sageza

Strategic Snapshot

Economical Disk-Based Recovery for the High-End: New Space Saving Pointer-Based Snapshots

By Charles King and Jim Balderston

The Sageza Group, Inc.
March 2004

The Sageza Group, Inc.
32108 Alvarado Blvd #354
Union City, CA 94587
650·390·0700 fax 650·649·2302
London +44 (0) 20·7900·2819
Milan +39 02·9544·1646

sageza.com
info@sageza.com

Economical Disk-Based Recovery for the High-End: New Space Saving Pointer-Based Snapshots

ABSTRACT

Just a few years ago, recovery point objectives (RPO) were typically measured in days. Increasing reliance and pressure on IT-based business processes and the resulting explosion of information storage requirements have reduced that time to a single day or less. As a result, many organizations have found that tape is too slow as a primary recovery option. Given evolving RPO and Recovery Time Objectives (RTO) and disk capacity requirements for supporting point in time (PIT) solutions, customers have been forced to make tradeoffs between cost and rapid recovery. Space Saving Snapshots offer economical alternatives to full-volume copies and require less disk space while significantly improving the RTO and RPO of tape-based backup. When Space Saving Snapshots are used in conjunction with full-volume copies, the combined solution provides levels of flexibility and security that neither technology offers alone. In short, Space Saving Snapshots can notably enhance the availability and economics of replication technology.

Economical Disk-Based Recovery for the High-End: New Space Saving Pointer-Based Snapshots

---

TABLE OF CONTENTS

Introduction

The value of business information in the 21st century is uncontested. Information is not just information but also a business critical component for enterprises to meet their goals. Virtually every organization shares the challenge of making its information easily accessible, manageable, and fully protected against system failure and other unforeseen disasters. Losing access to business critical information leaves an enterprise in a potentially dangerous position. To address this risk, new data replication solutions have come to market that offer customers additional options, namely space-saving point in time (PIT) copy solutions. In this paper, we will discuss these new solutions with respect to three key criteria: economics, recovery point objectives (RPO), the difference between current time and the most recent backup, and recovery time objectives (RTO), how quickly data can be restored and applications restarted.

Just a few years ago, RPOs were typically measured in days. Increasing reliance on IT-based business processes and the explosion of information storage requirements have reduced that time to a single day or less. As a result, many organizations have found that tape is no longer suitable as a primary recovery option. Given evolving RPO and RTO along with the disk capacity requirements for supporting PIT solutions, customers have been forced to make tradeoffs between cost and rapid recovery.

Space Saving Snapshots Improve RPO and RTO Economics

However, enterprises using high storage solutions no longer have to make these tradeoffs. New space-saving PIT copy solutions remove a key impediment to leveraging this technology for more affordable disk-based recovery and enhanced RPO and RTO.

Space Saving Snapshots do not require the 100% additional disk space requisite in traditional PIT copy solutions to accommodate copies. In most cases, these snapshots require between only five and twenty-five percent incremental storage based on the nature/profile of the data — this significantly changes the economics and process of data replication. As a result, organizations can replicate more of their data on disk and replicate key data more frequently while implementing Space Saving Snapshots that complement full volume copies.

Until now, businesses had primarily two options when addressing RPO and RTO challenges: full-volume copies or tape backup solutions. Full-volume copy solutions, such as EMC TimeFinder, IBM FlashCopy, and HDS ShadowImage, have been the preferred, if not the only local replication options for high-end storage solutions. Full-volume copy technologies are well established and well regarded, and with good reason. They offer 100 percent data replication locally for backup or parallel processing. They can handle virtually any workload and provide the best levels of performance (addressing the most stringent SLAs, assuming of course that the hardware has the power and performance to drive the replication process). However, full volume and traditional snapshot copy solutions require 100 percent additional disk capacity, which adds considerably to the overall costs of implementation and management.

Local Replication Strategies and Solutions for Rapid Recovery

How do traditional full volume and PIT copy solutions compare? First, some short definitions:

Full Volume Replication:

Full volume solutions are robust and flexible enough to handle the most demanding I/O workloads without affecting I/O response times to a host or application. The "cost" for this performance is the 100 percent additional storage capacity requisite to implement the replica.

Full Volume Pointer-Based Snapshots:

Full volume pointer-based snapshots are another option for creating replicas that need to be accessed immediately. They also allow the option for performing a full-volume copy upon each access of the data or through a behind-the-scenes copy process. Like full volume replicas, full-volume pointer-based copies require 100 percent additional storage capacity to implement the replica.

Space Saving Pointer-Based Snapshots:

Space Saving Pointer-Based Snapshots provide compact logical images of data, enabling users to make more PIT copies, reducing RPO cycles, while requiring less incremental disk capacity than full volume solutions. Snapshots, which offer fast, inexpensive (relative to full volume copy processes) pointer-based data replication, are solutions high-end storage customers can deploy to complement full-volume copy solutions as well as being a primary means of local replication.

When to use Space Saving Snapshots

Space Saving Snapshots offer fast, easy replication processes that consume less disk space than full volume solutions, but in some circumstances these snapshots may impact applications requiring the highest performance levels. If an organization needs predictable high performance in mission critical environments, it should continue to deploy full volume copies. However, organizations with differing or tiered levels of data classifications now have the option to explore Space Saving Snapshot solutions as a cost-effective means to meet data security, failover, and recovery requirements. This approach to backup becomes even more compelling as enterprises seek to maximize the price/performance of today's high-end storage arrays, which regularly support multiple applications and workloads within single systems.

When Space Saving Snapshots are used in conjunction with full-volume copies, new deployment models such as the three described in the following section offer recovery options and objectives not available by either technology alone.

