

Time-addressable and Event-addressable Storage: Combining Snapshots and Data Journaling for True Continuous Data Protection



© 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.

Original Release: 06/04/2007

Table of Contents

CDP- Definition, Characteristics, and Metrics	4
Time-addressable Data Storage	4
<i>Figure 1: Review and Selection of Time-stamped Snapshots</i>	<i>5</i>
<i>Time Addressable CDP with Data Snapshot</i>	<i>5</i>
<i>Event-addressable CDP with Journaling</i>	<i>5</i>
<i>Full Data Journaling</i>	<i>6</i>
<i>Metadata Journaling</i>	<i>6</i>
Summary and Conclusion	6
<i>Journaling in StorTrends iTX</i>	<i>7</i>
<i>Figure 2: Journaling Module Architecture in StorTrends iTX</i>	<i>7</i>
<i>StorTrends iTX Snapshot Technology</i>	<i>7</i>
<i>Figure 3: CDP Status Display in StorTrends iTX</i>	<i>8</i>
<i>Why AMI?</i>	<i>8</i>

CDP- Definition, Characteristics, and Metrics

Although there has been a sustained, high level of interest in the topic of Continuous Data Protection (CDP) in the storage industry for the last few of years, an authoritative definition of CDP remains a moving target. A survey of different vendors providing CDP solutions in the market today uncovers a wide spectrum of definitions. Some vendors intend CDP to mean the ability to restore data from any point in time. Others may describe it as the ability to get an application back on line quickly in a consistent form while minimizing the loss of data. Still others may define CDP as simply “snapshots plus data replication.”

A useful working definition is the one definition provided by the CDP Special Interest Group of the Storage Network Industry Association (SNIA). While the different notions of what CDP is and is not have prevented them to date from issuing a formal definition, they have been able to craft a working definition of CDP that is broad yet clear enough to be satisfactory:

“A methodology that continuously captures or tracks data modifications and stores changes independently of the primary data, enabling recovery points from any point in the past. CDP systems may be block, file, or application-based and can provide fine granularities of restorable objects to infinitely variable recovery points.”¹

Despite all of the differing ideas on the definition of CDP, there are a relatively clear set of conditions, characteristics, and metrics that set the boundaries for what a CDP solution must exhibit in order for it to be effective, or in other words, “true CDP.” First of all, it has to utilize RPO (Recovery Point Objective) as the key metric, by protecting data with a very fine level of granularity, since the CDP solution must be able to recover to any point in time. To do so, CDP solutions must be able to track data changes at the block, file, or application level, focusing down individual I/Os, data changes, or application events. As a secondary metric, the solution RTO (Recovery Time Objective) should also be low, meaning that the system can be restored to an uncorrupted point in time within a relatively short period of time.

Although there are many different ideas of what CDP is, it is clear that in order for a CDP solution to be most effective, it must protect data in a manner that is both **time-addressable** and **event-addressable**. Time-addressable and event-addressable capabilities are critical in order to realize the RPO and RTO targets that define CDP. This paper will give a detailed explanation of what time-addressable and event-addressable data protection are, and how the data protection and recovery technologies of **snapshots** and **journaling** incorporate these concepts into very effective CDP solutions.

Time-addressable Data Storage

In order to recreate data exactly as it existed at any previous point in time and ensure fast recovery of critical operations –in short, to provide nearly infinite RPO and low RTO- a CDP solution first needs to be time-addressable. This means that data changes or snapshots of volumes need to be “time stamped,” and associated with a specific moment in time of the system’s operation. This time stamping is necessary so that they can be reviewed relative to the point when disaster occurred, in order to determine whether or not they are suitable points for recovery.

