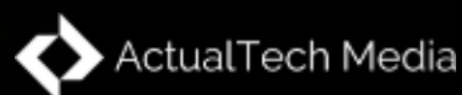


2015 Disaster Recovery as a Service Attitudes & Adoption Report



Scott D. Lowe
Partner and Co-Founder, ActualTech Media

David M. Davis
Partner and Co-Founder, ActualTech Media

Summer 2015

 **Infrascale™**

List of Figures

| | |
|--|----|
| Figure 1: Current disaster recovery solutions in place | 5 |
| Figure 2: Understanding the term “Disaster Recovery as a Service” | 6 |
| Figure 3: Application criticality breakdown | 7 |
| Figure 4: The primary methods for recovery in the event of a server failure..... | 8 |
| Figure 5: Company size breakdown..... | 8 |
| Figure 6: Understanding how company size impacts 15-minute recovery capability..... | 8 |
| Figure 7: Number of physical machines present in the organization | 9 |
| Figure 8: Number of virtual machines present in the organization | 9 |
| Figure 9: Amount of data that needs to be protected | 10 |
| Figure 10: 15-minute failover capability by amount of data being protected | 10 |
| Figure 11: Hypervisors in use | 11 |
| Figure 12: Gaining understanding for current application failover capabilities | 12 |
| Figure 13: Gaining understanding for current disaster failover capabilities | 12 |
| Figure 14: Failover capabilities for individual servers and applications | 13 |
| Figure 15: Failover capabilities in the event of a significant disaster | 13 |
| Figure 16: Failover capability broken down by industry vertical..... | 14 |
| Figure 17: 15-minute failover capability to a secondary site..... | 15 |
| Figure 18: Reasons for lack of on-demand failover capability | 15 |
| Figure 19: The number of critical outages experienced in the past six months | 16 |
| Figure 20: Time to restore key business applications | 16 |
| Figure 21: Average length of actual outages | 16 |
| Figure 22: Desired boot location for failing over critical applications..... | 17 |
| Figure 23: DRaaS deployment intention | 18 |
| Figure 24: Failover solution evaluation criteria..... | 18 |

Table of Contents

| | |
|---|----|
| List of Figures | 2 |
| Table of Contents..... | 3 |
| Executive Summary | 4 |
| Introduction | 5 |
| Understanding Disaster Recovery as a Service | 6 |
| Current Disaster Recovery Capabilities | 7 |
| Organizational Technical Characteristics | 9 |
| Understanding Peer Disaster Recovery Capabilities | 12 |
| Outage Handling Experiences..... | 16 |
| Purchase Intent for Disaster Recover as a Service..... | 18 |
| About ActualTech Media..... | 19 |
| About Infracore..... | 19 |

ActualTech Media © 2015. 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.

Executive Summary

Disaster recovery is a critical service that is not often performed at levels that truly protect the business. There are products available on the market that can bring to all companies – from the SMB to the enterprise – comprehensive disaster recovery capabilities that enable the gold standard in recovery – a 15-minute recovery period for *all* applications – while also being budget friendly.

As you will learn from this report, there is ample opportunity for solution providers in the Disaster Recovery as a Service (DRaaS) space to help customers better understand the real benefits that can come from such services and to significantly improve their overall recovery posture.

Here are some of the highlights from our report:

- 93% of respondents are either somewhat familiar or very familiar with the phrase “Disaster Recovery as a Service”
- 46% of respondents back up to a local appliance and then replicate to an offsite appliance; 22% still backup to tape and then store it offsite
- For respondents who do not currently have a DRaaS solution in place, 36% said that cost is the primary inhibitor to purchase; 26% said that they lack sufficient IT resources to deploy and manage. 17% of the respondents are currently evaluating DRaaS solutions
- Only 29% of respondents can currently recover key business applications within 15 minutes
- 43% of respondents would prefer to boot critical applications from their existing disaster recovery site/secondary data center; only 13% would prefer to boot applications from a public cloud
- When evaluating on-demand failover sites, cost was cited by 70% of respondents as a top 3 decision criteria; the next highest responses were security (48%), solution reliability (46%), infrastructure compatibility (39%) and solution complexity (34%)

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. Bear that in mind as you review the in-depth survey results.

