computing research

RESEARCH PAPER

When Plan B Goes Wrong: Avoiding the Pitfalls of DRaaS

June 2021

Sponsored by

iland Zerto

CONTENTS

Introduction

We all know that, from time to time, failures happen in data centers. Much talk of disaster recovery focuses on genuinely rare scenarios – terror attacks or freak weather events. The reality is rather more mundane – and surprisingly frequent according to the contributors to this exclusive research. Employees accidentally delete data. Floods and fires occur. People drill through power and network cables in nearby streets. Cooling systems in data centers fail and systems overheat. Hardware flakes out. Ransomware attacks land and encrypted data needs to be restored. An IT SysAdmin powers down a production database instead of the back-up. Mundane, workaday errors, that happen to enterprises public and private across the United States.

Despite the leaps in progress we've seen from AI, human beings are still very much involved in keeping the lights on, and the possibilities for human error are genuinely endless.

Fig. 1 : Looking at the last time your organization had to recover data and systems, which of the following most closely represents the cause?

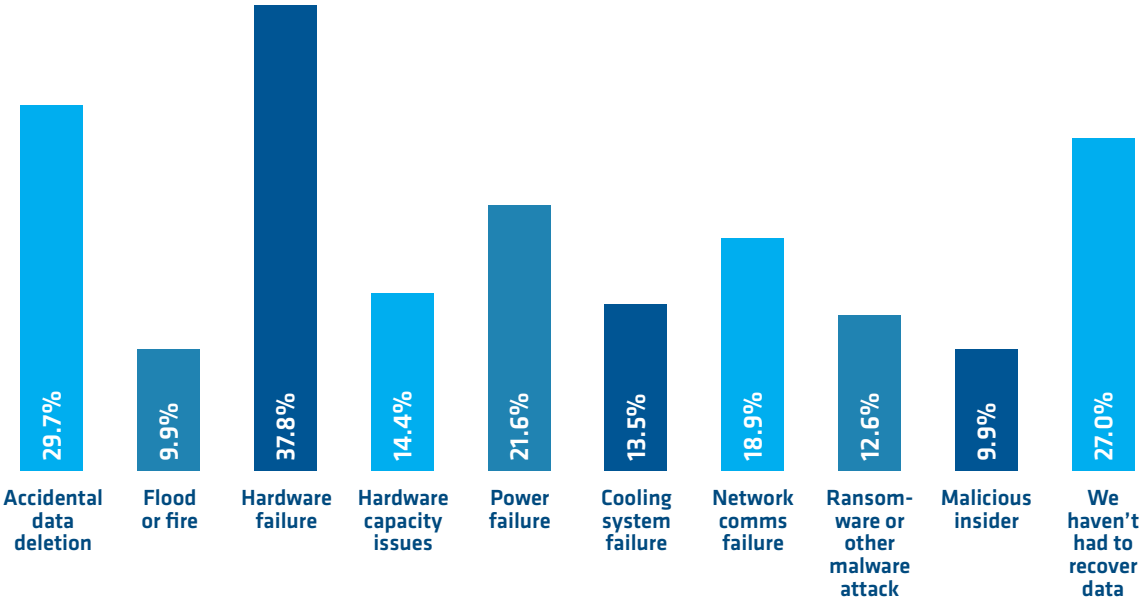

When did the outage occur?

