

WHITE PAPER

Disaster Recovery Management: An Emerging Approach to Avoiding Data Protection Disasters

Sponsored by: Continuity Software

Dan Yachin
June 2007

IDC OPINION

Disaster recovery (DR) and business continuity are a top priority for organizations of all sizes. In light of the increasing reliance on IT for business operations and the need to comply with regulations on the retention and protection of data, organizations' spending on data protection and recovery solutions is on the rise. Despite the increased spending on DR solutions, organizations are still facing difficulties in building and maintaining effective DR operations.

As DR implementations grow in complexity and size, it becomes increasingly difficult to make sure that critical data and business processes are successfully backed up and that required protection levels are constantly maintained.

Traditionally, organizations have addressed this problem by performing periodical testing of their DR environments to identify gaps between their main and backup sites. These tests are often performed by manually looking for configuration mismatches in specific applications, making fixes, running scripts, and implementing the required updates. But this approach carries several drawbacks such as the operational disruption as down critical systems are tested in failover scenarios; the inability to cover the entire environment (only specific applications are manually tested); and the lack of real-time information.

Given the dynamic nature of data center environments today, in which configuration changes occur frequently, the number of potential configuration gaps might create multiple DR vulnerabilities. As a result, organizations face uncertainty regarding their ability to successfully and in a timely manner restore operations in cases of disaster.

To address the shortcomings of traditional approaches, an automated DR management approach is needed that could maintain ongoing monitoring of complex DR environments and identify vulnerabilities in real-time rather than after the fact.

METHODOLOGY

IDC developed this white paper using a combination of existing market forecasts and direct, in-depth, primary research. To gain insight into the challenges of providing DR management solutions, and to learn how Continuity Software's RecoverGuard solution can help organizations mitigate risks associated with data protection gaps, IDC interviewed the company team on the issues of technology, product offerings, competitive landscape, and go-to-market strategy.

IN THIS WHITE PAPER

This IDC white paper discusses the need for DR management solutions to provide continuous monitoring and optimization of DR environments to identify and resolve data vulnerabilities that are the result of the frequent changes in enterprise data centers. It analyzes traditional approaches towards DR management and their shortcomings, and how emerging approaches can help organizations to continuously monitor their DR operation and its performance.

