

Continuity Software's High Availability and Disaster Recovery Monitoring and Management in the Field

March 2009



The enterprise routinely replicates data between primary and secondary environments, which theoretically provides strong levels of protection should the primary host or site go down. However, over time IT frequently makes changes in the primary environment without making identical changes to the secondary and/or localized high availability (HA) solution (i.e., clusters). This process can play havoc with replication systems, leading to a dizzying 75% failure rate within a year of a company's most recent disaster recovery (DR) test and mitigation.

Why not simply test more often to expose and fix gaps in the DR and HA systems? Since DR testing is expensive and disruptive, IT is understandably reluctant to take on a regular testing schedule. And even when companies do test, repairing gaps between complex systems and dependencies is a difficult undertaking. Protecting the critical DR and HA environment requires complex change management and comprehensive testing, but all too many corporations have failed to invest properly in these operations. This unfortunate state of affairs exposes their business-critical data to corruption and loss.

Continuity Software is answering this challenge by automating DR and HA testing and analysis across multi-vendor environments. Taneja Group classifies Continuity Software as disaster recovery management (DRM). DRM is an up-and-coming technology class that tackles the challenge of managing primary and secondary systems in complex DR/HA infrastructures. This Solution Profile will briefly discuss the DR and HA challenges facing corporations today, how Continuity Software enables IT to repair and optimize DR pathways and HA clusters, and how it works in the real world of enterprise datacenters.

The Challenge of Managing for HA and DR

The difficult challenge of manual change management in complex environments leads to serious problems. Different storage devices, different applications, and different replication software: changes made to any one of these in a replicated environment must be reproduced at the secondary cluster or remote site. The more complicated the

environment, the more of an issue this becomes.

For example, a corporate data center replicates its most critical data to a secondary site 100 miles away and also uses local clusters for failover. A critical SQL database stores its data on an HDS array and replicates to the secondary site using Hitachi Replication Manager. Back at the data center, Veritas Cluster Server replicates a high value

S O L U T I O N P R O F I L E

Oracle database to a secondary host within the cluster. Both primary and secondary systems represent a large investment of capital and ongoing expense and management. They must meet critical recovery time objectives (RTO) and recovery point objectives (RPO) with prompt failover and accurate data restoration.

Unfortunately, when the primary systems go down due to a city-wide power loss, IT contacts the hot site only to discover that the replicated system has lost recent transactions and that the data may be corrupted. There is no failover, and restoration will likely take days.

Once the power comes back on, a newly suspicious IT staff tests DR settings on its clusters. They find that although failover would have occurred, the secondary host was sluggish and could not come close to meeting service levels. When IT investigates the failure of its replicated systems, they find that a series of incremental changes made to the primary systems was never duplicated to the secondary ones. Over time, the replicated systems diverged to the point that RTO and RPO have become essentially meaningless. And now that the gaps have widened to the point of multiple failures, IT does not know where to start to mitigate the problem.

This is not an isolated example. Many large data centers have both hot sites and multiple clusters set up for transparent failover. In environment like this, a dozen or more clusters are replicating data in the event of a primary failure. But if IT has made configuration changes on the active parts of the cluster, they need to make the same

changes to the passive systems. Otherwise even minor inconsistencies will multiply over time, leading to data loss or corruption, dead in the water secondary servers, or extremely poor performance following a failover event.

The average cluster suffers unexpected downtime on the average of 8 hours a year. If failover is occurring as planned, this average may be acceptable. But all too often failover does not happen as it should – not because the secondary systems are failing, but because of serious configuration drift between systems that should be identical.

- **Challenge #1: Testing is dangerous, disruptive, and expensive.** HA/DR testing requires downing systems, which makes it impossible in the first place for 24x7 systems. When it can be done, it takes staff time and resources. The danger is present from the beginning, with the risk of being unable to bring the production systems back up again. The result is that companies might run HA/DR tests when they are deploying new software and equipment (emphasis on “might”), but frequently will not test after that. But when changes to the primary system are forgotten at the secondary level, the lack of testing cripples the change management process and risks the HA/DR environment.
- **Challenge #2: Configuration drift.** IT commonly makes ongoing minor changes to the primary systems, but often neglects to make the same changes to the secondary systems. Over time, divergence grows to the point that the entire failover process is threatened. The more mission-

S O L U T I O N P R O F I L E

critical the data, the worse the threat grows.

