FIVE COMMON MISTAKES IN DISASTER RECOVERY AND CONTINUITY OF OPERATIONS PLANNING

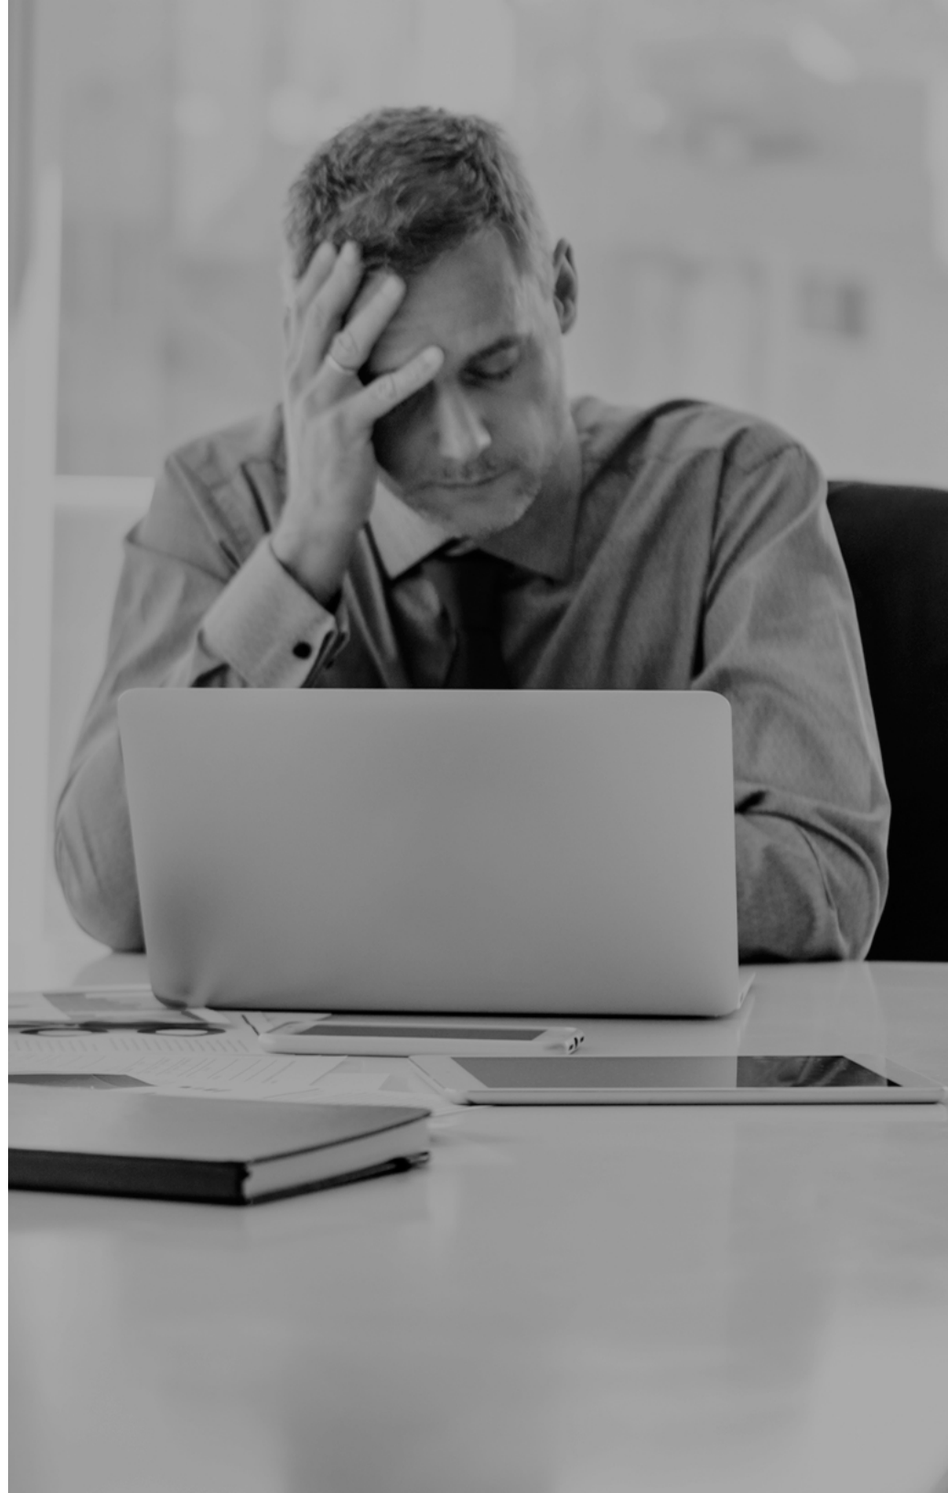

TABLE OF CONTENTS

INTRODUCTION

Rackspace has spent nearly two decades supporting federal agencies, and we've rarely come across two topics as widely confused as disaster recovery (DR) planning and the creation of a continuity of operations plan (COOP). Whether it's simple confusion about the difference between recovery point objective (RPO) and recovery time objective (RTO), or the game of hot potato involving the responsibilities of IT and mission owners, very few have a complete grasp of the differences between DR and COOP. In this white paper, we'll take a look at five of the most common mistakes made during DR and COOP planning and recommend techniques for avoiding them.

MISTAKE 1: IGNORING INEVITABLE FAILURES

The first and most common mistake is simply denying that there's a need for DR and COOP planning and, instead, attempting to design an architecture that "never fails." DR and COOP deniers seek to construct an impenetrable fortress, but there is no such thing. Instead of an impenetrable fortress, deniers end up with an architecture that cannot tolerate, or gracefully respond to, any failure.

As Amazon CTO Werner Vogels puts it, "Everything fails, all the time." You must plan for this inevitability. Ignoring it generally results in a system that more closely resembles a house of cards where one minor failure can bring the whole thing crashing down. Strive to build a system that automatically heals itself when minor problems do arise and can be easily reconstituted from nothing in the event of a catastrophic failure.

Count on failure and build plans for managing it. Automation can reduce the quantity of manual tasks needed to manage cloud workloads, and it plays a central role in planning for failure. Employ automation wherever possible and consider the implications of failure during each step of the design process. Automated management of common IT failures in cloud computing has completely changed how smaller system failures are handled, and many disaster scenarios can simply be avoided.

However, automation isn't always possible, and in many cases, especially when it comes to DR and COOP, a manual process may be preferable. You don't want a computer to make a binary decision to initiate a two-hour failover process in response to a failed package installation that could be backed out manually in five minutes.

In the age of artificial intelligence (AI) and machine learning (ML), sometimes a little human intellect can go a long way. Documented processes and decision matrices are invaluable for making tough decisions like activating a DR plan. If you don't have a formal DR plan and strategy, at least make sure you have the following documented, in multiple locations, for all of your IT systems:

1. A process to confirm that a disaster event has occurred
2. A means to categorize which failure mode you're recovering from or which DR scenario you're activating
3. A defined timeline to make a go/no-go decision
4. A specified and empowered decision-maker and lines of succession in that role
5. A list of key personnel and alternates
6. A communications plan

MISTAKE 2: CONFLATING DISASTER RECOVERY AND CONTINUITY OF OPERATIONS

The second biggest DR/COOP mistake is conflating these plans and assuming that if you have one, you also have the other. That couldn't be further from the truth. COOP is a U.S. government requirement focused on maintaining delivery of all essential aspects of government despite disruptive events. DR is focused on IT and comprises the policies, tools and procedures enabling recovery or continuation of IT systems following a disaster scenario. DR can be a subset of COOP, but COOP is much more broadly applicable to people, processes and functions beyond technology.

The National Institute of Standards and Technology (NIST) recommends that federal directives distinguish between COOP and DR plans to avoid confustion[2]. NIST also states that COOP focuses on restoring mission-essential functions (MEF) at an alternate site and performing those functions for up to 30 days, whereas a DR plan is designed to restore operability of one or more information systems at an alternate site after a major disruption.

As the Federal Emergency Management Agency (FEMA) explains, National Security Presidential Directive 51 (NSPD-51)/Homeland Security Presidential Directive 20 (HSPD-20) currently defines COOP as the "effort within individual executive departments and agencies to ensure that essential functions continue to be performed during a wide range of emergencies, including localized acts of nature, accidents and technological or attack-related emergencies."[3]

COOP is how your entire organization operates during a major disruption. DR is how you reconstitute an IT system after a major disruption. Each IT system likely has its own DR plan accounting for multiple disaster scenarios, while your entire organization may only have a single COOP.

MISTAKE 3: CONFUSING HIGH AVAILABILITY, FAULT TOLERANCE AND DISASTER RECOVERY

There are three closely related and commonly confused acronyms in IT: high availability (HA), fault tolerance (FT) and disaster recovery (DR). Too often, people attempt to "rank" these options, e.g., HA > FT > DR. They think, "HA is the most important, and I already have FT, so I don't need DR. Right?" This is the third most common mistake.