¹ Storage Networking Industry Association, SNIA DMF CDP Special Interest Group (Brian J. Olsen, Project Editor); *CDP Buyer’s Guide: An Overview of Today’s Continuous Data Protection (CDP) Solutions*. September 2006 – Second Edition

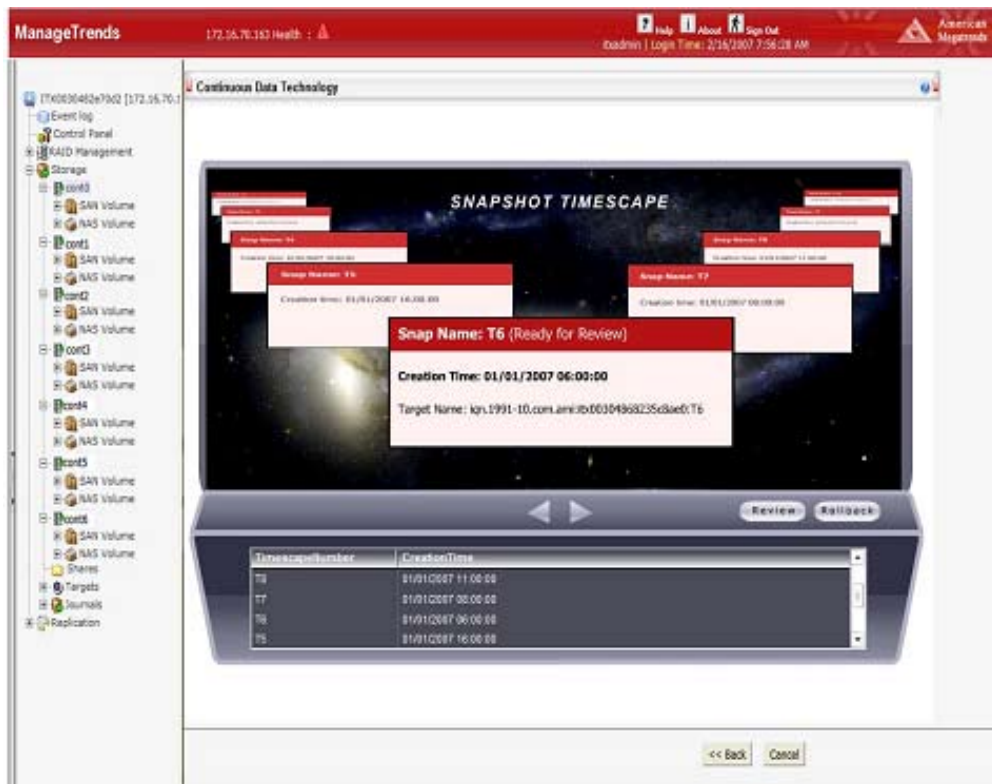


Figure 1: Review and Selection of Time-stamped Snapshots

Time Addressable CDP with Data Snapshot

When surveying the options available in the market today for time-addressable data protection, the most efficient and reliable method available is **data snapshot** technology. Briefly, snapshots are *point in time* images of the data contents of a volume that can be stored in a separate location (i.e., another volume or storage appliance), reviewed, and mounted to disk in a restore procedure. Snapshots are time-stamped at the moment they are captured, and can also be scheduled with a very fine granularity, on the order of minutes, hours, or days, depending on the snapshot scheduling software and the capability of the hardware to support snapshot scheduling.

Because of the granularity with which snapshots can be scheduled, often on the order of one snapshot per minute, the RTO can be extremely low when compared to other data replication technologies like tape backup. However, snapshots are generally considered “near-CDP” or “virtual CDP”, since, practically speaking, they can not provide an infinite number of recovery points. Regardless of the granularity with which they are scheduled, snapshots do leave a gap in coverage, however small, because of the small window of data exposure between the snapshots. That is why snapshots, as time-addressable data protection, should be combined with data protection that is event-addressable, to build a complete coverage solution across the board.