Challenge #3: Closing the gaps. IT simply cannot test all business services. Nor can IT test each business service end-to-end, such as simulating true load/operation on DR. And even when IT is aware that the gaps exist, it is very difficult to address them manually. The systems must match across a variety of concerns, including complete replicated data sets, service levels, accessibility of secondary data copies, and identical data deletions. And when DR testing uncovers serious issues, this complex environment makes it difficult to mitigate all incompatible configurations between the two systems.

Disaster Recovery Management and Continuity Software

This is where DRM comes in: by automating testing and change management, DRM can reliably and cost-effectively mitigate mismatches between primary and secondary replicated environments. DRM is not the same thing as data-protection management (DPM) or storage resource management (SRM). Instead it automates a high level of risk mitigation in replicated DR and HA cluster environments; a complex setting that is under-served by DPM and SRM as well as manual change management and test operations.



S O L U T I O N P R O F I L E

There are several vendors that practice DRM in their own replicated systems, but only Continuity Software applies DRM to multi-vendor replication systems from EMC, NetApp and HDS. This makes Continuity Software's RecoverGuard software a critical DRM tool for the real world of multi-vendor HA/DR environments in the corporate data center.

At a high level, RecoverGuard software continually scans the enterprise HA and DR infrastructure including storage, database, servers and replication configurations. When it discovers configuration gaps or vulnerabilities, it issues immediate, detailed alerts – allowing IT administrators to immediately remedy the situation. RecoverGuard is agentless and expands to heterogeneous DR environments across clusters and remote locations.

How does it accomplish this? RecoverGuard continually monitors characteristics such as failed dependencies, inconsistent data, incomplete data sets, and breached service levels. It scans the IT environment to collect critical configuration data from key IT assets, including storage, servers and databases. RecoverGuard then performs a comprehensive dependency analysis in order to build a detailed topology map, which serves as the foundation for future analysis. Based upon this topology, it can then identify and report on all gaps and vulnerabilities in the customer's environment (local and remote). Should the customer desire, they can also drill-down further in order to better understand the status of it SLA compliance and SLA exceptions, produce change and audit reports, review trend analysis of

configuration changes, further examine the risks detected in the infrastructure, as well as understand Business Services recoverability standing.

RecoverGuard protects multiple replication operations by sensing gaps and vulnerabilities throughout the end-to-end protection process. It automatically collects information from the IT infrastructure and scans for issues that will impact recoverability. Without this level of infrastructure discovery, inconsistent replication can lead to long failover times and even unrecoverable data, forcing IT to recover from backup and losing hours to days of changes to the production data. With large data centers commonly running as many as seven replication products, this enables IT to keep the replication tree consistent.

By identifying and mitigating critical gaps between hosts and their related resources, RecoverGuard solves the problem of complex change management in the HA/DR infrastructure. Businesses can set meaningful RPO and RTO objectives and confidently expect to meet them.

RecoverGuard does not add to the complexity of the HA/DR environment, but rather simplifies it by centralizing information and management across the entire replicated infrastructure. This makes the HA/DR process far more transparent and manageable, especially since RecoverGuard leverages existing configuration management databases. These RecoverGuard knowledge bases contain thousands of vulnerability signatures ranked by severity. Continuity

S O L U T I O N P R O F I L E

dynamically updates the knowledge base and automatically distributes it to customer sites.

In the Field: Customer Scenarios

Scenario #1: Financial institution

This financial firm invested heavily in HA/DR configurations and testing, including synchronous SRDF for its EMC DMX, and TrueCopy for its HDS disk arrays. In an effort to improve dependencies between the primary site and a warm secondary site, the company invited Continuity Software to demonstrate its gap finding and mitigation service. Over the course of 48 hours, RecoverGuard uncovered nearly two dozen serious dependency gaps that would have crippled the firm's RTO and RPO. These included an unprotected partition on an active transactional database, and failing SRDF replication on a critical production dataset. In both cases there was no major procedural or product fault, but rather a series of seemingly inconsequential changes to the primary arrays that had never been made on the secondary side.

