




ConversationalGeek®

Conversational Azure Backup Best Practices

By **Brien Posey** (Microsoft MVP, Commercial Scientist Astronaut Candidate)



**In this
book, you
will learn:**

- Why you need to focus your efforts on backup above all
- How performance, cost, and security play a role in your backup strategy
- 12 best practices to ensure your Azure investment is protected

4th
Edition

Sponsored by
veeam

Sponsored by Veeam

Veeam, the #1 global market leader in data protection and ransomware recovery, is on a mission to help every organization not just bounce back from a data outage or loss but bounce forward. With Veeam, organizations achieve radical resilience through data security, data recovery, and data freedom for their hybrid cloud.

The Veeam Data Platform delivers a single solution for cloud, virtual, physical, SaaS, and Kubernetes environments that gives IT and security leaders peace of mind that their apps and data are protected and always available. Headquartered in Columbus, Ohio, with offices in more than 30 countries, Veeam protects over 450,000 customers worldwide, who trust Veeam to keep their businesses running.



To learn more visit
www.veeam.com

Conversational Azure Best Practices (4th Edition)

Brien Posey

© 2024 Conversational Geek



ConversationalGeek®

Conversational Azure Backup Best Practices (4th Edition)

Published by Conversational Geek® Inc.

www.conversationalgeek.com

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Trademarks

Conversational Geek, the Conversational Geek logo and J. the Geek are trademarks of Conversational Geek®. All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. We cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or programs accompanying it.

Additional Information

For general information on our other products and services, or how to create a custom Conversational Geek book for your business or organization, please visit our website at ConversationalGeek.com

Publisher Acknowledgments

All of the folks responsible for the creation of this guide:

Authors:	Brien Posey
Project Editor:	Nick Cavallancia
Copy Editor:	Hope Crocker
Content Reviewers:	Shana Brewer Sam Nicholls

Note from the Author

Organizations leveraging Microsoft Azure often focus on the “using the cloud” part and forget completely about the “Oops... I put my data in the cloud” part. While critical workloads tend to get more focus than, say, your data in Microsoft 365, there is usually some question about how to best protect everything you’ve put into Azure.

This book focuses on what’s truly important – getting back to the basics of backup and using best practices to ensure the backup and data protection strategies you implement are truly effective.

We realize your use of Azure isn’t casual in nature; so the content in this book seeks to help you crystalize your backup concentration on the things that will keep those critical workloads, applications, and data recoverable should the situation arise.

Brien Posey



The “Conversational” Method

We have two objectives when we create a “Conversational” book: First, to make sure it’s written in a conversational tone so that it’s fun and easy to read. Second, to make sure you, the reader, can immediately take what you read and include it into your own conversations (personal or business-focused) with confidence.

These books are meant to increase your understanding of the subject. Terminology, conceptual ideas, trends in the market, and even fringe subject matter are brought together to ensure you can engage your customer, team, co-worker, friend, and even the know-it-all Best Buy geek on a level playing field.

“Geek in the Mirror” Boxes

We infuse humor into our books through both cartoons and light banter from the author. When you see one of these boxes it’s the author stepping outside the dialog to speak directly to you. It might be an anecdote, it might be a personal experience or gut reaction and analysis, it might just be a sarcastic quip, but these “geek in the mirror” boxes are not to be skipped.



Greetings. Within these boxes I can share just about anything on the subject at hand.

Read 'em!

Notes from the Authors

Organizations leveraging Microsoft Azure often focus on the “using the cloud” part and forget completely about the “Oops... I put my data in the cloud” part. While critical workloads tend to get more focus than, say, your data in Microsoft 365, there is usually some question about how to best protect everything you’ve put into Azure.

This book focuses on what’s truly important – getting back to the basics of backup and using best practices to ensure the backup and data protection strategies you implement are truly effective.

We realize your use of Azure isn’t casual in nature; so the content in this book seeks to help you crystalize your backup concentration on the things that will keep those critical workloads, applications, and data recoverable should the situation arise.

Brien Posey



Backups: What Are We Really Talking About Here?



So, how am I supposed to back THAT up???

