The Intersection of Hybrid Cloud, Workloads, Data Protection, and Disaster Recovery

SURVEY REPORT

ActualTech Media
RESEARCH

IN PARTNERSHIP WITH

Quest®

Contents

Introduction

Over the past decade, "the cloud" has evolved from a thorn in the side of the CIO to become a clear enabler of business, with incredible opportunity to help organizations change the economic landscape of IT, transform application development and deployment, and kickstart languishing digital transformation efforts. Along with all the good, however, have come some challenges. Some organizations have discovered that the hoped-for cost savings actually turned into cost overruns, while others discovered that "moving backups into the cloud" actually created unexpected recovery challenges.

In this report, Quest and ActualTech Media sought to discover how respondents feel about the cloud and how they make key decisions around cloud adoption and workload migration. In addition, we wanted to discover how organizations overall manage spend in both their private cloud environments. Finally, we questioned respondents about how they handle backup and disaster recovery in terms of the public cloud.

This report is broken up into two sections. The first discusses the public and private cloud landscape in general terms, while the second focuses on the role that public cloud plays in helping organizations protect their data assets.

EXECUTIVE SUMMARY AND KEY FINDINGS

There are a number of key findings from this research project:

- **81% of survey respondents are using the public cloud in some way** as a part of supporting their production workloads.
- **Even though the cloud is often seen as a solution to a problem, for 18% of organizations, the cloud actually presented a serious enough challenge** that they reverted to on-premises services.
- **Amazon Web Services (AWS) retains the lead in the cloud wars**, but Microsoft Azure is quickly catching up. Google Cloud Platform remains a distant third.
- **Just 6% of respondents use a purpose-built software tool or vendor solution to forecast their cloud spend.** The rest either don't do it at all or do it manually.
- **More than 50% of respondent organizations are using cloud in some way** as a part of their disaster recovery processes.

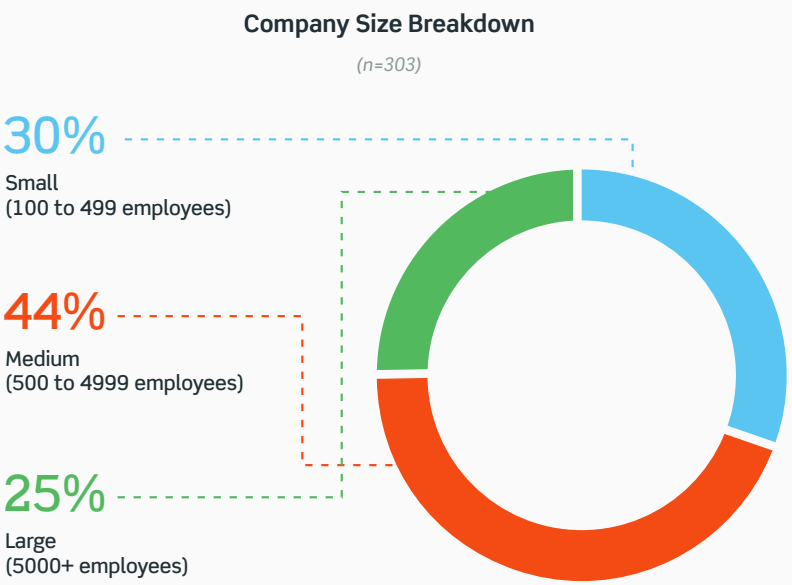

Figure 1: There was a wide range of companies represented by the respondents.

DEMOGRAPHICS

This survey focused primarily on decision-makers that may also hold responsibility for IT functional services, including data protection, cloud, virtual infrastructure, data center operations, monitoring, and applications. There were 303 valid respondents in the survey pool.

Survey respondents hailed from companies of all shapes and sizes, from 100 to 20,000+ employees. We broke responses down by small, medium, and large organizations, as shown in **Figure 1**. For the purpose of this analysis, company sizes break down as follows:

- **Small companies:** 100 - 499 employees
- **Medium companies:** 500 - 4,999 employees
- **Large Companies:** 5,000 or more employees

Throughout this report, where there are stark differences in statistics between companies of different sizes, we will provide a breakdown.

> Even though we've been discussing cloud for so long, many organizations are still taking tentative steps in that direction while others are jumping in with both feet.

Watching the cloud evolve has been an interesting experience as CIOs and IT pros have run through the hype cycle for over a decade now. Even though we've been discussing cloud for so long, many organizations are still taking tentative steps in that direction while others are jumping in with both feet. There are clear decisions that need to be made as workloads are migrated to the cloud. This section of our report outlines adoption patterns and workload migration decisions that respondents are grappling with as they undertake their projects.

Part 1: The Public and Private Cloud Landscape

CLOUD ADOPTION

One fact the data makes abundantly clear is that the pubic cloud has firmly enmeshed itself in the fabric of enterprise IT. Of the 303 responses analyzed, 81% indicated that they were running some type of production workload in the cloud (see **Figure 2**). In this context, cloud can be described as any public cloud platform such as AWS, Azure, Google Cloud Platform (GCP), or even a Software-as-a-Service (SaaS) solution such as Office 365 or Google Apps.

In terms of company size, mid-sized organizations are a bit more likely to be operating production workloads in the cloud (84%), but not by a significant margin over large companies (82%). Smaller companies appear to be a bit more hesitant, with 77% indicating that they're operating such workloads in the cloud (**Figure 3**).

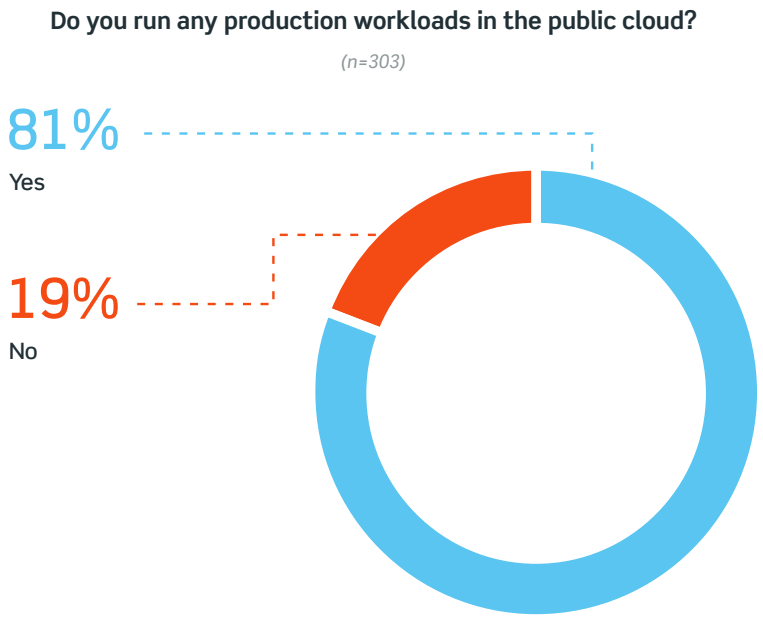

Figure 2: Most companies are running production workloads in the cloud.

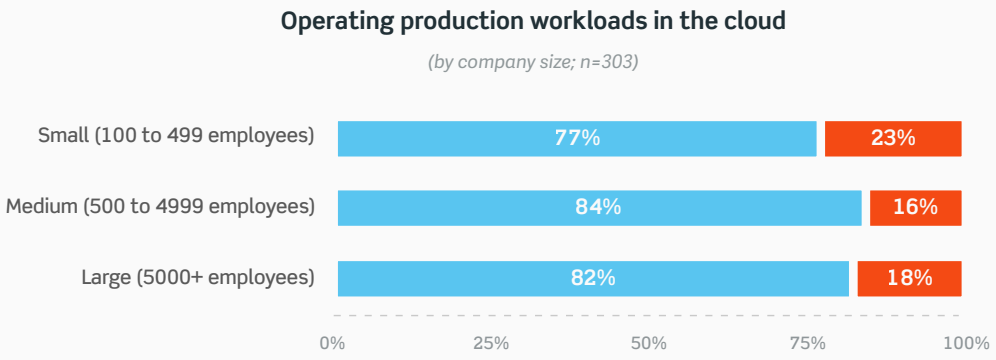

Figure 3: Medium-sized organizations have embraced the cloud more fully than others.