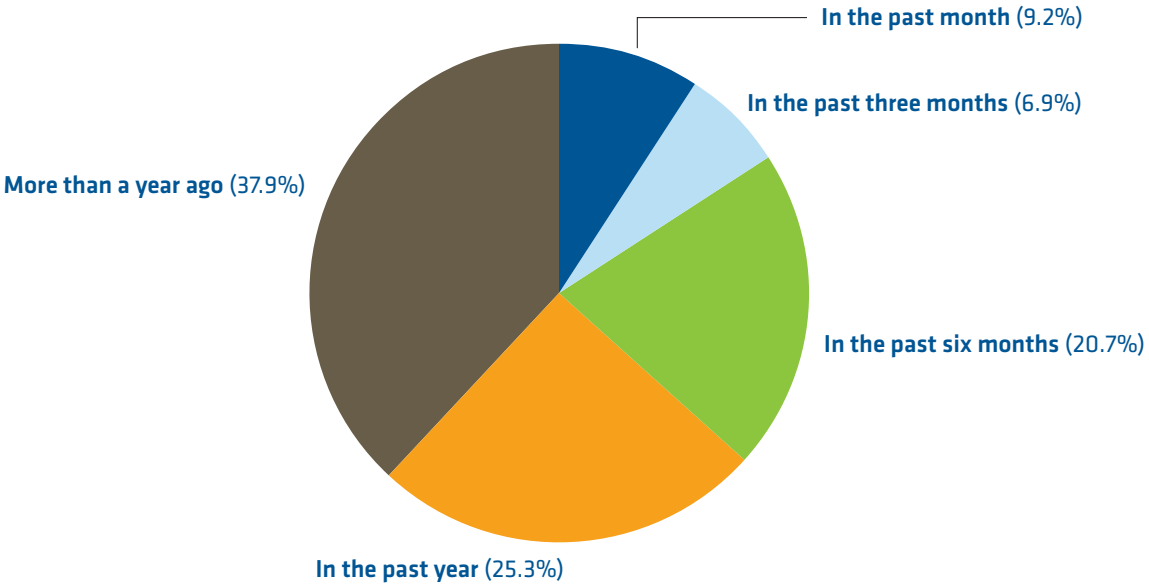

The diagrams above illustrate that whilst a fortunate 27 percent have never had to do a restore, the remaining 73 percent of contributors had experienced failure at some point. Almost two-thirds of that group had experienced the outage within the last 12 months and more than half of those within 6 months.

Enterprises know that they need to expect the unexpected and have a Plan B on standby in order that they can continue business operations in the event of a technical issue. Traditional DR has focused on replicating compute, storage and data systems in secondary data centers. Data is fed from the production site to the recovery site at a second data center which is (or ought to be) an exact match in configuration. Ideally this is in real time. One of the more obvious drawbacks of creating a second data center with all the hardware, compute, networking and connectivity it necessitates, is the sheer expense.

This is why Disaster Recovery as a Service (DRaaS) is such an attractive option for businesses. Running failover/replication with cloud service providers or hyperscale environments removes the need for a second data center and therefore can generate some significant cost savings. But are all DRaaS providers equal?

Computing surveyed 150 technical and business decision makers from organizations drawn from a wide cross-section of US enterprises, each employing a minimum of 500 people, to investigate the real-world experiences of enterprises when disaster strikes. Some 38 percent of contributors to our research were IT director level or above, and 45 percent at the manager level. All contributors were directly involved in disaster recovery planning and implementation.

The objectives of our research were to establish what DR systems were in place, how often plans are tested, and whether enterprises are confident in their ability to recover from disaster as easily and swiftly as possible. Is recovery smooth enough to ensure business continuity? Or are enterprises likely to find themselves mired in a situation where systems stay stubbornly down, configuration problems complicate things and unforeseen costs arise?

Key Findings Include:

- Some 54 percent of contributing organizations had a documented, company-wide DR plan in place.
- Failures and outages are more common than many people realize. Some 73 percent of contributors had experienced failure at some point. Almost two-thirds of that group had experienced the outage within the last 12 months and almost half of those within 6 months.
- Some 57 percent retained a second, separate on-premise data center for DR purposes.
- DR testing frequency is very low. Just 50 percent are testing only annually or at less frequent intervals, while 7 percent didn't test their DR at all.
- Of the organizations testing less frequently, the results of their last test led 50 percent them to believe that their DR may be inadequate, while 12 percent encountered issues that would have led to sustained downtime.
- Those testing more regularly are far more likely to have more ambitious objectives for data recovery than those who test less frequently. Some 32 percent expected to recover data in seconds or minutes and 23 percent actually did so. Some 47 percent expected recovery within hours and 43 percent achieved it.
- The most frequently encountered problem with failback to original sites is network problems, experienced by 62 percent.

Extent of Cloud DR

Our first port of call was to check whether contributors had documented DR plans in place and where on-premise applications and data were replicated. One slightly surprising finding was that only 54 percent of organizations contributing to this research had a documented company-wide plan in place. A further 35 percent had a plan in place for their specific part of the business and a very optimistic 11 percent had no documented plan in place at all.

