

Total Disaster Recovery in Clustered Storage Servers



StorTrends 3200i

© Copyright 1998-2007 American Megatrends, Inc.

All rights reserved.

American Megatrends, Inc.

6145-F Northbelt Parkway

Norcross, GA 30071

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.

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.

Total Disaster Recovery in Clustered Storage Servers

Enhanced Requirements for Disaster Recovery (DR) in Highly Available Clusters

Stringent availability requirements are oftentimes demanded in enterprise class Storage Servers. There are no specific standards followed in the Storage Industry today for this. However, there are two broad categories of implementations generally followed. Some vendors chose to implement this using a 'Dual Redundant' Storage Server, whereas StorTrends implements this using clustered nodes.

Dual-Redundant Storage Servers

Here, as may be clear from the term above, a Storage Server enclosure offers resilience against a single sub-component failure by providing multiple Hot-swappable subcomponents (this includes controllers) housed in the same physical enclosure. This offers redundancy at subcomponent levels, but if for some reason the entire enclosure is threatened then the availability will be compromised.

Clustered HA nodes

StorTrends implements High Availability using a Clustered Nodes. Here two Nodes are configured as a highly available cluster pair rendering uninterrupted data availability. This offers total redundancy, both at component and subcomponent levels. If, for example, somehow one enclosure gets totally unavailable, the other Node continues to provide the storage services without loss of availability.

Normally, the two Nodes in the clustered pair are kept close to each other (may be adjacent to each other or in different floors of the same building). So basically these Nodes are kept inside a Campus Area Network and now if the entire Campus gets affected then there will be loss of data availability. In order to enhance the DR strategy to deal with such situations some form of Asynchronous Remote Replication may be desirable. In the event of total failure of the HA Cluster (that is when both its Nodes are unavailable) the data in these servers would still be protected and available from the different geographical site to which replication has also been made.

High Availability:

Clustered pair allows two clustered StorTrends units to be grouped as Highly Available Cluster nodes for Microsoft application servers or clients. This solution provides protection against disk failures, path failures and node failures. In the event of a node failure, the failover action is automatic and seamless. This configuration also has built-in smarts for load balancing across the multiple available paths for performance scaling. To achieve this, a specially designed **Device Specific Module (DSM)** from AMI has to be installed in the MS clients.

Snap Assisted Replication:

Snap Assisted Replication is a form of asynchronous replication in which data is being replicated from primary to secondary periodically based on snapshots or point-in-time images of a volume or volumes configured in a consistency

Group (CG). These application consistent images of volumes in a CG are then periodically replicated out to Site C according to a conveniently set replication schedule.

Proposed Solution with StorTrends:

The volumes in StorTrends Unit located at Site A and Site B are configured as High Availability Clustered Pair. This HA configuration offers resilience against any single point of failure. Thus, the installation has to be carefully configured so that even the GigE switches offer redundancy against switch failures. It must be noted that the two Nodes of the cluster are normally located adjacent to each other or within close proximity of each other. So the Nodes A and B are really in the same Campus.

This Clustered Node then may be further replicated out to a remote geographic site to offer DR solution against site disasters. Here, the option is to use an Asynchronous form of Replication eg: Snap Assisted Replication (SAR). Thus, after setting up the HA cluster, the SAR replication has to be separately configured between the Primary node of HA cluster (Node at Site A) and the StorTrends Unit located at Site C.

This has been represented in the figure 1 below.

