Computer Associates®

White Paper

24x7 Operations: Snapshot Backup/Recovery Strategy, Benefits and Implementation Best Practices

Featuring BrightStor® ARCserve® Backup r11.1 and Microsoft Windows 2003 Server VSS Hardware Snapshot

Authors:
Steven A. Menges, Worldwide BrightStor® Product Marketing and Launch Manager
Paul Ignatius, Vice President, BrightStor Development
Matthew Dickson, Vice President, BrightStor Product Management
Steven Hwang, Manager, BrightStor Development
Ming Zhao, Software Engineer, BrightStor Development

October 2004

Introduction

Businesses and organizations of all sizes today rely on their applications and data to be successful and efficient. This data must be protected and available to the users who rely on it. Rapid growth in users and required storage capacity, coupled with constant changes in hardware and software technology, make it a challenge to keep critical data truly protected and available. Despite these challenges, the ever-present risk of data loss makes protection and availability a necessity.

To keep data protected and available, a backup and recovery solution is a necessity. The desired level of availability will differ depending upon the needs of each individual organization. In addition, there are various cost impacts to consider as an organization to move up this scale of availability. Therefore, the solution must be customizable and flexible to an organization's evolving requirements and constraints.

This paper examines the challenges associated with achieving 24x7 operations, availability and risk mitigation, along with strategies and solutions for meeting these challenges and fulfilling business needs.

Business Challenge

The need for 24x7 availability is fast becoming a reality for organizations of all sizes. Today's organizations cannot tolerate scheduled downtime (blackouts), nor can they afford extended periods of slow system response (brownouts) caused by traditional backup activities.

Another key element in availability is how quickly an organization can recover from data loss. Since human error and technical malfunctions cannot be 100% eliminated, they need to be able to recover data quickly when data loss inevitably occurs.

In the past, many organizations adopted "one-size-fits-all" strategies for backing up data center information. If the data was on or connected to a server, they would back it up once a week with a full backup, and conduct an incremental or differential backup each of the other six days.

However, with massive data growth in email, databases and other data becoming prevalent, a new strategy is emerging for cost-conscious organizations. This strategy involves classifying different types of data into categories based upon its importance to operations, and then applying appropriate levels of protection to each. Such a strategy enables organizations to more efficiently spend limited resources and mitigate risks.

Solutions

Tape-based backup has traditionally been a reliable, relatively low cost data protection solution. Yet three key factors have recently changed:

- Increased requirement for at or near 24x7 operations
- Drastic cost reductions for highly capable disk hardware
- Reduction in complexity of disk-to-disk (D2D), disk-to-disk-to-tape (D2D2T) and snapshot technologies

Today's organizations must target limited resources at solutions incorporating the most powerful possible performance technologies for critical applications and data. Lower-cost solutions can be used as appropriate for less-critical data to keep costs in line. Management can be involved in classifying the data and can sign off on the strategy, so the IT staff can be assured they are focusing in the right areas.

A strategy such as this one incorporates several key technology concepts which ideally can be part of one flexible solution:

Online ("Hot") Backup and Restore

Online or "hot" backup technology enables organizations to back up critical data (including applications such as Exchange, Oracle, SQL Server and so on) while it is up and available to users. Online backup solutions back up the files not in use (the same way as with a cold backup), and log any open files so the changes to those files are backed up as soon as they become stable. The backup operations do impact performance somewhat for users, causing slower than normal access speeds to the database or application being backed up; we call this a "brownout." During restores, the application knows how to bring all the data back efficiently. The key benefit is that it allows data to be accessible, even when it is being backed up.

Disk-to-Disk (D2D) and Disk-to-Disk-to-Tape (D2D2T)

D2D and D2D2T involve doing the backup/restore operations using disk as the backup media (or as a staging area for later off-load to tape). The key benefit is that the backups and restores can be done extremely quickly, and that even higher performance technologies such as snapshot copies can then be used.

Snapshot-Based Backup and Restore

Point-in-time or snapshot copy technology further reduces the brownout window to a few minutes or seconds. It involves briefly freezing the data set and taking a quick snapshot of it. This shot can be used as the backup copy or be backed up (migrated) to another disk or tape device. The key benefit is that the backups and restores can be done at an extremely fast pace with virtually zero impact on operations.