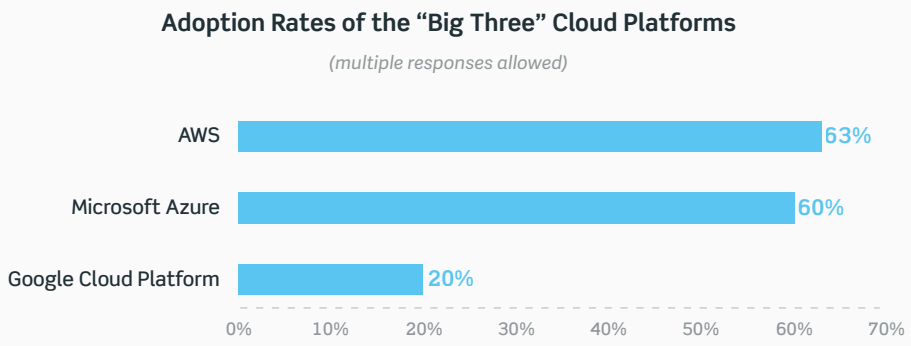

Figure 4: AWS remains in the lead in terms of adoption, but Azure isn't far behind.

CLOUD PLATFORM SELECTION

In terms of those that have formally adopted platforms such as AWS, Azure, or GCP, it's clear that Amazon remains in the lead, although Azure is growing quickly (**Figure 4**). GCP, a newer offering, remains a distant third in this race. Among respondents, 63% indicate that they're running AWS. Amazon's cloud fortunes have gained a boost from companies like VMware that are partnering with Amazon to help organizations bridge the private and public cloud gap with well-supported hybrid cloud solutions.

For Azure, the number is 60%, but with what we believe is a significant caveat. Office 365 may drive up this adoption figure, as organizations integrate Office 365 with on-premises environments through tools such as Azure Active Directory. For non-Office 365 users, we suspect that the Azure adoption rate is a bit lower than reflected in these results, but not likely by a lot.

As of today, GCP has been adopted by 20% of those polled. The platform is a bit younger than AWS and Azure, so it's no surprise that GCP adoption lags those services.

PUBLIC CLOUD DRIVERS

To understand the real driver behind cloud initiatives, we asked respondents to provide their primary reason for adopting cloud services (**Figure 5**). This was a single-answer question. A few years ago the top reason would have been cost reduction, as organizations sought to embrace the pay-as-you-go economic model of the cloud. Today, cost reduction comes in second place, with 25% of respondents identifying this as their key driver.

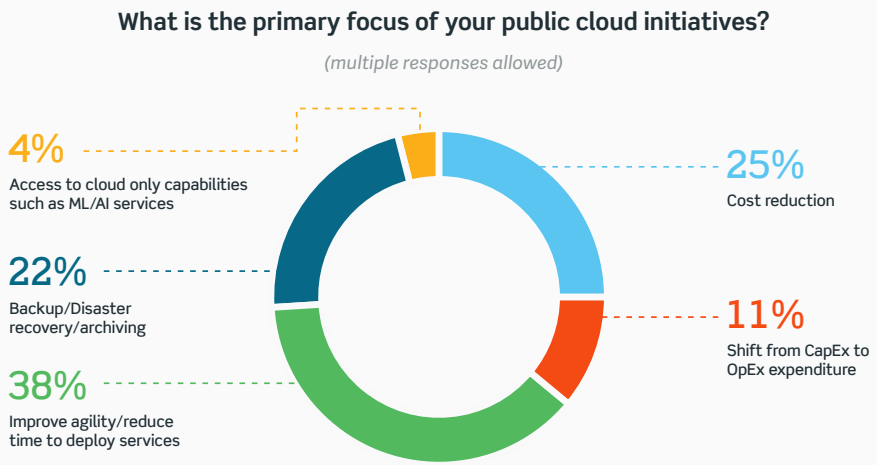

Figure 5: The primary drivers for moving to the public cloud have shifted.

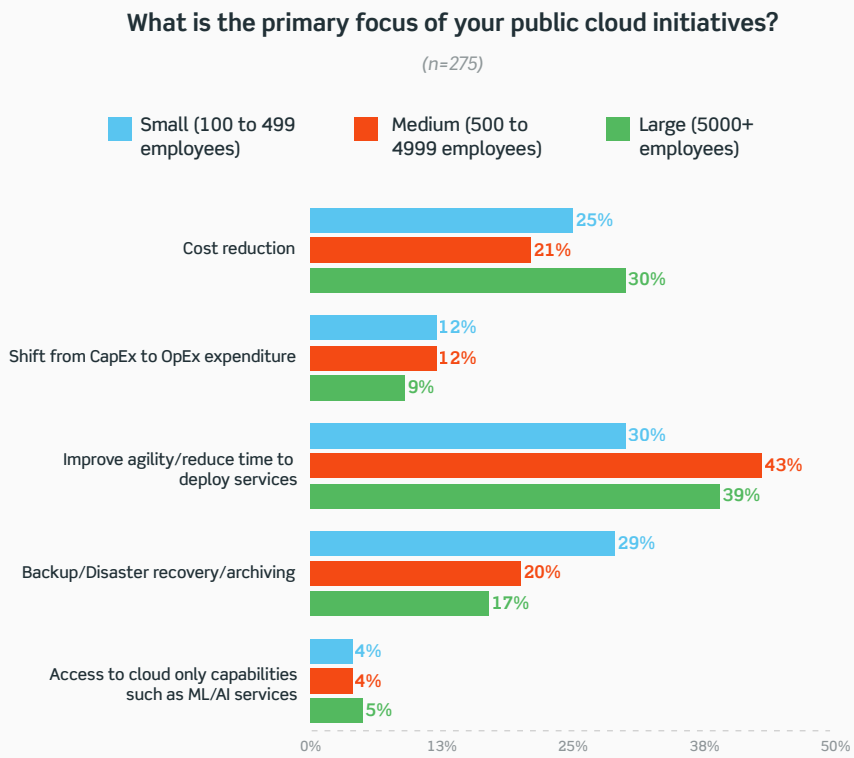

Figure 6: Primary focus of public cloud initiatives by company size.

Instead, the number one spot is now held by those seeking to improve overall IT agility and reduce the time it takes to deploy new services, which 38% of respondents cited as their key driver. In an era in which interest in digital transformation is increasing, it's not a surprise to see the agility and flexibility that cloud offers being brought to the forefront.

One of the original reasons companies were interested in cloud is still going strong. Coming in third place, 22% of respondents indicate that their primary cloud driver revolves around backup, disaster recovery, or archiving.

An emerging use case is top-of-mind for 4% of respondents, who indicate that their cloud plans are being driven by the unique capabilities that the cloud offers, such as quick access to machine learning and artificial intelligence capabilities.

There are clearly divergent drivers behind public cloud adoption when the data is sliced by company size. Cost is a relatively consistent factor, but even that varies a bit. For instance, 30% of large companies have adopted cloud primarily to reduce costs, while 21% of medium-sized companies have adopted cloud to save money.

What's interesting is what we see in the potential for increased workload agility and the use of the cloud for backup, disaster recovery, and archiving. Mid-size and large organizations see much more value than smaller ones in use of the cloud as a way to more quickly deploy new services. Smaller organizations, on the other hand, see a lot more benefit from the cloud in terms of data protection. This makes sense when you consider than smaller companies may not have secondary sites, so they view the cloud as their primary option there (**Figure 6**).

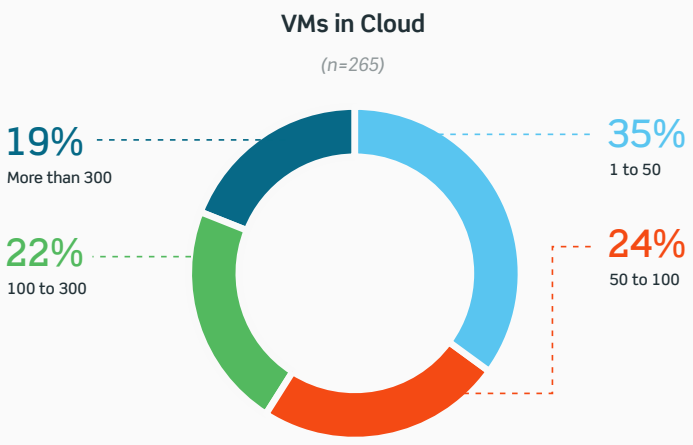

Figure 7: Organizations have been hesitant to place a lot of virtual machines in the cloud.

