

2016
**Disaster Recovery as a Service
Attitudes & Adoption Report**



Scott D. Lowe
David M. Davis
James Green
Partners, ActualTech Media

Spring 2016

 **Infrascale™**

Table of Contents

Table of Contents	2
List of Figures.....	3
Executive Summary.....	4
Introduction	5
Outage Handling Experiences	6
Disaster Recovery Adoption by Vertical	10
Understanding Disaster Recovery as a Service	11
Current Disaster Recovery Capabilities	12
Organizational Technical Characteristics.....	15
Understanding Peer Disaster Recovery Capabilities.....	17
Disaster Recovery Testing Processes	20
Purchase Intent for Disaster Recovery as a Service	21
About.....	23

ActualTech Media © 2016. All rights reserved.

Under no circumstances should this document be sold, copied, or reproduced in any way except with written permission.

The information contained with the document is given in good faith and is believed to be accurate, appropriate and reliable at the time it is given, but is provided without any warranty of accuracy, appropriateness or reliability.

The author does not accept any liability or responsibility for any loss suffered from the reader's use of the advice, recommendation, information, assistance or service, to the extent available by law.

List of Figures

Figure 1: Current disaster recovery solution in place	5
Figure 2: Outages experienced in the past six months	6
Figure 3: Average cost per hour of downtime	7
Figure 4: Time to restore a key business application	8
Figure 5: Preferred boot location for recovered workloads	9
Figure 6: DR solution capability by vertical	10
Figure 7: Understanding the term "Disaster Recovery as a Service"	11
Figure 8: Application criticality breakdown	12
Figure 9: Current disaster recovery solution type	13
Figure 10: Company size breakdown and ability to recover key business applications within 15 minutes	13
Figure 11: Company size breakdown and ability to recover key business applications within 15 minutes	14
Figure 12: Number of physical and virtual servers under management	15
Figure 13: Amount of data that needs to be protected	16
Figure 14: Hypervisors in use	16
Figure 15: Understanding failover capabilities	17
Figure 16: Failover capabilities with regard to RTO	18
Figure 17: Reasons for lack of on-demand failover capability	19
Figure 18: Disaster recovery testing frequency	20
Figure 19: Intent to deploy DRaaS	21
Figure 20: Top evaluation criteria for on-demand failover solutions	22

Executive Summary

Disaster recovery might just be the most overlooked responsibility in IT departments around the globe. That's especially unfortunate, considering that there are products available in the market that can provide all companies comprehensive disaster recovery capabilities that enable a 15-minute recovery period for all applications.

As you'll learn in this report, the market for Disaster Recovery as a Service (DRaaS) providers is ripe with opportunity to educate customers on the potential business benefits of a well thought out disaster recovery plan. Perhaps one of the most critical findings is that 80% of respondents continue to use yesterday's technology for disaster recovery or do not leverage cloud. At the same time, demands on the business necessitate quick recovery. There is a fundamental disconnect between tape – yesterday's technology – and quick failover, particularly given the emergence of affordable DRaaS offerings.

Here are some of the highlights from our report:

- 20% of businesses still lack a disaster recovery solution.
- Tape-based backups still account for 42% of the existing disaster recovery strategies while appliance-based backups account for 38%. These companies still rely on yesterday's technologies for DR, which don't actually provide failover services and get their users back to business fast.
- 22% of respondents report experiencing more than a single outage in the past 6 months. As you'll learn later in this report, the self-reported cost of downtime might be underreported and may not include the full breadth of cost.
- Not surprisingly, cost was the most cited criteria for not having an on-demand failover solution and the most important factor when evaluating an on-demand failover solution.
- 37% of respondents can't even speculate what an outage costs their business. On the harmless side, this could just be due to the respondent's position in their organization. On the more dangerous side, however, is the possibility that the respondents who are decision makers aren't even aware of their risk in the event of an outage. Given the financial challenges and public relations issues that can result from an outage, this is a pretty scary situation.
- 22% of respondents test their DR plans less than one time per year – or never.

Please note that we limited survey results to US-based companies with 100 to 5,000 employees only and requested that only those responsible for disaster recovery complete the survey.

Introduction

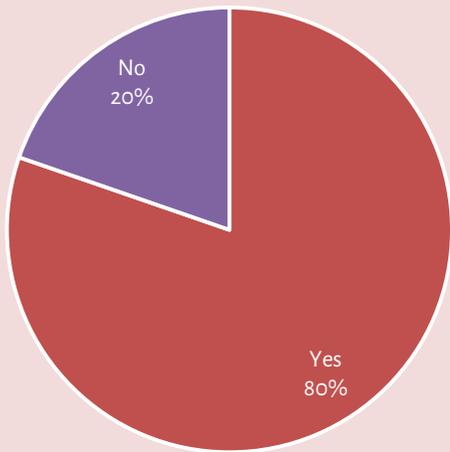
Although disaster recovery is not likely to be listed by many IT professionals as their favorite data center topic, the fact remains that it's a vital consideration in the greater data protection strategy. A robust disaster recovery strategy can be the difference between surviving the roller coaster of modern business and closing up shop when disaster strikes.

For a variety of reasons – the primary ones being cost and complexity – an unsettling number of businesses opt out of investing in disaster recovery solutions. Instead of ignoring disaster recovery altogether, many organizations make the difficult decision to protect a subset of their entire application landscape in the interest of balancing risk mitigation and cost. Clearly, it would be preferred to protect *every* workload in the data center.

In 2015, Infracore commissioned ActualTech Media to undertake a study of how organizations were handling their disaster recovery needs and to understand the adoption patterns of disaster recovery as a service. This report – the 2016 edition – is the second version of that effort, and shows how the market has changed over the last year and reveals new insights about the reality of the disaster recovery landscape.

It is our hope that this report will equip you with information about how your peers are handling disaster recovery and be able to compare your own processes and procedures.

Do you currently have a disaster recovery solution in place? (N=274)



To help set the stage, we first sought to understand the respondents' current disaster recovery status. As seen in Figure 1, 80% of respondents do have some form of protection from a disaster. This is roughly the same result as last year, which suggests one of two possibilities:

- The urgency of data protection is still lost on 20% of organizations, or is perceived as too expensive or challenging.
- The 20% without a disaster recovery solution have yet to discover a solution that fits both their needs *and* their budget.

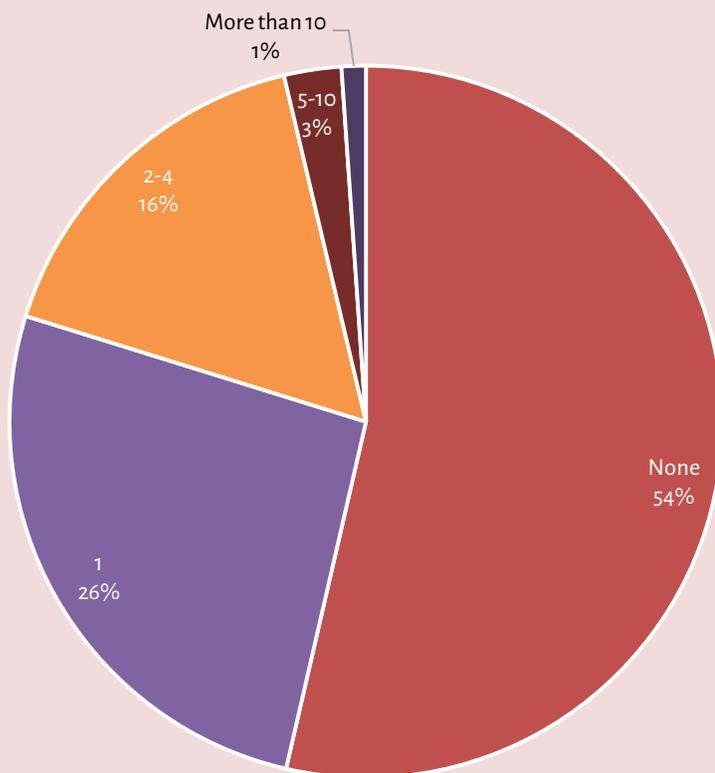
Figure 1: Current disaster recovery solution in place

Outage Handling Experiences

Even the best-designed data center experiences outages. It's an unavoidable certainty. What matters is not whether it will happen, but how it is handled when it does happen. How commonly a company experiences outages and how quick the IT staff are able to recover from the situation can directly correlate to overall profitability of the business. Downtime is like washing money down the drain.

In light of this, we sought to understand how frequently respondents are experiencing outages. Figure 2 shows that 54% of respondents either didn't experience any critical outages in the past 6 months or are not sure. The other 46% did experience some sort of disruption to a vital part of their business.

Figure 2 also shows that 20% of respondents report experiencing more than a single outage in the past 6 months. In the context of Figure 3, this reinforces the idea that many companies may be throwing away money by not investing in an on-demand failover solution to protect them.



**Over the last 6 months,
how many critical application
outages has your company experienced?**

Figure 2: Outages experienced in the past six months

IMPORTANT NOTE: WE'VE ASKED PEOPLE TO SELF-REPORT THEIR DOWNTIME STATISTICS. PEOPLE OFTEN FORGET TO INCLUDE SOME INFORMATION, SO IT'S VERY LIKELY THAT PEOPLE HAVE EXPERIENCED MORE OUTAGES AND MORE INDIRECT EXPENSE THAN WERE REPORTED.

