

StorTrends® iTX: Application-aware Scheduled Snapshots with Microsoft® VSS and SAFE



© 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.

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.

Original Release: 07/24/2007

First Revision: 10/07/2008

Introduction

In any storage system, the one feature that is of paramount importance to the user is data protection. While the user takes for granted that data will not be lost due to a system crash or a design bug, the user also expects a storage system to protect him from inadvertent deletions, unwanted modifications, malicious agents and the like.

In response to this most basic need, snapshots have become an industry-wide standard for protecting data at a very fine granularity. Whether commissioned according to a schedule or under the direction of an administrator, the data snapshots taken by a storage system represent point-in-time pictures of the contents of a volume in a storage system, similar to an incremental backup of the system's data. What distinguishes snapshot implementations from vendor to vendor is the manner in which the stored data is backed up, as this can greatly impact both space utilization and performance. The StorTrends® line of storage appliances from American Megatrends, which utilize the innovative StorTrends® iTX storage software, offer industry-leading support for snapshots, with an implementation that optimizes performance and storage capacity utilization.

While snapshots present a consistent state of the volume to the user in terms of the outstanding I/Os to the storage server, this state may not be consistent with respect to the application using the volume, particularly in terms of its recoverability. Moreover, it is often the case that an application is using multiple volumes on the storage server for various purposes. For example, Microsoft Exchange Server uses both a log volume and a data volume to maintain, manage and protect its data. Thus, a snapshot of both volumes must be synchronized in order to achieve application recoverability, since these two volumes can not function independently of each other. Also, often the I/Os are being cached on the initiator, which can be doing delayed-writes to the volume, meaning that the snapshot taken from within the storage server does not always assure application-consistent data recoverability during a disaster.

To mitigate this potential problem there needs to be an application- and host-aware agent on the application server that communicates with the storage server for the purpose of creating application-aware and consistent snapshots. This host agent must also be able to quiesce the application and trigger the snapshots on all of the application-specific volumes (also known as the **consistency group**). StorTrends iTX, which comes preinstalled with every StorTrends storage appliance, utilizes its Microsoft VSS Hardware Provider and Application-specific agents such as Microsoft Exchange and Oracle® snapshot agents, to resolve this challenge.

Microsoft Volume Shadow Copy Service (VSS)

Volume Shadow Copy Service (VSS), introduced by Microsoft in Windows® Server 2003 and featured in Windows® Storage Server, is the infrastructure that enables built-in high-fidelity snapshot (shadow-copy) capability. VSS is designed to support both host-based and hardware based storage solutions. Host based solutions, as the name suggests, are host-resident and either need special software providers or can use the default system software provider of VSS. Note that for purposes of this discussion, StorTrends iTX is a hardware-based storage solution external to the host machine. Therefore, in order to leverage the features of VSS, a Host Attach Kit (HAK) is provided by AMI that contains a StorTrends iTX VSS Hardware Provider as one of its components.

For any VSS-based shadow copy infrastructure there are 3 main components that need to interact with the VSS Framework in the manner specified by VSS. These components are:

- **Requestor** – This is the component that invokes the Shadow Copy (snapshot) operation. This could either be a backup application creating a VSS-enabled backup, or a management application such as the StorTrends iTX Scheduler service, which is the portion of the StorTrends iTX HAK that handles the creation of scheduled snapshots.
- **Writer** – As mentioned above, in order to create application consistent snapshots all outstanding application writes must be quiesced and further new writes queued for the duration of the snapshot process, until the point when the snapshot operation

