Summary Insights: Delivering End-to-End Disaster Recovery

RFG Perspective: In the hybrid cloud world, end-to-end Disaster Recovery (DR) will have to be re-invented to cope with the accumulated technology changes. It's no longer a matter of setting RTO and RPO policies in a closed network of large enterprise data centers – and sticking to them. What constitutes a disaster has changed – and the ability to recover from these damaging forces needs to be part of the DR plan.

Disasters can be caused by conventional causes – natural catastrophies, power outages, or network or system outages – or, as is common in the New Normal environment, by malware and other cybersecurity attacks.

The shift to the various cloud manifestations, containerization, DevOps, and the end-users' "always on" demands has increased the disaster recovery (DR) challenges that organizations face. In addition, the COVID-19 pandemic has changed the business model for many organizations, increasing downtime risks.

To preserve high levels of uptime and consumer confidence – and to preserve service level agreements (SLAs) – business and IT executives must make sure that they have thoroughly tested communications, disaster recovery and business continuity (DR/BC) plans, as well as comprehensive testing procedures.

To achieve minimal downtime, DR/BC procedures must be automated wherever possible and tested multiple times per year. If these plans are not tested, executives should not expect their DR/BC plans to restore normal operations quickly, as planned. Executives must make sure that their organization's DR plans and procedures are well-documented, widely understood – and practiced across the entire organization.

INTRODUCTION

The well-established disaster-recovery (DR) procedures of traditional enterprise data centers must be extended to support the organization's new hybrid multi-cloud environment. Today's computing environment has undergone substantial change in terms of what can cause a catastrophe and where the incident may occur.

In the world of data-center-centric computing and highly distributed, multi-site cloud computing, the damage may occur across multiple sites and regions. In fact, it may not be easy to determine the exact location of the outage – even before the work begins to repair it. This means that IT executives and business unit executives have to "Think Different" – the mantra of the late Steve Jobs of Apple Computer. Visualizing what types of problems could occur, at some time in the future, will go a long way toward ensuring business continuity in the hybrid cloud.

It's no longer a matter of setting RTO and RPO policies in a closed network of large enterprise data centers – and sticking to them. In the hybrid cloud world, end-to-

end Disaster Recovery (DR) will have to be re-invented to cope with the accumulated technology changes.

RFG 100 Panelists and Insights from Participating RFG 100 Executives

On Mar. 24, 2021, Robert Frances Group facilitated an RFG 100 video conference entitled *"Delivering DR Automation, End-to-End, Including DevSecOps Governance Architecture in an Agile Context."*

The panelists on the call were:

- Shawn Butler, Vice President, Architecture and Analytics, BNS
- Don Molaro, Chief Technology Advisor at Worldwide Technology.
- Brook Schoenfield, Master Security Architect, IOActive Inc.
- Evan Bauer, Founder and CEO, OpStack Inc.

POLL RESULTS

RFG 100 executives were polled about the challenges of documenting, testing and proving a DR/BC plan in a hybrid cloud world. We found the following results in the polling data:

- 60% of the respondents stated they test quarterly; the percentage grows to 70% when adding in the semi-annual and annual testing.
- 47% of respondents perform DR for their data centers and private cloud, while 27% of respondents report that they run DR for all sites
- Surprisingly, a number of those polled said they regularly failed to meet their stated organization's RPO and RTO requirements
- 20% of respondents claimed they store all data off-site and in the cloud
- 7% have fully automated DR testing

DR in the Age of End-to-End, Multi-Site Computing

Enterprises need to think holistically about DR. In fact, there would be value in the organization going back and defining disaster all over again to make sure that it fully covers all the failure categories now being encountered. Then, enterprises would need to develop a single, consistent communications plan for DR/BC that covers all data center, cloud, and edge sites and all potential disaster types one can expect in the 21st century.

That means that the "mindset" of modernizing applications needs to connect with the "mindset" of more traditional operations. Security failure testing, DR/BC planning and data-recovery practice sessions are all essential to business continuity in a New Normal world.

Complicating matters is that recent years have seen a significant change in application development and testing practices for traditional and cloud-native systems. This bifurcation of App Dev and DevOps practices is making end-to-end DR/BC solutions more difficult to "harmonize" – and work together smoothly – than ever before.

In the COVID-19 pandemic, DevOps is a great accelerator for modernizing applications and migrating them to the hybrid cloud, but the need to protect corporate data has never been higher. Continuous security scanning and penetration testing, DR practice sessions, and repetition of data-recovery procedures are essential to business continuity in today's digital economy.

In the midst of this computing revolution, service disruption has more causes than ever: power outages, network outages, natural disasters, service interruptions from CSP providers and SaaS providers – and the threat of cybersecurity breaches.

Making the situation even worse is that some SaaS providers are no longer providing end-to-end backup and recovery services, preferring to work with partners to provide back-up and restore capabilities. This shift can result in the third-party RPO and RTO procedures failing to meet corporate objectives.

Due to the rapid rate of change, organizations should increase the frequency of testing exercises within the organization. The passage of time has led to incomplete, and possibly inaccurate, documentation and playbooks, which can result in a weakened ability to quickly recover following an outage. Scheduled testing should not only address failover and recovery, but it should also address a planned fallback to a previous, known, good configuration, which has proven to be a tougher challenge for most enterprises.

Data Protection in the Hybrid Cloud

Protecting data remains a top-priority for end-to-end disaster and recovery strategy. IT organizations must replicate and back up critical data in anticipation of outages. They must practice the art of data recovery from all data sources, whether they're on-prem or off-prem data resources – including cloud service provider (CSP) infrastructure.

The traditional datacenter hosts enterprise workloads (e.g., ERP, SAP, Oracle and most "crown jewel" applications) that are anchored within the data center. In today's hybrid cloud world, most new applications are containerized and managed by Kubernetes or other orchestration software and located somewhere in the cloud or in an edge location.