DR plans are created to address major disasters, not service interruptions. FT and HA address localized losses of individual system functionality. DR is about widespread losses of functionality, potentially even external to your HA or FT system. Your may be thinking, "External? Then how is that my problem?"

Consider what happens when your organization's core router goes down. In most well-architected networks, the core router is configured as part of an FT or HA pair. But when they both fail, your entire network goes down. So how would that happen? It could be caused by something that impacts your entire headquarters building such as a fire, flood or an overzealous backhoe operator. This is a

legitimate disaster, nothing that FT or HA design could really account for. You might be thinking: "Well, I'm good — my application runs in the cloud!" Sure, your application continues chugging along, but nobody in your office, or anywhere on your corporate network, can get to the cloud.

In this scenario, you would probably need to activate your COOP processes to enable alternate work locations until your HQ was restored to normal operations. You would also want to activate your DR plan for the core network to restore access to IT systems, as well as your DR plans for any other systems residing at HQ.

MISTAKE 4: FAILING TO ACCOUNT FOR ALL FAILURE MODES

Planning for failure means considering all potential failure modes. A failure of the imagination is the fourth-most common mistake. Don't put all of your focus on protecting against one specific type of failure while the door is left wide open for countless others. When it comes to DR planning, it's critical to consider scenarios that are obscure and unlikely — but entirely plausible. Some disruptive events happen more frequently than you expect. Think through outages of an entire cloud region and/or service, cloud credential compromises, insider threats, data corruption (both malevolent and accidental) and encryption snafus.

Cloud technology provides many native mechanisms that easily add redundancy by distributing components across multiple availability zones and geographic regions. If one data center fails, your system remains available in another. If a disaster affects an entire region, a fully functional replica of the application will continue operating in a geographically distant region. However, what happens if an application upgrade in your primary location overwrites the entire master database, line by line? Ironically, your HA design works flawlessly! Each of these transactions is fully replicated in milliseconds to each of your alternate locations. Unfortunately, because of the volume of these transactions, your transaction logs have been looped several times. You now have a perfectly synchronized but completely useless system in your primary and alternate locations.

Faults and failures usually result in data loss or loss of availability. The extent of the damage depends on the type of failure and your organization's mission, but failures can be detrimental on many fronts, and failures can take many forms. To truly understand how each potential failure will affect your organization, you must thoroughly analyze the direct and indirect costs, which accrue every second an application is down.

Decreased productivity occurs when employees can't access applications needed for work or abandon their normal work to fix the problem, and lost data can be expensive or impossible to replace. Failures often significantly damage an organization's reputation, especially when customer data is lost. Planning for all plausible failure modes will help you protect your application data and avoid the costly damage of system failures.

MISTAKE 5: COMMINGLING DISASTER RECOVERY WITH PRODUCTION

The final mistake we see regularly, even among cloud-savvy customers, is commingling DR and production environments. "Commingling" means running DR right alongside production and attempting to separate the environments by using pre-cloud concepts like network isolation. Unfortunately, using multiple network segments for environment segregation and isolation does not actually provide tangible separation, especially in the cloud. True isolation is only guaranteed at the cloud account level. Network isolation often isn't a real isolation at all, and it confers almost no benefit when you're attempting to isolate entire environments, which by definition must be able to communicate. Instead, it offers a false sense of security, makes management and operations more difficult, and has the potential to increase the blast radius of a single failure.

The ideal approach would be isolating DR environments in separate cloud accounts, potentially in separate regions, and maybe, for the most extreme protections, with an entirely separate cloud service provider (CSP). This provides several governance and security benefits by guaranteeing permissions enforcement as well as offering increased protections from multiple failure modes in each environment. In the cloud, as in any other IT environment, managing permissions and credentials is extremely important. Creating an

WITHOUT A DISASTER RECOVERY PLAN, YOU'RE LEFT TO WING IT:

Where are the backups? Who maintains the backups? Do we even have backups? Looks like we do have backups, but the last successful backup was over a week ago. (Guess we lost everything for that new customer we launched on Monday!) Where should we restore the backups, primary or secondary? Can we even perform a restore on a replicated database? Uh-oh, this backup was prior to the schema upgrade for the latest version. Now we have to stand up a third environment and install the previous version … but we don't even have the fresh install procedures for that version. We have to install two versions back, then upgrade, then restore. Wait — is this table encrypted?! Where's the private key?!

account for each environment makes it easy to explicitly guarantee that assigned permissions are restricted to a specific environment by default. Segmenting by account also enables management to have more visibility into the costs of each environment, making it easier to optimize spend on IT resources by clearly delineating DR and production.