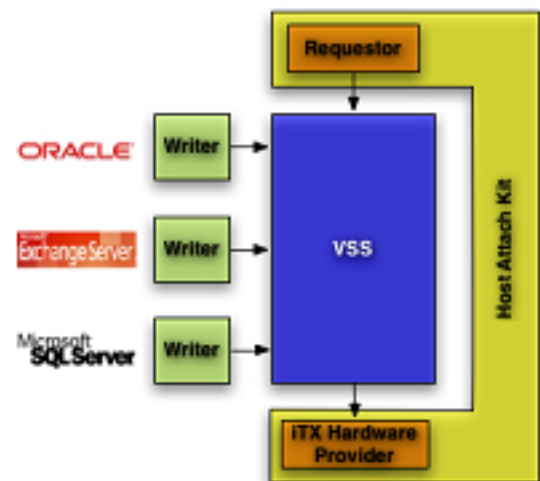


Figure 1: VSS-based Snapshot Architecture

is completed. Writers are application-specific modules that do precisely this function.

- **Provider** – Providers are interfaces to snapshot capabilities, as described earlier. The StorTrends iTX Hardware Provider interacts with the snapshot monitor in the StorTrends iTX boxes in order to create a snapshot. As a result, snapshot creation, deletion, and rollback in StorTrends iTX are fast operations with very little overhead. The powerful and feature-rich character of StorTrends snapshots will be discussed in greater detail in a subsequent section.

Snapshot Creation Sequence – the VSS way

In order to take snapshots using VSS, the Requestor, Writer and Provider must interact with the VSS coordinator in a specific way. The sequence of operations involved in this process¹ is outlined below.

- 1) The requestor (in this case, the StorTrends iTX Scheduler Service) notifies the VSS coordinator service to prepare for shadow copy creation.
- 2) The VSS coordinator notifies the application-specific writer, e.g., Exchange or Oracle, to prepare its data for making a shadow copy.
- 3) The writer prepares the data in whatever way is appropriate for that particular application, such as completing all open transactions, rolling transaction logs, and flushing caches.
- 4) Once the data is prepared for shadow copy creation, the writer notifies the VSS coordinator.
- 5) The VSS coordinator relays the message to the requestor (iTX Scheduler), and the requestor initiates the “commit” snapshot phase.
- 6) The VSS coordinator temporarily halts (quiesces) application I/O write requests (I/O read requests are still possible) for the several seconds required to create the shadow copy of the volume or volumes.
- 7) The StorTrends iTX hardware provider next informs the Snapshot Monitoring Agent in the StorTrends iTX appliance to create the snapshots. Note that there could be multiple volumes in the consistency group, and the snapshots of each volume must be taken. According to the VSS Specification, a maximum of 10 seconds is allowed for this entire operation. Fortunately, StorTrends iTX can commission snapshots extremely quickly, and multiple volume snapshots can be created simultaneously with a single call. It is important to point out that the scheduler can specify the number of snapshots to retain per volume, and in cases where the number of snapshots has exceeded this amount, the oldest snapshot in the volume is deleted before the new snapshot is created. Again, with StorTrends Advanced Snapshot technology this is not a concern, as all these operations are finished well within the 10 second cut-off time stipulated by VSS.
- 8) After the shadow copy is created, the VSS coordinator releases the writer from its temporary quiescent phase.
- 9) All queued write I/Os are completed.
- 10) VSS queries the writers to confirm that write I/Os were successfully held during shadow copy creation.
- 11) If the writes were not successfully held (meaning the shadow copy data is potentially inconsistent), the shadow copy is deleted and the requestor is notified. The requestor can retry the process (loop back to 1) or notify the administrator to retry at a later time.

As a direct result of steps 6 and 11-12, the VSS coordinator service ensures that the shadow copy is a highly consistent copy of the original data. Moreover, because all of the necessary preparation for making the shadow copy is completed prior to halting the application write I/Os, production performance is not affected.