The IT industry, perhaps even more so than any other industry, seems to love using buzz words. Sadly, backups are by no means immune from this trend.

The problem with the way that buzzwords and technical jargon get thrown around is that they can cause terms to become ambiguous, or to lose their original meaning. Take the term *Disaster Recovery* for example. In some circles, disaster recovery is synonymous with backup and restoration. For others, disaster

recovery refers to having the ability to failover a workload to the cloud in the event that an organization's data center becomes incapacitated.

In an effort to cut through the ambiguity, we want to say up front that this book focuses squarely on backing up resources that are being hosted in the Azure cloud. There might be brief mentions of other clouds, assets that reside on-premises, or some of the more exotic protective mechanisms, but backup and restoration are really at the heart of this book. In fact, our goal is to take a back-to-basics approach and have a very frank discussion of backup best practices and shortcomings as they relate to the Azure cloud.



Several years back, IBM poked fun at the IT industry's pervasive use of buzz words when they made a commercial featuring a group of conference attendees playing Buzzword Bingo (goto.cg/BuzzBingo).

The One Buzzword that Matters

Even though we have tried to stay away from using buzzwords as much as we can, there is one that we have to talk about. That buzzword, or buzz phrase rather, is "location agnostic".

For years, public cloud providers were engaged in an all-out marketing assault in an effort to get their customers to move anything and everything to the cloud. The marketing message was simple: "Cloud good, data center bad". We even once heard a speaker at a conference say that if by now you are not operating 100% in the cloud, then you are a dinosaur.

Thankfully, cloud providers have backed off, at least somewhat, and seem to have accepted the idea that their customers are going to be operating hybrid environments. Some workloads are best suited for the cloud, while others are probably a better fit for on-premises data centers. This is where the phrase “location agnostic” comes into play.

Location agnostic means that a workload could be running practically anywhere. Its geographic location is becoming increasingly irrelevant.

The reason the idea of location agnostics is so important is that it makes it possible to begin once again managing our resources in a cohesive way. Back when everything existed on-premises, there were universally accepted best practices pertaining to how those resources should be managed and protected.

However, as the public cloud started to gain mainstream acceptance, people quickly realized that they had to adopt two different ways of doing things. One set of best practices – and tools – applied to resources that were running on-premises, and an entirely different set applied to cloud resources. Now the industry is starting to get back to the point where a common set of best practices can be applied regardless of where an organization’s IT resources physically reside.

The one big caveat to this, however, is that these best practices aren’t necessarily the same as those that were previously used. Remember, they were intended for an either/or situation. Either a workload was running in the cloud, or it was on-premises.

12 Best Practices for Backing Up Resources Residing in Microsoft Azure

Because the best practices for backup and restoration have evolved significantly in recent years, we wanted to talk about some of these, especially as they relate to protecting resources that are stored in the Azure cloud.

1. Actually Back Up Your Data

This one might sound like a bit of a no-brainer, but it's incredibly important. You need to back up your data. Yes, seriously. This is the number one best practice.

So why on Earth are we telling you to be sure to back up your data? After all, the simple fact that you're reading this book means that you probably already know that backups are important.

The reason why we're including the need for data backups among the other best practices is because it directly ties into a couple of the things that we have already talked about. Namely, the cloud providers marketing efforts, and the concept of location agnostics. Let us explain.

If you were to ask most people what message the big cloud providers' marketing teams were trying to convey 10 years ago, they would probably tell you that the message was that the cloud makes everything cheaper and easier.

In conveying this message, cloud providers would often stress the idea that operating in the public cloud is far more convenient than running workloads on-premises. This was due to the fact that the cloud provider handles all of those maintenance tasks that you, as an administrator, don't want to be bothered with.

Depending on the type of workload, these maintenance tasks might include things like keeping the hardware running, planning for the organization's future capacity requirements, or installing security updates.

The problem with this is somewhere along the line, people started getting the perception that cloud providers handle all of the maintenance for you. Since backups can be thought of as a maintenance task, it probably seemed only natural to assume that the cloud providers were handling backups on their customers' behalf. As it turned out though, they weren't.