Introduction

Like insurance, disaster recovery is one of those IT functions that every company knows they need to buy, but that they hope they never have to use. In addition — and also like insurance — many companies don't invest in disaster recovery solutions, and they leave themselves in a vulnerable position in the event that an incident or event takes place that results in the need for full business and application recovery.

Traditionally, disaster recovery solutions have been relatively expensive to procure and complex to maintain. Companies have had to make difficult choices and tradeoffs around which applications deserved protection, even while knowing that they would prefer to protect *all* of their applications.

Infrascale is working to enable companies of all sizes with reliable and affordable DR services. By providing a comprehensive disaster recovery as a service (DRaaS) product, Infrascale wants to bring disaster recovery to everyone in a way that is budget friendly while making it feasible to protect all workloads. To this end, Infrascale commissioned ActualTech Media to undertake a study of how organizations are currently handling their disaster recovery needs.

The results of this survey are compiled and analyzed in this report. Here, you will learn how your peers are handling disaster recovery and be able to compare your own processes and procedures. You will also learn how Infrascale can help you tackle even the most challenging disaster recovery needs and achieve 15-minute enterprise-class disaster recovery for your own organization.

Do you currently have a disaster recovery solution in place? (N=358)




Figure 1: Current disaster recovery solutions in place

Before getting fully started with the survey results, it's important to understand context around responses. To that end, perhaps one of the most critical questions asked respondents to indicate whether or not they currently have a disaster recovery solution in place. A full 83% of respondents indicated that they do as you can see in *Figure 1*.

There are two key questions that arise from this response:

1. Why are there still 17% of respondents without such a solution?
2. Are the 83% fully satisfied with the solution they currently have in place?

Read on to learn more.

Understanding Disaster Recovery as a Service

Just over one-quarter of respondents indicated that they are *very familiar* with the term “disaster recovery as a service.” This suggests that there is significant opportunity for vendors in this space to continue to educate customers on the potential benefits of DRaaS. Also as shown in *Figure 2*, 66% indicated *some* familiarity with the term, while the rest have never heard the term before. In short, 93% of respondents have at least some understanding of the term.

How familiar are you with the term “Disaster Recovery as a Service” (DRaaS)? (N=352)




Figure 2: Understanding the term “Disaster Recovery as a Service”

Current Disaster Recovery Capabilities

We asked a number of questions around respondents' current use of disaster recovery and application availability tools in order to gain an understanding for how things look today. By establishing a baseline for current capabilities, it becomes possible to determine a path forward to improve disaster recovery services.

Before that, though, it's important to understand the primary application support goals that respondents have as they consider disaster recovery needs. To that end, we asked respondents to tell us which two applications are most critical for protection. By far, databases (54%) and email (44%) were identified as the applications considered most mission-critical and that require the most comprehensive disaster recovery features. *Figure 3* provides you with an overview of the applications considered the most critical by survey respondents. Please note that respondents were allowed to select two applications for this answer, so the percentages do not add up to 100%.

Which category of applications do you consider most mission-critical? (Select two answers)

