

Data Loss in a Virtual Environment— An Emerging Problem

Solutions to successfully meet the requirements of business continuity.



An Altegrity Company

- 2 Introduction
- 3 Common Virtual Data Loss Scenarios
- 4 Recent Virtualization Data Loss Case Studies
- 5 Data Loss Consequences


Introduction

The terms “business continuity” and “disaster recovery” have been blurred over the years and this has created confusion for businesses who are attempting to protect their business operations. A business continuity plan is a comprehensive policy that ensures all of a business’ departments can successfully operate with minimal or limited impact during a disruptive event.¹ The disaster recovery plan and emergency response procedures are generally part of a larger business continuity plan.

The advent of virtualization technology has enabled business continuity planning and execution for many organizations. However, virtualization technology is complex and requires specialized skill and knowledge from both IT staff and management. In fact, if deployed or managed carelessly, virtualization can itself create business disruptions or data disasters as outlined by the graphs below.

If deployed or managed carelessly, virtualization can itself create business disruptions or data disasters.


Traditional System



Source of Failure: Human vs. Machine

- Human error
- Lack of education
- System and hardware failure
- Environmental (power outage, over-voltage, etc.)

Virtual System



(figure 1)

According to Forrester Research’s report on the business state of disaster recovery preparedness, a joint effort with the *Disaster Recovery Journal*,² many organizations have improved their disaster recovery capabilities over the past few years. Despite a slow economy, survey respondents reported an increased confidence in being prepared for a data center disaster or site failure.

Seventy-six percent of survey respondents reported no disaster or major disruption in the past five years, yet Forrester Researcher reports that companies should not take comfort in this statistic. Instead, it should serve as a wake-up call because a whopping 25 percent of companies are likely to declare a disaster. Furthermore,

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.

² Forrester Research’s 2010 report on the business state of disaster recovery preparedness, a joint effort with the *Disaster Recovery Journal*, http://www.drj.com/images/surveys_pdf/forrester/2011Forrester_survey.pdf

business disruptions are much more common than “declared disasters.” Getting an organization to declare a disaster can be a matter of perspective, according to Don Stewart, director of professional services at Ongoing Operations, a non-profit business continuity service provider for U.S. credit unions. “In some events, IT is so focused on fixing the problem that they don’t inform senior management of the disaster event,” Stewart relates. Some organizations have not defined what a business disruption is, therefore senior management will hesitate to declare a disaster if the event is perceived to be minor; for instance in the case of a phone system failure, or delays in e-mail messaging.

Fifteen percent of respondents knew the cost of their business’ downtime; it averaged nearly \$145,000 USD per hour.


Staying prepared requires more than having a documented business continuity plan; it requires teamwork from all stakeholders. Having a stake in planning at this level ensures that business operations would be maintained in the event of a disruption. Stewart recommends that a good plan starts with a risk impact analysis. Most companies, according to Stewart, will purchase an in-depth risk assessment and then do nothing about it; “the report just sits there with no further actions being taken.” This is as effective as making a list of essentials to pack in a kit in case of a house fire but never assembling the kit.

Common Virtual Data Loss Scenarios

When data loss happens within a virtual data center it is usually due to human error. Other virtual data loss events result from hardware failure and are exacerbated by the lack of a disaster recovery plan. Disaster recovery plans that are weak or not regularly tested force IT staff to focus on un-tested repairs based on faulty troubleshooting. Obviously, nobody wants a data loss or business disruption on the systems they are responsible for. Too often, a serious data loss or business disruption results in unemployment for all responsible or thought to be responsible.

Other data loss scenarios have been due to overconfidence in a SAN’s redundancy. Precious recovery time is lost when it is discovered that important backups are corrupted or not readable during the middle of a disaster. This is the worst possible time to learn that backups have been failing or that backup software has not been reporting media failures during backup sessions. Reference the pie chart provided by Kroll Ontrack that provides a breakdown of 2010 virtual data loss types:

Virtual Data Loss Types



(figure 2)

Recent Virtualization Data Loss Case Studies

A virtualization data loss can be catastrophic for an organization. Determining the financial impact of a business disruption is difficult because there are both tangible factors, including productivity loss, missed sales opportunities and staff's hourly time, but also less tangible factors to downtime such as potential non-compliance penalties, damage to corporate image and weakened customer confidence. The previously quoted Forrester-DRJ survey noted that 15 percent of respondents knew the cost of their business' downtime; it averaged nearly \$145,000 USD per hour. That is an estimate that would make any director or CIO take notice of the readiness of their business continuity plan. Virtualization technology can compound those losses as illustrated by the following cases.