Event-addressable CDP with Journaling

Beyond RPO and RTO requirements, which focus on time, CDP users would be wise to choose a CDP solution that protects data in a format that is amenable to the storage environment or the application itself. This is done through utilizing an event-addressable storage methodology. Event-addressable storage (EAS) is founded on the idea that data storage and hosted applications have certain telltale events (I/Os, write data in certain patterns, and so on). Being able to piece together these events from a record or journal that is logging them and leaves a clear trail is a very useful mechanism for protecting and restoring data. Event-addressable storage adds detail to the reference points used for selecting the optimal recovery point, by contributing information about application events to timestamp information for many failure scenarios.

To complement the time-addressable data protection of snapshots, event-addressable storage should be added to a CDP solution; the most efficient and reliable method of event-addressable data protection is **journaling**. Journaling is

essentially a software-based approach to EAS that is a “rolling database” of changes to a system, keeping a record of all I/Os. The beauty of journaling is that it creates an “exact virtual disk” replicate of the system - which can be block, file-based, or application-aware CDP. For block data, journaling works across structured, semi-structured, or unstructured data, and is exceptionally useful for database application recovery because of its granularity and flexibility across multiple platforms and application environments.

Full Data Journaling

When discussing journaling as a CDP solution, a distinction must be made between “full data” and “metadata” journaling, because of the performance overhead and suitability of each for different environments. In full data journaling, everything is written to the journal, including data blocks. This provides a strong guarantee that in the event of data corruption, the system can be restored to its previous uncorrupted state. At the block-level, journaling works by copying block data as it is changed, and tracking the changes. For file data, it works at the physical file level, and is able to achieve more granular recovery with applications since it can recover at the file level rather than the entire volume.

When disaster strikes, a journaling system can combine the blocks that have not been overwritten since the restore time with copies of the blocks that have changed, and present the combined image to the system. The RPO of journaling is infinite, since the granularity can be tuned to the I/O level; furthermore, the data size that can be protected is quite large.

Metadata Journaling

A potential concern with full data journaling, however, is the storage overhead created by enabling it. Since data must be written twice in full data journaling, it can demand anywhere from 5% to 40% additional storage overhead, so that protecting 100 terabytes (TB) of data may require another 5 to 40 TB of disk space. In response to this concern, **metadata journaling** also plays an important role in event-addressable approaches to CDP, and offers tremendous capacity and potential benefits to users. In metadata journaling, only the changes to file metadata are recorded, a practice that is generally two to three times faster than full data journaling. Metadata journaling provides the guarantee that no matter when the computer crashes, the file system metadata can easily be made consistent, and data restore can proceed from an examination of system metadata.

Summary and Conclusion

As the name implies, CDP is a method of uninterrupted safeguarding against data loss or corruption, by using one of a number of possible services or methodologies. Generally speaking, CDP services operate by continuously capturing all changes to data and storing these changes in a separate storage location. CDP services also provide the capability for virtually instantaneous restoration of data to any point in time prior to a catastrophe or data loss.

This paper discussed how replication that is both time-stamped and event addressable through the combined technologies of data snapshots and journaling is true CDP that offers complete system data protection. By combining the time-addressable and event-addressable data protection technologies of data snapshots and journaling, an application recovery-oriented solution with optimized RPO and RTO characteristics for any type of failure is now a possibility.

Until now, CDP has only been deployed in enterprise systems due to the complexity and cost of implementing a reliable solution. But by capitalizing on the unique features of snapshot-capable iSCSI storage appliances, and placing the CDP functionality within the software rather than hardware, CDP is now within the reach of small and medium-sized businesses (SMBs). The key benefits of CDP that allow organizations to reduce data loss (higher RPO) and perform quicker recovery of short-term data in the event of a data corruption or loss (lower RTO) are all presented at an attractive price in the innovative StorTrends® iTX storage software stack from American Megatrends (AMI).