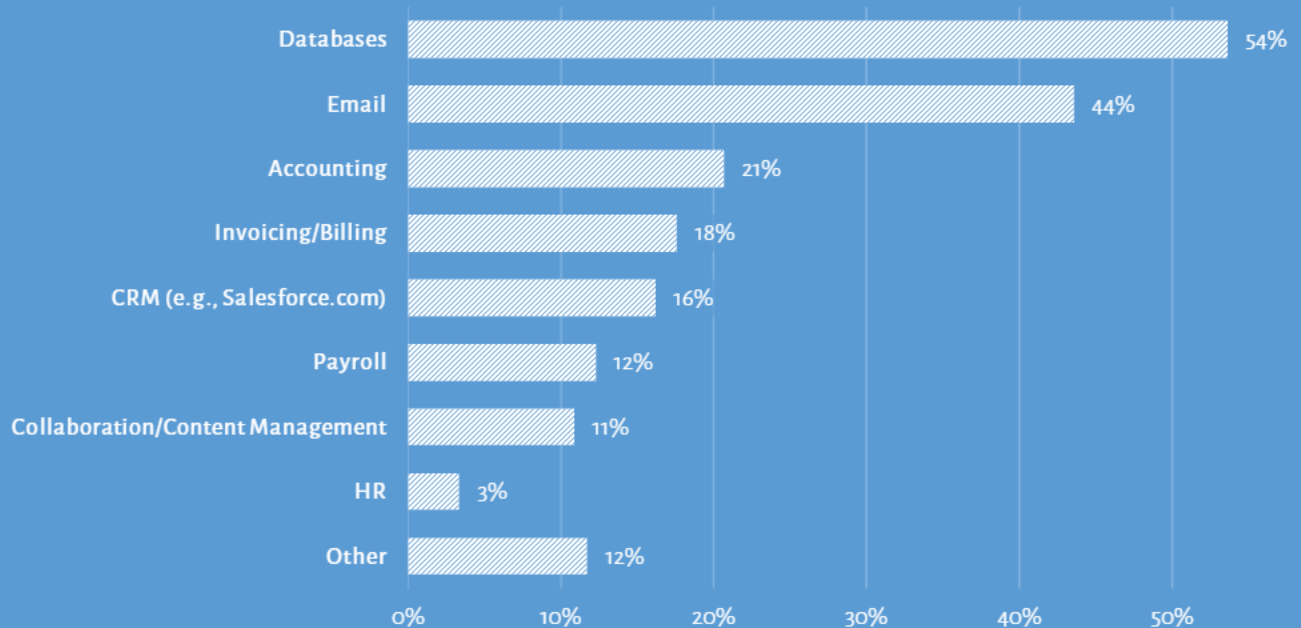


Figure 3: Application criticality breakdown

There are a number of ways by which applications can fail, with server host failure being a common occurrence. We asked respondents to provide feedback concerning their current recovery processes as they pertain to server host failure. 38% of respondents indicate that they would perform a typical restoration process from either a disk- or appliance-based recovery service. *Figure 4* provides you with the full set of responses, which also indicate that there that there is still a significant number of respondents (16%) that continue to rely on tape as a primary recovery mechanism.

In the event of a server failure, with your existing processes, how would you restore the applications? (N=348)

