Buyer Case Study

Taking a Proactive Approach to IT Infrastructure Resiliency

Dan Yachin
Sponsored by: Continuity Software

IDC OPINION

A multibillion-dollar global financial services company was facing challenges in meeting its IT service availability objectives after making substantial changes to its IT services resiliency program. The rapid growth of the company's IT and storage infrastructure, the continuous implementation of new technologies, and the outsourcing of key IT operations to a third party had resulted in greater exposure to risk, as well as increased difficulty in effectively planning and conducting audits and disaster recovery (DR) tests. As a result, a decision was made to adopt a more proactive approach to infrastructure resiliency management. As part of this change, new service availability targets were established for the entire business, namely a 24-hour recovery time objective (RTO) and a 30-minute recovery point objective (RPO).

After experiencing difficulties in meeting the new service availability objectives, the company identified cross-domain configuration drifts as a major cause of the risks and problems that occurred during production and annual DR tests. To tackle this issue and improve its ability to monitor and control the datacenter environment, the company chose Continuity Software's AvailabilityGuard solution to analyze and audit infrastructure configurations. Continuity Software helped the company achieve the following benefits:

- Identify configuration risks that were previously overlooked or not anticipated
- Significantly shorten the time spent on identifying and fixing availability and performance issues
- Exceed strict RPO and RTO objectives
- Reduce the time to provision and deliver new workloads while improving reliability and stability
- Enforce and audit the use of best practices by employees and outsourcing partners
- Cope with employee retention issues by preserving organizational knowledge
- Reduce OPEX by proactively identifying risks

IN THIS BUYER CASE STUDY

This IDC Case Study discusses the use of Continuity Software's AvailabilityGuard by a global insurance company to address infrastructure resiliency and high availability risks. It analyzes the company's challenges in meeting its service availability objectives and how AvailabilityGuard helped tackle these challenges by enabling a proactive approach to identifying and resolving problems.

SITUATION OVERVIEW

Organization Overview

Continuity Software's customer is a large US-based global financial services provider with more than 10,000 employees and over $500 billion in total assets under management. The company relies on an extensive IT infrastructure supporting multiple mission-critical applications and services hosted in two primary datacenters with approximately 2,000 physical servers and three 10Gb network links connecting the two locations. Around 75% of the open systems workloads (Windows and Linux) are virtualized, using approximately 20,000 VMs.

The company has a large storage infrastructure handling approximately 12PB of raw capacity across the two sites, distributed between Tier 1 and Tier 2 storage systems. A fully redundant full-mesh storage area network (SAN) topology is used in each datacenter. Network-attached storage (NAS) is used for storing data from some Tier 1 and Tier 2 applications, as well as for general purpose file shares.

For data protection and recovery, the company uses a wide range of solutions deployed across both storage tiers, including replication software for moving data between the two datacenters (with a default RPO of 30 seconds), and for providing a consistent recovery point across different applications – both across open systems applications and between open systems and mainframe applications. Some critical applications and workloads run in active-active environments between the two datacenters. All Windows, Linux, and VMware production environments utilize load balancing and path failover. Dedicated object-based storage arrays are used for storing compliance archive data, which is replicated between the two sites.

Within each datacenter, the company has deployed various high availability and clustering solutions. In terms of backup, the company is tapeless in both datacenters and uses purpose-built backup appliances (PBBA) as targets for backup data, as well as for data coming from different backup applications. Backups are not intended to be part of the primary data source in a DR scenario, but all data is replicated in case it is needed, either during or after the DR process.

For business continuity and disaster recovery (BCDR) purposes, the company uses dedicated DR infrastructure for some of its Tier 1 applications. For most environments, when a DR event is initiated, the model environments are shut down or paused, and the production environments from the other site are then brought online using the replicated data residing on the local arrays. To assess the readiness of the DR infrastructure, the company performs a full site DR test failover and failback each year.

Challenges and Solution

According to the company, a few years ago, it experienced a major IT outage – the biggest in the organization's history – when an entire storage array went down for three days. The primary cause of the outage was a hardware/software bug in the array. At the same time, failure to implement industry and vendor best practices substantially complicated and prolonged the event. The impact of the outage was substantial in terms of business interruption, resulting in data loss as well as financial loss.

Following the outage, and due to its increasing dependence on the IT infrastructure to deliver critical services to a large number of customers, three years ago, the company established new challenging targets of a 24-hour RTO and an RPO of 30 minutes for the entire business. These objectives are based on a full site failover. A small number of applications, based on an active-active architecture, can failover in a more granular manner – an initiative the company is currently in the process of extending. There are currently no quantified objectives for resiliency and high

availability within each datacenter, but critical systems are all expected to provide recovery within seconds to minutes.

In addition to identifying gaps in complying with the new service-level agreements (SLAs), the IT team was facing several other challenges. Being highly focused on driving standards across its IT environments, the company has been making significant efforts to test and audit its IT infrastructure. However, ongoing and annual tests are extremely time-consuming to plan, prepare, and run. Running effective tests has become an even more challenging task, as the company is constantly modernizing its IT environments with new, rapidly evolving technologies, implementing strategic initiatives such as virtualization, consolidation, active-active storage architecture, private and hybrid clouds, and other new technologies, which introduce further risk and exposure.