StorTrends® iTX combines its low-latency snapshot technology with efficient I/O journaling to offer a powerful, dependable CDP repository. The CDP services layer is very tightly integrated with the Journaling and Snapshot modules as shown in the illustration below. The snapshot technology in StorTrends iTX provides near-continuous data protection by creating and retaining hundreds of low-latency snap shots per volume at scheduled intervals, to produce numerous hourly, daily, monthly and yearly snapshots. When combined with journaled storage, StorTrends iTX offers **true CDP** capabilities.

Every write I/O is journaled, and entire volumes and application consistency groups can be rolled back to any point in time (APIT), to provide an instantaneous Recovery Point Objective (RPO). The upcoming iTX v2.7 integrates this further with 'Visual Snapshot' technology to offer a very intuitive and user friendly Snapshot Recovery Wizard for the selection, review, and rollback of snapshots. This near instantaneous rollback feature, paired with Redirect-on-Write (ROW) snapshots, enables StorTrends iTX to offer the industry's fastest Recovery Time Objective (RTO), on the order of milliseconds in response time.

Journaling in StorTrends iTX

The journaling module in StorTrends iTX registers with the caching/locking module above the DVM layer to provide I/O logging and support consistency groups. This module also tightly interfaces with the snapshot service to keep track of I/Os and events, forming the basis of event-addressable storage in StorTrends iTX.

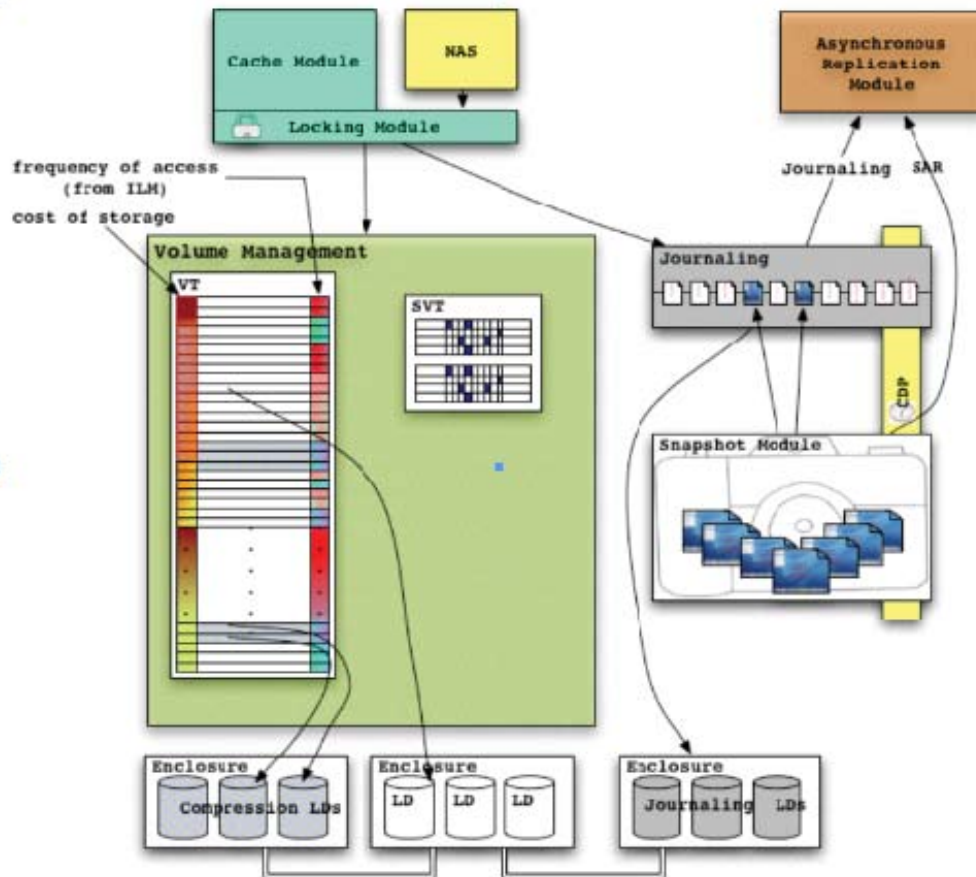


Figure 2: Journaling Module Architecture in StorTrends iTX

Journal files can be configured from up to 32 logical drives of any size to serve as event-addressable storage. Additional features of journaling CDP in StorTrends iTX include:

