

Planning for and Surviving a Data Disaster

Solutions to successfully meet the requirements of business continuity.

- 2 Introduction
- 2 Managing Host Storage for Virtual Environments
- 5 Evaluating the Common Precursors of Data Loss Events
- 7 Tips on How to Improve Data Recovery Service Success

Introduction

Business continuity was conceived as a solution to protect mainframe computer systems and early data centers. The profession was originally called “disaster recovery” and consisted of project planning and support from equipment vendors. As the profession matured, disaster recovery planning became a subset of an organization’s business continuity plan. The business continuity plan has become the umbrella-like policy that ensures all of a business’ departments can operate successfully with minimal or limited impact during a disruptive event.¹ The disaster recovery plan and emergency response procedures all fall under the business continuity plan. What started as a formal procedure to protect expensive computer equipment has crossed over to protect all elements of a business organization.

The advent of virtualization technology has enabled business continuity planning and execution for many organizations. Virtualization technology is complex and requires proficiencies from IT staff and management to gain a complete return from an organization’s investment. Unfortunately, if not deployed or managed carefully, virtualization can itself create business disruptions or data disasters.

Managing Host Storage for Virtual Environments

This article examines the state of virtualization within the corporate world. Also, data recovery service providers give practical tips on virtualization to optimize (and perhaps even lower the cost of) data recovery.

Asset Identification

Asset identification and management for physical hardware systems has always been fairly straightforward. Naming conventions and identification within the virtualized environment, on the other hand, tend to be cryptic and are complicated by explosive system growth. (See Figure 1. Find the development domain controllers—if you can.) Virtual machine identification taxonomy that is cryptic or inconsistent between physical and virtual systems often leads to human error when there is a business disruption.