The Impact of Adding Space Saving Snapshots

Sample Continuity Plan A: Protecting Data via Daily Tape Backup Only.

For organizations not currently deploying any form of local disk-based replication, implementing Space Saving Snapshot technology is a cost-effective way to reduce RPO from 24 hours (as in typical tape backups) while requiring a relatively small investment in additional disk capacity. Space Saving Snapshots can also enhance recovery times since the restore operation will be from disk rather than tapes (which are typically stored offsite). The

incremental capacity required depends on the application's write ratios and interval of snapshots, but is typically between five and twenty-five percent of the source capacity. Thus, for an organization with 4TB of source data, the typical incremental capacity needed for Space Saving Snapshots would be ~1TB.

Results:

- RPO reduced from 24 hours to ~1 hour.
- RTO reduced from hours or possibly days, to minutes or possibly hours.
- Deploying Space Saving Snapshots significantly reduces RPO and RTO

Sample Continuity Plan B: Protecting Some Production Data with Full-Volume Copies

Organizations that currently deploy full-volume copies to protect a small percentage of their production data can deploy snapshot technology as a cost-effective way to improve protection levels for the balance of their information. By taking snapshots at regular intervals against the production data that is not protected by full-volume copies, organizations can reduce RPO for volumes that were only backed up to tape every 24 hours. Space Saving Snapshots can also enhance recovery times since the restore operation will be from disk rather than tapes (which are typically stored offsite). For an organization with 2TB of data currently protected only by tape, the incremental capacity needed for snapshots would be ~500GB.

Results (for previously unprotected data):

- RPO reduced from 24 hours to ~1 hour
- RTO reduced from hours or possibly days, to minutes or hours
- Deploying Space Saving Snapshots in conjunction with full-volume copies helps meet mixed service levels while reducing RPO and RTO

Sample Continuity Plan C: Protecting Most Production Data with Full-Volume Copies

Organizations that are currently protecting most of their data with full-volume copies can deploy snapshot solutions as a cost-effective means of further reducing RPO while requiring an incremental investment in additional disk capacity. Typically, organizations will maintain two full-volume copies per source volume, cycling each one every 12 hours. One copy remains attached and actively mirrored while the other copy is in a *split state* (inactive unless needed as the source of a restore operation). Snapshots are taken at regular intervals against active full-volume copies, allowing organizations to maintain the high performance associated with full-volume copies while reducing their recovery point objectives from 12 hours to 1-2 hours (depending on requirements). For an organization with 4TB of source data protected by two full volume-copies, the incremental capacity needed for Space Saving Snapshots would be ~1TB. For the same organization's balance of unprotected data (e.g. 2TB), the incremental capacity needed for snapshots would be ~500GB.

Results:

- RPO for currently protected with full-volume copies is reduced from 12 hours to 1 hour
- RPO for currently unprotected data is reduced from 24 hours from tape to 1 hour
- RTO reduced to minutes
- Deploying Space Saving Snapshots in conjunction with full-volume copies helps meet mixed service levels while reducing RPO and RTO

The Landscape of High End Data Replication from Various Vendors

Several storage vendors offer solutions that meet the basic full-volume replication needs of high-end data storage customers, though some products offer more options than others do (*refer to table 1*). To complicate the issue, different vendors use the terms "Snap," "point-in-time" copies, and related solutions differently. At this time EMC appears to be the only vendor offering a true pointer-based space-saving snapshot product (TimeFinder/Snap) that is commensurate with high-end storage environments such as EMC's Symmetrix.

Table 1: Comparative Replication Capabilities of High End Storage Solutions

Vendor Product	Space-Saving Pointer-Based Snapshots	Full Volume Mirrors	Full Volume Pointer-Based Copies	Full Data Level Pointer-Based Copies	Single Array / Multiple Array Consistent Splits
EMC Symmetrix DMX	Yes	Yes	Yes	Yes	Yes/Yes
HDS Lightning	No	Yes	No	No	Yes/No
IBM ESS	No	Yes	Yes	Yes	Mainframe Only

What Does It All Mean?

Just a few years ago, RPO were typically measured in days. However, given evolving RPO and RTO and disk capacity requirements for supporting PIT solutions, customers have been forced to make tradeoffs between cost and rapid recovery. While full volume mirror copies will likely continue as the replication standard for mission critical, high-end, high performance data environments well into the future, Space Saving Snapshots offer high-end storage customers a new option.

Space Saving Snapshots can notably enhance the availability and economics of replication technology. We believe this is helping fuel the growing realization that not all data is equal. In addition, Space Saving Snapshots further ignite interest in these snapshot offerings for high-end storage environments that support mixed storage and data tiers. Equally, if not more important is the fact that blending full volume copy technology and snapshot offerings within high-end data environments provides levels of flexibility and security that neither solution offers alone. We believe the complementary features of full volume and snapshot replication solutions provide a best practices approach that can securely and cost-effectively meet the ever-growing demands of an enterprise's value chain relationships with partners, suppliers, and customers.

Virtually all storage vendors offer solutions that provide meet basic full volume replication needs, however, at this time EMC is unique in offering its Space Saving Snapshot product – TimeFinder/Snap — that is commensurate with high-end storage environments such as EMC's Symmetrix. For enterprises that are suffering from the proliferation of their secondary and tertiary data stores, EMC's Space Saving Snapshot offering may prove to be a viable solution to a vexing data replication need. As such, enterprises that seek to blend full mirror and snap data replication technologies to support high end, tiered storage environments are well advised to consider EMC's TimeFinder/Snap offerings.