To get a better understanding of exactly what outages cost respondents, we asked them to quantify the amount of money their business loses during every hour of downtime (Figure 3). The responses tell us several things:

- 37% of respondents can't even speculate what an outage costs their business. On the harmless side, this could just be due to the respondent's position in their organization. On the more dangerous side, however, is the possibility that the respondents who are decision makers aren't even aware of their risk in the event of an outage.
- Of those who are able to calculate their financial exposure, 44% of respondents' businesses *lose more than \$10,000 every hour that an outage continues*. This means that a single multi-hour outage could conceivably cost more than a reliable, on-demand failover solution. When compared to information shared by other outlets, such as Gartner, these figures are quite low. Bear in mind that these are self-reported figures that may not include all potential costs – compliance fines, overtime, public relations – incurred as the result of an outage.
- Larger companies experience *far more downtime cost* than smaller companies. This makes sense when you consider that it impacts more employees than in a smaller company.
- 61% of those that indicate that DRaaS solutions are too expensive are also unable to quantify the financial impact of an outage.

If you experience an outage, what is the average cost per hour of downtime (lost revenue productivity, brand damage, etc.)?

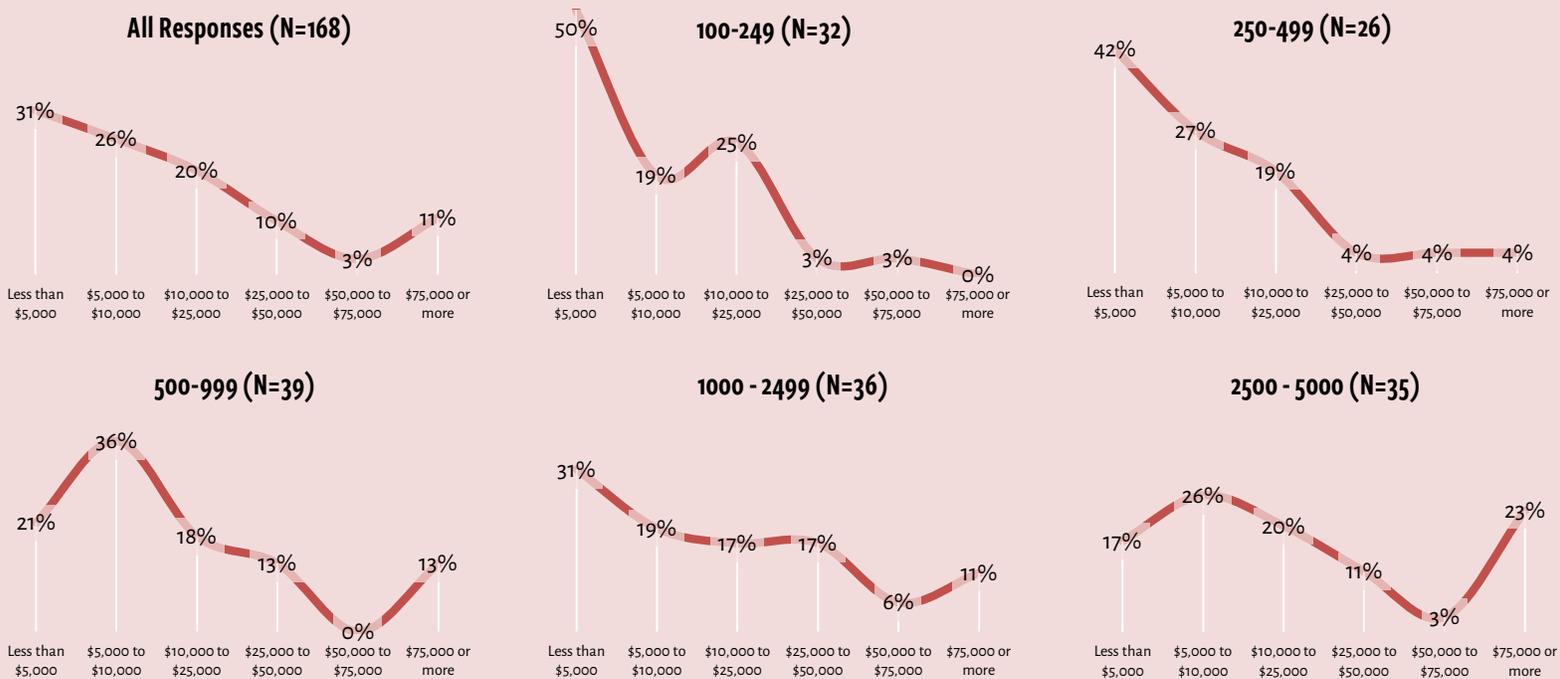


Figure 3: Average cost per hour of downtime

To get a sense of what the ongoing reality is in the experience of our respondents, we asked them to tell us how long it would take to recover a key business application; 89% of respondents indicated that they would be down for more than 15 minutes. Since all downtime translates to a loss of revenue and productivity in some form, recovery times measured in hours and days can impact the bottom line and IT's credibility.

If a key business application crashes, how long would it take to get it back online? (runbook) (N=254)

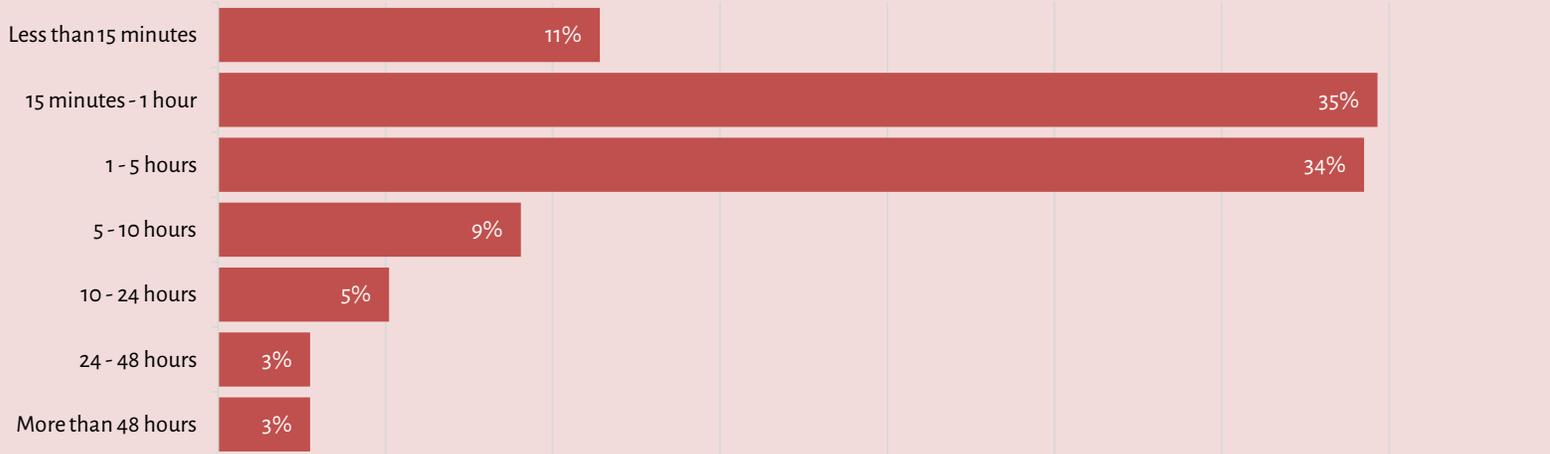


Figure 4: Time to restore a key business application

Recovery options vary greatly depending on the disaster recovery solution that is in use and the type of disaster that has occurred. We wanted to know what respondents prefer in terms of the platform for their recovery. As you can see in Figure 5, most respondents are most comfortable recovering the workload to their existing DR site. However, 26% prefer to recover into a cloud infrastructure and 24% would like to boot the workload on the local backup appliance. These figures almost exactly match those from the 2015 version of this report.