SITUATION OVERVIEW

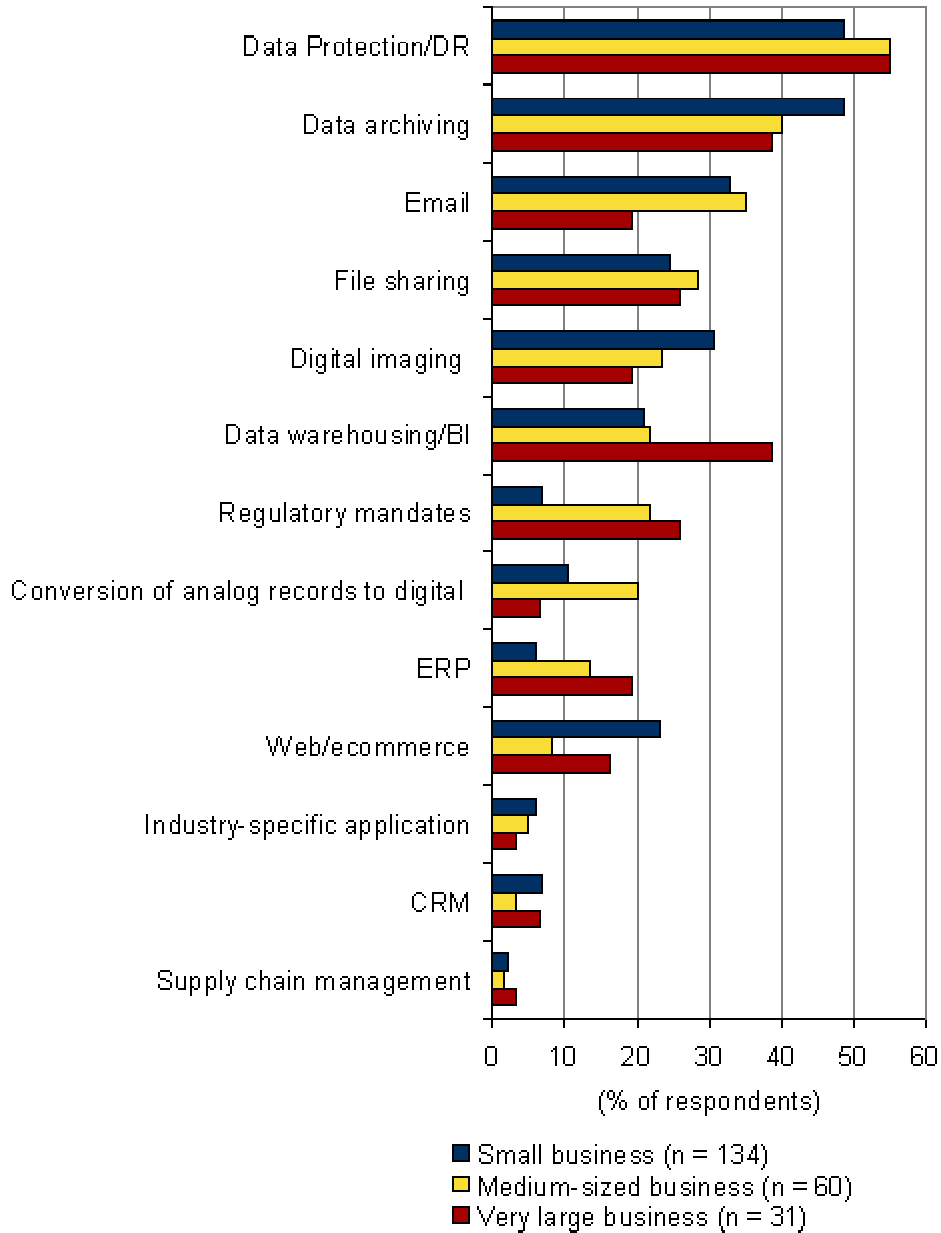
Introduction

Disaster recovery remains a top concern for organizations. Fueled by real-world disasters such as 9/11 and Hurricane Katrina, regulatory compliance and corporate governance requirements, awareness of the need to protect corporate data assets and ensure business continuity is rising steadily among companies of all sizes and across various industries.

Given the growing dependency on IT to support day-to-day business operations, and the risk of losing data that is critical to the functioning of a business, the importance of disaster recovery cannot be underestimated. An IDC survey found that data protection/disaster recovery is the top storage priority of very large and medium-sized businesses in the U.S. (see Figure 1) – and their spending on data protection and recovery (DPR) solutions, including backup, data availability, replication, and mirroring, among others, is on the rise. The survey also revealed that over 50% of all U.S. firms — small, medium-sized, and very large — rated data protection/disaster recovery as their number one driver for planned spending on additional storage capacity.