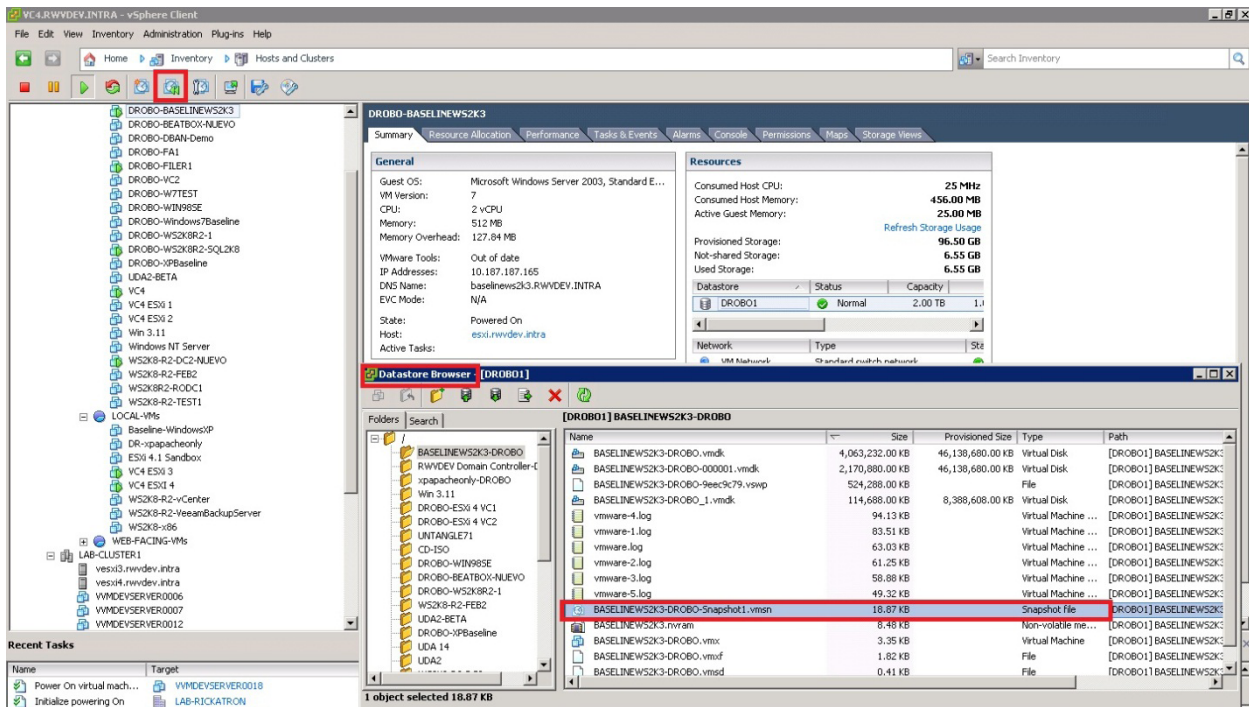


Figure 1 – “Identifying Virtual Machine Snapshots in vSphere,” Rick Vanover

Source: www.techrepublic.com

¹ For purposes of this article a business disruption is anything that prevents day-to-day work from being done, including power disruption, downed phone lines, and so forth. A data disaster occurs when data is corrupted. Hence, a data disaster is a subset of business disruption.

Data recovery service providers report that victims of IT disasters have little or no basic documentation of failed storage systems. Data recovery engineers suggest that if an organization were to keep a simple, current screen shot of each host server's virtual machines it would accelerate recovery efforts to identify priority systems.

Snapshot functionality within hypervisors was never meant to replace backup solutions; instead it is a method of preserving the virtual machine's state and disk data at a specific point in time.

Virtualization Snapshot Management

Snapshot functionality within hypervisors was never meant to replace backup solutions; instead it is a method of preserving the virtual machine's state and disk data at a specific point in time. For example, a routine snapshot before system maintenance should be committed to the primary virtual disk after successful maintenance procedures. (See Figure 2.)

Multiple snapshots slow down access to virtual disk data due to the fact that snapshot files contain a subset of the data stored in the primary virtual disk file. Snapshots also use up valuable disk space and if left unchecked can fill up a data-store and cripple a virtual environment. When a disaster occurs, all of the virtual disk files need to be recovered to meet recovery point objectives.

Data recovery service providers cite multiple snapshots as a hindrance to a successful recovery when virtual disk files have been deleted or if a data-store volume has been reformatted.

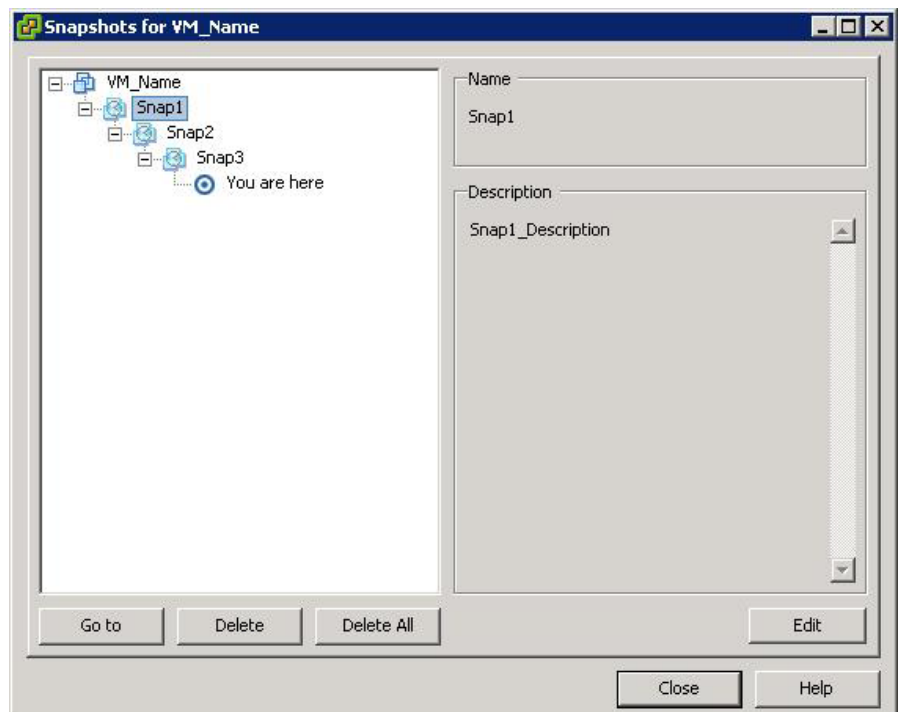


Figure 2 - "Troubleshooting Virtual Machine Snapshot Problems," Ruben Garcia, February, 2010