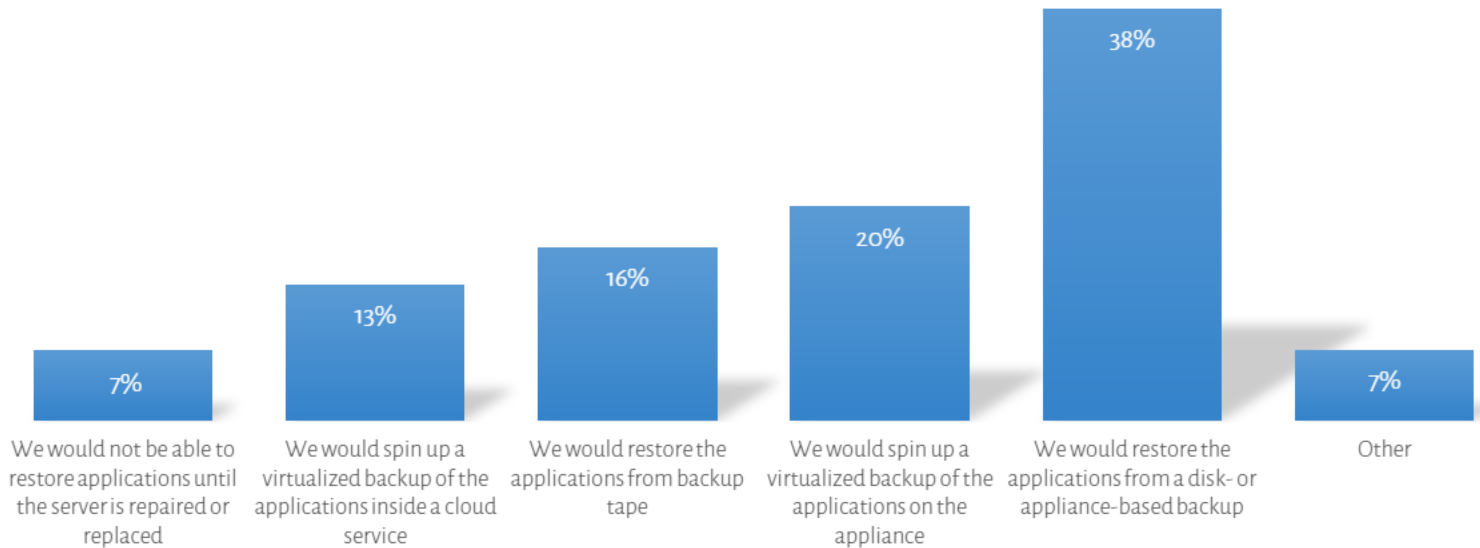


Figure 4: The primary methods for recovery in the event of a server failure

How many people work in your company? (N=358)




Figure 5: Company size breakdown

The survey results don't show that much of a difference based on the size of the company (*Figure 5*). In other words, all businesses, regardless of size, suffer equally from an inability to failover key business apps in the wake of a disaster. As shown in *Figure 6*, there is very little difference in the ability for large companies to recover in 15 minutes when compared to small companies.

15-minute recovery capability by company size (N=292)




Figure 6: Understanding how company size impacts 15-minute recovery capability

Organizational Technical Characteristics

Understanding respondent organizations' technical characteristics is important in order to gain insight into what is actually being predicted and to be able to provide some context around technical requirements. With that in mind, our survey queried respondents about a number of elements inside their data centers.

The number of employees in an organization does not always determine the actual technical size of that company. We asked respondents to tell us how many physical and virtual servers they're running in their organizations. These results are shown in *Figures 7 and 8*. As you can see, respondent data center sizes run the gamut from very small to large.

Approximately how many physical servers exist in your environment? (N=352)




Figure 7: Number of physical servers present in the organization

Approximately how many virtual servers exist in your environment? (N=352)




Figure 8: Number of virtual machines present in the organization

The quantity of data that needs to be protected directly impacts the size of the disaster recovery solution that needs to be procured. In *Figure 9*, you can see that close to 70% of respondents are protecting less than 50 TB of data. However, while the size of the protected data may inform the size of the overall disaster recovery solution, small organizations increasingly have similar recovery requirements as their larger brethren. As such, while protected data size is important, it's really more of an interesting data point for vendors rather than customers. Vendors need to address a broad range of market needs spanning both small and large customers alike.

How much data does your company need to protect? (N=358)




Figure 9: Amount of data that needs to be protected

Can you fail over your key business applications to a secondary site within 15 minutes?



Figure 10: 15-minute failover capability by amount of data being protected

Earlier in this report, you learned that company size is not a good predictor for whether or not the company has the capability to fail over to a secondary site within 15 minutes. Instead, it's interesting to note that the amount of data being protected seems to provide some insight. In fact, for companies with less than 2 TB of data, only 3% can failover critical business applications within 15 minutes.

What type of virtualization environment do you have in place in your organization? (N=358)




Figure 11: 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% virtual to all virtual all the time — 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 that virtualize workloads, VMware vSphere is the clear market leader.

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 — a step short from full disaster recovery — 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 12*, close to 60% of respondents currently have some kind of failover solution in place that allows them to recover critical applications. Of course, this also means that over 40% of respondents don't yet have this capability.

The results are a bit worse when looking at what happens after a major disaster (*Figure 13*). Barely half of respondents currently have a failover solution that can protect them in the event of a natural disaster. That's potentially a business-ending event.

Do you have a failover solution if a critical application goes down?
(N=351)




Figure 12: Gaining understanding for current application failover capabilities – for critical applications

Do you have a failover solution that can protect your organization from a disaster such as a hurricane, tornado, or other event that results in the loss of a data center? (N=348)




Figure 13: Gaining understanding for current full disaster failover capabilities – to protect against natural disasters

For those that have a failover solution (*Figure 14*), 43% can protect all of their mission-critical applications while 29% can failover just *some*. On the good-news front, 17% can fail over everything, but that left 12% as either not certain about what they could fail over or having different capability. The message here is that there is major room for improvement when it comes to failover capabilities. Even non-mission-critical systems exist for a reason, and maintaining availability for those systems will be increasingly important over time. *Figure 14* shows you results for single applications and servers while *Figure 15* displays capabilities for what happens when real disaster strikes. The results maintain a similar order of magnitude.