FIGURE 1

U.S. Business Top 2 Storage Priorities in the Next 12 Months by Company Size

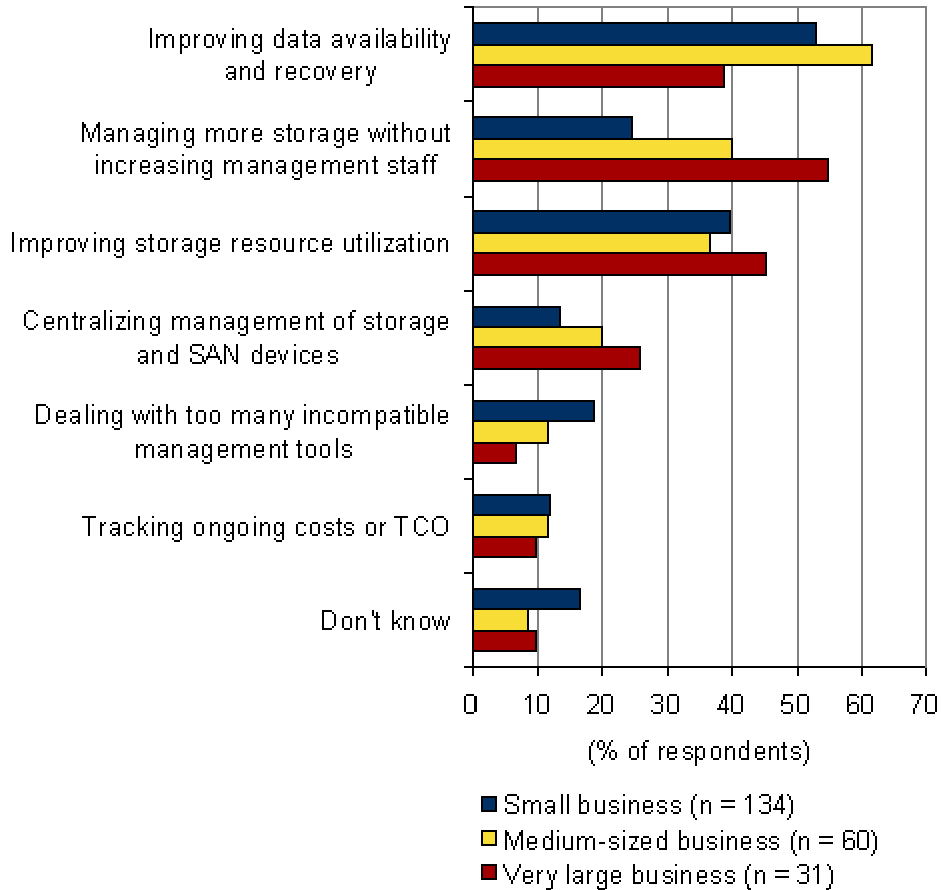


Source: IDC/Infoworld 2005 Trends in Storage Survey

Despite the increased spending on various DPR solutions, organizations are still expressing a need to improve data availability and recovery. As seen in Figure 2 below, 50% of small firms and 60% of medium-sized firms mentioned this issue as their primary challenge.

FIGURE 2

U.S. Business Top 2 Challenges Relating to Storage Management by Company Size



Source: IDC, 2007

The need to improve data availability and recovery stems from the fact that disaster recovery operations are usually poorly managed. For the most part, DR is about allowing an organization to restore operations with as little disruption as possible and as quickly as possible in case of a disaster. To accomplish that, many DR strategies are based on identifying the business processes and functions that are critical for ongoing operations, and the systems and applications that support them. After defining different SLA metrics such as Recovery Time Objective (time required to restore operations, RTO) and Recovery Point Objective (point in time by which data must be restored, RPO), different DPR solutions can be used to replicate data to a secondary or remote site for recovery purposes.

The main problem associated with DR implementations is the lack of sufficient means to continuously monitor and ensure that all critical data is successfully and properly backed up, and that the predefined protection levels are consistently met. As a result, many organizations are uncertain as to whether all data will be restored, in a timely manner, in case of disaster.

The challenge of maintaining ongoing visibility into data protection levels stems from the complexity of a typical DR implementation, which consists of multiple different DPR products. Coordinating and managing this heterogeneous environment involves high costs and significant technical challenges, which in most cases cannot be solved in a comprehensive manner. This problem is compounded by the ever-growing amounts of data that need to be protected.

Data Protection and Recovery Management (DPRM) solutions have emerged in recent years to address this problem. These solutions are aimed at providing a global, unified view of different backup products, configurations, jobs, and assets, and have provided organizations with some new capabilities. For example, DPRM products can provide important statistics on DR performance and utilization, as well as on success and failure data and, in many ways, provide functionality or information missing from the native backup application itself. In addition, DPRM products are increasingly being used to provide proof of data protection processes and controls for the purposes of compliance and IT governance (see *Worldwide Data Protection and Recovery Management 2007-2011 Forecast and Analysis: Who's Who and Why It Matters*, IDC#205571).

DPRM solutions can significantly improve the ability to keep track, measure, and improve upon current data protection and recovery processes, as well as to prove compliance. But these solutions were not designed to address another critical problem, which serves as a key inhibitor to obtaining visibility and control over DR environments – the inability to analyze configuration changes in the surrounding IT environment and their impact on the organization's DR readiness on an ongoing basis.