**Where would you prefer to boot critical applications in the event of a server failure?
(N=269)**

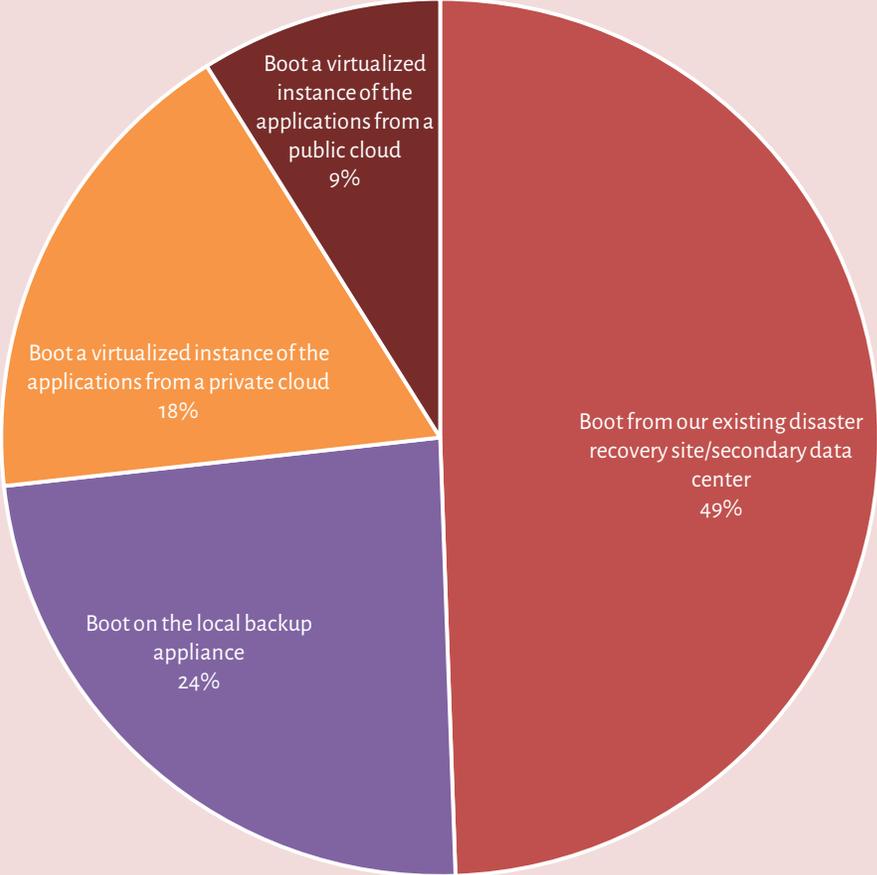


Figure 5: Preferred boot location for recovered workloads

Disaster Recovery Adoption by Vertical

It's interesting to look at the organizations that do or do not have a disaster recovery solution in place with regard to industry verticals. Figure 6 shows that according to respondents, *Government* and *Finance, Banking, or Insurance* are the most at-risk. Many government and financial companies will be required to possess a certain level of disaster recovery capability to be compliant with industry standards. On the other hand, *Retail* and *Energy or Oil & Gas* seem to enjoy living life on the edge. Ironically, VAR/MSP's—the customers that we would expect to be the benchmark for DR – are the companies least equipped to do so for themselves, with 43% lacking a DR solution.

Do you currently have a disaster recovery solution in place? (By vertical, N=274)

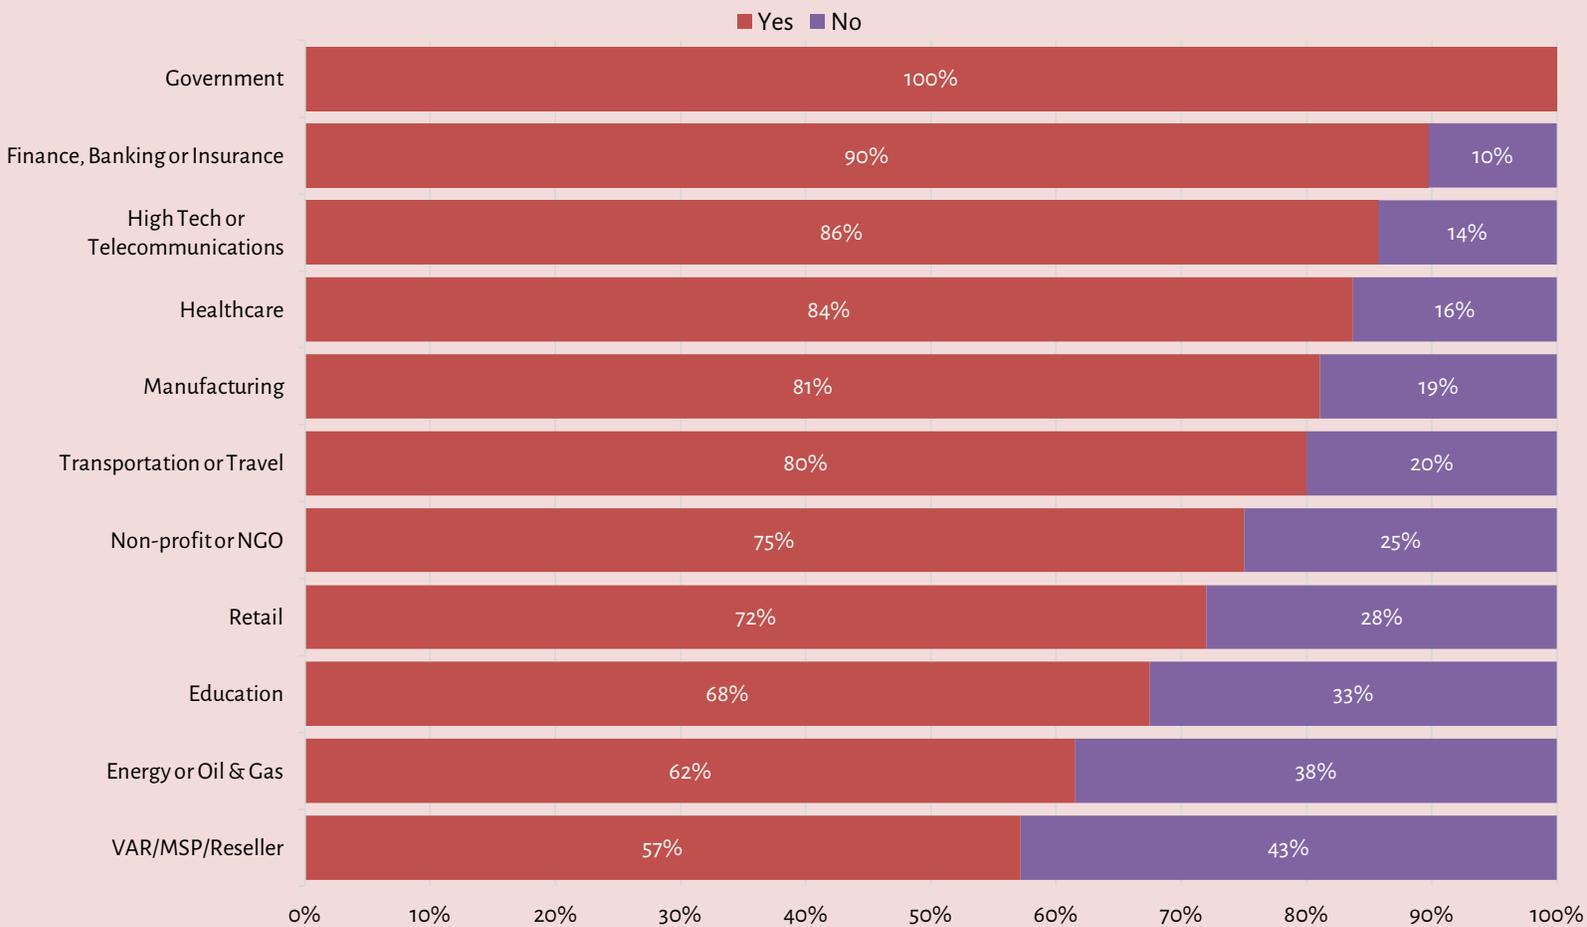


Figure 6: DR solution capability by vertical

Understanding Disaster Recovery as a Service

As shown in Figure 7, 20% of respondents indicated that they are *very familiar* with the term “disaster recovery as a service.” It seems plausible that they are either existing DRaaS users or their organizations have recently been doing active research into disaster recovery solutions. Consistent with last year’s survey, there remains significant opportunity for vendors to continue educating customers on both the concept and the value of DRaaS. As you can see in Figure 7, 67% of respondents have *some* familiarity with the term; 20% are very familiar with the term, and the remainder are completely unfamiliar. Cumulatively, 87% of respondents are at least somewhat familiar with the term.

How familiar are you with the term “Disaster Recovery as a Service” (“DRaaS”)? (N=274)

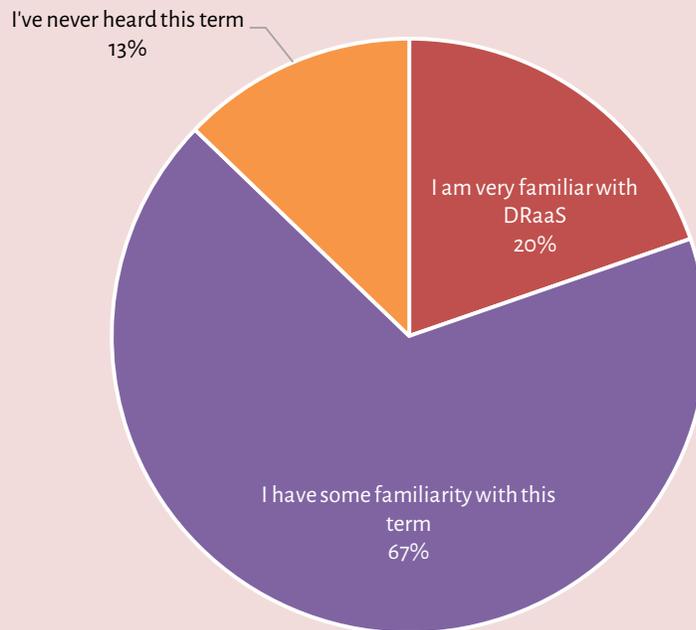


Figure 7: Understanding the term “Disaster Recovery as a Service”

Current Disaster Recovery Capabilities

We asked a number of questions intended to gauge the maturity of existing disaster recovery implementations. Understanding the current state of the environment from an availability standpoint helps make sense of decisions that are being made and also helps to clarify organizational priorities.