Most cloud providers, including Microsoft, don't back up their customers' data on their behalf. Microsoft uses a shared responsibility model. This states that Microsoft is responsible for ensuring the wellbeing of the Azure infrastructure, but customers are responsible for protecting their own data.

Although the shared responsibility model has been around for a while, it is still very relevant today. In fact, Microsoft is releasing its own backup tool for Microsoft 365, thereby underscoring the idea that it is important for organizations to back up their own data. After all, if Microsoft were performing backups for their customers, then there would be no need for Microsoft to create a backup tool. The same can be said for Azure Backup.

The bottom line is that the need for backups doesn't go away just because a resource resides in the cloud. Data loss can, and certainly still does happen in cloud environments.

This all ties back to the concept of location agnostics. Your data needs to be backed up, regardless of where that data resides.

2. Use the Right Tool for the Job

A second-best practice is to make sure that you are using the right tool for the job. Your backup software needs to be aligned with your backup requirements.

We will be the first to admit that this seems like yet another odd thing to include in a list of best practices. At best, it seems a little bit over simplistic. At worst, it sounds like a pitch designed to sell backup software. In reality though, there are three very important reasons why we are including the concept of using the right tool for the job among our list of best practices for backing up resources in Azure.

The first reason why it is so important to use the right tool for the job is that backup software that was designed for on-premises use might not do such a good job backing up resources in the cloud. That isn't to say that a legacy backup product can't

back up resources in Azure. We have seen it done. What we are saying though is that trying to repurpose an on-premises backup product for use in the cloud can be problematic. In some cases, the software may work, but fail to provide an optimal experience (for example, it might not utilize native Azure snapshots). In other situations, the software might partially work, but leave you with gaps in your coverage.



It's important to be completely confident that your backup solution is going to work when needed. There is no worse feeling than thinking that your backup solution was doing its job, only to discover during a crisis that your data wasn't actually being backed up properly.

A second reason why it is so important to use the right tool for the job is that a lot of the legacy backup products require the use of agents. As IT pros, we all know that agents can be a pain. They can be difficult to deploy and manage, and sometimes they seem to stop working for no reason. We have even seen situations in which someone re-imaged a computer and forgot to reinstall the backup agent when they were done, leaving that machine unprotected.

Even if you put all of these potential annoyances aside, there is a more important reason for staying away from the use of an agent-based approach. Backup agents are often designed to communicate across an obscure TCP or UDP port. That might not be a problem if you are using that agent exclusively on-premises, but there is a good chance that the agent's port requirements may keep it from functioning across cloud boundaries.

Finally, reason number three has to do with support. As a general rule, software vendors support their software only when it is used in accordance with their recommendations. If a backup

vendor designed a particular piece of software to be used on-premises, then they probably don't support using that software to protect Azure resources, even if the software seems to work with Azure.

As a matter of self-preservation, you never want to put yourself in a situation where you might one day end up having to explain to your boss that you can't recover the organization's data because you were using backup software in an unsupported manner. The lesson here is to use a backup product that avoids the use of agents, and that is designed for use with Azure.



Using backup software that was designed for on-premises use to protect Azure resources may in fact be a violation of the software's licensing agreement.

Of course, this raises the question of how you can tell if a backup product is designed for Azure, or if it just happens to work with Azure. Generally speaking, the backup vendor's website should tell you if their solution is designed for Azure. However, there is an easier way. If you want to make sure that your backup product of choice truly is designed to work with Azure, then check to see if it is available in the Azure Marketplace. Not only will acquiring a backup application from the Azure Marketplace guarantee Azure compatibility, it can also simplify the deployment process.

3. Run Your Backup Software in a Way that Makes Sense for Your Organization

There are two main options for backing up Azure resources. The first option is to deploy a backup server, install backup software,

and then configure that software based on your backup requirements.

The other high-level option is to use backup software as a managed service. This approach is often referred to as Backup as a Service or BaaS. The advantage to using a BaaS provider is that you don't have to worry about maintaining the backup infrastructure. The BaaS provider handles all of that for you. The provider makes sure that backup servers remain online, and they also handle tasks such as patch management and capacity planning. The customer needs only to configure the software to back up their Azure resources.