Fig. 2 : Where are your on-premise data and applications replicated for disaster recovery?

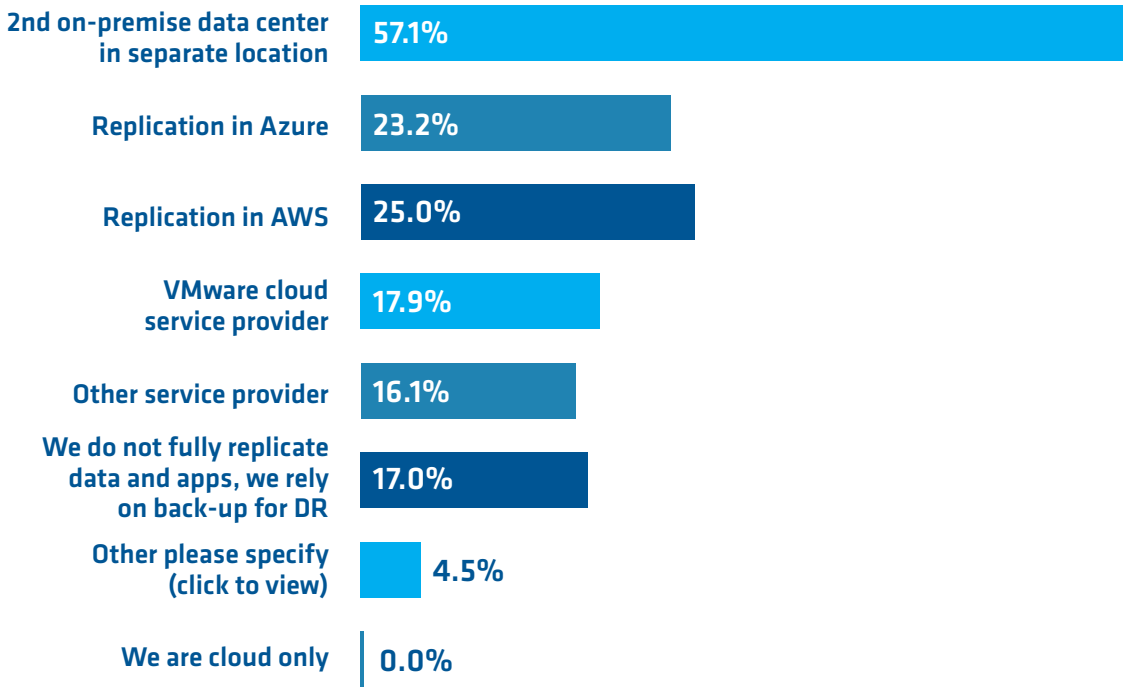

The illustration above indicates that even some organizations with documented plans are not practicing full replication of apps and data. Some 17% of contributors rely solely on back up for DR. However, the remaining contributors are using a range of methods, including replicating their data in cloud service providers like VMware, AWS, or Azure. The majority, almost 60% of these organizations, still have a second on-premise data center in a separate location.

One of the more surprising aspects of these findings is the comparatively small number of respondents who were relying on the cloud for their replication. This suggests that the move to the cloud – and DRaaS – has been slow among the larger businesses which were the focus of our research. The reason for this may be reasonably straightforward – long multiple year contracts were common in DR up until relatively recently so it may be that for many of our contributors these contracts are only just coming up to the point of renewal.

Lingering concerns about security, compliance and also cost control and visibility may also have contributed to a slowness to embrace cloud DR. The extent to which concerns like these can be overcome depends greatly on the cloud providers involved. Not all providers are equal – a subject we will be discussing in greater depth later on.

However, despite the persistence of the secondary data center, there are counter forces which are driving cloud DR. Technologies such as Zerto and VMware for example, which are proven to work effectively in the cloud, are providing tangible benefits such as real-time replication to increasing numbers of businesses. Technology leaders value the opinions of other technology leaders when making long term strategy decisions, and the maturing of these technologies is likely to be reflected in the process.

The pandemic is also likely to be accelerating the transition to DRaaS even as this paper is being written. The lockdowns have vastly increased the importance of remote access generally – which is one of cloud's biggest selling points.