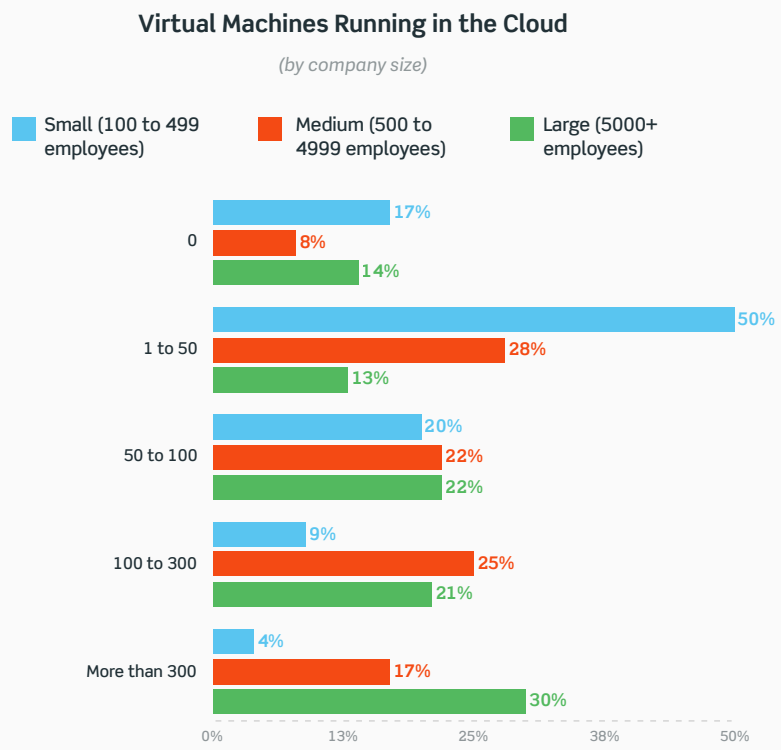

Figure 8: Perhaps surprisingly, mid-sized companies are the leaders in putting virtual machines in the cloud.

VIRTUAL MACHINES IN THE CLOUD

Not every company operating in the cloud is running virtual machines in the cloud. Many are operating SaaS applications and the like. However, for those that do run virtual machines, the quantity varies, not surprisingly, by company size.

In total, more than one-third of respondents (35%) are operating just a few – 1 to 50 – virtual machines in the cloud. Just 19% are running more than 300. For the chart broken down by company size, the full data set of 303 respondents is included (**Figure 7**).

> In total, more than one-third of respondents (35%) are operating just a few – 1 to 50 – virtual machines in the cloud. Just 19% are running more than 300.

It's interesting to note that mid-sized organizations are significantly ahead of both their smaller and larger brethren in terms of moving virtual machines into the public cloud (**Figure 8**). Just 8% of mid-sized organizations have no virtual machines in the cloud. In contrast, 17% and 15% of small and large organizations, respectively, have yet to move virtual machines into the cloud.

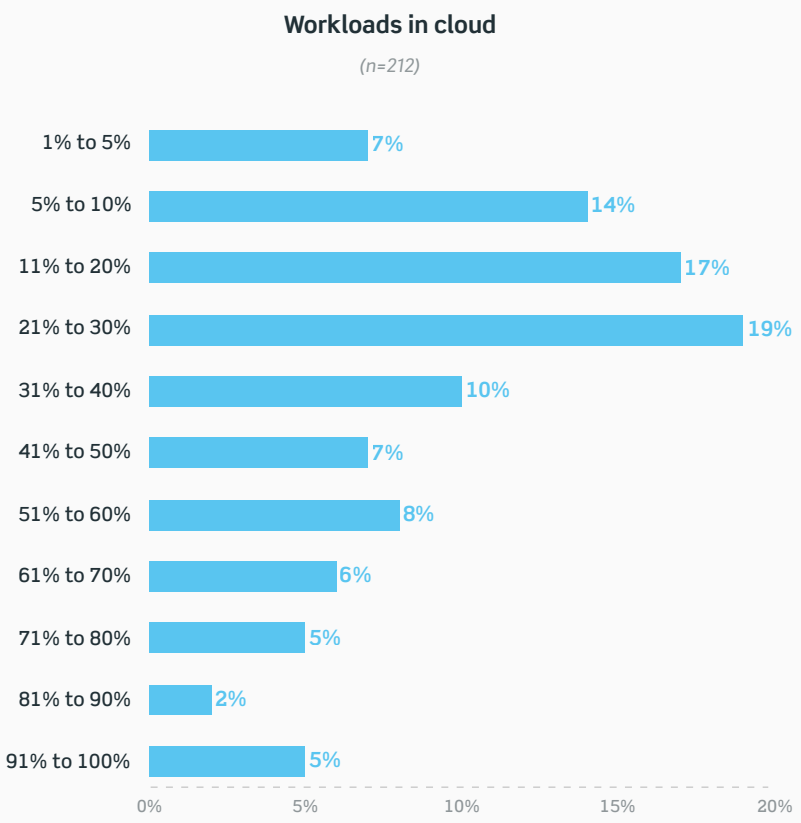

Figure 9: Organizations are still mostly keeping workloads in their own data centers.

GENERAL WORKLOADS OPERATING IN THE CLOUD

For anyone predicting the demise of the data center, consider those predictions curtailed (**Figure 9**). Just 5% of respondent organizations have 91% to 100% of their workloads running in the public cloud. In fact, 74% of organizations have fewer than 50% of their workloads in the public cloud. And 57% have fewer than 30% of their workloads operating outside the data center walls.

> By analyzing the data in this way, it becomes clear that smaller companies can more easily adopt cloud to a greater degree than large ones. This likely reflects the greater inertia to overcome in a large company vs. a smaller one.

There are two interesting ways to analyze this data point.

The first is to consider the sizes of companies that have migrated a certain mass of workloads into the cloud. In **Figure 10** below, note that, for companies that have moved 91% to 100% of the workloads into the cloud, 40% are small companies and 60% are medium ones. None of the large company respondents in this survey work in organizations that have moved 91% to 100% of their workloads into the cloud.

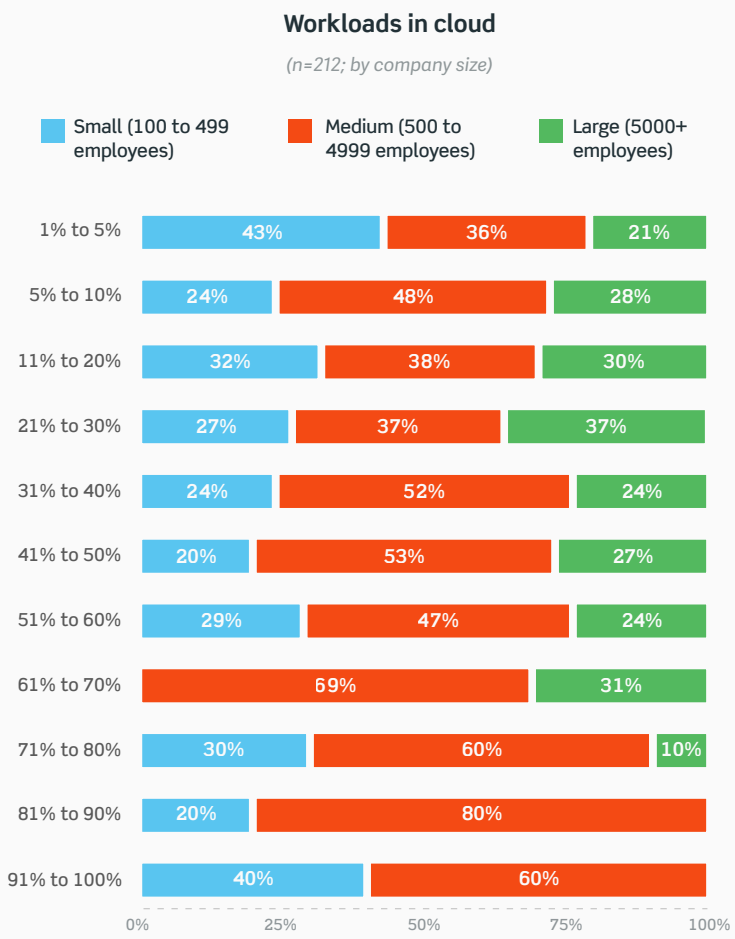

Figure 10: Percentage of workloads operating in the public cloud (by company size; breakdown 1)