As easy and convenient as BaaS may be, it isn't a good fit for every situation. Because BaaS is offered as a managed service, the software can be somewhat rigid. Those who need super granular control of their backups or who need control of the underlying backup infrastructure might be better off using a more traditional backup.

Interestingly, large organizations sometimes find that BaaS is a good fit for some of their workloads, but not others. While there is no rule saying that you can't mix backup types, doing so can lead to backup silos and it can complicate backup management. The easiest way to avoid these types of problems while still capitalizing on the advantages of using a mixture of backup types is to find a vendor who offers both traditional backup software and managed backup services. Doing so will likely allow you to have the best of both worlds, without forcing you to adopt two completely different approaches to backup management.

4. Use Automation to Work Smarter, Not Harder

The third best practice for protecting your resources within the Azure cloud is to use backup automation whenever you can.

Earlier, we brought up the idea that one of the things that initially made public clouds like Azure so appealing was that the cloud providers handle many of the maintenance tasks that IT

pros would have otherwise had to do themselves. While it is true that there are still some maintenance tasks that are left to the IT professional, it is possible to narrow the gap by leveraging automation.

Automation is actually a really good fit for backup-related tasks. In fact, it's one of those things that has been used in one way or another for what seems like forever. Consider for example, that the backup software used in the 90s was capable of running a backup job at a scheduled time. This was a form of automation.

But job scheduling isn't the only way that automation can be used to assist with the backup process. Modern cloud backup applications make it possible to automate things like VM snapshot creation, and intelligent tiering of both snapshots and backups across different cloud storage tiers based on organizations retention policies.

Admittedly, it is tempting to think of backup automation as a convenience feature. Some may consider it to be one of those things that is nice to have, but not necessarily essential. However, backup automation can help to improve backup reliability by making sure that backups are created and maintained in a predictable – and compliant – way.

Automation also helps IT pros to be more effective than might otherwise be possible. We all know that this age of ever-shrinking IT budgets has resulted in staffing resources being stretched thin. Automating mundane tasks such as those related to data protection helps to decrease the administrative workload, thereby giving IT professionals a bit more free time that they can use to work on other things.

Another benefit to using backup automation is that doing so can help to eliminate the possibility of human error. If for example, the backup software is automating snapshot creation and handling your snapshot lifecycles based on a retention policy,

then it lessens the chances that a human error will result in a data loss event.

Simply put, set-and-forget simplicity through policy-based automation is a key capability to look for in a backup solution. It can help to reduce costs and eliminate human error while also helping backup operators to make better use of their time.

5. Be Aware of Cloud Costs

Early on, public cloud providers relentlessly marketed themselves as being the cheap alternative to running business workloads on-premises. Over time though, many of us have learned the hard lesson that operating in the cloud can be just as expensive (if not more so) as keeping workloads running in-house. This isn't to say that you can't save money by hosting a workload in the cloud. Under the right circumstances, migrating a workload to Microsoft Azure can yield a significant cost saving.

The key to realizing actual cost savings by running a workload in the cloud is to understand that cloud providers such as Microsoft, Amazon, and others are not running a charity. Like any other business, their goal is to make money. Not surprisingly, it can be quite expensive to host a workload in the cloud.

If you want to avoid being surprised by the costs you'll incur, then you will need to spend a little bit of time learning how Microsoft bills its Azure customers.

A deep dive into Azure billing is well beyond the scope of this book, but we do want to take a moment and talk about data egress fees. These fees can have an enormous impact on the cost of your Azure backups.

Simply put, data egress fees are charges for data that leaves the cloud. These fees are not unique to Microsoft. Although the data egress fee amounts vary from one cloud provider to the next, most cloud providers do charge their customers a fee any time data leaves *their* cloud. Presumably these fees were put into

place as a way of discouraging customers from moving their cloud workloads to another provider or moving those workloads back on-premises. Regardless of intent, data egress fees come into play when an organization creates or restores a backup.

Designing your environment with backups as a consideration can help avoid data egress fees. Utilizing direct connect networks, keeping data within a region, etc. – in essence, staying within one cloud provider – will lower the risk of incurring egress fees.