The level of exposed risk shocked the company, but with Continuity's detailed analysis, the firm was able to quickly and efficiently mitigate those gaps. RecoverGuard now runs automatically at scheduled intervals to identify and help close any subsequent gaps created by changes to either environment.

Scenario #2: Mobile communications company

With a culture of strict procedures and careful oversight, a mobile phone company

was convinced that their replication environment was rock solid. They reluctantly agreed to test RecoverGuard on PB-scale databases running on EMC Symmetrix DMX systems in three different locations. RecoverGuard located multiple gaps and inconsistencies, several of them immediately serious and others that threatened to become so. The cell phone company was able to quickly prioritize and mitigate the problems. They have since deployed RecoverGuard across their replication infrastructure, and have automated daily testing. The ongoing analysis pinpoints existing and potential problems, enabling the company to effectively counter any configuration gaps between their primary and secondary systems.

Scenario #3: Utilities company

A large utilities company employed several sophisticated HA/DR data centers across three states, making for a complex web of interdependencies. The infrastructure sports multiple platforms and databases, as well as critical commercial and in-house applications. The company very much wanted to submit its HA/DR environment to accurate testing, but disruption from downtime was not an option. In lieu of regular testing, the best thing the company could do was to observe best practices. Then it was introduced to Continuity. RecoverGuard runs on the HA/DR environment and every night analyzes the replication systems without downing services or systems. If it detects issues, it immediately delivers an actionable alert to IT. The company also chose to use Continuity's disaster assurance service, which sends alerts

S O L U T I O N P R O F I L E

and analysis to Continuity Software's HA/DR Specialists Team as well. Continuity Software also enables the utilities company to leverage RecoverGuard's topology map to detect under- and over-utilized assets.

Continuity Software Benefits

As we wrote earlier, the enterprise HA/DR environment is replete with challenges. Chief among them are the challenges of regular testing, configuration drift between theoretically identical systems, and gap mitigation between said systems. Continuity Software addresses these challenges even in multi-vendor HA/DR environments.

- **Comprehensive HA/DR testing without downtime.** RecoverGuard regularly and non-disruptively detects vulnerabilities and threats in complex HA/DR environments. If primary system configurations diverge from secondary configurations – as they frequently do -- RecoverGuard will detect it. The software runs at scheduled times and delivers alerts and analysis for immediate and prospective issues. Since the process pulls from the dynamically updated gap knowledge base, information and remediation is immediately available.
- **Mitigate configuration gaps.** Even when an administrator becomes aware of a configuration mismatch, it can be hard to repair in complex environments. Once RecoverGuard detects gaps, it issues detailed alerts and tickets to aid administrators in fixing the problem quickly. Administrators can always refer to up-to-date displays of RPO and RTO

readiness, and can take advantage of Continuity Software's unique Disaster Recovery Assurance (DR Assurance) plan for professional mitigation services.

- **Optimizing the IT Infrastructure.** Optimization requires more than closing configuration gaps. RecoverGuard also detects less than optimal infrastructure settings and underutilized storage. Companies can use the RecoverGuard topology map to optimize their DR resources. Continuity Software's DR team is available to help with sophisticated troubleshooting and optimization throughout the DR infrastructure.

Taneja Group

Disaster recovery has been a hot topic for more than two decades, and companies continue to spend huge amounts of money on high availability clusters and remote secondary sites. But a major challenge is ensuring that failovers and data restores are going to happen within critical timeframes. Regular DR tests are highly disruptive and costly, and even when IT discovers a mismatch between primary and secondary systems; it can be difficult to repair. The reality is that even when IT manually repairs a gap the system will not stay repaired. Within hours, the slightest change might not replicate properly and the whole crippled DR scenario begins again.

That is why disaster recovery management (DRM) as a technology class is growing fast. System and storage vendors offer their own product-specific testing mechanisms, but only Continuity Software offers DRM and

S O L U T I O N P R O F I L E

optimization for EMC, NetApp, and HDS environments. We find that Continuity Software's RecoverGuard software and DR Assurance services are increasingly valuable options for companies who cannot afford to

miss their DR, data protection, and business continuity objectives, and who want to bullet-proof and optimize their failover and replicated environments.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.