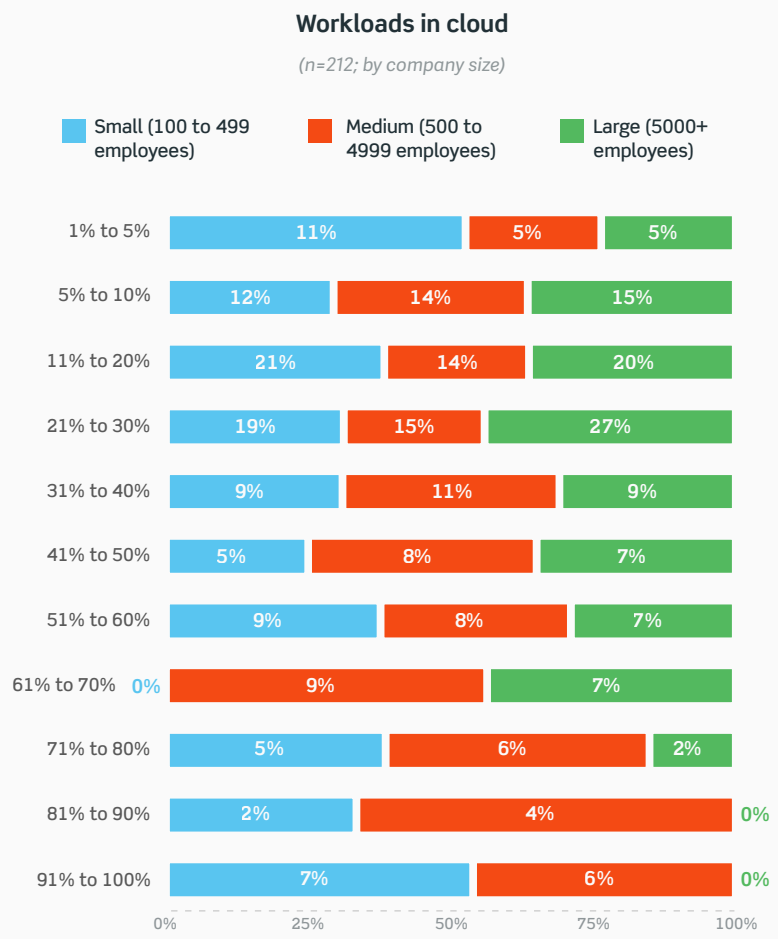

Figure 11: Percentage of workloads operating in the public cloud (by company size; breakdown 2)

By analyzing the data in this way, it becomes clear that smaller companies can more easily adopt cloud to a greater degree than large ones. This likely reflects the greater inertia to overcome in a large company vs. a smaller one.

The second way to analyze this data point is to consider it by company size. The chart in **Figure 11** is a 100% stacked bar chart that shows order of magnitude differences in the data. In this case, 11% of small companies have moved 1% to 5% of their workloads into the cloud, while 7% have moved 91% to 100%. This breakdown reinforces the hypothesis that smaller companies have an easier time moving more workloads into the cloud.

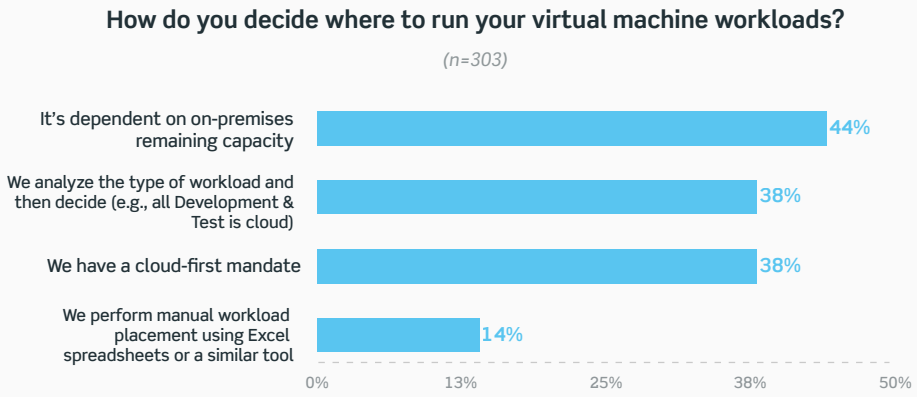

Figure 12: Workload placement decision processes.

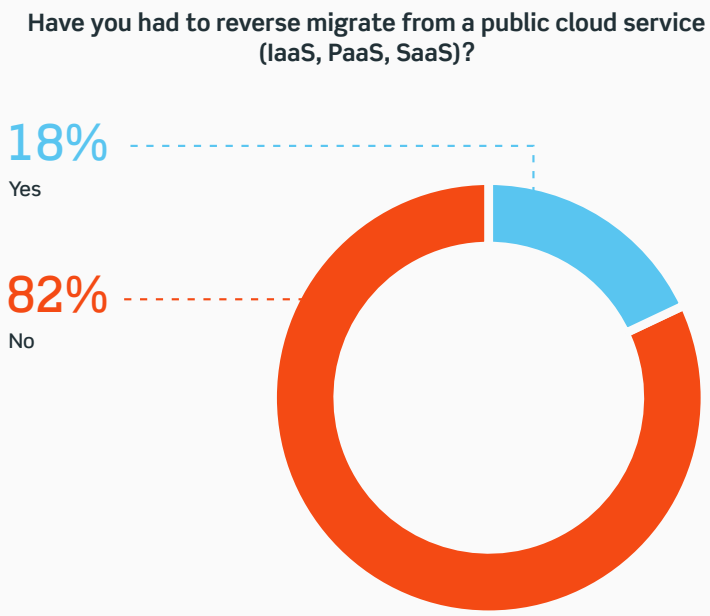

Figure 13: It's not uncommon for companies to move workloads back on-premises from the cloud.

WORKLOAD PLACEMENT DECISION PROCESS

As new workloads are placed into production, decisions have to be made around where to operate them. Do they run locally, or do you place them in the cloud? We wanted to understand how people make these decisions, so we asked a simple question about their decision process. Please note that this question allowed multiple responses, since multiple factors may play into the workload placement decision.

In short, 38% of respondents say that they have a cloud-first mandate. This doesn't mean that every workload is simply moved to the cloud, but it does mean that placing a workload in the cloud is a higher priority than placing it in an on-premises data center.

Another 38% indicate that they have a workload breakdown matrix that helps them decide where to run things (**Figure 12**). For example, all test and development workloads may be required to operate in the cloud. For many, the biggest choice comes down to on-premises capacity. Forty-four percent (44%) of respondents said that they consider current on-premises capacity before making a decision. Finally, 14% indicate that they use manual workload placement tools such as Excel to help them decide where to place workloads.

MIGRATING AWAY FROM THE PUBLIC CLOUD

A number of organizations that have moved workloads into the cloud have discovered that, at some point, they needed to move it out of the cloud and back into an on-premises environment. In fact, of those that have moved data into the cloud, 18% have had to bring some or all of it back out (**Figure 13**). There are a number of reasons this happens:

- **Security/personally-identifiable information (PII).** The company suffered a security event, was worried about the potential for a security event or concerned about the safety of PII operating in the cloud.
- **Data loss.** The company suffered some kind of data loss event in the cloud that convinced them to move data back onsite.
- **Perceived simplicity compared to their on-premises environment.** Right or wrong, people perceive simplicity in what they know. For many, public cloud is an unknown, which adds to the perceived complexity of the environment.

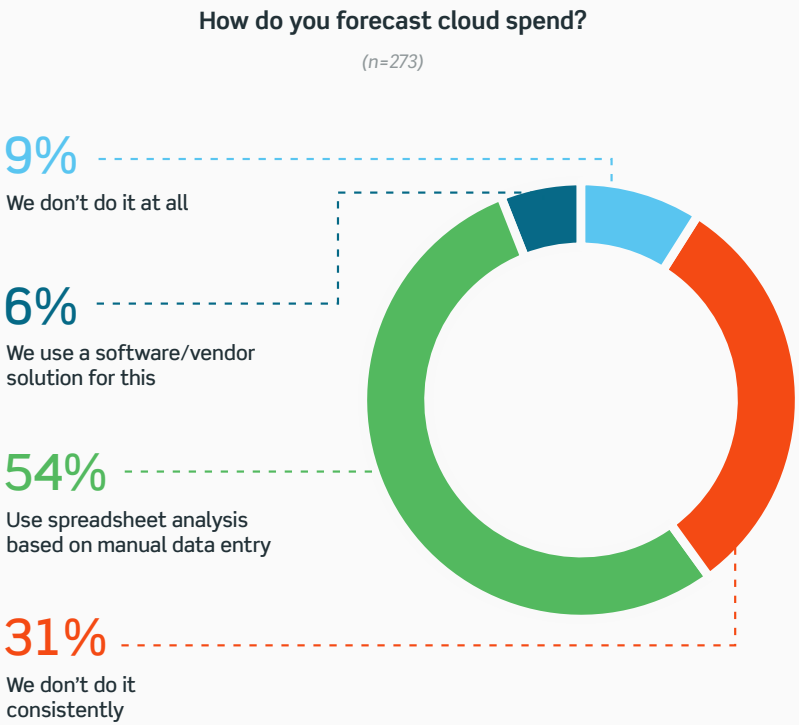

Figure 14: Most companies don't focus properly on forecasting their cloud expenditures.