¹ For more information please refer to the whitepaper “Storage Management Using Windows Server™ 2003 and Windows® Storage Server 2003 Virtual Disk Service and Volume Shadow Copy Service”, Microsoft Corporation, August 2003

Transportable Snapshots

Transportable snapshots are a feature of VSS that only works with hardware providers; this feature enables the transport of data between servers that access the same storage array. By design, storage allocated to one server cannot be used by a second server simultaneously, as this might otherwise result in the two servers managing and writing the same data, an obvious potential conflict. However, the ability to use a snapshot of the same data in the other server provides a realistic solution, as data protection is maintained while simultaneously easing data management restrictions. In this way, snapshots taken with StorTrends iTX Hardware Provider on one server can then be exposed as targets to a different server. Since support for writable snapshots is a key feature of in StorTrends iTX, these snapshots can also be used in writable mode, paving the way for a number of advanced storage management solutions.

StorTrends iTX Snapshots

Snapshots in StorTrends iTX are Redirect-on-Write (ROW) Snapshots, meaning that all new writes to the volume to portions that were previously written prior to the snapshot are redirected to new disk space, and managed using mapped tables. The implication of ROW snapshots is that in comparison to conventional Copy-on-Write (COW) Snapshots, which read old data, write them to a different area, and then execute new writes, the number of I/Os saved by ROW snapshots is considerable. Moreover, any portions that remained unwritten to during the previous snapshot lifecycle can be tracked (via provision tracking) and reused. The use of ROW snapshots in StorTrends iTX allows it to deliver much better performance, both in terms of I/O and space savings. Comparisons of StorTrends iTX I/O performance with other contemporary storage stacks support this point. In fact, testing has shown that contemporary storage stacks can suffer a performance drop of 50% or more of original performance for a snapshot, whereas the performance drop of StorTrends iTX is negligible (on the order of a 5 – 10 % decrease), even after multiple snapshots.

StorTrends iTX offers support for up to 1022 Read-Only and 1022 Writable snapshots, and the capability to create snapshots in both Read-only and Writable modes is just one powerful feature of the StorTrends snapshot implementation. Writable snapshots are typically used in testing disaster recovery points, as well as in CDP, boot server consolidation and the like. Testing of StorTrends iTX's snapshot capability has also confirmed that creation, deletion and rollback operations with StorTrends iTX snapshots generally requires no more than 100 ms to completion. The combination such high performance and extensive unique features make the snapshot technology in StorTrends iTX the best that is currently available in the market.

Scheduling Snapshots

Administrators of enterprise-scale storage systems typically back up the data on their production servers to other disk- and tape-based storage servers on a scheduled basis; this schedule could be nightly, weekly or monthly depending on the nature of the data involved. In order to do so they employ custom scripts and daemons to schedule such tasks. StorTrends iTX makes this painstaking job infinitely easier by allowing administrators to generate and manage schedules for snapshot creation using its Web-based storage management interface, ManageTrends®, as well as through the command-line interface (CLI). Once a schedule is established, the administrator can determine the number of snapshots he or she wishes to retain; if and when this count overflows, the StorTrends scheduler will automatically delete the oldest snapshot and create a new snapshot in its place. Additionally, a backup target can be selected and the backup utility will use this target for backing up the snapshot. Because of the potential need to create multiple schedules and establish criteria for snapshot retention, StorTrends iTX allows up to 8 levels of independent snapshot schedules for each volume or consistency group.