Test, Test, Test

The level of confidence that enterprises have in their Plan B is determined largely by the extent and frequency with which they test it. Many organizations have SLAs to their own customers for services or contracts to deliver a certain quality of service. A tried and tested Plan B is crucial if they are to be sure that those SLAs and contractual obligations can be met. Given how dynamic – and complex – infrastructure typically is in larger organizations, the frequency of DR testing needs to mirror the pace of change.

However, in a majority of organizations, it doesn't. Some 50 percent are testing only annually or even less so, while 7 percent didn't test at all.

Given the frequency and likelihood of outages, the infrequency – or absence – of testing among our contributors is an eyebrow-raising finding, particularly given the continual message that comes through from much of our research on cybersecurity, that the strategy and solution focus is very much on remediation of breaches rather than prevention. If your strategy for undoing the damage from a ransomware attack is to recover copies of locked data, regularly testing the way in which you plan to do so is probably a good idea – particularly as more clever examples of ransomware target back-ups. Recent history is packed with an abundance of cautionary tales of those who failed to do so.

Fig. 3 : Do you test your disaster recovery plan or run drills to check what would happen in the event of an outage requiring failover?

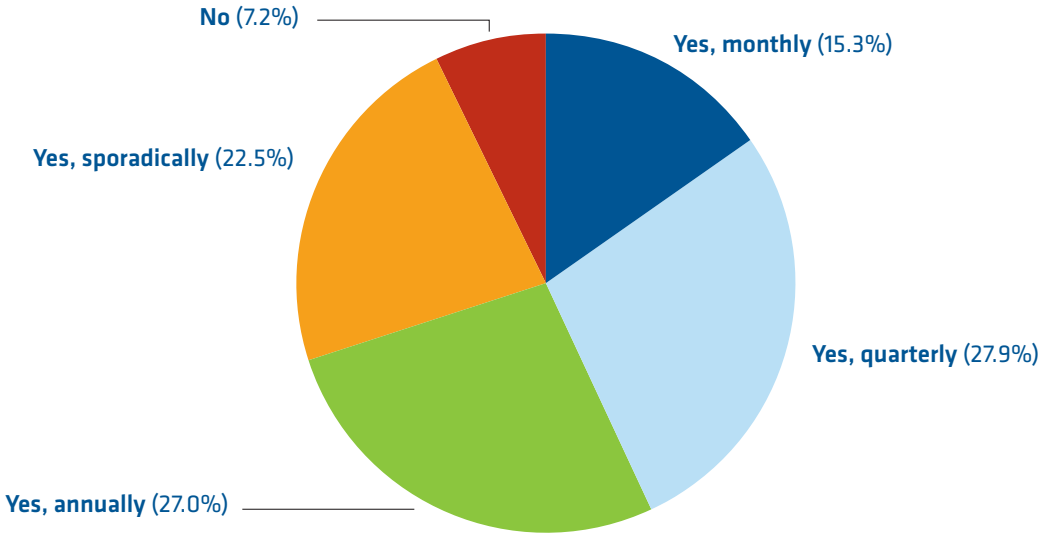

Few would argue that regular testing is sensible, so why are fewer organizations than you might think actually doing it? One explanation is fear of what the tests will show up. There are probably a substantial number of IT professionals who suspect strongly that their DR isn't fit for purpose, but they haven't got the means necessary to replace it with one that is. It's a time-limited approach to risk management, but probably not as uncommon as it should be.

Time is important when it comes to testing. Almost every organization is subject to budgetary and resource constraints. Technical skills shortages are endemic and much of this resource is expended on day-to-day production and application availability. DR testing simply gets pushed down the never ending to-do list. Also, some DR approaches, be they secondary data centers or DRaaS can be quite tricky to test. If you have to bring production systems down to test or schedule this out of hours there is extra cost involved which compounds the problem of low prioritization.

There may be some nuance contained within these responses that isn't immediately obvious. It is possible that business critical applications are being tested more regularly than less important ones. An intelligent and focused testing strategy like this requires business input in order to ensure that everyone agrees on the applications and data which are operationally critical and those which are less so. There are likely to be multiple levels of prioritization. From this process a DR test schedule which reflects the criticality of different applications and services can be built.