With respect to priorities, it all comes down to applications. Since the goal of disaster recovery is to ensure the availability of mission-critical business applications, we asked respondents to tell us their tolerance for downtime on certain common business applications. Note that, for 2015, we asked respondents to simply identify their most mission critical applications. For 2016, we requested that respondents tell us their downtime tolerance for each application. It becomes clear in Figure 8 that databases and e-mail are the most critical applications to respondents and require the shortest Recovery Time Objective (RTO), which is consistent with last year's results.

What is your tolerance for downtime for each of the following classes of applications?

15 minutes or less 15 minutes to 2 hours More than 2 hours

■ ■ ■

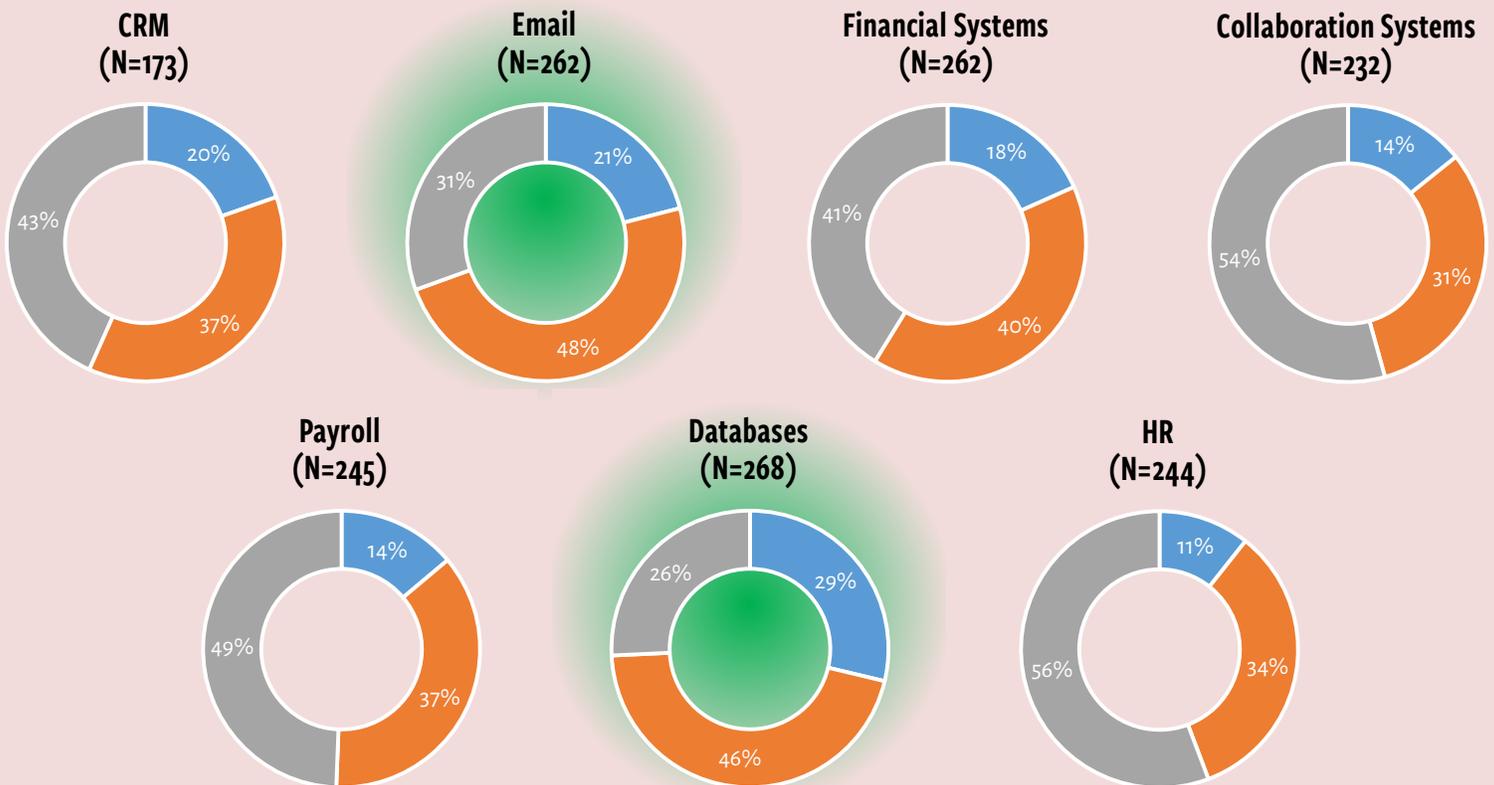


Figure 8: Application criticality breakdown

We want to understand the way that organizations are currently protecting their data. It's no surprise that "tried and true" tape backups – which are stored offsite – remain the most popular method (42%), closely followed by replicating local backups to an offsite appliance (37%). For those looking for rapid recovery or quick RTO, these technologies are not generally ideally suited for meeting this promise. Please note that respondents were allowed to choose multiple answers to this question.

Which statement best characterizes your current DR solution? (multiple selections allowed)

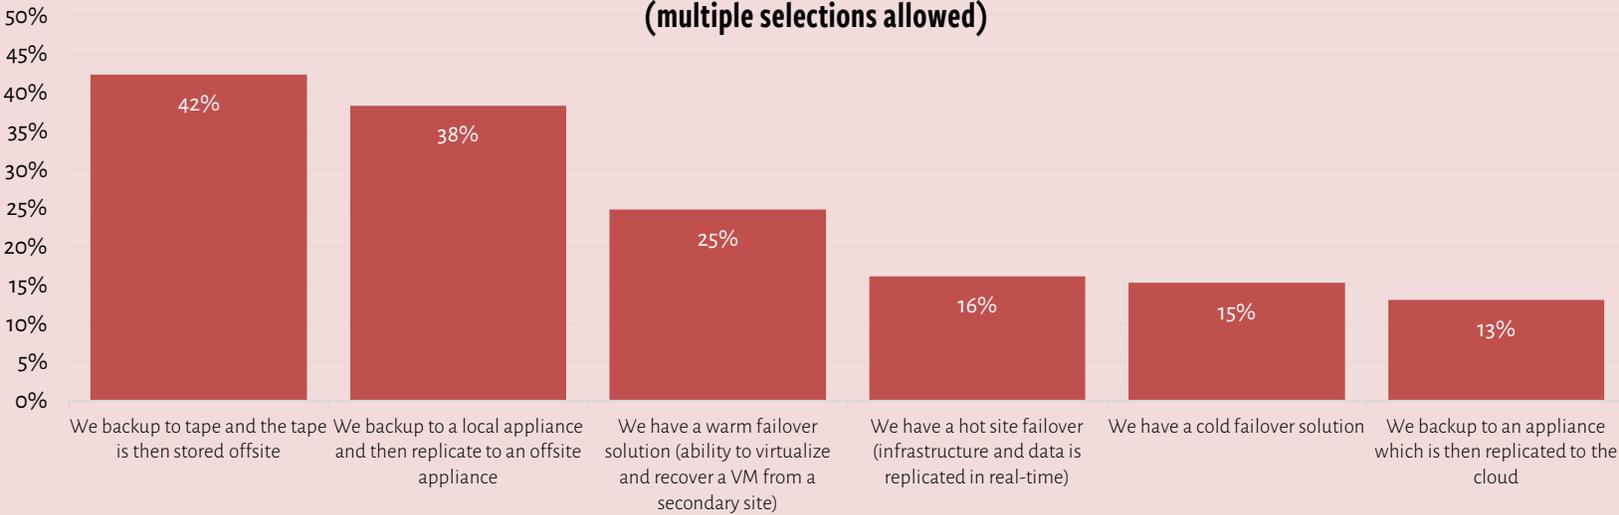
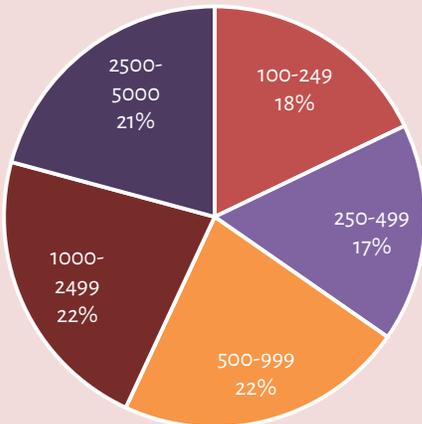


Figure 9: Current disaster recovery solution type

The survey sampling represents businesses from small to quite large in almost equal proportions. Figure 10 shows the breakdown. This data, when viewed in conjunction with the the ability to recover key apps within 15 minutes, suggests that larger companies are generally better able to perform quick recovery, although this is not necessarily a universal truth.