Hardware-Based Snapshot

Hardware snapshot technology is an advanced, powerful technology that can enable very high levels of data availability for 24x7 operations. Snapshot-based backup and restore can be complex and include the coordination of many elements, including the:

- Application to be protected (such as Exchange, Microsoft SQL Server, Oracle and so on)
- File system and operating system
- Creation and maintenance of the snapshot (changes from vendor to vendor and technology to technology)
- Performance of a backup and restore operation with the snapshot data

Hardware-based snapshot technology is provided by a disk array hardware vendor or storage virtualization provider (HDS, EMC, HP and so on), and will work only with an organization's own arrays. It creates a full, point-in-time snapshot copy of the data set stored in another one of the vendor's disk arrays (or another section of the same disk array). This snapshot copy can then be used as the backup copy for restores, or it can be backed up/migrated to another disk or tape device for longer-term retention or archival use. Since the snapshot is an exact full copy of the original data set, it will double the required disk space for the first snapshot copy, and so on. The interface used to generate hardware snapshots varies from vendor to vendor.

Furthermore, to fully leverage this type of technology for applications such as Exchange in a backup and restore solution, backup software vendors must integrate, test and certify their products with the application, operating system and the hardware snapshot technologies of the disk array vendors. The key benefit is that backups and restores can be done quickly, both locally and remotely, with maximum reliability and little to no impact on operations.

Microsoft Windows Server 2003 Volume Shadow Copy Service (VSS)

This feature of Microsoft Windows Server 2003 enables organizations to create very fast point-in-time copies (Microsoft refers to them as shadow copies) of applications and their associated data volumes while maintaining complete data integrity. The backup system, independent of the application server, manages archiving of the shadow copy; specifically, it uses a technique known as transportable shadow copy to make the point-in-time copy available to the backup server. The key benefit is that this technology helps makes hardware snapshot technology more accessible to organizations of all sizes, with limited budgets and administrative resources.

The concepts of high-performance D2D, D2D2T and snapshot based backup/recovery have been around for a while, but complexity — combined with the hardware, software and administrative costs — made these technologies realistic for only a small percentage of organizations. The recent frustration many IT managers have felt is largely due to the gap between what was predicted (namely, "simple, low-cost D2D solutions") and the reality, namely, "still complex, highly people-intensive D2D solutions with lower-cost hardware."

In the past, disk vendors were able to develop and perfect hardware snapshot technology for their own arrays, but the coordination did not exist yet between the disk arrays, operating systems, and backup and recovery solutions. Today, organizations such as Microsoft, CA and leading disk array vendors such as Hitachi (HDS), HP and EMC work together to integrate and test solutions, therefore fully leveraging the capabilities of these technologies.

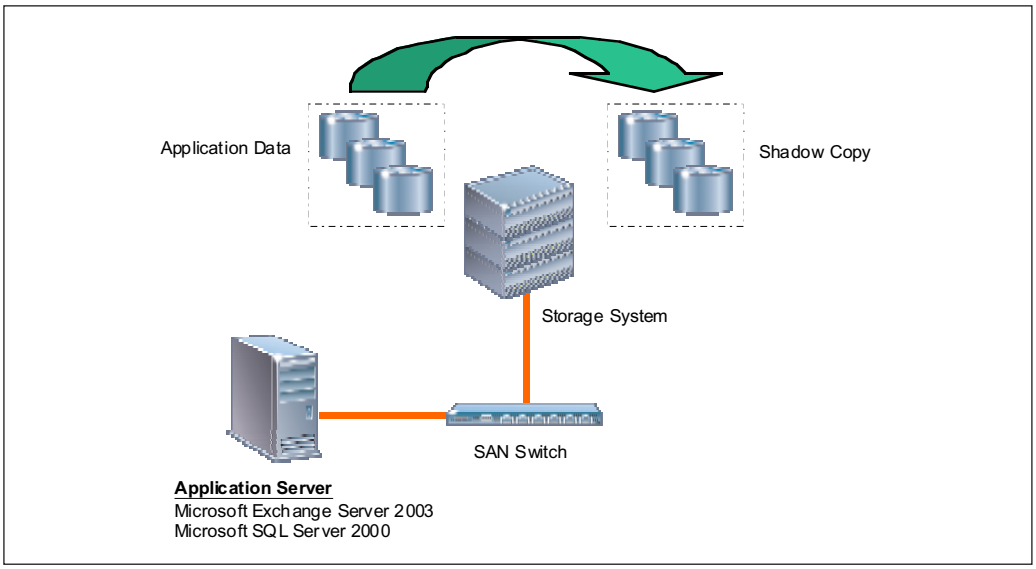

Figure 1: Applications or databases like Exchange, SQL Server or Oracle are snapshot copied and then immediately separated from the original data set

The benefits of using a risk mitigation strategy that incorporates high-performance D2D, D2D2T and snapshot-based backup/recovery as an integrated part of the overall solution are numerous. They include:

- Logical, defendable data protection strategy
- Efficient 24x7 operations
- Ability to continue leveraging existing hardware assets
- Focusing budget dollars on critical data, while minimizing costs elsewhere
- Reducing administrative cycles

The Role of CA Storage Technology

The Microsoft Windows Server 2003 VSS framework contains technology that makes the required three-way integration seamless and efficient for applications such as Exchange 2003 and others. Leading hardware vendors including HDS, HP and EMC have included the needed integration to make their technology work within the VSS framework. CA has developed the functionality to manage it all as part of a backup/recovery, data availability and/or overall storage management or IT management solution.