The DR Testing Conundrum

Today's enterprise data centers are highly dynamic environments, in which multiple production configuration changes occur frequently. As any modification or upgrade of the infrastructure or the applications can affect an organization's DR readiness, making sure that all changes are properly reflected in the secondary site, and that there are no configuration gaps between the two sites, is becoming a significant challenge.

Traditionally, organizations have been dealing with this problem by performing periodic testing and manual audits to identify potential gaps between the main and secondary sites. These tests are often performed by manually looking for configuration mismatches in specific applications, making fixes, running scripts, and implementing the required updates.

The periodic manual testing method has several key drawbacks:

- ☒ **Operational Disruption:** Manual DR testing is highly resource-consuming process that often involves operational disruption as critical systems need to be taken down to test failover scenarios. In addition, DR testing requires the participation and support of many employees from multiple IT groups (e.g., operating system, middleware, database, storage, application). To minimize business disruption, organizations are scheduling DR testing in advance, and usually perform them during weekends. Thus many organizations are not testing their DR environment frequently (in many cases only once or twice a year, if at all). Nor is it tested in real-life scenarios when the business environment is active.

- ☒ Costs: Performing a manual DR testing involves significant time costs for IT and business staff, IT and communication costs, and others. Given the high costs of building and maintaining a DR environment, organizations are often left with limited budgets for testing it, which means they test less frequently.
- ☒ No Knowledge of Configuration Changes: Continuous detection of configuration changes to the different elements in an enterprise data center is a challenging task on its own, let alone verifying that changes were replicated to the secondary site. Obtaining a comprehensive view of configuration changes and their reflection on the secondary site is virtually impossible using the manual testing approach, as each application should be tested and analyzed on an ongoing basis. The more complex the IT environment becomes, the more difficult it is to keep track of configuration changes across multiple components from multiple vendors.
- ☒ Lack of real-time information: As mentioned, modern data centers are highly complex and dynamic entities in which configuration changes occur very frequently. Therefore, periodic DR testing cannot provide a consistent, accurate view of the DR environment and its readiness, as the information obtained is outdated almost immediately.
- ☒ Incomplete Coverage: A large organization may have hundreds if not thousands of different applications running in production environments and thus manually testing each and every application is practically impossible. As a result, some applications of lower priority might be tested only once every few years. In addition, testing the entire environment as a complete service may provide different results because of dependencies between business services.
- ☒ No Alignment between Business and IT: DR testing is often considered an IT operation, in which specific applications and IT infrastructure components are checked for DR readiness. As DR should by definition revolve around business continuity, it is important to map and correlate IT components to the business service they support to make sure that the entire business service is constantly protected – a goal that is difficult to achieve using manual practices.

As data centers grow larger and more complex, and business requirements frequently drive changes to the applications and IT infrastructure, the shortcomings of periodical testing and manual audits are further exacerbated. In this dynamic environment, manually identifying configuration gaps for each server, storage device, database and so forth, is practically impossible. In addition, due to regulatory compliance pressures, the move to 24x7 environments, and the increasing reliance of customer-facing online services, among other issues, organizations are requiring shorter RTO and RPO to minimize data loss risks.

As a result, there is increasing demand for new solutions to maintain ongoing monitoring and control of DR environments, and optimize their performance. It should be noted that technology cannot entirely replace the need to test the DR environment to make sure it performs as expected. But an automated means to monitor the health of DR operations, coupled with tools and approaches that can reduce the effort involved in DR testing and increase its effectiveness, can significantly improve organizations' DR readiness.

The Need for Disaster Recovery Management

Simply performing more tests is one way to mitigate the risks of data protection gaps in dynamic environments. But given such issues as the high costs and business disruption, organizations are reluctant to allocate more resources to periodic testing, especially as in most cases the money for DR testing comes from the same budget as that for DPR products. As a result, the number of unidentified and unfixed data protection gaps may increase substantially and leave critical applications and data unprotected.

In light of this situation there is an increasing need for solutions to automate DR monitoring and testing processes, as well as to reduce attached costs and resource consumption. In addition to optimization of DR operations, however, there is a need to make sure that configuration changes in the production environment are successfully applied to the secondary data center, that no configuration gaps exist between the two sites, and that protection levels meet SLA requirements.