How many people work in your company? (N=274)



With your existing DR solution, can you failover your key business applications to a second site within 15 minutes? (N=244)

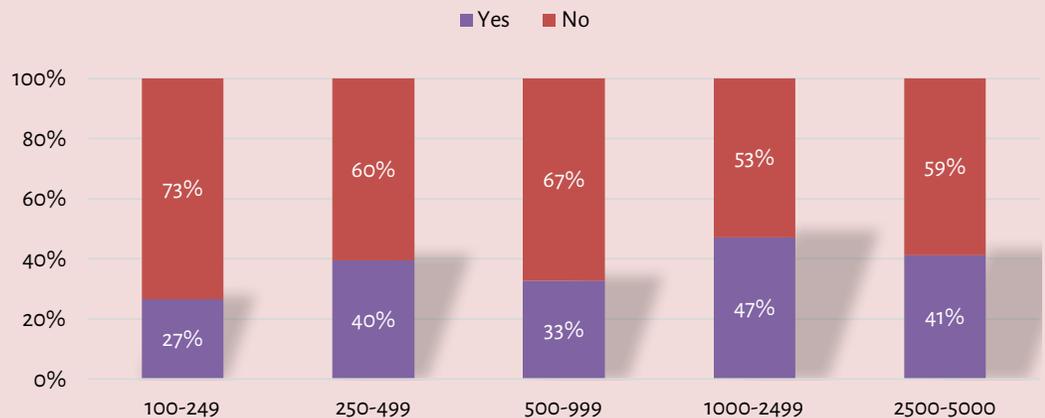


Figure 10: Company size breakdown and ability to recover key business applications within 15 minutes

Different kinds of companies have different needs and requirements when it comes to failing over critical business applications. In Figure 11, you can see that certain verticals are more prepared than others when it comes to the ability to recover key business applications within 15 minutes. Those that provide services to others – VARs and MSPs – seem to be the best prepared, with 57% of respondents in that vertical indicating their ability to recover within 15 minutes. Back in Figure 6, you saw that VARs and MSPs were at the bottom of the list when it comes to the number that have disaster recovery plans, so seeing them at the top of Figure 11 might be a bit of a surprise. Basically, we interpret this to mean that, while not all VARs and MSPs have solid DR plans, for those that do, they get it right.

Coming in second and third places are government and financial institutions, with 53% and 47%, respectively. From there, things get a little less positive, with all other verticals falling below the 40% mark.

**With your existing DR solution, can you failover your key business applications to a second site within 15 minutes?
(N=244, by vertical)**

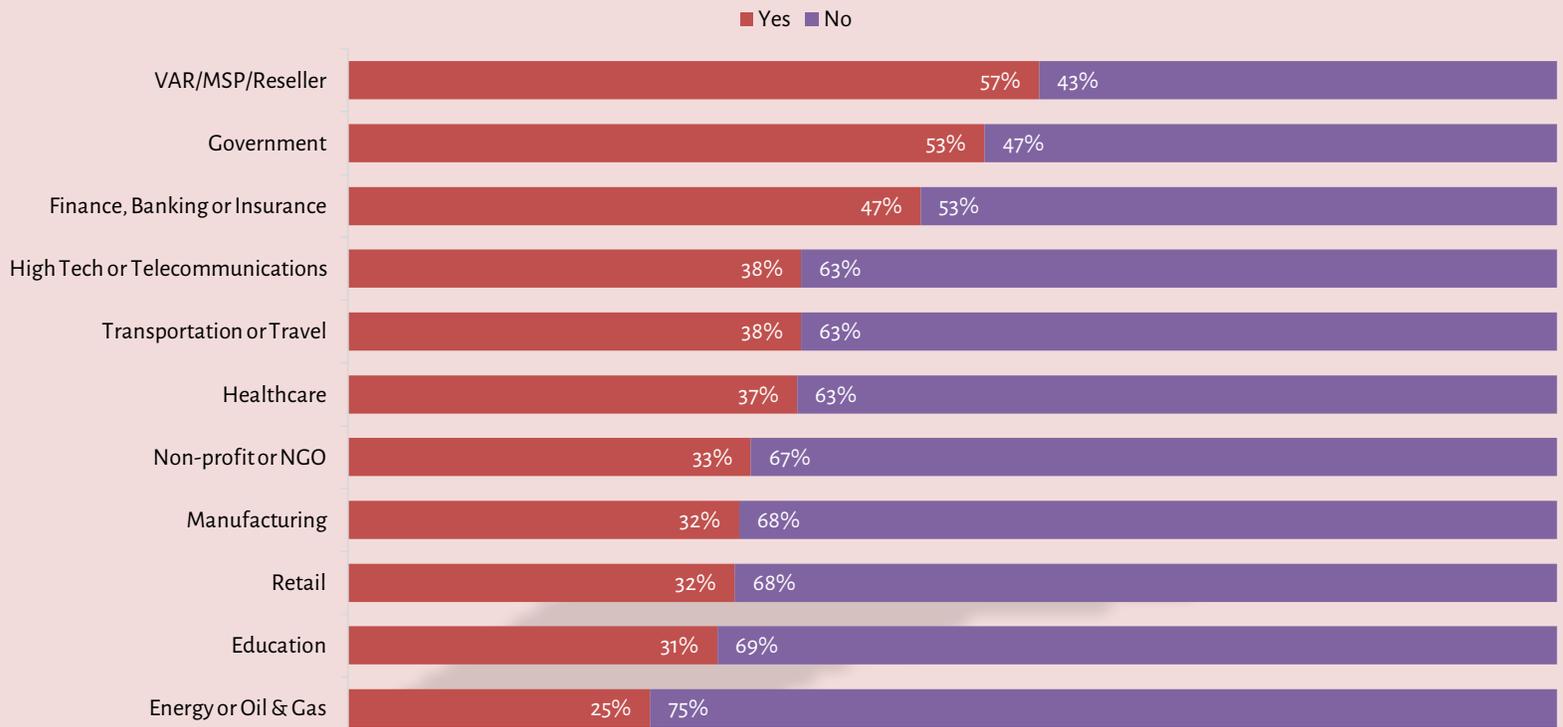


Figure 11: Company size breakdown and ability to recover key business applications within 15 minutes

Organizational Technical Characteristics

The scope of data center systems under management affects decisions regarding disaster recovery. Further, since the size of a company from an employee perspective doesn't necessarily correlate to the number and scale of applications requiring protection, we asked respondents about the systems that run their business applications.

As you can see in Figure 12, there are far more virtual servers used in respondent organizations than there are physical servers. This makes sense; after all, a single physical server can run dozens of virtual machines.

For companies that are slow to virtualize their workloads, many DRaaS solutions can't help you since they only protect virtualized systems.

Approximately how many servers exist in your environment across all sites and locations? (N=274)

Physical Virtual

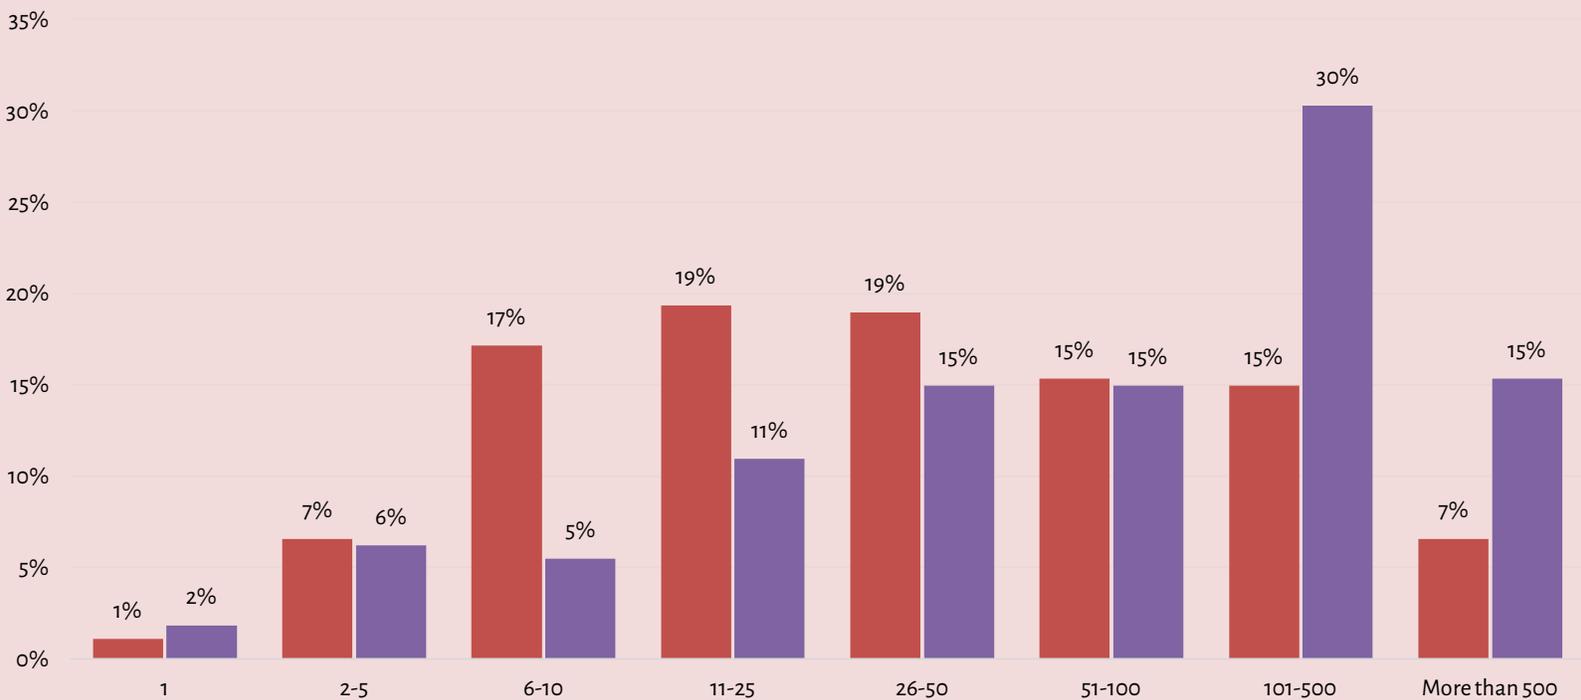


Figure 12: Number of physical and virtual servers under management

Although companies of all sizes need to protect data, it's interesting to understand the amount of data that organizations need to protect. In Figure 13, 44% of respondents are managing more than 50 TB of data. Interestingly, this is a 13% increase from the 2015 survey (31% with more than 50 TB), suggesting that the amount of data under management in the organizations surveyed has *grown substantially* over the last year.

