also faithfully replicated leaving both the sides equally vulnerable. The best way to protect this is to deploy CDP on the recovery site so that in the event of such blemishes, data can be quickly restored to a pre-disaster state.

As can be seen from the following diagram, implementing CDP on the recovery site using conventional methods generate many 'internal' I/Os per replication I/O. The incoming replication I/O is first staged, the original data is put aside in an 'undo log' and then the replication I/O makes its way to the replicated volume. This stretches the latency further threatening the viability of such a SR operation.



**Figure 6: Contemporary CDP with SR**

StorTrends N-Way router logic in the presence of a CDP Journal automatically routes the replicated I/O to the journal and completes the I/O. Multiple journal readers supported in StorTrends allow creating synchronous image for failover as well as APIT image of the primary volumes in the event of logical data errors. The following diagram shows how the I/Os are drastically reduced assisting in delivering a very viable and cost-effective SR solution along with CDP.



**Figure 7: StorTrends CDP with SR**

## StorTrends Advantages

To summarize, the advantages of StorTrends Synchronous Replication technique has the following advantages:

- Very Affordable: allows SR in heterogeneous environments
- Low I/O latency even with snapshots
- Fast resynchronization
- Low RTO
- Support Path Redundancy and Load Balancing
- Permits Replication to fewer and slower spindles
- Supports Coalescing of outstanding I/Os.
- One to one, one to many, many to one, bidirectional and Round Robin topologies supported
- Supports recovery sites in write-back mode for low latency
- Seamlessly integrates recovery site CDP

## Summary

Where Business needs dictate Synchronous Replication as the only viable DR strategy, implementing an affordable solution becomes a difficult challenge. The costs incurred to replicate synchronously almost invariably restrict synchronous replication to be within the bounds of Metropolitan Area. The other demands on performance and latency further push towards having an expensive recovery storage server with many fast and costly disk drives. With StorTrends, the affordable recovery server which lies in close geographical proximity to the production sites can further be asynchronously replicated out to one or many geographically remote locations. StorTrends, with its many architectural innovations, provides a simple and elegant solution to this apparently formidable task.

### *Why AMI?*

AMI offers a wide array of disaster recovery and high availability solutions for your business needs. We provide services that range from storage needs analysis to the design and implementation of a custom disaster recovery solution. We can help your business plan for when things are at their worst, while reducing the cost and complexity of your storage environment. For more information on AMI StorTrends solutions, visit [www.StorTrends.com](http://www.StorTrends.com), email to [sales@ami.com](mailto:sales@ami.com), or call (800) U.Buy.AMI.