In this dynamic environment, communication and orchestration of changes among multiple stakeholders has become acutely important in order to keep track of the rapid changes. In addition, the company is currently in the midst of outsourcing its IT operations, implementations, and helpdesk to a third party. All this in turn has driven the need to develop best practices for outsourcers to work with and to audit and verify that the outsourcing company is maintaining these standards and meeting the company's expected SLAs.

In light of these challenges, the company decided to adopt a more proactive approach to service resiliency management, including through the implementation of new technologies that could help achieve the abovementioned SLAs. Another key driver for product selection was the need to ensure that best practices are implemented for local high availability within each datacenter. To address these needs and others, the company was looking for a solution that could improve its ability to constantly monitor the entire IT environment and assess its operational readiness.

The company already has a large number of tools to analyze various aspects of the IT infrastructure. However, these tools are generally isolated to their specific domain (e.g., network), and very few of them are focused on reducing overall IT risk. According to the company, "We could see performance-related issues or capacity-related issues, but nothing that would focus on best practices or configuration risk detection — especially across multiple technology domains or silos." Furthermore, the company noticed from production problems and, during the annual DR tests, that many of the identified risks and complications were a result of cross-domain configuration issues. Tackling this problem was therefore a key consideration in selecting the appropriate solution.

Results

Designed specifically for the purpose of detecting availability risks and preventing outages, Continuity Software's AvailabilityGuard was selected by the company to continuously validate infrastructure resiliency and DR readiness. According to the company, it could not find "any other tools that are dedicated to analyzing and auditing infrastructure configurations."

In essence, AvailabilityGuard is an IT operations analytics solution for preventing infrastructure outages. It operates by performing daily cross-domain scans of infrastructure configurations, covering servers, storage arrays, databases, and the virtualization layer and calculating the dependencies and relationships between them. The configuration data is collected from the pre-production, production, and DR environments in a nonintrusive manner using read-only commands with no agents installed on systems. AvailabilityGuard analyzes the data against a knowledgebase of more than 6,000 known gaps and vulnerabilities to detect configuration drifts and deviations from best practices and vendor recommendations. When potential data protection, availability, or DR risks are detected, a ticket containing the details regarding the severity of the problem and its business impact is automatically issued to the appropriate team, along with a suggested resolution.

AvailabilityGuard was implemented across the company's two sites, covering the entire host and storage-related hardware and software infrastructure. It is currently set to run weekly, looking for risks or drifts in configurations in the production/DR environment.

By using AvailabilityGuard, the company was able to "reduce the time required for daily firefighting and fixing risky issues, especially ones that we never thought to look for." When AvailabilityGuard was first enabled across the entire production environment, it identified an extensive number of risks that had accumulated over several years. Once the initial backlog was addressed, the company notes that "the various technology teams started to routinely depend on AvailabilityGuard to identify things that had been overlooked or not anticipated."

By using AvailabilityGuard, the company was also able to exceed the defined SLAs during a company-wide DR test for the first time. It comments that, "Since the company is divided into three business units, corporate IT's ability to demonstrate to the business that, for the first time ever, it could achieve the stringent business-defined SLAs for IT resiliency had a tremendous impact on its reputation." The company also reports definite OPEX and revenue savings that it has observed from using AvailabilityGuard to proactively identify risks, as well as improved corporate reputation.

Oher benefits the company has gained from using AvailabilityGuard include:

- Significantly shorter time spent on identifying problems and provisioning and delivering new workloads that are more reliable and stable, all in a timely manner
- Enforcing and auditing best practices (As outlined above, this has become a critical issue, as the company now depends on outsourcing to operate its entire IT environment)
- Maintaining stability and organizational knowledge as the company undergoes employee churn following the move to outsourcing

ESSENTIAL GUIDANCE

When it comes to problems facing datacenter operators today, downtime tops the list. IDC's research is consistently showing that downtime due to system failure is the top issue in the datacenter, as reported by enterprises and colocation providers alike. Other top concerns include downtime due to human error, latency and performance issues, delays in IT deployments due to power or space constraints, and downtime due to natural disasters.

Avoiding downtime is more critical than ever before, as organizations have become dependent on their IT infrastructure to support critical, revenue-generating business operations. However, modern IT environments are growing increasingly dynamic and complex, with multiple moving parts that must be synchronized, which increases the potential for downtime events. In addition to reducing the number of opportunities for human involvement in the datacenter to lower the likelihood of downtime due to human error, organizations should also identify existing infrastructure that needs updating to avoid downtime due to system failure.

This case study shows that preventing or minimizing the impact of downtime — whether due to human error or system failure — requires a proactive approach to infrastructure resiliency. Such an approach should be largely based on the ability to continuously monitor the entire IT environment, identify potential risks, and assess its operational readiness, while enforcing the use of best practices. These capabilities, as this case study demonstrates, are essential in order to identify and address issues before they escalate and meet even the strictest RPO and RTO objectives.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com