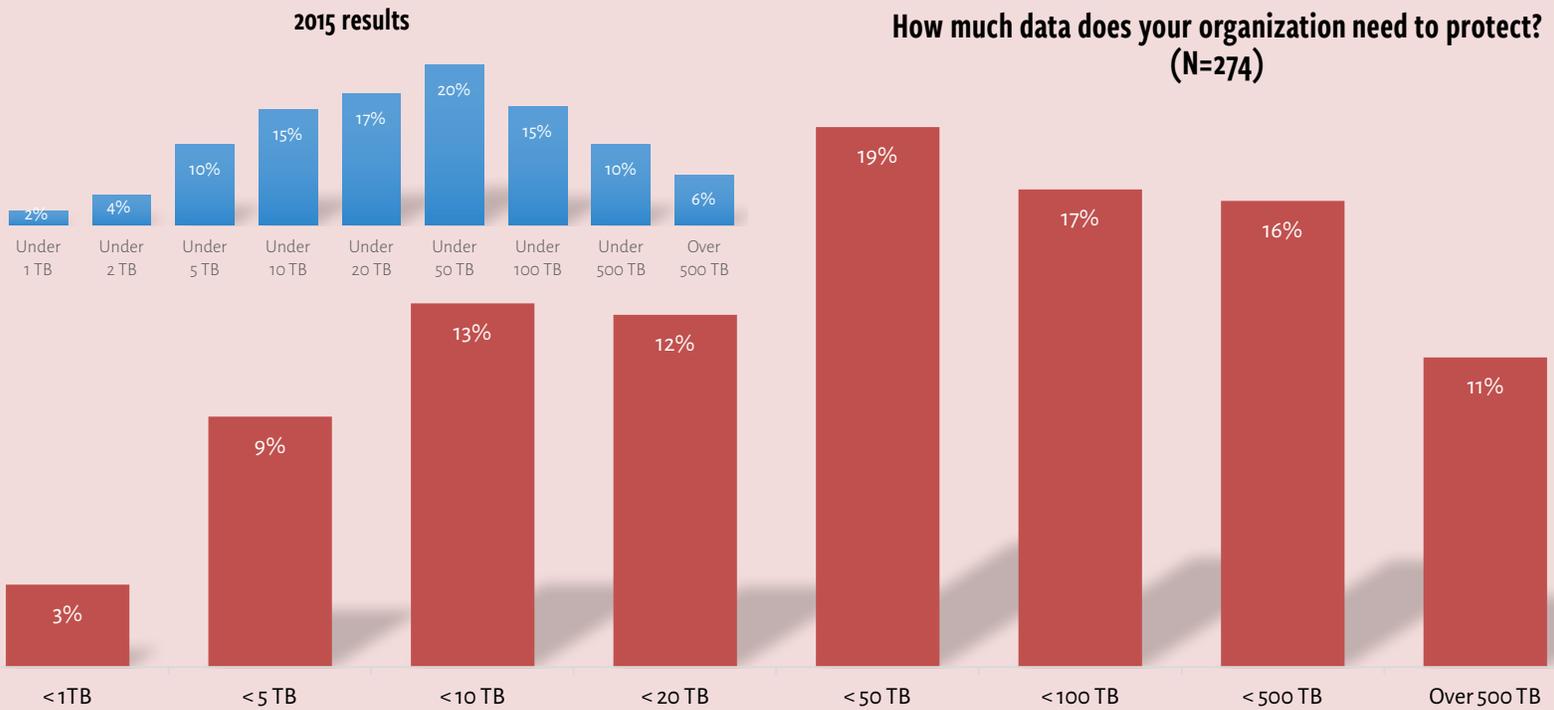


Figure 13: Amount of data that needs to be protected

What type of virtualization environment do you have in place in your organization?
(N=274)

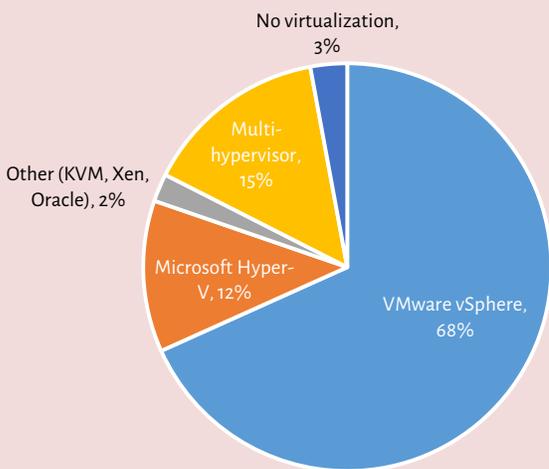


Figure 14: Hypervisors in use

The vast majority of businesses do virtualize some or all of their applications, though the overall percentage of virtualized workloads varies from company to company. This variation in virtualization penetration — from 0% to 100% virtual — means that customers need disaster recovery solutions that can support a wide breadth of virtual and physical systems and that can support the operating systems and applications that run inside those environments. It's no surprise that, for those who virtualize workloads, VMware vSphere is the clear market leader. It is notable, however, that a full 15% of respondents are operating a multi-hypervisor data center today.

Understanding Peer Disaster Recovery Capabilities

Everyone wants to be able to edge out the competition in some way and, believe it or not, disaster recovery capabilities are important enough that they can become a strategic differentiator. After all, if you and your biggest competitor both suffer disasters at the exact same time, but you can recover in 15 minutes, while it takes your competitor 48 hours, the advantage to you is clear.

So, where do you fall when it comes to disaster recovery? We asked respondents a series of questions in order to gauge their current status.

There are multiple services that need to be protected in the data center and, traditionally, companies have had to prioritize which services deserved protection. Disaster recovery for all services was considered too expensive or too complex. However, failover services have started to become more commonplace as some of these kinds of services are built into the hypervisor and as myriad failover solutions have come on the market in recent years.

Results of our survey support these observations. As shown in Figure 15, more than half of respondents have some capability to survive a disaster by failing over to an alternate resource or facility. While this is good, it also means that a disturbing 45% of respondents will be left in a less desirable position and will have to resort to time-consuming and cumbersome procedures, such as manually restoring from backups.

Do you have a failover solution that can protect your organization from an event such as a server failure or natural disaster?
(N=274)

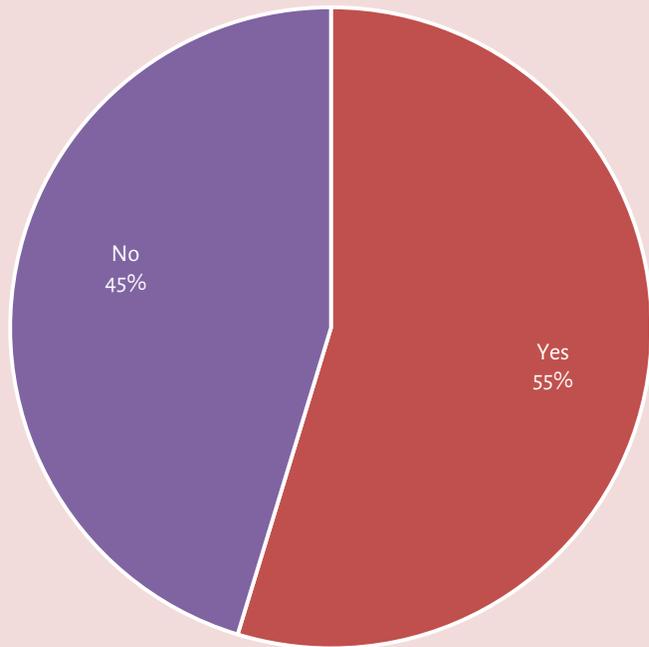


Figure 15: Understanding failover capabilities

As shown in Figure 15 the degree to which respondents are actually prepared varies greatly. To further clarify the exact capabilities respondents have with regard to failover, we asked them to describe their response to a failure, specifically with regard to RTO.

Figure 16 shows that of the respondents who said they had some level of failover capability, only 49% have a 15-minute RTO, and only 18% of those are able to get all of their applications back online within that 15-minute window. This suggests that even among respondents with some failover capability, there's substantial room for improvement in both the number of applications that can be failed over and the speed with which all protected applications can be recovered.