Results Day

Computing asked those who test less frequently what the outcome was of their most recent test – however long ago that happened to be. Several facts jump out from the diagrams below. The first is that the option for perfect, seamless recovery was ticked by - literally – nobody. The second is that 50 percent fear their DR may be inadequate. The third is that a further 12 percent encountered issues that would have led to sustained downtime – a scenario that if it occurred in the real world would be operationally damaging in the extreme. Whilst the sample size of respondents to this question was small, it remains clear that failure to test is quite likely to lead to failure to recover from a disaster when one strikes. It's also clear from the second diagram that when recovery does fail there are a wide range of issues contributing to that failure.

Fig. 4 : Thinking of your most recent test/drill, which of the following statements most closely resembles the outcome?

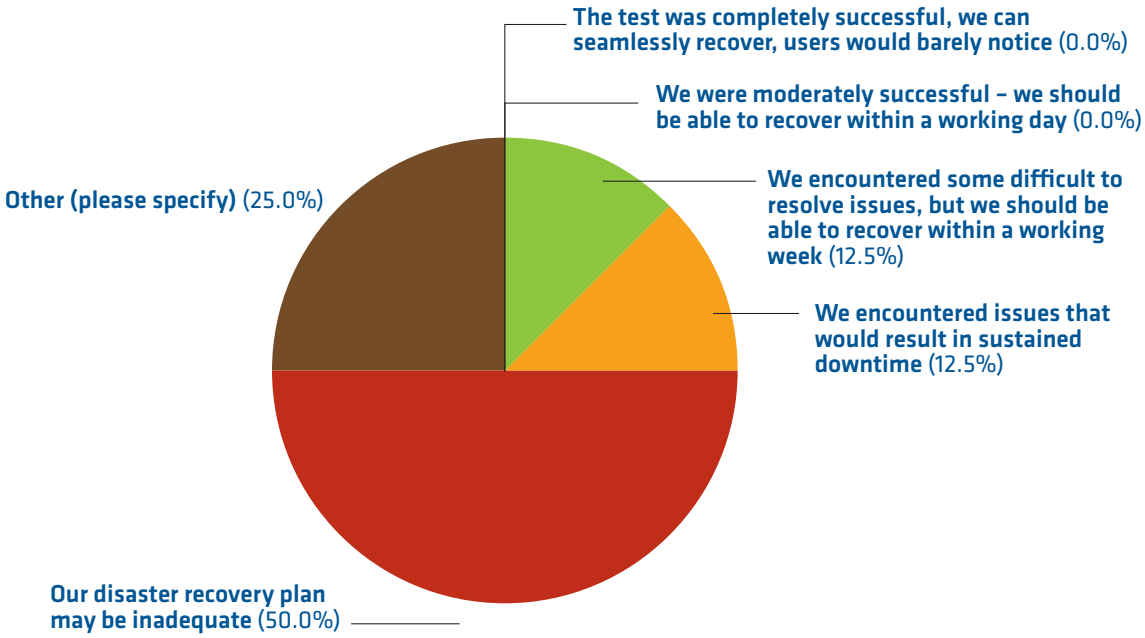

Fig. 5 : Which of the following issues did you experience?