In case you're wondering, data egress fees vary widely based on the provider and on the volume of outbound data. As of the time that this book was written, Microsoft allows up to 100 GB of outbound data per month before any Azure data egress fees kick in. Once the 5 GB threshold is exceeded, the price for outbound data is \$0.087 per gigabyte for the first 10 TB of data each month (Microsoft offers a discount for transfers larger than 10 TB). The pricing structure is documented at gogo.cg/3JCb43k.

Admittedly, \$0.087 per gigabyte doesn't sound like all that much money, but let's do some math. Suppose for a moment that you needed to restore 1 TB of data from Azure to an on-premises VM. One terabyte is equal to 1024 gigabytes. Just to make things fair, let's assume that 100 gigabytes of data are going to be restored for free. This means that you would be billed for transferring 924 GB of data. At \$0.087 per gigabyte, that works out to \$80.39.

While an \$80.39 charge probably isn't going to deplete your entire IT budget, remember that this figure is based on restoring a single terabyte of data. A large-scale data restoration operation can cost considerably more. Never mind the fact that there may be other charges related to the operation, such as fees related to storage IOPS.

Reducing backup-related costs involves much more than just avoiding unnecessary data egress fees. The frequency with which backups are created has an impact on cost, as do other

factors such as storage tier selection and data redundancy. One of the best ways to help make Azure backup costs more predictable is to leverage a cost calculator. There are various Azure cost calculators available for download, but some of the better Azure backup tools include an integrated cost calculator that can not only estimate your backup costs, but also help you to figure out how to minimize those costs.

6. Use Storage Tiers Effectively

There are costs tied to the use of Azure storage. If you are going to be storing your backup data in the Azure cloud, then it is important to factor those costs into your backup strategy.

Aside from the previously discussed data egress fees, there are two main costs that you need to be aware of. First, there are capacity related costs. When you store data in the Azure cloud, you aren't purchasing storage, you're renting it. As such, Microsoft charges a fee each month for every gigabyte of storage that is in use. Suppose for example, that you were to write a 100 GB file to the Azure cloud. You would be charged for 100 GB of storage each month, for as long as the file remains in the cloud.

The other type of charge to be aware of is a usage charge. Microsoft bills its customers for storage IOPS related to reading or writing data. Hence, the more frequently a piece of data is accessed, the higher the cost of keeping that data in Azure.

One thing to keep in mind is that Microsoft Azure storage is not one size fits all. There are actually several different types of Azure storage. Azure block blob storage, for example, is classified into tiers that include the Premium Performance tier, Hot tier, Cool tier, and Archive tier. Each of these tiers has its own unique performance characteristics, and its own pricing structure. The general rule is the "colder" the storage, the lower the cost/GB and the slower the restore time. Even so, it's much less expensive than maintaining on-premises storage arrays. The best

way to keep your cloud storage costs in check is to make use of the storage tiers, balancing cost and performance.

As you can no doubt imagine, having so many different types of storage at your disposal can make it tough to figure out the best type of storage for a particular use case. This is where automation comes into play. Some Azure backup solutions have the ability to automatically tier snapshots to lower cost classes of storage based on retention and compliance policies.

7. Isolate Your Backup Data

One of the most important things that you should do with regard to your Azure backups is to keep your backup data completely isolated from everything else. This is the only way to guarantee the integrity of your backups.

Early on, ransomware infections tended to target the Windows library folders (Documents, Pictures, etc.) for encryption. Over time though, ransomware evolved into something much more damaging. Modern ransomware variants still encrypt the Windows libraries, but they usually also encrypt the data found on network shares. There are even a few types of ransomware that are specifically designed to attack backups.

The one saving grace is that ransomware cannot encrypt data that it cannot access. That's great news if an end user happens to trigger a ransomware attack. However, if an administrator were to accidentally trigger a ransomware attack then the damage could be massive. After all, the ransomware will have access to everything that the administrator has access to.

That's why it is so incredibly important to keep your backups isolated. Backups often represent the only viable tool for recovering from a ransomware attack, short of paying the ransom. If the backup were to become a casualty of the attack, then the organization may be left with no other option to restore data.