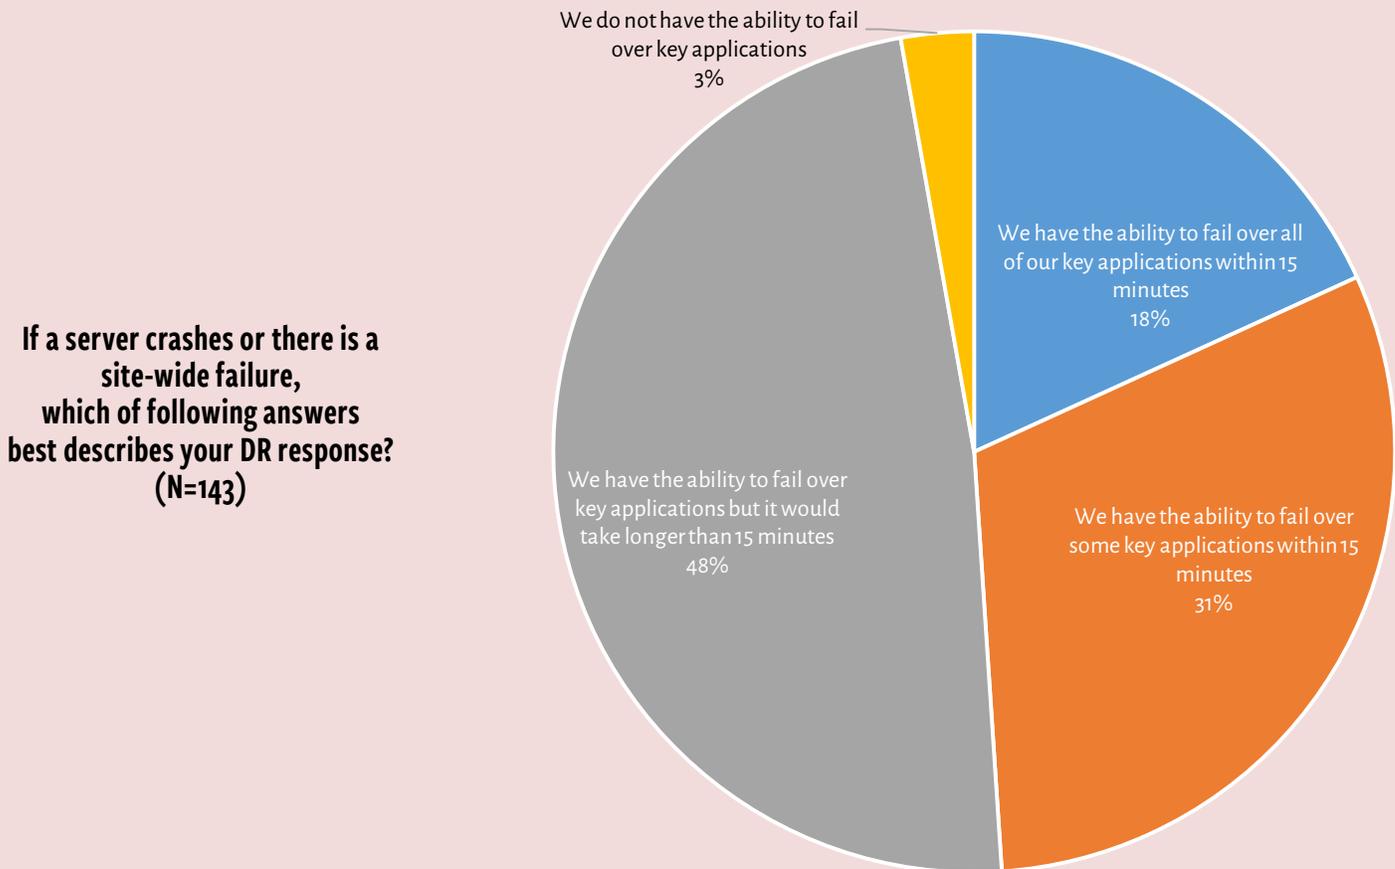


Figure 16: Failover capabilities with regard to RTO

As mentioned earlier in this report, disaster recovery is often neglected and two of the common reasons are cost and complexity. With all of the potential business benefits of an on-demand failover solution, we wanted to learn about the barriers to adoption. As expected, cost is the primary factor (Figure 17). Interestingly, however, the second most common reason cited for not having an on-demand failover solution is that it just hasn't been a priority. This lends substantial credence to the thesis presented in the introduction that the *urgency* of quality disaster recovery capabilities may not be apparent to some organizations.

If you don't currently have an on-demand failover solution, why not?
(N=124, Multiple responses allowed)

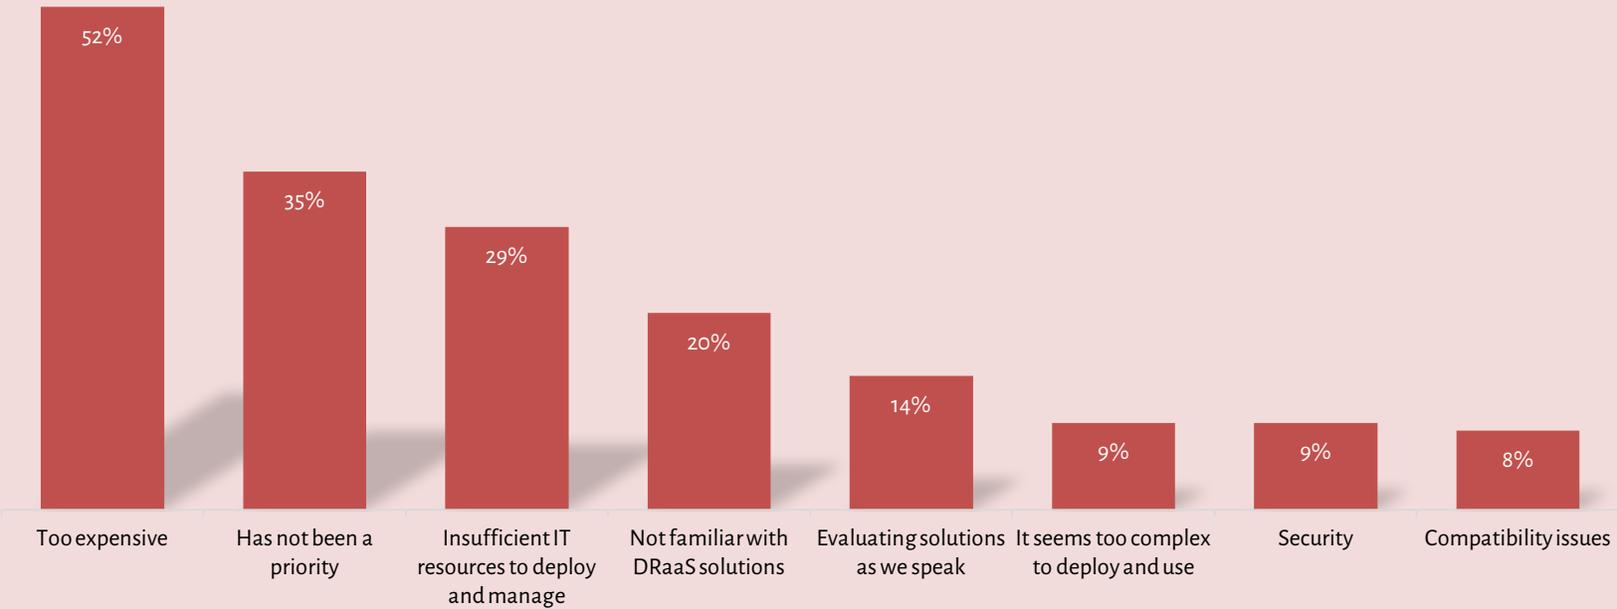


Figure 17: Reasons for lack of on-demand failover capability

Disaster Recovery Testing Processes

We know that DR plans often get tested less frequently than would be ideal, and in some cases they don't even get tested at all. In order to find out just how bad the problem is, we asked our respondents how frequently they actually test their DR plan (Figure 18). We found that 11% of respondents say that they have never tested their DR plan, and 11% more say that they test less often than once per year. We also found that 16% of respondents have DR providers that charge for testing, which may be why many rarely perform testing.

**How often do you actually test your DR plan?
(N=220)**

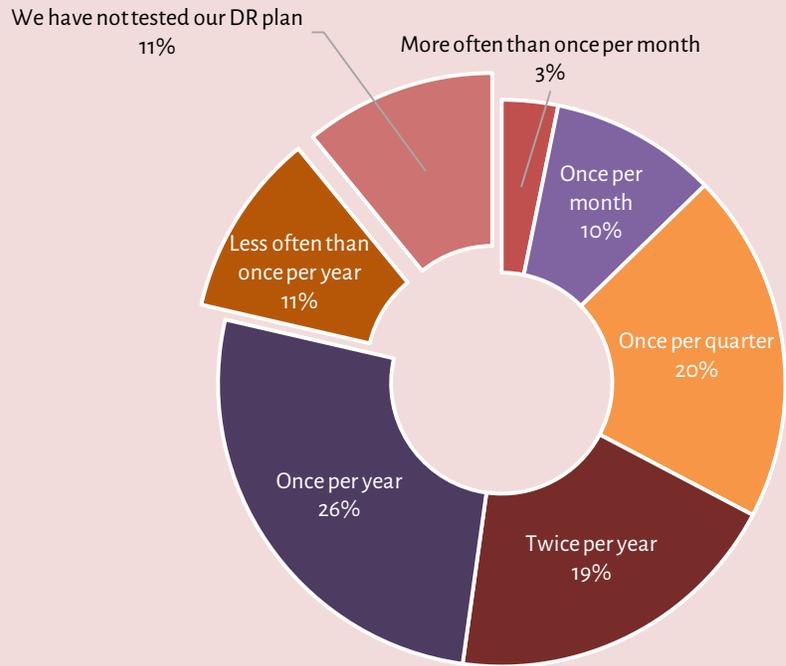


Figure 18: Disaster recovery testing frequency

Purchase Intent for Disaster Recovery as a Service

A whopping 46% of respondents indicated that disaster recovery is a top priority for their organization in 2016. In many cases, this will mean a capital purchase. So besides asking how they are prioritizing DR, we asked specifically whether they were evaluating on-demand failover solutions to meet their DR needs. You'll see in Figure 19 that a full 37% are either already using a DRaaS solution or plan to deploy one. Of those who plan to deploy one, however, 20% of them are still more than 6 months away from deployment.