- **Higher than expected costs.** There are countless stories of companies moving workloads into the cloud and then receiving surprise invoices from their provider. Often, these issues can be avoided with the right level of visibility and proper workload configuration, but many do not have the right tools or skills to make that assessment.

CLOUD SPEND FORECASTING METHODS

Cloud cost opinions have dramatically shifted over the years. Early on, people looked at cloud as a panacea in terms of costs. However, as time has passed and companies have become more savvy about cloud costs, many have begun to understand that cost savings may not be there; but, as shown in **Figure 5**, organizations are seeing benefits in terms of agility.

Cost is vitally important, but you wouldn't know it by the way respondents forecast their cloud spend, a task that will be increasingly important as cloud adoption continues to grow. Forty percent (40%) of respondents either don't forecast at all (9%) or don't do it consistently (31%).

> Cost is vitally important, but you wouldn't know it by the way respondents forecast their cloud spend, a task that will be increasingly important as cloud adoption continues to grow.

For those that do some level of forecasting, 54% use a spreadsheet or other manual tool; just 6% use a purpose-built software tool or vendor solution to forecast their spend. There is clearly room for improvement (**Figure 14**).

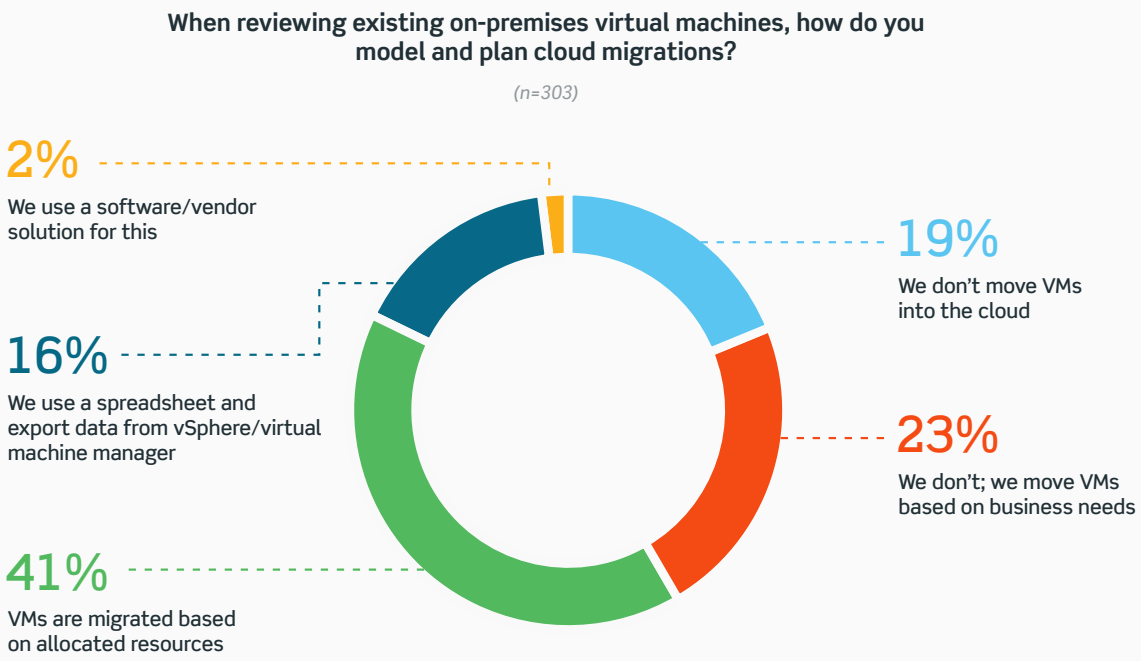

Figure 15: Cloud spend forecasting practices.

VIRTUAL MACHINE CLOUD MIGRATION ANALYSIS

Careful analysis of workload migrations from on-premises into the cloud is important. There are a variety of ways by which organizations model and plan these migrations, with a breakdown shown in **Figure 15**. Note that 19% of organizations don't move virtual machines into the cloud at all, which makes sense given that 81% of companies have adopted cloud.

> Only 16% use a spreadsheet with data gathered from virtualization management tools, and just 2% use a dedicated tool for this purpose. As is the case elsewhere, there's a lot of opportunity for increased maturity.

Beyond that, from a formal modeling perspective, 23% don't really model and plan migrations. Instead, they react to current business needs. For an additional 41%, it's a simple resource-based decision: which environment has the right resources for the application? Only 16% use a spreadsheet with data gathered from virtualization management tools, and just 2% use a dedicated tool for this purpose. As is the case elsewhere, there's a lot of opportunity for increased maturity.

Do you have detailed insight into your on-premises data center spend?

(n=303)

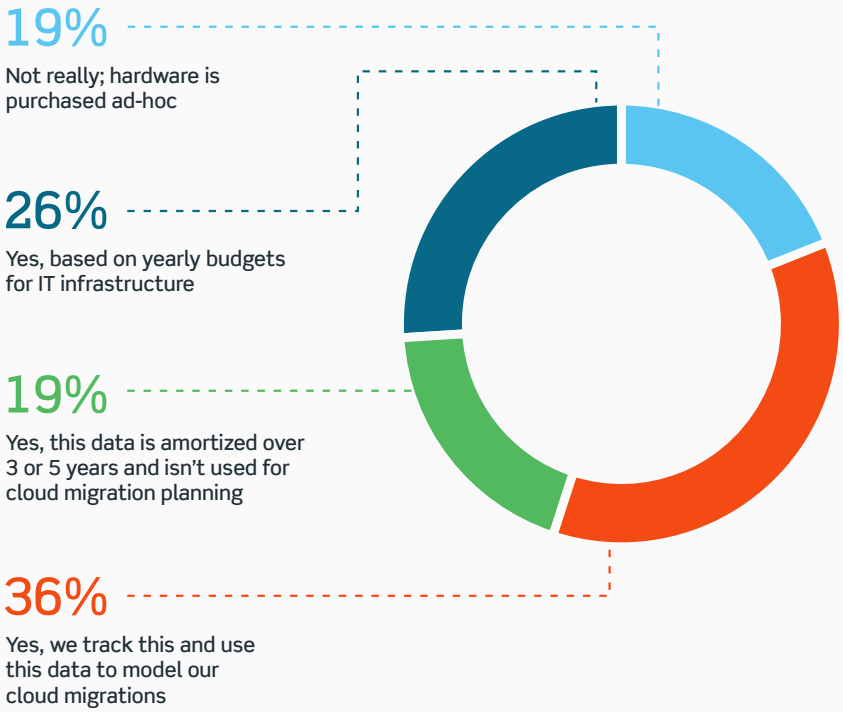

Figure 16: On-premises data center spending practices.

ON-PREMISES DATA CENTER SPEND ANALYSIS

The cloud is clearly a priority for many, but that doesn't mean on-premises data centers are being ignored. To the contrary, they're more important than ever as companies rush to deploy hybrid clouds to support digital transformation efforts. And it's important that the spend associated with these local environments be well understood.

The good news is that most organizations—81%—do a good job tracking their data center spend. Just 19% say that they purchase equipment ad hoc. An additional 26% have established annual budgets for IT infrastructure, and they stick to those as their sole planning methodology. In reality, these numbers are likely far higher (as you'll see), since most organizations do some level of annual budgeting; but this was a single response question that requested that the respondent make a clear choice.

The data from **Figure 16** shows that 36% of organizations track their local spend and use some of this information to model cloud migrations to determine where best to operate workloads from a cost perspective. An additional 19% amortize their data center spend over three to five years, but don't use this data for cloud migration planning.