To accomplish that, DR management should provide automatic, continuous discovery of configuration gaps, as an alternative to manually detecting and fixing gaps. In addition to the reduction of operational disruption and the potential cost savings, a key benefit of the automatic approach is the ability to maintain a consistent view of the DR environment, which cannot be achieved with manual periodic tests. With automation capabilities, on the other hand, configuration gaps are detected immediately and alerts can be provided in real time.

In addition, DR management should take a business-centric view to enable organizations to make sure that an entire business activity or process is properly protected and is compliant with data protection policies and legislation.

CONTINUITY SOFTWARE DR MANAGEMENT SOLUTIONS

Founded in 2005, Continuity Software is a provider of DR management solutions. The company's RecoverGuard is an agent-less discovery and monitoring solution that automatically detects data protection gaps in the DR infrastructure by continuously scanning the enterprise's storage, replications, servers, and databases. Thus, it allows organizations to optimize their DR operations, ensure that critical data and systems are protected, and avoid corporate policy violations and SLA breaches.

RecoverGuard operates by gathering information from the DR environment and providing a topology of DR dependencies. It then automatically detects and analyzes gaps and unprotected areas in the production environment based on a vast knowledgebase of DR gap signatures. After analyzing and checking the current DR status of critical applications versus corporate requirements, the product provides recommendations on specific remediation actions that can be taken based on best practices and recovery objectives, and according to the level of risk posed by each of the discovered gaps. This process is documented for auditing and compliance purposes.

RecoverGuard consists of the following components:

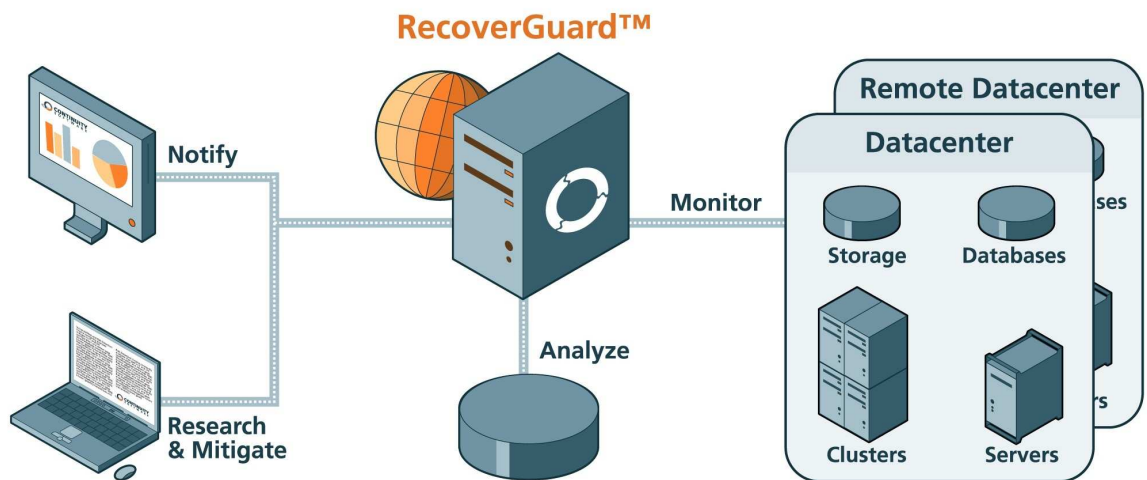
- IT Discovery and Scanning Engine: Performs a periodic scan of the IT infrastructure to collect configuration data from the storage infrastructure,

servers, clusters and databases, using protocols such as WMI, SSH, and Telnet, as well as proprietary vendor APIs. RecoverGuard also plugs into Storage Resource Management products to obtain information on storage systems and their replication environment.