There are no guarantees that you will be able to recover from a ransomware attack by paying the ransom. There have been many stories of people who have paid a ransom but were still not able to decrypt their data. There are also stories of secondary attacks occurring days or even hours after a ransom is paid. The attacker knows that the victim has already paid the ransom once, and therefore encrypts their data a second time in an effort to extort additional money.

One trick to keeping your backups safe is to make sure that the backup data is not accessible from any of your standard user accounts. Instead, create purpose-built accounts that are used only for backup and restoration tasks. These should be the only accounts that have access to the data.

Another thing that you can do to help to guard your backups against ransomware is to leverage role-based access control. A good backup application should allow you to delegate roles and permissions in a way that keeps any one single backup operator from having access to absolutely everything. This not only helps to limit the damage that ransomware can do, it may also help an organization to more easily meet its compliance mandates.

Another way that you can help to protect your Azure backups against ransomware is to use the Azure Key Vault to encrypt your backups. Backup encryption is essential for any number of security and compliance related reasons, but let's talk about why this is so important from a ransomware perspective.

Because ransomware has become such a pervasive threat, many organizations have put an increased emphasis on their backups.

Ransomware authors know that if an organization can just restore a backup following an attack, then they won't pay the ransom. That being the case, ransomware authors needed an additional threat to help entice victims into paying the ransom. This additional threat is exposure. There are several ransomware infections that not only encrypt data, but also threaten to expose the data on the Dark Web unless a ransom is paid. Organizations that use the Azure Key Vault to encrypt their backups can help to prevent the backup contents from being accessible to ransomware authors or other cyber criminals who want to steal data.

8. Utilize Multifactor Authentication

While we are on the subject of keeping your backups secure, be sure to take advantage of multifactor authentication. Even if you don't want to require multifactor authentication on an organization-wide basis, you should enable it for any accounts that have access to your backup data. This will keep an attacker from being able to gain access to your backup data by performing a brute force attack against an account that has access to the data.

9. Use a Single Backup Solution

Throughout this book, we have stressed the concept of location agnostics. Location agnostics is extremely important when it comes to backup applications. If a backup application is truly location agnostic, then it will be able to back up your data regardless of where it physically resides. There are massive benefits to having a location agnostic solution.

The most obvious of these benefits are that using a single solution to back up all of your data, regardless of its location, reduces costs and complexity. It's always going to be less expensive to manage a single solution than to juggle an entire collection of disparate backup applications – hence the earlier

recommendation to use a vendor who offers both BaaS and traditional backup software.

What is more important, however, is that having a backup application that is truly location agnostic gives you the flexibility to run business workloads in the location that makes the most sense. You might, for example, have some workloads running on-premises in a VMware environment, other workloads running in a Nutanix environment at a secondary datacenter, a few workloads running elsewhere in a Kubernetes cluster, and still other workloads running in Azure or another public cloud. Having a single cohesive backup solution that can work both on-premises and in the cloud, regardless of geographic location, hypervisor or cloud means that you can run your workloads where it makes the most sense to do so, without having to worry about how you are going to back them up.

Another way that such a solution helps organizations to be flexible is because it greatly simplifies workload portability. You can restore or move on-premises workloads to the cloud or bring cloud workloads back on-premises. Having this ability gives you a degree of future proofing, because you are no longer locked into using a specific cloud or platform. You can be agile and flexible enough to move workloads as your business needs dictate.

There are several backup applications available that can back up a variety of workloads, regardless of where they are located. As you evaluate the various backup solutions however, there are some critically important things to think about.

First, does the backup solution allow for true data portability? In other words, can you backup a workload in one location and easily (easily being the operative word) restore it elsewhere?

Another consideration is whether the solution will work with all of your Azure workloads. Remember, Azure is a collection of services and most organizations that use Azure leverage a variety

of Azure services. As such, it is important to be able to back up things like Azure virtual machines, Azure files, and Azure SQL.

A third consideration is whether the solution supports multi-tenancy. Larger organizations often operate multiple Azure subscriptions. Sometimes having multiple subscriptions can make compliance easier because multi-tenancy creates natural isolation boundaries. Multi-tenancy is also sometimes used in a way that gives each department its own subscription, so that the Azure resources that a department uses are being paid from their own budget. Whatever the reason, backup and restoration will be a lot easier if your backup tool supports multi-tenancy.