Part 2: The Role of Public Cloud in Backup and Recovery

Backup and disaster recovery are common cloud use cases, and often the first services that make their way out of the local data center and into the cloud. There are a number of reasons:

- **Using the cloud negates the need to stand up secondary data centers for disaster recovery.** It's a significant cost savings with serious time-to-value and CapEx advantages, not to mention the ongoing operational benefits, including reduced maintenance, fewer patches, and a whole lot more.
- **Using the public cloud as a part of a backup or disaster recovery strategy allows organizations to more quickly adopt such services since the cloud**, unlike a secondary data center with adequate infrastructure, is immediately available.
- **The natural geographic segregation from on-premises data centers** provides organizations, particularly small ones with only regional footprints, the opportunity to better protect their data.
- **Cloud-based disaster recovery often allows organizations that suffer disaster to quickly recover in the cloud,** even if the primary data center is no longer available.

DISASTER RECOVERY AND THE CLOUD

Most respondent organizations—60%—are using the public cloud in some way as a part of their overall disaster recovery strategy (**Figure 17**). Bear in mind that this doesn't mean that respondents have complete mirror architectures in the cloud; instead, they may only protect a subset of their applications.

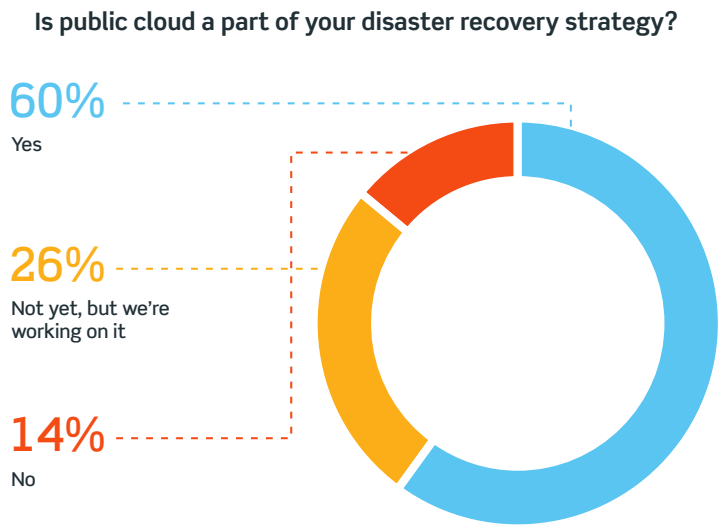

Figure 17: Most organizations use the cloud as part of their disaster recovery strategy.

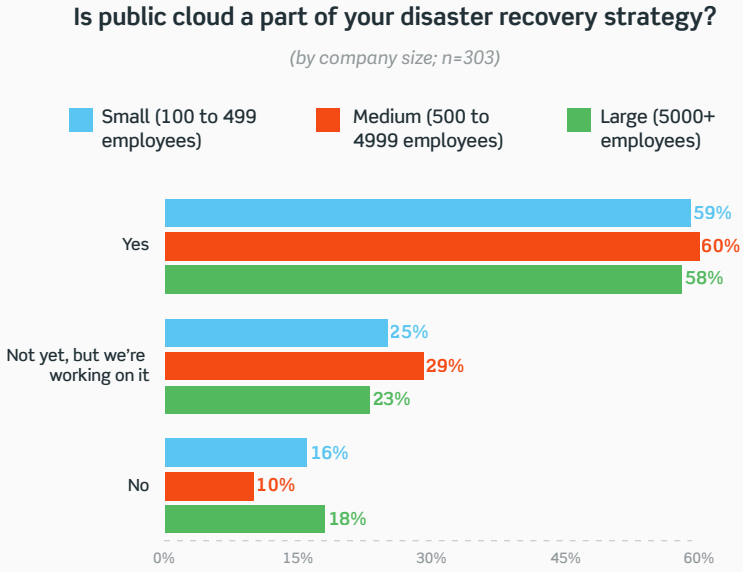

Figure 18: Organizations in which cloud is a part of the disaster recovery strategy, broken down by company size.

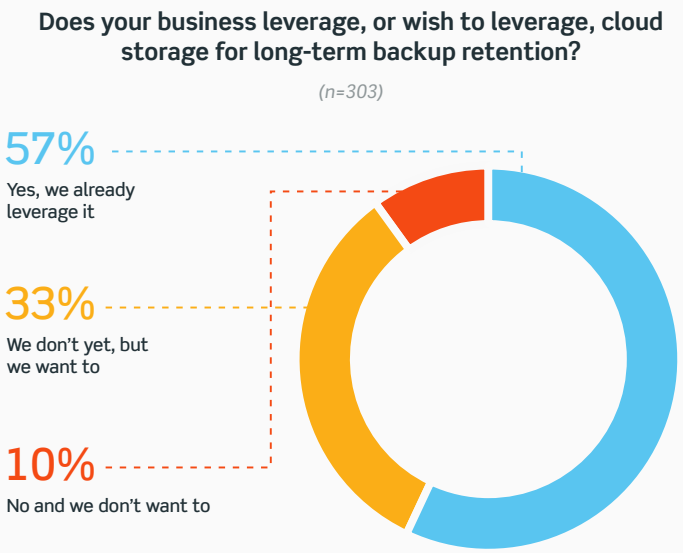

Figure 19: Most organizations see the cloud as a long-term backup retention repository.

An additional 26% of respondents are at work on such plans, while just 14% indicate that they currently have no presence or active projects toward establishing a disaster recovery presence in the cloud.

Remember that this question is very broad. It may mean that the respondent organization is backing up directly to the cloud. It may mean that the respondent has moved a mission-critical application to a SaaS provider. It could mean that a single application bursts to the cloud and runs there in the event of a failure of the primary data center.

Mirroring other findings in this survey, mid-size organizations are ahead of small and large companies in their use of public cloud, but not by a lot (**Figure 18**). For actual production use, 60% of mid-size companies indicate that public cloud is a part of a disaster recovery strategy. The difference comes when you consider companies that are working on getting to this state. Here, mid-size companies are ahead of others by a handy margin (29% for mid-sized vs. 25% for small and 23% for large companies).

CLOUD STORAGE FOR LONG-TERM BACKUP RETENTION

Another early use cases around cloud involved data protection. The cloud's limitless resources and geographic diversity make it an attractive option for long-term backup retention. In fact, it's such a popular option that just 10% of survey respondents say that they have no interest in it. In other words, 90% are either doing it or want to do it.

Of those 90%, 33% have yet to take the plunge. They don't currently use the cloud for backup retention, but they want to (**Figure 19**). The remainder—57%—are doing something around cloud with regard to backup retention. It could be as simple as copying backup files to S3, or as comprehensive as having implemented a direct-to-cloud backup solution that uses the cloud for its long-term repository.

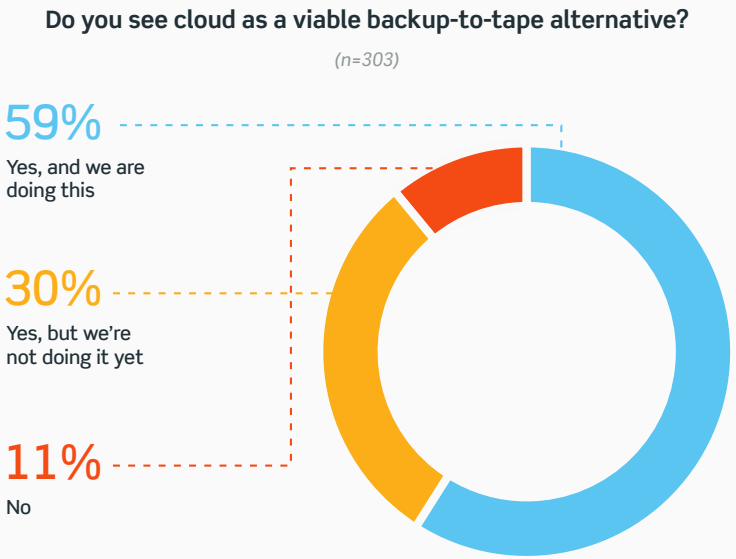

Figure 20: A strong majority views the cloud as an alternative to tape.