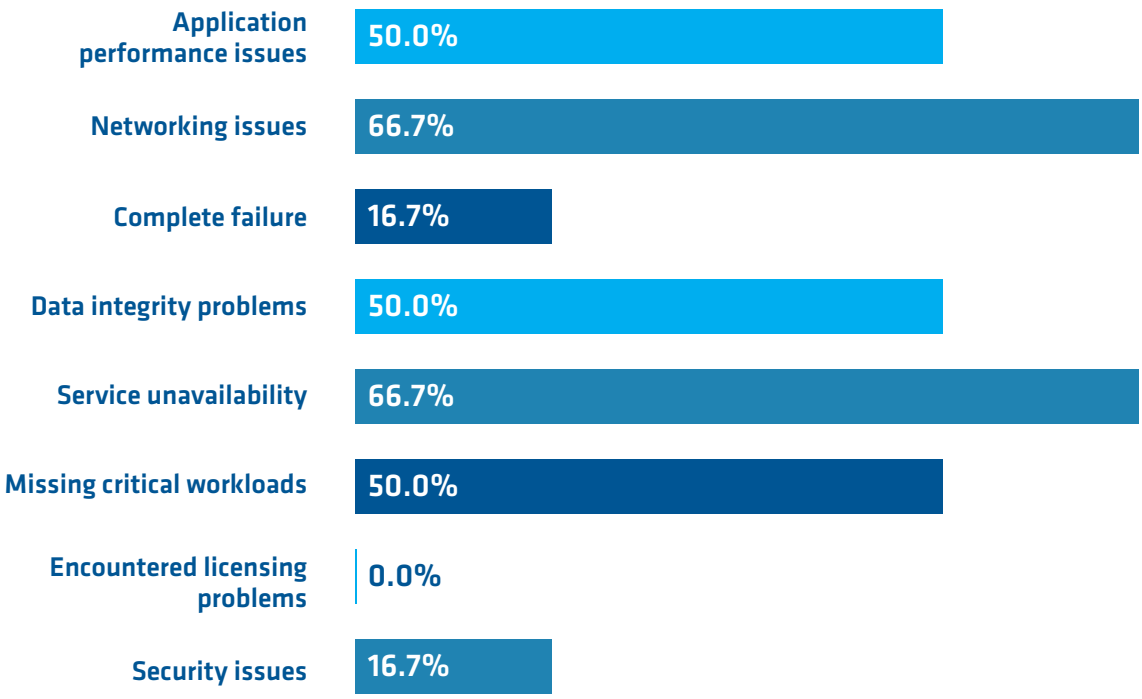

The final diagram below illustrates the responses of those who do test their DR more regularly. It is clear that those testing more regularly are far more likely to have more ambitious objectives for data recovery than those who test less frequently, with a majority expecting to recover data within hours at the most, and a majority of those doing just that.

Fig. 6 : What was your recovery time objective and did you meet it?

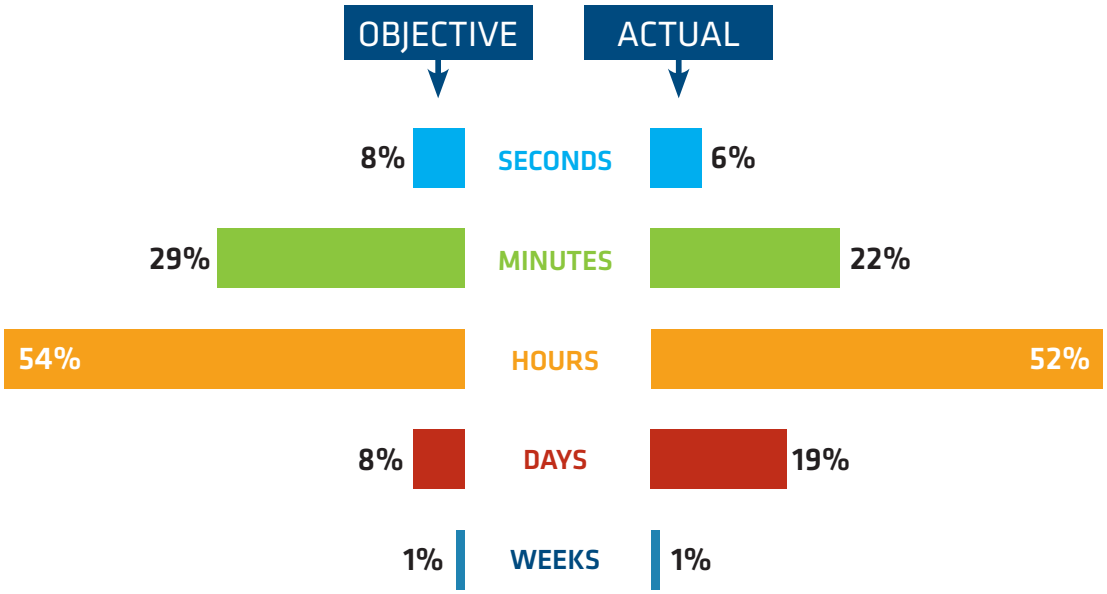

The Importance of Failing Back

Of course, failing over to a secondary location is only half the story. At some point, you have to failback – return the data and apps to their original location. The diagram below illustrates some of the issues that contributors experienced with failback.