CA, Microsoft and these leading hardware vendors have tested and certified the complete solution and are the first to make it available today. CA today offers management of the complete, tested and certified solution as part of BrightStor® ARCserve® Backup r11.1 and the entire BrightStor® r11.1 integrated, automated suite of Intelligent Storage Management solutions.

BrightStor® ARCserve® Backup r11.1 for Windows

BrightStor® ARCserve® Backup for Windows delivers leading backup and restore protection for all Windows server systems as well as Windows, Linux, Mac OS X and UNIX client environments. Certified by Microsoft on all of its server platforms, this product is easy to use, reliable and scalable. Through deep integration with other BrightStor® solutions, it provides comprehensive data protection to ensure data availability.

BrightStor® ARCserve® Backup r11.1 Enterprise Module for Windows

The BrightStor ARCserve Backup Enterprise Module for Windows enables BrightStor ARCserve Backup customers to easily add enterprise functionality when and where they need it. By installing this module on top of BrightStor ARCserve Backup, organizations can use enterprise-class functions such as multi-streaming, portal, vaulting, managing other BrightStor ARCserve Backup environments and the ability to use Enterprise Options (such as the Enterprise Option for VSS Hardware Snapshot). The key benefit is that it is part of the solution that makes everything work together.

BrightStor® ARCserve® Backup r11.1 Enterprise Option for VSS Hardware Snapshot

The BrightStor® ARCserve® Backup Enterprise Option for VSS Hardware Snapshot enables BrightStor ARCserve Backup customers to manage and use hardware snapshot-based copies of data sets to deliver instant point-in-time recovery from disk. It also supports zero-impact backups and staging to disk, D2D2T backup, vaulting and laboratory testing. The BrightStor ARCserve Backup Enterprise Option for VSS Hardware Snapshot further supports the Microsoft Windows Server 2003 Volume Shadow Copy Service (VSS) to leverage technology available with leading disk arrays. The key benefit is that it is part of the solution that makes everything work together.

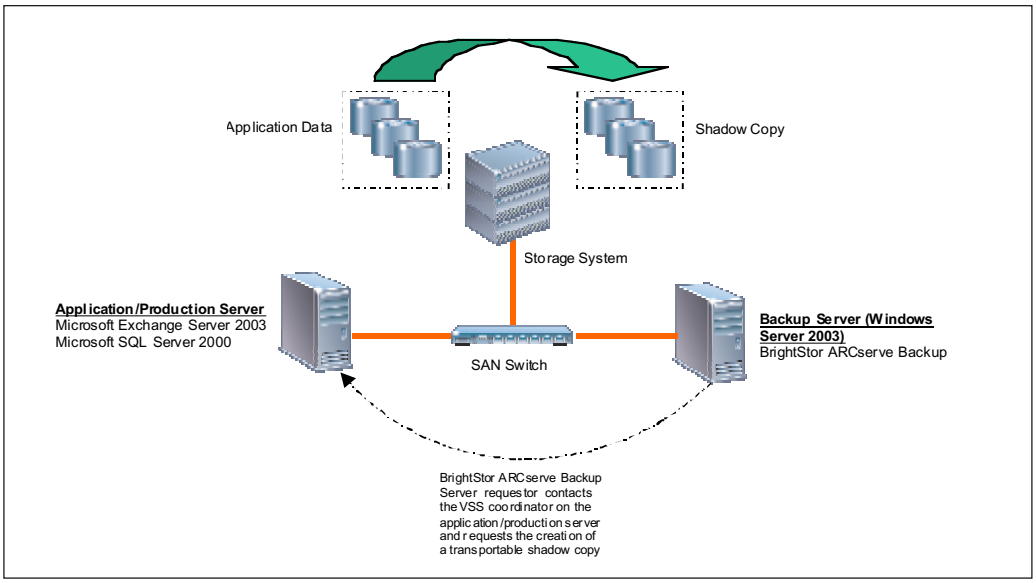

Figure 2: BrightStor ARCserve Backup r11.1 controlling the hardware snapshot process