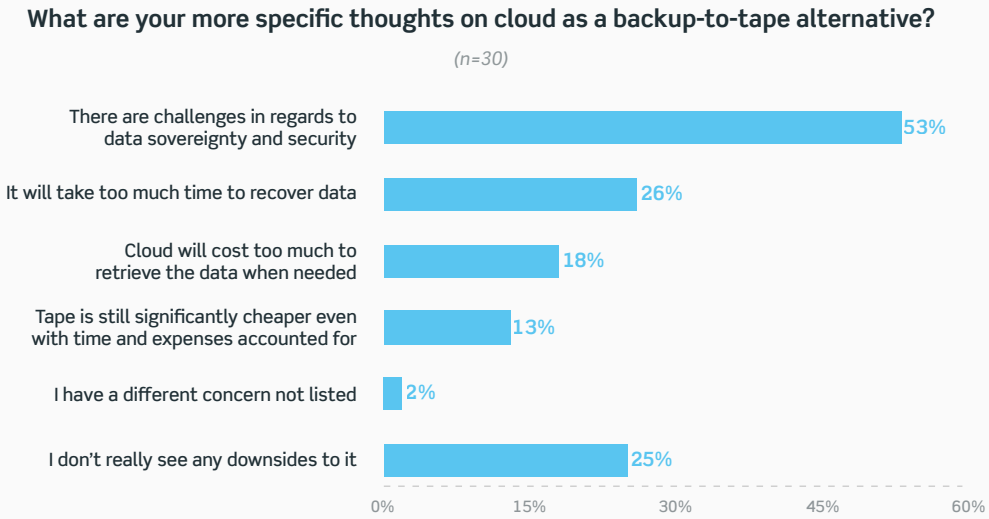

Figure 21: Concerns around use of cloud as compared to tape.

THE VIABILITY OF CLOUD AS A BACKUP-TO-TAPE ALTERNATIVE

Anyone in this industry long enough has heard the cry: "Tape is dead!"

Except it's not. A lot of companies still use tape, and for a lot of good reasons... and some bad ones, of course. On the plus side, tape is cheap when compared with most other data storage solutions, and it's well understood. On the negative side, tape can be flaky; you'd better hope you don't have a bad backup tape when you need to do a recovery. And you may need to retain old tape hardware even when you upgrade so that you can continue to recover data in the event of a data loss incident.

Given the downsides of tape and the popularity of cloud as a backup target, it's no surprise that 59% of respondents see cloud as a viable replacement for tape, and have started to move in this direction. Another 30% have yet to take the first step but want to, and just 10% indicate that they don't see cloud as a viable replacement for their trusty tape drives (**Figure 20**).

For those that aren't interested in cloud as an alternative to tape, there are a lot of reasons (**Figure 21**). More than half—53%—of respondents worry about data sovereignty and security in the cloud. Others (26%) worry about recovery time with the cloud, which is a valid concern if the data protection strategy doesn't include standing up workloads in the cloud in the event of a disaster.

Still others worry about the cost of cloud in terms of how much it will cost to retrieve data, and also indicate that the cost of tape is very low. Twenty-five percent (25%) indicate that they don't see a lot of downside to the cloud, but they just don't seem to have an interest in making the jump.

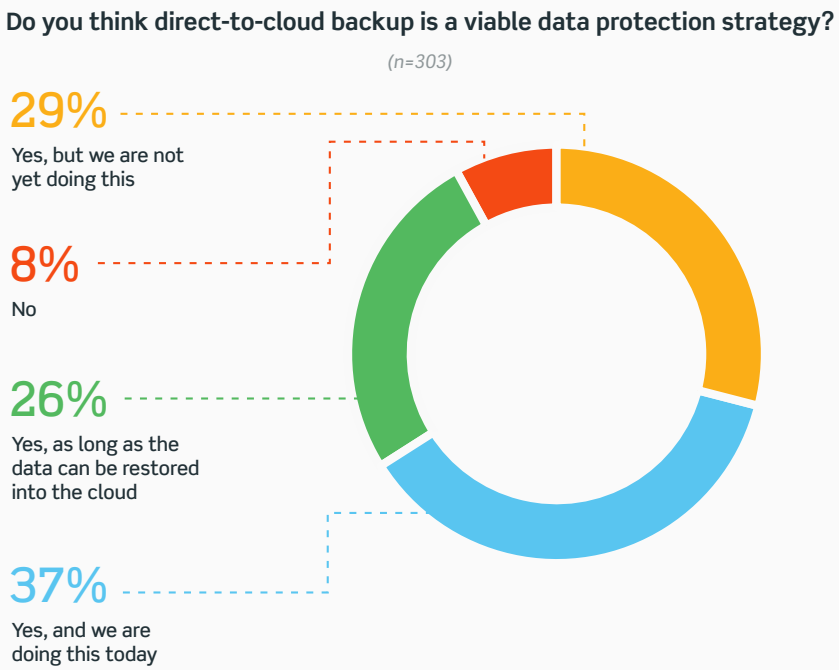

Figure 22: The viability of direct-to-cloud backup.

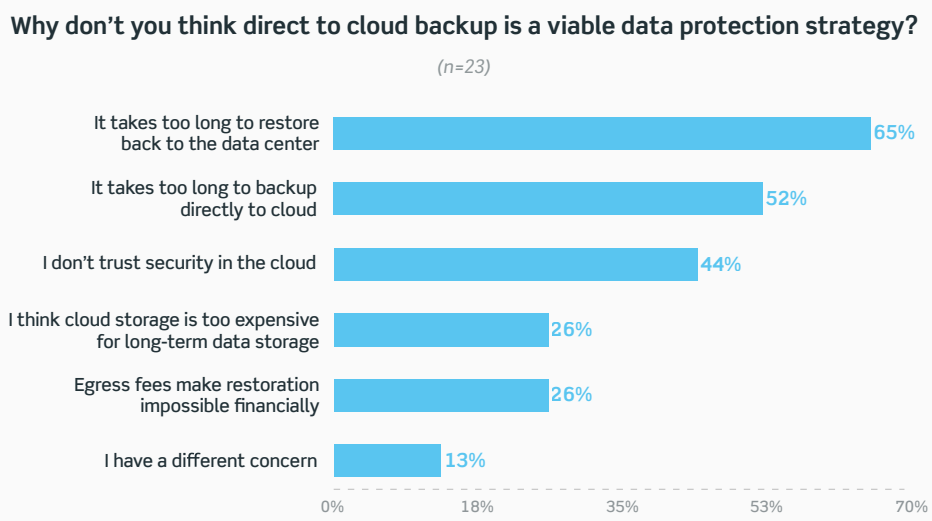

Figure 23: Concerns around use of cloud as compared to tape.

THE VIABILITY OF DIRECT-TO-CLOUD BACKUP AS A DATA PROTECTION STRATEGY

You don't need to back up data onsite. You can skip this step and back up data right to the cloud. This can be a viable strategy that provides a great deal of simplicity, but also includes challenges, including bandwidth needs and restore latency.

Among respondents, 26% agree with this assertion (**Figure 22**). They're comfortable backing data up to the cloud directly, but only if those workloads can then be restored into the cloud, too. With this capability, an organization can recover workloads immediately and avoid high recovery time objectives (RTOs) for their recovery process. They can get the benefits of the cloud and very low RTOs, and the perform on-premises recovery at a pace that makes sense for them.

Other respondents are already backing up to cloud (37%), while still others agree that this is a viable strategy but aren't doing it yet (29%). Just 8% of respondents indicate that they don't believe that direct-to-cloud is a viable data protection strategy.

There are a number of reasons that the disbelievers in direct-to-cloud backup feel the way they do. It's a small number—23 in this survey—but 65% of them have serious concerns around the time it takes to recover data back to the data center (**Figure 23**). Moreover, 52% have reservations around how long it takes to back data up to the cloud. Forty-four percent (44%) have little faith in cloud security.

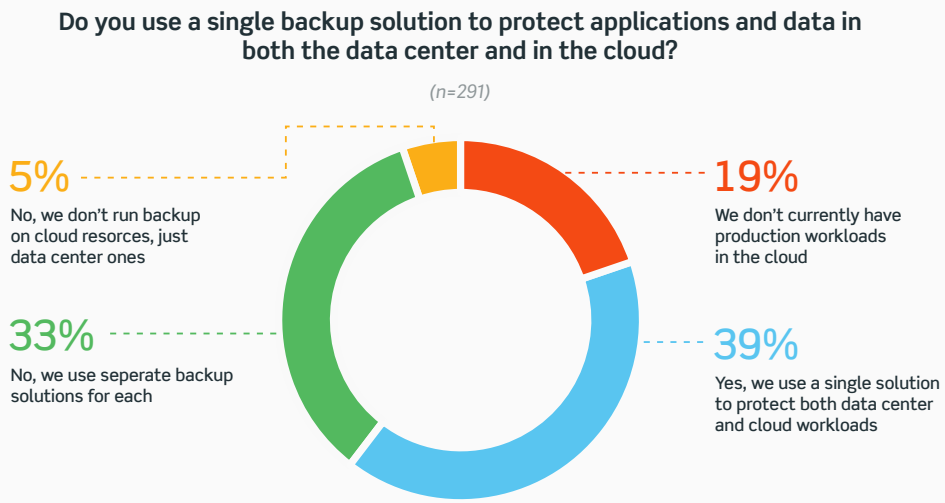

Figure 24: Application backup strategy.