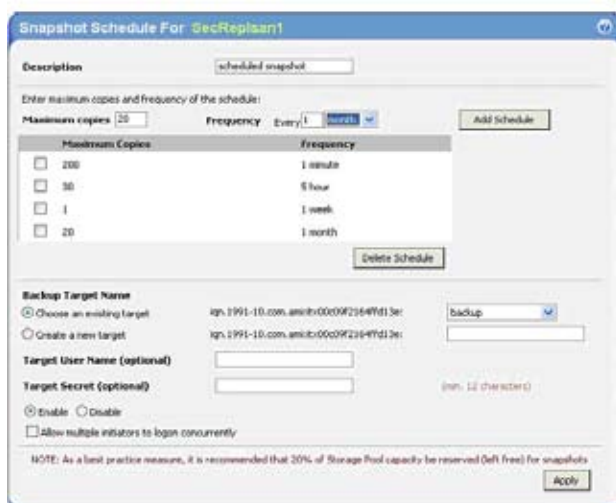


Figure 2: Snapshot Scheduling with ManageTrends®

Host Attach Kit (HAK) and Agents for non-VSS environments

Volume Shadow Copy service (VSS) was introduced by Microsoft in Windows® 2003 Server; Windows® 2000 and Windows® XP do not have native support for VSS, as both were released prior to Windows® 2003 Server. However, StorTrends iTX offers a solution for this potential obstacle through its Snapshot Agent Framework Environment (SAFE). SAFE is a component of the StorTrends iTX Host Attach Kit, which also includes the StorTrends iTX Scheduler service and the StorTrends iTX Hardware Provider for VSS

environments. SAFE automatically verifies if the environment is VSS or non-VSS, and acts accordingly.

The StorTrends iTX HAK also contains agents for Microsoft Exchange and Oracle that assist in suspending the respective applications during the process of creating snapshots in non-VSS environments. This function is essentially identical to the role of writers in the VSS environment. SAFE loads the application agents dynamically, depending on the configuration present in the *safe.ini* file in the HAK install path (*drive\Program Files\Ami\iTXHardwareProvider*). By default, all the agents are configured in *safe.ini* in a disabled state; to enable a particular agent, the user must change the “enable” field from “no” to “yes” in the *safe.ini* configuration file for the corresponding agent. All the enabled application agents configured in the *safe.ini* are managed by SAFE prior to the taking of snapshots. SAFE uses the StorTrends iTX Scheduler Service to poll StorTrends appliances on a minute-by-minute basis to obtain the snapshot schedule for the volume, as it is defined by either ManageTrends or the CLI.

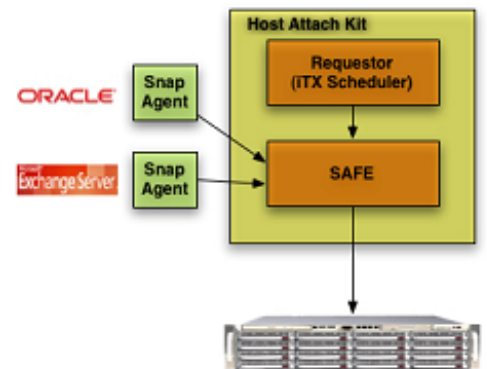


Figure 3: Non VSS-based Snapshot Architecture

If the Scheduler Service finds that a particular volume is scheduled to take a snapshot, using the calls defined by SAFE, it next identifies whether the system is VSS- or non-VSS based. If the system is a VSS-based one, SAFE initiates the VSS snapshot sequence explained earlier. If, however, the system is non-VSS based, SAFE determines whether or not any of the loaded agents are using the StorTrends iTX volume. If it is determined that any of the supported applications are using the StorTrends iTX volume for storage, SAFE will then communicate with the corresponding application agent to suspend the application and close all the file handles before executing the snapshot. The communication between SAFE and the agents are standardized calls, and new application agents can be created as plug-ins to work with SAFE. Currently, application agents are available to work with Oracle and Microsoft Exchange Server.

Summary

The snapshot implementation featured in StorTrends iTX is undoubtedly one of the best, most feature-rich in the industry. Its host of powerful and attractive features combined with stellar performance pushes StorTrends iTX to stand out from its contemporaries. The benefits of scheduling snapshots using an intuitive GUI and the ability to create application-specific snapshots that preserve consistency-group relationships make StorTrends iTX an indispensable tool for storage administrators.

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

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.

© 2008 American Megatrends, Inc.

All Rights Reserved



American Megatrends Inc.

5555 Oakbrook Parkway, Building 200

Norcross GA 30093 | t: 770.246.8600

Sales & Product Information

sales@ami.com | t: 800.828.9264

Technical Support

support@ami.com | t: 770.246.8645

www.ami.com