- Extensive tools for monitoring and control
- Support for both continuous and periodic replications
- Seamless migration between various modes of replication
- Optional compression of data for efficient link utilization
- Link-level scheduling
- Flexible path configuration for performance enhancements

StorTrends iTX Snapshot Technology

Time-addressable CDP is also offered by StorTrends iTX by utilizing several of the snapshot technology innovations available in the software stack, including:

- Snapshot technology that utilizes the Redirect on Write (ROW) method

- Near-unlimited snapshot capacity, allowing snaps to be taken every few minutes
- Minimal performance degradation when snapshots are enabled
- Easy, quick rollback to any snapshot

AMI's unique snapshot technology offers near-zero performance degradation by reducing the number of I/O operations by a factor of three. StorTrends iTX features the capability to rollback to any snapshot in less than 100 ms, along with an unrivalled feature called **Snapshot-assisted Replication**, where snapshots can be reviewed, and then mounted as part of the recovery process. **Delta Snapshots** are another unique feature of StorTrends iTX that copy only the part of the data store which has changed since the last snapshot, to save on writes (saving I/O overhead) as well as storage space.

The StorTrends 3100 and 1100 Series storage appliances are capable of virtual-CDP with solely the advanced snapshot capability found in StorTrends iTX; combined with its powerful Journal Assisted Replication, it is an unbeatable combination for data protection and value in a single product.

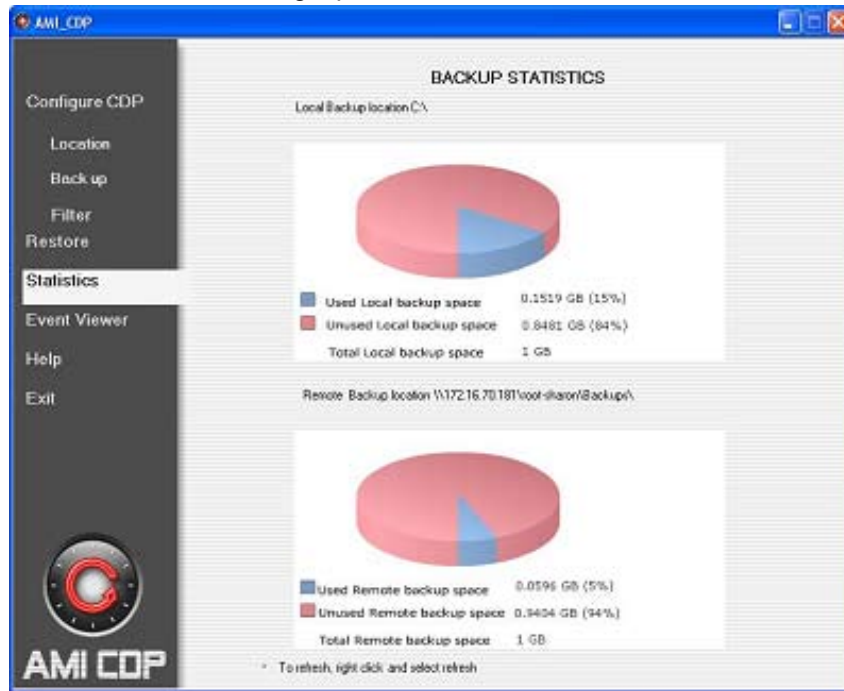


Figure 3: CDP Status Display in StorTrends iTX

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.



American Megatrends Inc.
6145-F Northbelt Parkway
Norcross GA 30071 | 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

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.

© 2007 American Megatrends, Inc.

All Rights Reserved