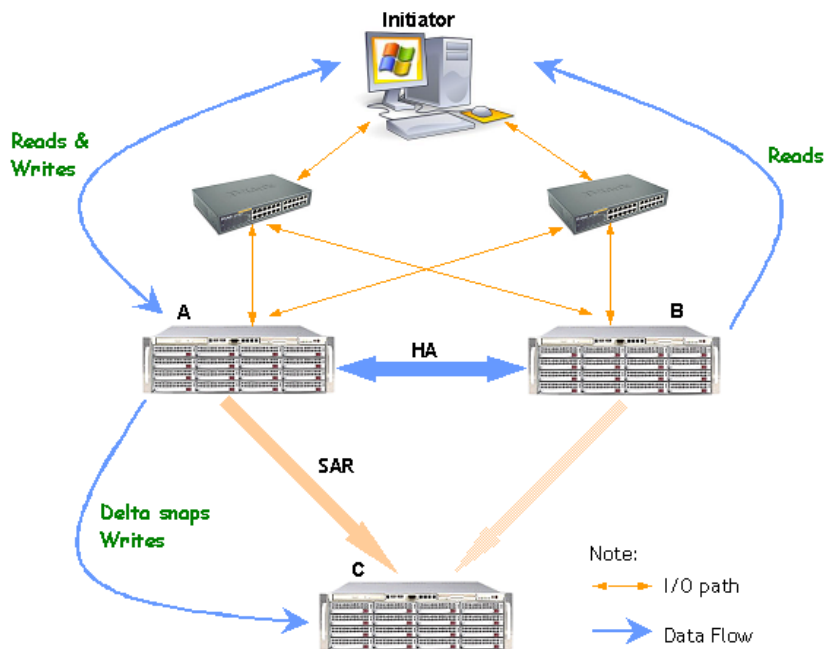


Figure 1: Optimal Clustered Nodes replicating to remote site

How the Solution works:

The High Availability pair is created using the Site A volume(s) as the primary and Site B volume(s) as the secondary. Data is replicated across these HA pairs synchronously and continuously.

Now the Snap Assisted Replication pair is created using the Site A volume(s) as primary and Site C volume(s) as the remote secondary.

When the primary volume in site A receives IO's from the Initiator, the data will be replicated to the secondary volume in site B (HA) and the same data will be replicated to alternate secondary device i.e. Site C (Snap Assisted Replication). Of course the replication to site C will occur asynchronously according to the snapshot schedule and replication schedules configured.

This is illustrated in the figure 2 below.

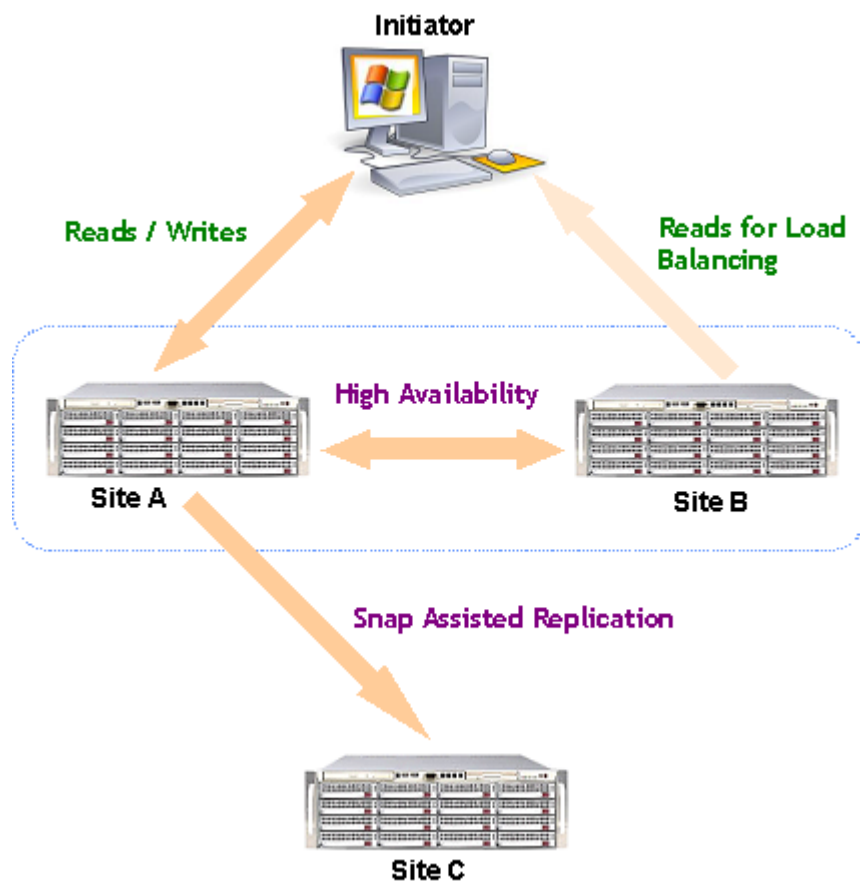
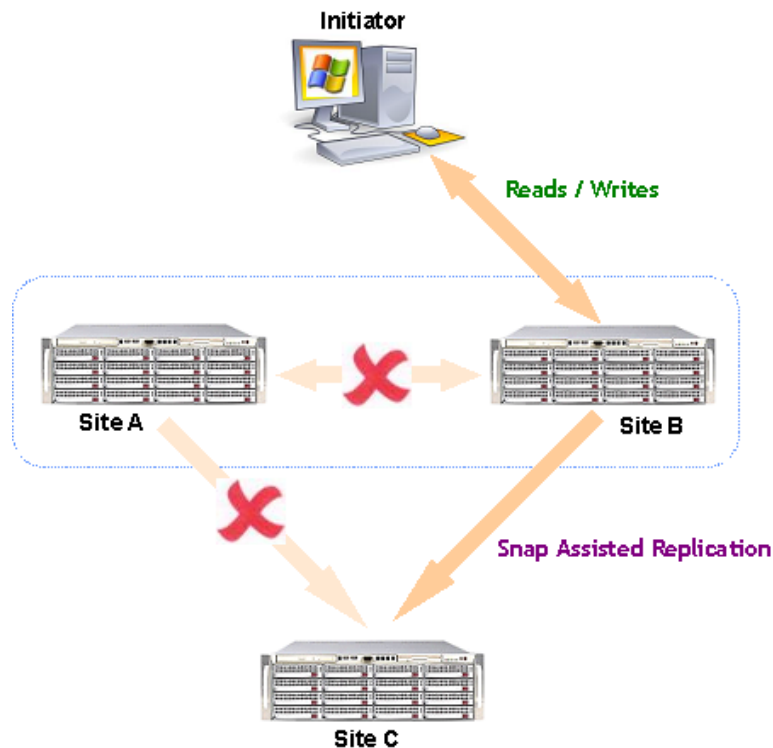