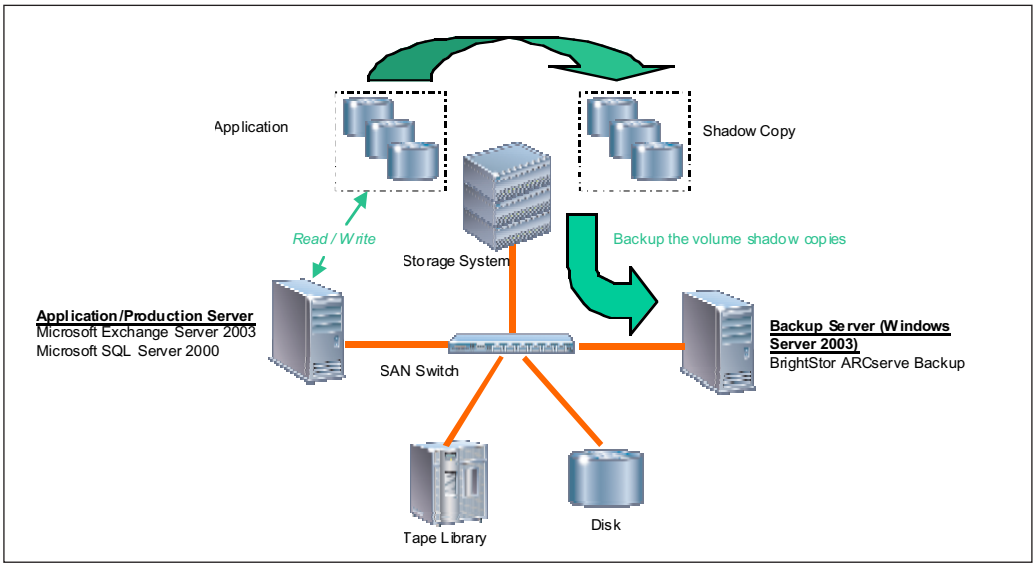

Figure 3: BrightStor ARCserve Backup r11.1 using the snapshot as part of a D2D2T data availability solution

It is important to note that not all third-party (non-CA) backup and restore solutions are designed to leverage these technologies as a part of an automated, integrated solution. Therefore, if an organization is considering various options, it should be sure to verify what is supported by the vendor, how the integration (if any) supports all the needed technology, and how easy the entire solution will be to use by administrators. (For detailed information and best practices on how to efficiently implement and use the solution, please refer to the implementation guide entitled: ***Best Practices For Using BrightStor ARCserve Backup r11.1 and Microsoft Windows 2003 Server VSS Hardware Snapshot.***)

Conclusion

The quest for 24x7 operations has made maximum availability of critical data a necessity for organizations that want to remain efficient and competitive. However, many organizational budgets have remained flat. Therefore, the storage solutions chosen must meet an organization's particular needs and enable complete protection, the highest possible availability, rock-solid reliability and minimal total cost of ownership (TCO).

The first step is to clearly identify recovery, backup and archive requirements so that an organization can focus on the right set of technologies to meet those requirements. An organization should then carefully look at the solutions that meet their toughest requirements and that also offer the flexibility to combine lower cost technologies for less critical data. They can then compare the solutions in light of the entire TCO, including required consulting, components, vendor management costs and so on.

Managing and keeping available critical applications and growing data is a tremendous job, but it is one that has been made easier with an integrated snapshot-based backup and recovery solution from CA, Microsoft and leading disk vendors such as HDS, HP and EMC. The fact that these advanced technologies can now be leveraged as part of an integrated data availability solution — including the ability to use more traditional and lower-cost technologies like tape — makes the new technologies even more attractive to cost-conscience organizations.

High-performance, relatively low-cost D2D and snapshot technologies are readily available and may already be sitting on your data center floor or loading dock, but fully leveraging their capability has been complex. Now, fully leveraging these technologies to meet your business objectives has been made simpler for organizations of any size. In addition, using these technologies as part of an integrated, automated, fully tested solution such as CA's BrightStor r11.1 can help reduce TCO by cutting back on required administrator cycles, consulting services and vendor management costs. By leveraging other BrightStor r11.1 Intelligent Storage Management solutions as well as CA's integrated eTrust™ security and Unicenter® enterprise infrastructure management solutions, IT managers can gain additional efficiencies while getting best-of-breed IT solutions from a single trusted source.

The bottom line is that the vendors have worked together to make their latest, most advanced technologies efficiently work together as part of a data availability solution virtually any type or size of business can now efficiently implement and manage. The goal of data availability for 24x7 operations on a limited budget is now attainable for virtually everyone.

Computer Associates®