If you have a failover solution, to what level are you able to fail over your servers and applications in the event of an application failure?




Figure 14: Failover capabilities for individual servers and application

If you have a failover solution, to what level are you able to fail over your servers and applications in the event of a disaster such as a tornado or hurricane or other event that results in the loss of the primary data center?




Figure 15: Failover capabilities in the event of a significant disaster

The market vertical in which a respondent works provides some additional insight into the importance that is placed on failover and disaster recovery in different organizations. Some verticals, such as high tech and finance, seem to do a pretty good job at being able to fail over at the application level (*Figure 16*). Education and the Oil & Gas verticals, on the other hand, are quite challenged.

Please note, a longer green line signifies that an organization can't protect all of their mission critical systems.




Figure 16: Failover capability broken down by industry vertical

With your existing solution, can you fail over your key business applications to a second site within 15 minutes? (N=340)




Figure 17: 15-minute fail over capability to a secondary site

15-minute recovery capability is quickly becoming a business necessity for most applications. Many of today's disaster recovery solutions do a very good job at actually getting systems and applications back after recovery, but they often cannot achieve the goal within this 15-minute window. Respondents verify that this is true: only 29% indicate that they can fail over key business applications to a secondary site within a 15-minute window (*Figure 17*); 57% do *not* have this capability; 14% are uncertain, which we take to mean that they either simply don't know and that they don't have full confidence that their existing solution can meet this hurdle.

Given the importance of failover, the question then becomes one of why. In other words, why can't companies achieve a 15-minute failover goal?

One word: *Money*.

It's far and away the primary reason that on-demand failover is lacking in respondent organizations (36%; *Figure 18*). Priorities also play a part: 26% of respondents have other needs that are more pressing than on-demand failover. Sometimes, the here-and-now has to be taken into account before what many people consider to be solutions for handling "what if" situations. Another 25% simply don't have technical resources to deploy and manage. 11% believe that on-demand failover is too difficult to achieve. Vendors providing DRaaS services still have work to do when it comes to helping potential customers understand the ease by which such services can be deployed and managed.

17% of respondents indicate that they are currently evaluating solutions.

If you don't currently have an on-demand failover solution, why not? (Multi-select)




Figure 18: Reasons for lack of on-demand failover capability

Outage Handling Experiences

Suffering outages is a part of the game when it comes to managing data centers. This fact is the primary reason that companies are seeking to improve their overall recovery capabilities.

The more outages that you experience, the more downtime you suffer. It's a linear relationship. For most businesses, downtime equates directly to lost revenue and, sometimes, increased expenses. It's certainly not a financial win! Most respondents, 58% as shown in *Figure 19*, have not experienced any critical outages in the past 6 months and 2% are unsure about how many they had. That means that a whopping 40% of respondents did, indeed, suffer outages with 5% claiming 2 to 5 outages, and 1% experiencing *more than 10 outages*.

Over the last 6 months, how many critical application outages has your company experienced? (N=340)




Figure 19: The number of critical outages experienced in the past six months

Along with outages comes the need to recover from those outages. There are two metrics here that are important:

- How long do you expect it to take to recover?
- How long does it *actually take* to recover?

As you can see in *Figures 20 and 21*, people have a pretty good handle on their recovery times as expectation fairly aligns with reality. However, only around 10% to 12% of respondents can actually recover in less than 15 minutes. More than half of recoveries actually take 15 minutes to 2 hours and the remainder can take 5 hours to 48 hours to complete. That's the loss of 1 to 2 business days' worth of effort.

How long would it take you to restore a key business application if it crashed? (N=340)




Figure 20: Time to restore key business applications

What has been the average server or application outage duration? (N=144)




Figure 21: Average length of actual outage

You may notice that there appears to be a discrepancy between the results shown in *Figures 17, 20 & 21*. In *Figure 17*, note that 29% of respondents indicate that they can recover key business applications within 15 minutes. In *Figures 20 and 21*, that number drops to 10% to 12%. We suspect that, while people may be able to recovery *some* applications in 15 minutes, when they were pressed to put an actual timeframe around their answers, that they started to think deeper and get more granular by targeting approximate actual times.