Are you currently evaluating on-demand failover solutions? (N=274)

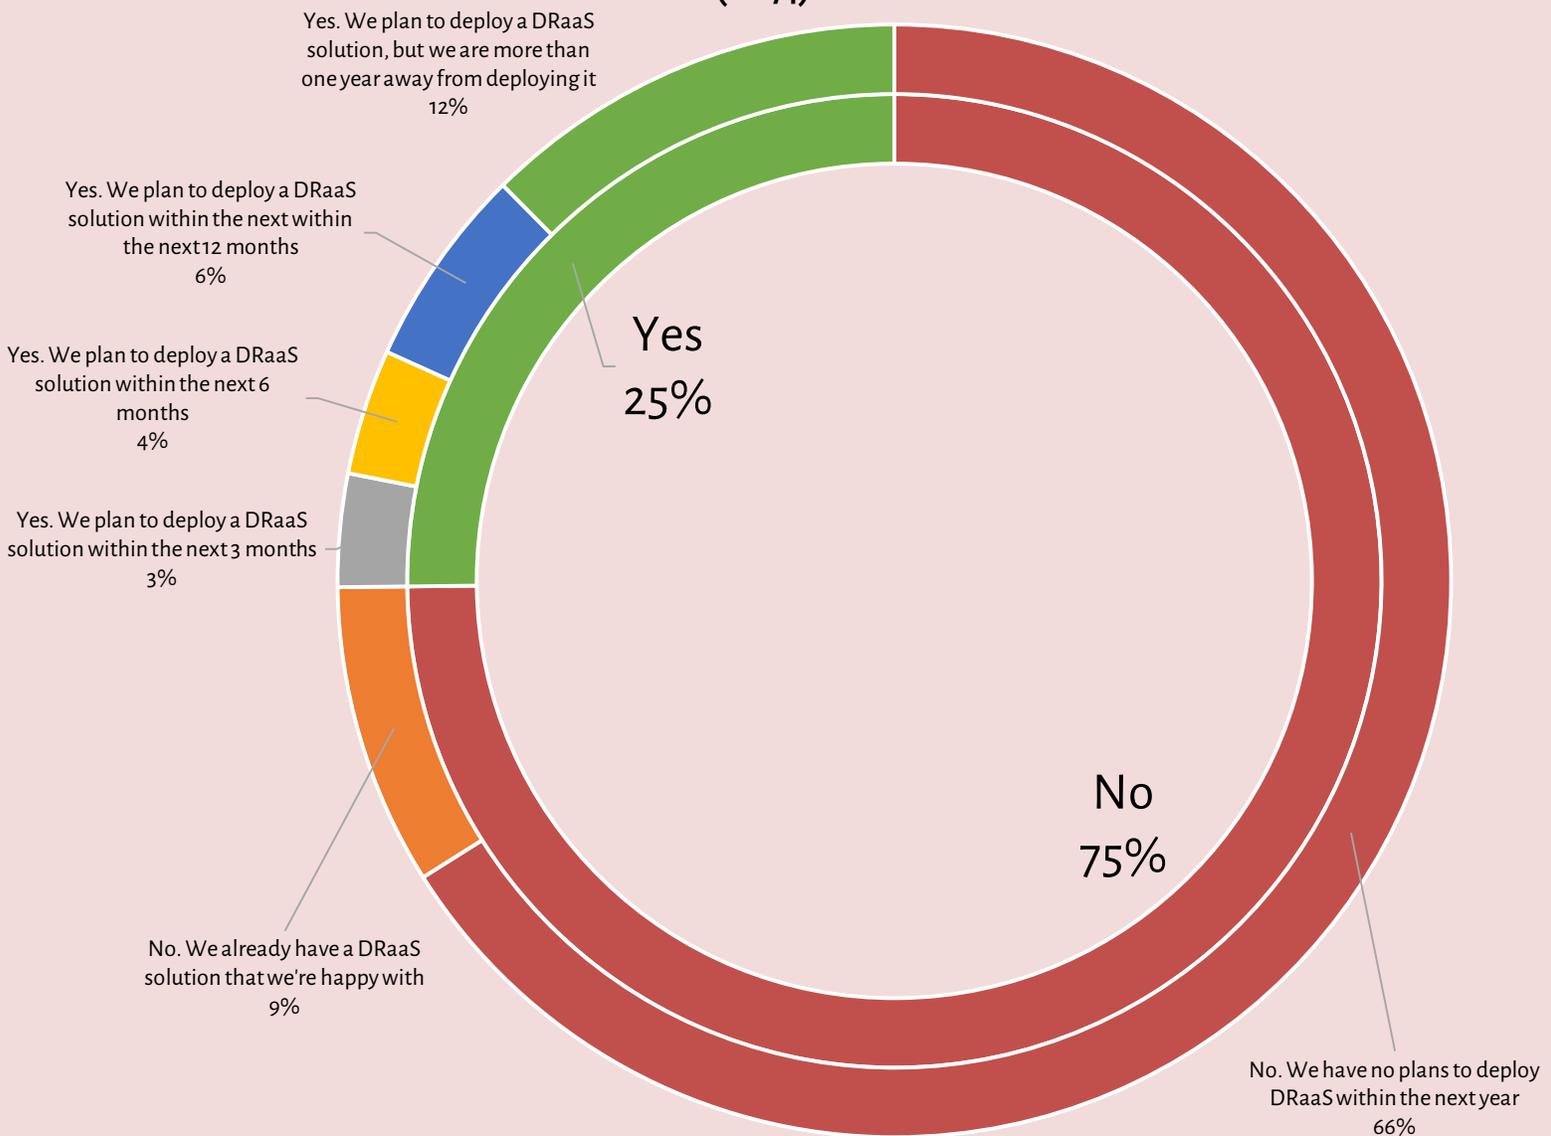


Figure 19: Intent to deploy DRaaS

We asked all respondents – even those not currently in the buying process – which criteria are most important to them when evaluating on-demand failover solutions. Unsurprisingly, the overwhelming leader is cost, as you can see in Figure 20. Reliability, security, and compatibility with existing infrastructure are also ranked highly. (Each item in the figure below is ranked on a scale of 1 to 10)

When evaluating on-demand failover solutions, what are your top decision criteria?

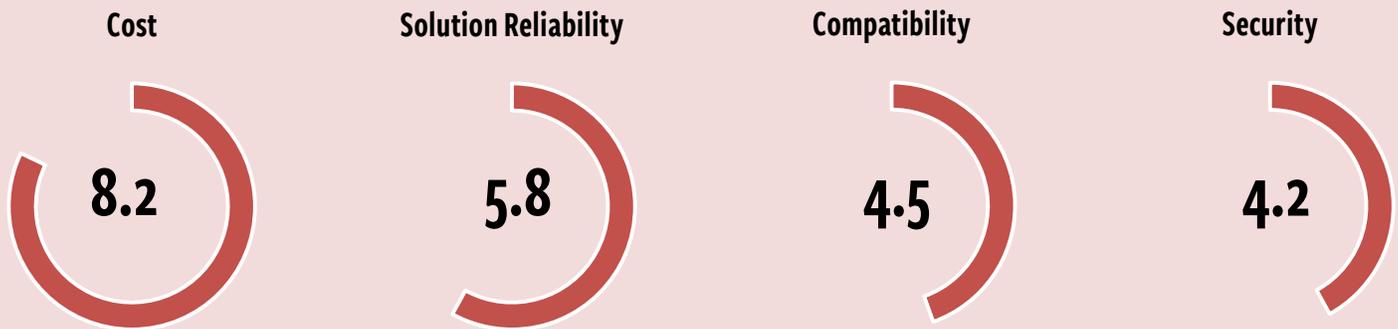


Figure 20: Top evaluation criteria for on-demand failover solutions

About

About Infracale

Infracale is a provider of the most powerful disaster recovery solution in the world. Founded in 2006, the company aims to give every company the ability to recover from a disaster – quickly, easily and affordably. Combining intelligent software with the power of the cloud is how Infracale cracks the disaster recovery cost barrier without complex, expensive hardware, enabling any company to restore operations in less than 15 minutes with a push of a button. Infracale equips business with the confidence to handle the unexpected by providing less downtime, greater security, and always-on availability.

About ActualTech Media

ActualTech Media delivers authoritative content services and assets for top IT companies across the globe. Leading IT industry influencers Scott D. Lowe, David M. Davis, James Green and partners develop trusted, third-party content designed to educate, convince and convert IT buyers. ActualTech Media helps its clients reach the right technical and business audiences with content that resonates and leads to results.