DevOps is a great accelerator for modern containerized applications, but unfortunately, the developers are working on quick agile-development cycles and are pressuring others to meet their schedules. This can cause a lack of data-protection hygiene, which can be seen in terms of poor practices regarding data architecture, data governance, data management, and data quality.

One more obstacle also exists: the abstraction being built into the software stack, allowing developers' code to run anywhere in a cloud. This code is subject to deployment issues, which can create a sudden crisis due to software flaws or, occasionally, to intentional interference.

Ideally, a crisis would lead to intensive consultation with the teams focused on data, governance, operations, and security. However, that frequently does not happen. Why? In many places, there are distinctly separate corporate cultures for traditional data centers, which have these types of teams in place, and born-on-the-web, cloud-native solutions, many of which are still operating with immature processes.

Where Do We Go From Here?

What can be done to raise awareness of these issues across IT organizations and business units? The first step is a change in mindset.

A little paranoia is a good thing, as former Intel CEO Andy Grove once remarked. Everyone must start with the assumption that if something has not been tested, then it might not work as planned. With a "testing" mindset, employees and managers must prepare for many types of disruptions – and for the need to recover mission-critical workloads when disruptions occur. That is why organizations are exploring partnering with companies to accelerate adoption of digital transformation – and to find ways to "scale up" their testing of data for business continuity (BC) purposes.

Testing must be emphasized, to ensure business continuity, no matter what causes data services to be interrupted. The acid test for adequate DR/BC preparation is this: "Pull the plug" and find out what does – and doesn't – work when finding and restoring data services. Then, fix the problem and consider increasing the frequency of testing until executives are satisfied that the tested disruption is considered adequately addressed – and that it is not remedied with a short-term patch.

In the New Normal environment, the multi-cloud world of microservices and containers is being combined with the on-premises world of scalable enterprise applications and databases. Historically, those two worlds – of cloud-native workloads and datacenter workloads – have operated with different sets of expectations regarding availability and resiliency. Therefore, a full plan for business resiliency will likely involve the company's cloud services and its SaaS resources, including applications built with open-source code and externally controlled APIs.

What You Can Do – Right Now.

- **Identify your organization's critical revenue streams.** The code, data, and infrastructure for these streams must be protected, above all else, in your end-to-end DR/BC planning.

- **Continuously scan your applications code for security gaps and vulnerabilities.** Good, known copies of applications code cause fewer security problems than code in repositories that are not well-managed. This applies to code that was developed or acquired from a third-party. Scanning containers before they are checked out of a repository will prevent replication of serious security flaws, which could lead to lengthy outages and downtime.

- **Wrap IaC (Infrastructure as Code) around aging applications.** This will ensure consistency of code, and a uniform approach to restarting systems, resulting in for faster re-starts on recovery.

- **Keep tabs on your organization's CI/CD software life-cycle management.** Agile organizations use CI/CD software development to reduce lengthy application development cycles. Monitoring CI/CD closely often leads to more efficient and effective operations deployments.

- **Test your back-up and recovery procedures more frequently.** Most enterprises have back-up and recovery procedures that have been well-documented over the years. But "pulling the plug" can reveal weaknesses in your documented DR/BC plan – and it can reveal security gaps. One best practice is to failover to the backup site, remain there for a week, and then fallback to the original site. That approach not only proves the procedure works, but it also provides the operations groups with consistent and constant DR training.

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; *Contact: inquiry@rfgonline.com*

- **Review your access management practices.** In the age of work-from-home (WFH), software keys and tokens that travel home with an organization's end-users might lead to security breaches later on.

- **Communicate effectively:** Develop and clearly enumerate your expectations about DR/BC and SLAs to all of the key owners of DevOps and deployment processes and gain buy-in. Then, communicate your expectations about availability – and your priorities about enforcing them – throughout the organization. That will help the DR people in the field to become comfortable with the DR/BC plans and satisfied that they will work optimally under the challenging circumstances of a real-life disaster. Having good communications within the organization also means that the DR planning team is prepared to accept all the objections that may be thrown at them by one or more business units. Finding a way to arm them with prepared answers to frequently asked questions (FAQs) will create more "buy-in" for DR/BC practices across the enterprise.

RFG POV: Executives must ensure their organization's planned disaster recovery (DR) procedures address their end-to-end hybrid cloud. The adequacy of existing RPO and RTO requirements must be re-evaluated – and executives must ensure that their DR plan satisfies the requirements of the New Normal computing world.

Today's IT infrastructure has Core, Cloud and Edge components – and each must be tested separately – and then tested together with the other applications that span the data center and the Cloud. All of this takes time, but it will be time well-spent if the organization can return to its work in case of natural disasters, outages (e.g., networking, power outages) – and downtime due to security gaps and cyberattacks.

Deep inspection of all systems supporting enterprise workloads – and all connections to cloud service providers (CSPs) and managed service providers (MSPs) – will help to strengthen the resilience of the infrastructure. Importantly, one cannot project that all of the CSP connections will perform seamlessly under the pressure of a major DR incident.

Consistent testing will reveal whether an organization's DR/BC plan is effective – or not. Once the various components are locked down, and all stakeholders are satisfied that their concerns are addressed, executives should schedule repeated testing, and leverage automation software wherever possible. This will ensure that recovery times are minimized and that human error is reduced as much as possible from the equation.

Additional relevant research and consulting services are available. Interested readers should contact Client Services to arrange further discussion or interview with Mr. Cal Braunstein, CEO and Executive Director of Research. Jean S. Bozman, President of Cloud Architects Advisors LLC, co-authored this report.

---

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; *Contact: inquiry@rfgonline.com*