A commonly overlooked problem with using a single cloud account for production and DR environments is cloud account compromise. Cloud credentials get accidentally compromised all the time. Maybe you've done it. You threw some credentials in a config file to test something. It worked, so you push your code to a public source code repo and celebrate. You quickly forget about the credentials sitting in the config file, now shared with the world. It just takes one bad actor to find those credentials before you remember, and your cloud account is compromised. This happens every day.

Now let's say you've done everything else right. You've built an HA production environment that is auto-healing in the event of minor failures. You have a solid backup strategy. You have a DR plan that you've not only tested numerous times but actually activated once before during a regional service outage. Everything worked marvelously and your system remained online with no appreciable data loss or downtime. High-fives all around.

However, that bad actor just used your credentials to destroy your entire production system. Your security team identified the breach within an hour and worked with the CSP to disable the offending credentials. You perform a damage assessment and discover that there is literally nothing left of the production environment. Frustrating? Yes, but there is a road back, starting with the activation of the DR plan for catastrophic failure.

"Error, not found?" Unfortunately, you made just one simple misstep in your planning: Your DR environment runs in the same cloud account as your production environment — the same production account whose credentials you checked into that public repo. The bad actor wiped out your entire DR configuration and all backups as well.

This scenario actually happened. In 2014, a fairly successful company called Code Spaces was forced out of business in less than 24 hours when its cloud account was compromised.[4] Its entire production environment, DR environment and all its backups were deleted. The company was born in the cloud, and its only footprint was digital, but almost every asset it had was gone. Code Spaces had staked its brand identity on data protection and recovery capabilities:

"[Code Spaces offers] full redundancy, duplicated and distributed among data centers on three continents. Backing up data is one thing, but it is meaningless without a recovery plan … and one that is well-practiced and proven to work time and time again."

Code Spaces did almost everything right.

CONCLUSION

COOP and DR plans remain vitally important, even in the cloud. They are not the same, and you probably need both. Cloud makes it easy to do things like HA, but that can't replace DR. When defining your DR plans, make sure you think about all possible failures. Then go the extra mile and separate your production resources from DR. At Rackspace, we take a security-first approach, working with each of our government and commercial customers to institute defense in depth and develop a comprehensive disaster recovery plan.

Learn more about Rackspace DR and COOP solutions.

ABOUT THE AUTHOR

Brad Schulteis is currently Director of Government Solutions at Rackspace. He has over 15 years of enterprise IT experience across the public and private sectors. He previously worked at AWS Worldwide Public Sector, supporting some of its largest U.S. government customers. Prior to that, he worked at the Department of Defense (DoD), managing all aspects of its private cloud computing platform. His DoD work earned him a Value Engineering Award, Special Act Award and Joint Meritorious Unit Award. Concurrently, he participated in many government IT transformation working groups and study committees. He has a proven track record of working with complex requirements and driving change in large IT enterprises.

https://www.linkedin.com/in/brad-schulteis

REFERENCES

1. Congressional Research Service, "Executive Branch Continuity of Operations: An Overview," R. Eric Peterson, Feb. 2, 2005, http://congressionalresearch.com/RL31857/document.php
2. National Institute of Standards and Technology, "Contingency Planning Guide for Federal Information Systems," Marianne Swanson et al., May 2010, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-34r1.pdf
3. Federal Emergency Management Agency, "Continuity of Operations: An Overview," https://www.fema.gov/pdf/about/org/ncp/coop_brochure.pdf
4. Threatpost, "Hacker Puts Hosting Service Code Spaces Out of Business," Michael Mimoso, June 18, 2014, https://threatpost.com/hacker-puts-hosting-service-code-spaces-out-of-business/106761/

ABOUT RACKSPACE

Rackspace is modernizing IT in today's multi-cloud world. By delivering IT as a service, we help government institutions realize the power of digital transformation without the complexity and expense of managing it on their own. Our comprehensive portfolio of managed services across applications, data, security and infrastructure on the world's leading public and private cloud platforms enables us to provide unbiased expertise. Rackspace has been honored by Fortune, Forbes, Glassdoor and others as one of the best places to work.

Learn more at www.rackspace.com or call us at **1-800-961-2888**.

SEC-CWP-Federal_Govt_5_Common_DR_COOP_Mistakes-11285

July 9, 2018