If and when an application does happen to fail and you're in the process of recovering it in some way, you need to decide on how you want to actually accomplish recovery. There are a number of different options from which you can choose and the variety of options at your disposal is directly related to the disaster recovery capabilities available in your solution. We asked survey respondents to tell us where they would prefer to boot critical applications in the event of a server failure, with the respondent results shown in *Figure 22*.

More than 40% of respondents would prefer to boot servers from an existing secondary site. However, there are a number of people that would prefer to leverage either a public cloud (13%) or a private cloud (25%). 18% indicated that they would like to boot workloads right on the backup appliance.

Where would you prefer to boot critical applications in the event of a server failure? (N=340)




Figure 22: Failover service desired boot location for critical applications

Purchase Intent for Disaster Recover as a Service

It's surprising that 51% of respondents tell us that they do not have any plans to consider DRaaS within the next year. With that said, there are 49% of people that *do* plan to deploy DRaaS, but 19% of those plan to do so more than 12 months out. This leaves 31% of respondents considering DRaaS deployments within the next 12 months (*Figure 23*).

Are you currently evaluating on-demand failover (sometimes known as DRaaS) solutions? (N=340)

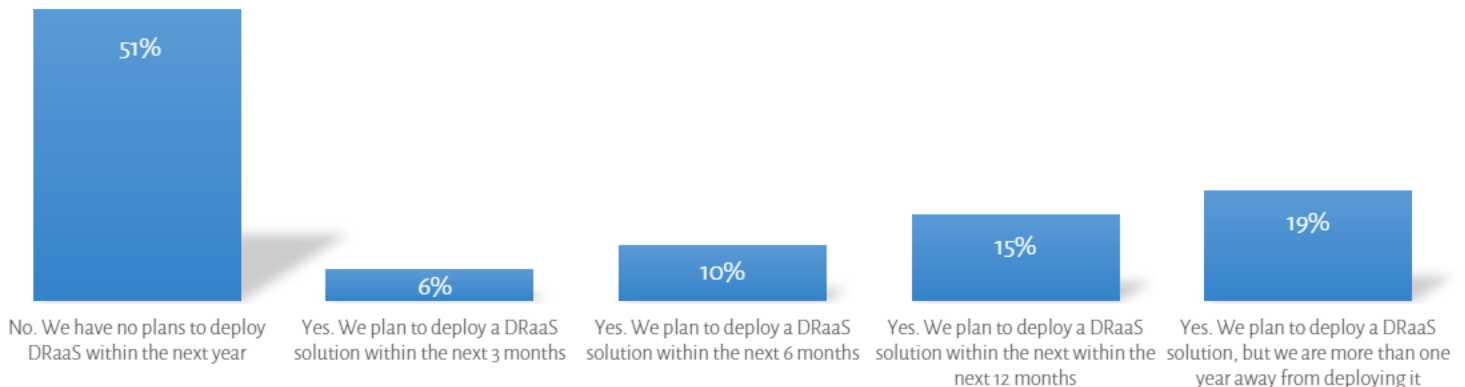


Figure 23: DRaaS deployment intention

In what probably comes as no surprise, solution cost is of primary concern, considering the purchase of an on-demand failover solution is the key driver (70%; *Figure 24*). Security takes second place with 48% of respondent votes and solution reliability comes in third place with 46%. Note that respondents were allowed to choose up to three responses for this question.

When evaluating on-demand failover solutions, what are the top 3 decision criteria?




Figure 24: Failover solution evaluation criteria

About the Sponsors

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 and partners develop trusted, 3rd-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.

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 The Authors



Scott Lowe is a vExpert and partner in and co-founder of ActualTech Media. Scott has been in the IT field for close to twenty years and spent ten of those years in filling the CIO role for various organizations. Scott has written thousands of articles and blog postings and regularly contributes to such sites as TechRepublic, Wikibon, and virtualizationadmin.com.



David Davis is a partner in and co-founder of ActualTech Media. With over 20 years in enterprise technology, he has served as an IT Manager and has authored hundreds of papers, ebooks, and video training courses. He's a 6 x vExpert, VCP, VCAP, & CCIE# 9369 and his blog can be found at www.VirtualizationSoftware.com