Historically, many DR solutions simply didn't consider failback – and some of the issues experienced by our contributors might well reflect this. Often, significant periods of downtime would occur because DR systems would need to be powered down in order to be replicated back to the original site. DRaaS is easier to failback from because specialist DR software such as Zerto, which is purpose built for cloud environments, provides native failback capability that allows you to seamlessly failover and use the same process to fail back.

One of the key determiners of the success of failback is the degree of compatibility with the platform it failed over to. If an organization is running VMware locally, the failover location most likely to meet the objectives of a Plan B is a VMware based DRaaS provider. If you choose a different platform, whilst data can be fairly easily replicated, the failback is likely to prove tricky because of the conversion from one hypervisor type to another. This may well be why so many contributors experienced network issues. If network adapters pick up a hardware change or hypervisor change the relevant VMs are left without network access and any data or application on them will remain unavailable. A 100% compatible platform neatly swerves this problem as all components recognize each other.

Fig. 7 : Which, if any, of the following issues occurred when you 'failed back' returning from the recovery site to production?

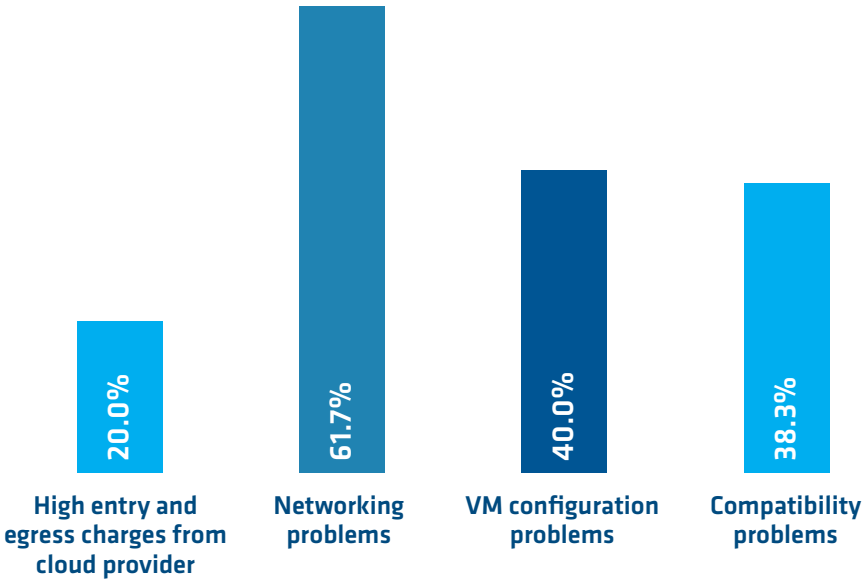

If the networking, configuration and compatibility issues can be resolved with these solutions what of the more commercial issues our contributors experienced? Data entry and egress charges have cropped up repeatedly in *Computing* research on cloud. Hyperscalers tend not to charge a great deal to push data into their clouds but make up the cost when organizations want to take data out. This may not figure too much in DR planning if budgets are based on disasters being huge but vanishingly unlikely. However, as we've seen, the reality is likely to be less spectacular but rather more frequent. This is likely to make recovery from these smaller scale events a lot more expensive than anticipated.

Conclusions – In DRaaS, One Size Does not Fit All

Our research suggests that DR is treated, in many organizations as a relatively low priority and that this is likely to be based on some misconceptions of what constitutes a disaster. When determining budgets, executives assess the likelihood of data centers being hit by a plane or falling into a sinkhole and determine, correctly, that the chances of these events actually occurring is very low indeed. However, our research has shown us that smaller scale disruptions generated by human error are regular events. Some 73 percent of those contributing to our research had had to restore data or services within the last year.

The three-quarters of contributors who replicate their data and applications in a secondary environment use a variety of methods. The majority still retain a secondary data center, but a significant proportion of respondents also replicated to hyperscale environments or VMware DRaaS providers, and this is a trend that is likely to gather pace as data center infrastructure ages, contracts expire, and the economic and societal impact of the pandemic continues to reverberate.