Figure 2: Optimal HA cluster and Remote DR

When the Node at Site A becomes unavailable or is not accessible, Site B takes over automatically as the Primary Node, that is failover will happen automatically, and the Cluster becomes degraded but still remains available. Now the Asynchronous replication has to be manually configured and started between Nodes B and C. Please refer to the Figure 3 overleaf.

Operation during degraded cluster Mode

When Node A becomes available for service, automatically this event is detected and data starts getting replicated to Node A from Node B. During this process, Node A is also Re-synchronized to bring it current with Node B. Once this is completed, the cluster Node becomes completely operational and optimal. Now if Node B goes out of operation, Node A seamlessly and automatically takes over as the Primary Node. Following similar procedure, the Asynchronous



[Figure 3: Degraded HA cluster with Remote DR](#)

replication has to be manually started between Node A and Node C to resume and restore remote DR operations.

Conclusion

StorTrends, with its extensive suite of DR features, offers a very robust solution with total redundancy in enterprise class Data Centers. Two StorTrends Nodes can be configured as HA pair in a Campus Area Network providing a very resilient IP-SAN Server. This architecture protects against any single point of failure including component, sub-component, path and Node failures. Intuitive GUI tools make configuration and management of such a node a breeze. Coupled to this, a remote storage server can be configured to offer complete protection against physical and site disasters. Further, snapshots can be scheduled at each of the Nodes to offer added insurance of protection against any Logical or Data disasters.

Why AMI?

Since 1996, AMI has been a leader in the data storage industry, with the inception of the MegaRAID® RAID Controller. AMI was the largest third party RAID Controller manufacturer in 1997, and by 2001 had reached complete market share. At that time, AMI sold its MegaRAID® Division to LSI Logic in 2001. This asset sale allowed AMI to develop the StorTrends® IPStorage product family. The entirety of AMI's many years of expertise in data storage has been rolled into the StorTrends® product. AMI clearly understands the industry's needs and requirements for data storage and has used this understanding to develop a complete line of "best in class" IP-Storage products. AMI looks forward to entertaining any questions regarding the suitability of the StorTrends® IP-Storage Array for the Microsoft® Exchange environment.

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, email to sales@ami.com, or call (800) U.Buy.AMI.