Business Continuity Plans Not Tested

Business continuity experts and data recovery service providers consistently report that companies are not completely testing their business continuity or disaster recovery plans. According to Don Stewart, director of professional services at Ongoing Operations, a non-profit business continuity service provider for U.S. credit unions, “many business continuity plans are not exercised fully. Recently, when testing one company’s plan, the recovery coordinator said that the first step of the action plan was to find the 500-plus page business continuity plan; he asked for time to find the plan.” Stewart reports that many organizations perform a yearly tabletop exercise of their plan but that it mostly turns into a review and updating of participant’s phone numbers. Stewart says that the only way to verify continuity and recovery plans is to conduct a real-life exercise; acting out the disaster provides insight to management on how to contain a business disruption and how long it will take to restore services (i.e., how much the disruption will cost).

Stewart recommends that instead of measuring the maximum allowable outage during an exercise in days or hours it should be measured by economic numbers, that is, dollars. Once the financial impact of a disruption is determined, priorities are driven by realistic goals.

Mercy Medical Center, Cedar Rapids, Iowa, provides a success story of having a business continuity plan in place for the entire organization. They successfully put their business continuity plan into action during the Midwest floods of 2008, and according to the hospital’s website, after three weeks the hospital returned to full operations. The *Wall Street Journal’s Health Blog*² has a compelling interview about the plan’s evacuation and recovery process.

Data recovery service providers report that organizations are doing better with business continuity planning but have observed that the choice of a data recovery service is rarely a formal part of the disaster recovery plan. If selected during the midst of a disaster, the criteria used to evaluate a data recovery service provider will be less robust than the criteria for choosing which coffee maker to install in the breakroom. Since an inexperienced or underequipped data recovery provider can make a bad situation worse, the best time to choose a provider is before a data disaster. Here are five criteria for evaluating data recovery providers:

- Identify companies that have the technology and resources to solve a wide array of data loss challenges. A data disaster may affect several platforms. For example, a disruption may affect UNIX, Linux, and Windows®-based systems that are all running on one virtual server.
- Identify companies that provide data recovery solutions to fit your specific needs
- Identify companies that will provide you with the information required to make an educated purchase decision
- Identify companies that offer professional customer service whenever and wherever you need it
- Identify companies that have well documented and established procedures for maintaining the security and confidentiality of your data. Few data recovery providers have submitted themselves to and passed an IT security compliance audit. If it is critical that your data remains secure, find a company that has passed a third-party security audit.

² “When Hospitals Fall Victim to Disaster,” Shirley S. Wang, Wall Street Journal, June, 2008

Evaluating the Common Precursors of Data Loss Events

Information technology staff and data recovery professionals list the following precursors as the most frequent causes of data loss:

- Human error
- Storage hardware failure
- IT disaster recovery plan that is weak or not exercised regularly
- Overconfidence in a SAN storage redundancy
- Corrupted or unreadable backups or archives of missing data

Here are some comments and observations from IT professionals about the importance of the projects they are working on, the magnitude of those projects, and the impacts data disasters have.

“We’ve been in the planning stages for three months now. I can’t tell you how many scoping and business impact analyses I have done. I don’t trust any storage, SSD, Cloud, or tape, which is why my data is stored in multiple locations. I plan for failure and have a solution to protect the data.”—From an IT architect who is ready to start an e-mail migration affecting forty thousand users.

“It’s a six year business intelligence project with data aggregates in the 100TB range. There’s a lot of time being spent on creating metrics and mapping the data. The raw data is going to have thirty- to forty billion rows in a single table. There’s no room for error for the team I’m working with.”—From a retail sector business analyst.

“Nobody tests their backups; they’re either incomplete or untested. They have this sense of security with virtualized storage and don’t backup! And get this, backups are made to the same SAN that holds the original data; when the SAN goes down, everything is inaccessible and this situation delays the recovery effort.” —From a data recovery engineer.