The optimum location of a secondary environment is likely to be unique to each organization and indeed, individual applications, workloads and data. What is arguably of greater importance than location is the frequency of testing. Any reader of this paper is likely to be hyper-aware of the dynamic and fragile nature of technical ecosystems – which makes the finding that half of organizations represented in this research are testing annually at best rather worrying.

The results of infrequent testing are predictable. Half of those with infrequent testing fear their DR may be inadequate. Failure to test DR will, at some point, lead to a recovery failure. It really is only a matter of time.

There are many reasons why DR testing doesn't happen as frequently as it needs to, but most relate back to the importance accorded to DR within the organization itself. Do executives truly understand the definition of a disaster? Like cyber security, the most successful DR strategies involve DR being considered at the inception of new applications and services rather than as an afterthought once they've gone live. Only if DR is integral to an application, can solutions be aligned with confidence, and a consensus that testing has to be as frequent as change can be established. Business input is also vital to ensure a consensus on what is business critical. DR test schedules should reflect the varying levels of importance of different applications, data and services.

The importance of testing is paramount but another finding of our research was that whether or not contributors were restoring data under test conditions or failing back after a real world restore, some of the issues encountered – such as networking issues – were shared. Many of these issues can be traced back to the DR site or service. Part of the process of integrating DR into an application or service is the alignment of platforms and outcomes. For example, for customers running VMware-based infrastructure, a VMware-based DRaaS provider is going to provide the best possible experience related to application performance, data integrity etc. The more closely matched configurations are, the fewer hitches are likely to occur.

It is no surprise that DR has the potential to be every bit as complex as the infrastructure it is protecting. Our research also demonstrated that a significant minority of businesses are looking for DRaaS solutions to help them manage this complexity – as well as reduce the costs inherent in running a second physical data center. This is leading them towards hyperscalers. However, DR in a hyperscale environment does not always bring the promised cost reductions.

There are a number of reasons for cost expectations not being met. Firstly, the "instances" that compute and storage are sold in by hyperscalers means that their customers end up having to bend their configuration to fit. This is a highly error-prone exercise given the transience inherent in a great deal of infrastructure. At best, businesses end up overpaying for instances they don't need. At worst, their failover doesn't work because they have underspecified their instances to keep costs down. Related to this are data egress costs. Data entry costs from hyperscalers are usually quite attractive. Data egress costs tends to be less so and failback from a disaster means that all the relevant data has to move in one go back to the original production site. The costs can be much higher than anticipated.

There is also the question of security – particularly as a feasible event to be recovering from is a ransomware attack. Hyperscalers do not integrate security and compliance into their service offerings and depending on the level of service that has been purchased, there may be a time lag between a disaster occurring and the point where you can activate the service.

The inflexibility inherent in the commercial packages offered by hyperscalers means that for many organizations they will not provide the most cost effective DRaaS platform. More flexible, specialist providers may well offer a complete match on the configuration of your on-premise

data center – which means that a lot of the technical issues with failover and failback that our contributors have experienced are less likely to occur. They are also likely to be far more competitively priced in the longer term because specifications can be completely tailored rather than picked from a menu of choices and, crucially, security and compliance support are already included rather than optional extras.

Enterprises are advised to give very serious thought to the status of DR in their organization. Once that is understood they are in a better position to assess their environments and decide which DR solutions are best aligned to business-specific applications. Finding the right platform for critical applications and services enables technical teams to meet expectations of business continuity without being tripped up by compatibility issues or hidden costs.

In DRaaS, one size does definitely not fit all.

About the sponsors, iLand and Zerto

iland is a global cloud service provider of secure and compliant hosting for infrastructure (IaaS), disaster recovery (DRaaS), and backup as a service (BaaS). Industry analysts recognize iland as a leader in disaster recovery. The award-winning iland Secure Cloud Console natively combines deep-layered security, predictive analytics, and compliance to deliver unmatched visibility and ease of management for all of iland's cloud services. With headquarters in Houston, Texas, London, UK, and Sydney, Australia, iland delivers cloud services from its cloud regions throughout North America, Europe, Australia, and Asia.

For more information:

Visit: www.iland.com