The Case of the Reformatted Server

A business in Italy recently experienced a business disruption when their 4TB virtual host server lost access to the storage system. The virtual environment contained 40 virtual machines in a mixed operating system environment; some were Linux, a few were legacy UNIX systems and the remainder consisted of Microsoft® Windows® servers. These supported business application servers, web servers and database servers.

The virtual host server was operating a Linux-based hypervisor with two 2TB LUNs attached. At some point, the storage LUNs were reformatted. The reason for the reformat wasn't divulged, but the damage to the existing file system structures was severe and extensive. During the reformat process, the Linux storage manager writes EXT file system metadata in predefined areas throughout the volume. This metadata contains only a couple of thousand bytes of information, yet the impact upon the virtual host server's file system and virtual disk files was devastating.

Each virtual machine had four to six virtual disk files totaling 70-90 virtual disk files stored by the host server. Some of the virtualized Microsoft Windows servers employed "dynamic disk volume" configurations (i.e., "logical volume manager" in Linux parlance) between multiple virtual disk files, further complicating recovery efforts.

After the organization's IT department exhausted all internal resources, a professional data recovery firm was engaged to recover the data. Despite the damage, virtual disk files were found and critical data was restored.

It is industry best practices combined with IT management procedures that ensure data protection.

The Case of the Disastrous Data Merge

A United States business merger suffered a disaster while the two company's IT departments were merging their data. Evidence suggests the disaster was caused by employee sabotage and the cause is still under investigation by computer forensic investigators.

The first company's virtual host server held over 400 virtual machines across 20 storage LUNs. During the data merge, someone with administrative access to the virtual host server systematically deleted the 400 virtual machines and their virtual disk files, causing the loss of over 440 virtual disk files and over a thousand snapshot files.

The merging company quickly engaged emergency data recovery services and prioritized core servers that provided essential services. In three days, those systems were up and running. For the next two weeks emergency recovery efforts continued on the rest of the storage system. This required extensive recovery engineering efforts to search the unallocated areas of the storage LUN for potential virtual disk files, identifiable only by their file system attributes.

Through a combined effort of backup restoration and original volume recovery, data was recovered. Most of the virtual disk files were complete, while other virtual disks required the file contents to be extracted due to file system damage.

The Case of the Off-Site SAN Reformat

Disaster recovery efforts went from bad to worse for a company in Luxembourg. During routine maintenance on the company's SAN storage that housed its virtual machines, the SAN was presented to a different physical server by accident. When the SAN storage was identified as "unknown," the volume was automatically reformatted. Initially, some staff panicked due to the potential for data loss. They were relieved when they were reminded of the identical SAN storage located off-site which employed the SAN equipment's automated site replication technology. It was thought that this would be a minor business disruption.

Upon logging into the remote SAN, the IT team discovered that the remote SAN was an identical copy of the primary site; the SAN's automated site replication technology had not been disabled prior to the maintenance. Thus, when the reformat occurred at the primary site, the secondary SAN was reformatted as well. Through the efforts of experienced data recovery engineers, the virtual machines and virtual disk files were successfully recovered.

This organization did not have any backups because it was assumed that dual storage architecture and site replication mechanisms provided complete data and system redundancy. This case is especially compelling because storage equipment features provided a false sense of security. In reality, it is industry best practices combined with IT management procedures that ensure data protection.

Data Loss Consequences

Virtualization technology has revolutionized the IT industry and has delivered on the promise of a reduction in facility expenses and equipment costs. 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 requires 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.

As more storage is consumed by virtualization technology, additional attention must be given to the management and protection of virtual assets. Maintaining business continuity by having well planned and tested disaster recovery plans is essential.

Successful organizations realize that any disruption within the virtual infrastructure, regardless of how small, will have an amplified 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 caused by a data disaster.

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