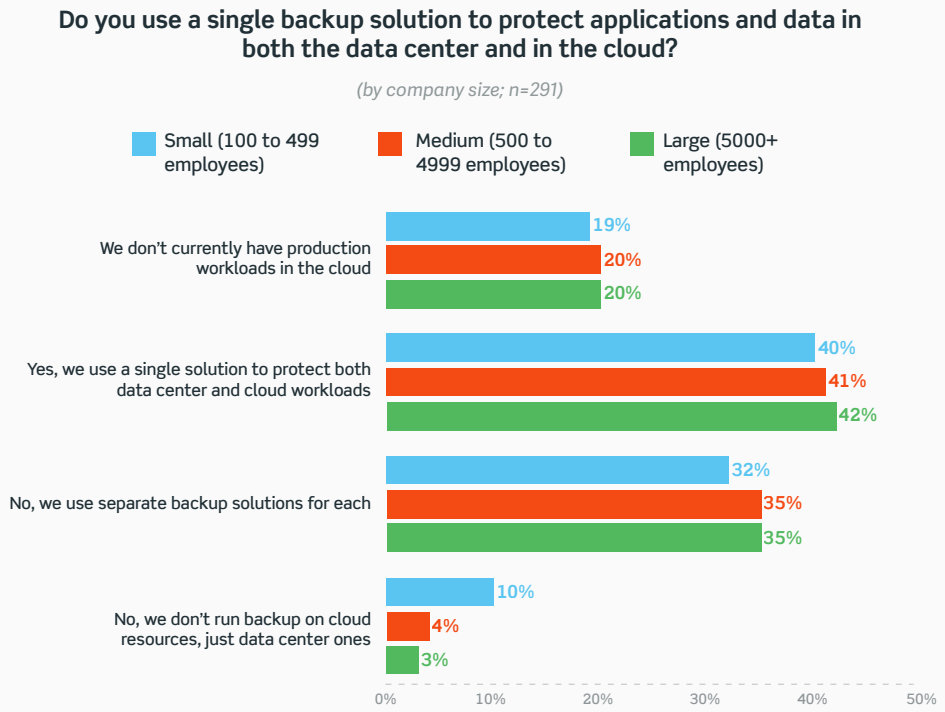

Figure 25: Application backup strategy by company size.

And then there are costs. On this front, 26% believe that cloud storage is too expensive for long-term data storage and 26% are (rightly) worried about egress fees that would be imposed during a recovery operation. Remember, taking data out of the cloud imposes a fee for each gigabyte. When you perform a full recovery, that's a lot of bytes!

BACKUP APPLICATION STRATEGY

Some companies use different applications depending on where an application lives. If it's onsite, one tool supports it; if it's in the cloud, it's protected another way. Among respondents, 19% don't need to worry about the cloud side since they have no applications in the cloud. For the rest, there are varying procedures (**Figure 24**).

Thirty-nine percent (39%) of respondents use a single solution for both on-premises and cloud applications. That's the ideal scenario, since it maintains a level of simplicity in backup and recovery. But not everyone is there yet; 33% use different tools depending on where applications reside. And, finally, 5% say they're not providing separate backup for cloud applications, instead protecting only those workloads running in the data center.

When considered by company size, smaller companies are the outlier—to a point (**Figure 25**). Here, you can see that 10% of small companies don't back up cloud resources, compared with just 4% and 3% of medium-size and large ones, respectively. Further, small companies are somewhat less likely to use a separate backup application for their on-premises and cloud resources.

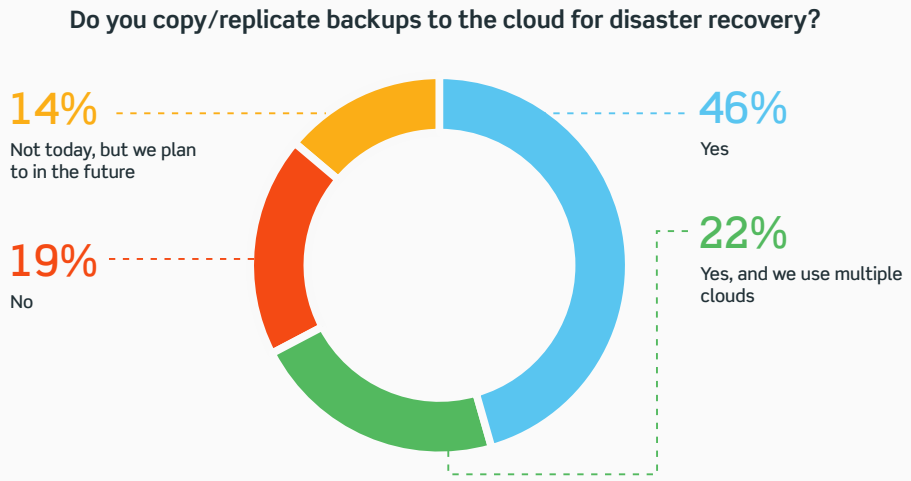

Figure 26: Backup replication to the cloud as a part of disaster recovery.

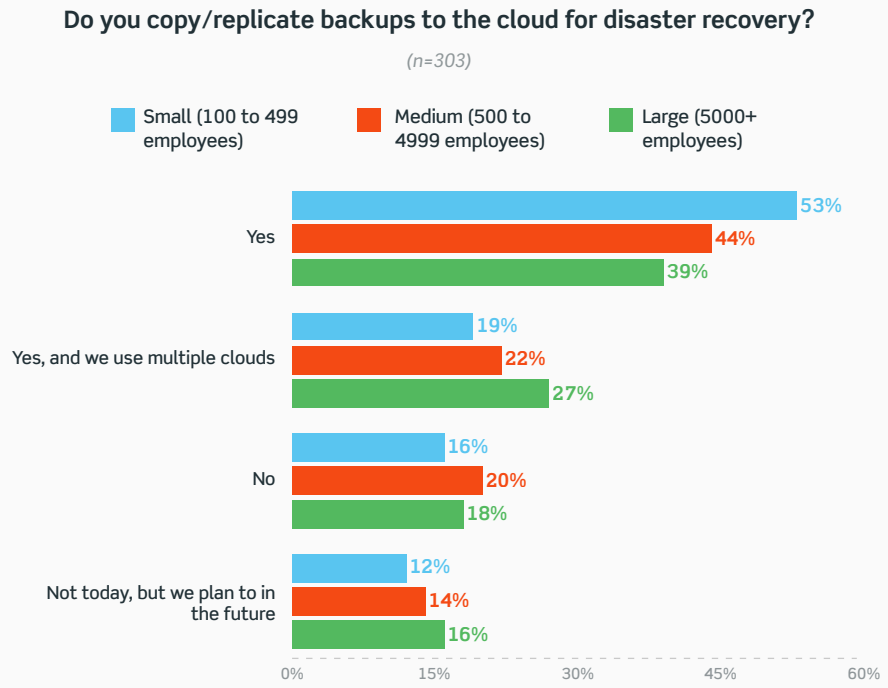

Figure 27: Backup replication to cloud as a part of disaster recovery, by company size.

DISASTER RECOVERY DATA REPLICATION STRATEGY

In keeping with the theme of data protection and the cloud, there was a desire to learn about how respondent organizations place copies of backups. To that end, they were asked about whether they copy their backups to the cloud (**Figure 26**). Forty-six percent (46%) of respondents are doing this, but we caution the reader that the definition of "cloud" in this question could be construed as either public or private cloud.

An additional 22% indicate that they replicate to multiple places. Fourteen percent (14%) say that they aren't doing this now, but they want to; finally, 19% aren't replicating to the cloud and don't currently have such plans.

We question the high number of people doing this today, and believe that public and private clouds were considered in how this question was answered.

Smaller companies that want to protect their data are more likely than others to copy their backups to the cloud (53%). However, large companies are more likely to copy such backups to multiple clouds (27%). See **Figure 27**.

About Quest Data Protection Solutions

More than 70,000 customers worldwide depend on Quest to optimize IT infrastructure, plan cloud migrations and maximize system, application and data protection and availability in the datacenter and in the cloud. With Quest, you can tear down the IT silos that are keeping valuable resources locked away and putting your data at risk.

To learn how Quest will help you get the most from the cloud, visit: **https://www.quest.com/solutions/cloud-management/**