“RAID controller failures are the biggest support calls we deal with. These types of failures are slow to identify and big on disaster. Most IT admins do not have a plan to handle these types of events until the entire system crashes. When we support these types of calls, we do not go to the backup right away. We analyze the I/O event logs to see when the problems started. Then, through a combined effort of our replication solution and portions of other backups, we selectively restore the missing data. It’s a planned recovery execution so that recovery time objectives are met. This also helps us to meet recovery point objectives that business owners have established.”—From a business continuity service provider.

These observations illustrate that as storage and data increases in size, there is an ever increasing need to prepare for business disruptions and data loss. IT projects are getting bigger, planning is taking longer, and enterprise data is closing in on the petabyte range. Backups are not providing enough coverage, and data loss events have a devastating impact on companies in competitive markets.

Planning For and Surviving a Data Disaster

According to IDC's worldwide tracking of external disk storage systems, total disk storage capacity shipped was over 5,100 petabytes—a 55.7 percent increase over the previous year.³ This continued growth will require IT management to maintain disaster recovery documentation and to exercise recovery plans regularly. This will minimize or eliminate business disruptions due to data loss within virtualized environments.

Business disruptions caused by data disasters present a challenging situation. IT staff are scrambling to get priority systems up and running while senior management worries about the larger impact to the organization and its clients.

Successful organizations realize that any disruption, regardless of how small, will have an impact on the business as a whole. This has led IT leaders and business continuity planners to proactively include data recovery services in their contingency plans. Choosing a data recovery service vendor before a disaster occurs prepares the IT team for a successful survival of a business disruption.

According to IDC's worldwide tracking of external disk storage systems, total disk storage capacity shipped was over 5,100 petabytes—a 55.7 percent increase over the previous year.³

Tips Cheat Sheet

Virtualization brings an extra layer of complexity to a host system and when data loss occurs, it is critical to select a data recovery provider experienced in recovery from virtualized systems. Below are tips on how to safely recover data from virtualized environments:

- Restore backups to a different volume. This ensures that all important files are good on the backup before possibly overwriting data on the active volume.
- If there is a RAID problem, image each drive from the RAID before attempting a rebuild. Sometimes a RAID rebuild does not work correctly and can make the problem worse. Test backups by restoring to a different location before overwriting the RAID array.
- Do not create any new files on the disk needing recovery or continue to run virtual machines until the important data is recovered. New files can overwrite the files that need recovery if restoring the backup fails. Virtual machines using snapshots and thin provisioned virtual disks that are still in use after the data loss can also overwrite files that need recovery.
- Do not run FSCK or CHKDSK or other file system repair tools on a virtual disk unless a good backup has been validated by restoring it to a different volume. These repair tools assume that there is a good backup of the data and can overwrite file pointers to make a file system consistent. If desired, these tools can be run in read-only mode to find any major corruption before repairs are made.
- If one virtual disk needs recovery but others are still running from the same volume and cannot be shut down during the recovery, clone or migrate them to another volume. If a deleted virtual disk or snapshot needs recovery, it is best to copy or clone the virtual machines instead of migrating them so they are not found as part of the deleted recovery.
- Shutdown or clone/copy any other active virtual machines on the same volume that are thin provisioned or are using snapshots. Any writing to the new blocks on the volume can overwrite recoverable data.

³ "Worldwide Disk Storage Systems Finishes 2010 with Double-Digit Growth on Strong Fourth Quarter Results," IDC, March, 2011



For more information, call or visit us online.

800.872.2599 in the U.S. and Canada

+1.952.937.5161

www.krollontrack.com

Copyright © 2011 Kroll Ontrack Inc. All Rights Reserved.
Kroll Ontrack, Ontrack and other Kroll Ontrack brand and product names referred to herein are trademarks or registered trademarks of Kroll Ontrack Inc. and/or its parent company, Kroll Inc., in the United States and/or other countries. All other brand and product names are trademarks or registered trademarks of their respective owners.

XXXXX