- ☒ Data Dependency Engine: Identifies the dependencies between servers, clusters, storage devices, their databases, and their replications in the production environment.
- ☒ Gap Analysis Engine: Analyzes the data map compiled by the Data Dependency Engine to automatically detect different types of gaps (e.g., data completeness, data inconsistency, SLA breaches, data accessibility problems, incorrect deployment or maintenance actions, data tampering, wastes, general faults, and customer specific problems) and make suggestions for their resolution. This way, under-protected and over-protected data areas can be identified, as well as under-performing components (in terms of performance or availability) in the DR environment.
- ☒ Report Generator – automatically produces detailed reports on detected configurations and gaps, wastes, and general DR audit information.

FIGURE 3

Continuity Software RecoverGuard



Monitor	Analyze	Notify	Research & Mitigate
Constant Agent less data collection and events detection of IT Configuration (WMI, SSH, Telnet, CLI, Device API)	Vulnerabilities and Risk Detection Engine constantly analyzes changes in IT and DR configuration and detect gaps (Knowledgebase covers over 1000 known gaps)	Threats Dashboard, SLA, Data Protection & Availability views, Remote Notification System	IT Topology Visualized Tickets: • Threat handling recommendations • Incident tracking Optimization opportunities Waste findings Audit IT changes Reports & Queries

Source: Continuity Software, 2007

Based on these capabilities, RecoverGuard can be used to evaluate and mitigate risks and infrastructure vulnerabilities that may compromise the protection level of IT assets. In addition, Continuity Software's Data Protection Risk Assessment solution provides periodic reports on data protection risks and SLA violations, as well as recommendations on how to optimize the performance of the DR environment. In this process, RecoverGuard can also detect other optimization opportunities in the IT infrastructure, such as unallocated and unused storage elements, misconfigured replications that are abusing bandwidth, unused old database copies that consume storage resources, and so on.

RecoverGuard can also be used to improve DR planning, testing, and processes by allowing organizations to focus their DR resources on the applications and infrastructure areas that have changed. In addition, it can be used for regulatory compliance purposes by verifying that configuration changes do not affect the organization's DR readiness and providing alerts on violations. On the same note, RecoverGuard can be used to address section 404 of the Sarbanes-Oxley Act, which requires enterprises to map the systems that support their controls and financial reporting processes to their financial statements.

CHALLENGES/OPPORTUNITIES

While DR is expected to remain a top priority for organizations of all sizes, traditional DR approaches must meet the increasing challenge of providing visibility into the true status of DR environments on an ongoing basis. The more this environment grows in complexity, the more challenging this task becomes. As a result, organizations are looking for new solutions to better manage their DR implementations.

As an alternative to the manual approach of testing DR environments, Continuity Software is providing a DR management solution that offers important capabilities such as the automatic detection of known, unknown, and suspected vulnerabilities, and the ability to mitigate risks immediately by identifying root-cause problems and their impact, while reducing the manual labor required for performing these tasks. The RecoverGuard product can also be used to address the high costs associated with DR operations by allowing a cost-effective budget allocation based on the identification of specific business services that do not meet their service levels, and pinpointing which resources are required to achieve the required level of protection.

To capitalize on the market opportunity in the DR management space, Continuity Software should continue developing interfaces to third-party storage platforms, databases, operating systems and other infrastructure elements in enterprise data centers, as well as proprietary, homegrown systems. Integration with major system management vendors is also key to addressing the requirements of large enterprises. In addition, the company should seek to broaden its database of known protection gap signatures, which it claims currently contains more than 1,000 signatures. Continuity Software should also seek to establish partnerships with other technology vendors with complementing capabilities. For example, the company could benefit from partnering with DPR and DPRM vendors that can add the ability to fix gaps as they are discovered on top of RecoverGuard's gap detection capabilities. Another important channel could be DR hosting service providers, which the company is currently working to partner with.

CONCLUSION

Despite the growing awareness of the need to protect corporate critical business processes and data against loss due to disasters, and the increased spending on different DPR products, organizations are still facing difficulties in making sure that their DR objectives are met. To a large extent, this situation is due to the lack of means to monitor the health of DR environments on an ongoing basis, and the reliance on periodic manual testing to find and resolve data protection vulnerabilities. IDC believes that given the dynamic nature of enterprise data centers, in which production configuration changes occur frequently, an automated approach for identifying data protection gaps is becoming increasingly necessary.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2007 IDC. Reproduction without written permission is completely forbidden.