One last thing to think about is extensibility. Larger organizations often have custom applications that are not natively supported by any backup application. The best way to protect those applications is to use an API provided by the backup vendor to add backup functionality to the application. As such, any backup application that you choose should support extensibility through industry standard APIs, such as the RESTful API.

10. Avoid Complexity Wherever You Can

The entire concept of data backups is really straightforward. You are creating a duplicate copy of a protected resource so that you have a way of getting your data back following some sort of catastrophe.

It really doesn't get much simpler than that. Even so, there are some extraordinarily complex backup solutions on the market today. While such solutions presumably work, it is usually better to go with a less complicated solution. This holds true whether you are leveraging a BaaS provider or if you want to continue to operate your own backups.

One reason for this is that a simple backup solution can be implemented far more quickly than a complex backup that requires extensive architectural planning. The reason why this matters is that the simple solution allows you to begin protecting

your data right away, while a more complex solution may leave your data unprotected for a period of time during the implementation process.

A second reason why it is better to go with a simpler backup solution if at all possible is that excessive complexity is often the root cause of administrative mistakes. When it comes to something as important as your backup, you really don't want to be making configuration errors. The simpler a backup product is, the less chance there is that you will make a mistake when setting it up or using it.

11. Having the Ability to Restore Matters

Obviously, the entire point of backing up your data is so that you can get that data back if something bad should happen. Not every data loss event is the same. Your backup solution needs to be able to restore data in a way that aligns with the problem that you are trying to recover from. Imagine for instance that a user accidentally deletes a file from a file server. You shouldn't have to restore the entire file server just to get that one file back.

Make sure that your backup solution has granular restoration capabilities so that you can restore the minimum amount of data needed to recover from the situation at hand. Ideally, you should be able to perform a restore job targeted at the following levels:

- Host Server (if operating on premises)
- Virtual Machine
- Application
- Infrastructure Component (such as the Active Directory)
- File

While the idea that there are different types of restorations may seem really obvious, some of the backup and data recovery tools

that are available today force you into performing one specific type of recovery. Unless you have a backup solution in place that offers granular restoration capabilities, you may end up having to restore an entire system just to recover a file or an application.

12. Use Immutable Storage

One of the most important best practices for Azure backups is to write your backups to immutable BLOB storage. Immutable storage refers to cloud storage that does not allow data to be deleted or modified once it has been written.

Before we discuss the benefits of using immutable storage, let's take a moment and talk about the elephant in the room - cost. Earlier we explained that Microsoft and other cloud providers bill you each month based on the amount of storage space that you are consuming. Using immutable storage that does not allow data to be deleted might sound like a recipe for ever increasing storage costs. Thankfully, using immutable storage does not lock you into keeping your backup data forever.

Immutable storage supports the use of data lifecycle policies, meaning that you can create a policy that will cause Azure to automatically purge old backups after a specific length of time.

Microsoft Azure currently supports two different types of immutability policies. The first is a time-based retention policy, which renders data immutable for a predetermined length of time. Once the retention period ends, an administrator can delete the immutable data, but Azure will not allow that data to be overwritten or modified.

The other type of immutable data retention policy is a legal hold policy. This type of retention policy is generally used in response to litigation and data is held indefinitely until an organization's legal department chooses to release the hold.

Although immutable storage is probably best known as a document retention solution, it is ideally suited for use as backup storage. In fact, there are two major benefits to storing backups on immutable storage.

The first of these benefits is that using immutable storage makes it impossible for ransomware to attack a backup. While ransomware could theoretically encrypt a backup that is stored on immutable storage, that encryption would be treated as a new version of the data. The original, unencrypted backup is not overwritten because the underlying storage will not allow any changes to be made to the existing backup data.

The second benefit to using immutable storage for backups is that doing so can greatly simplify regulatory compliance. Requirements vary by regulation, but it is relatively common for compliance mandates to require backups to be retained for a specific length of time. Using immutable storage (along with a suitable retention policy) can help to ensure that such a requirement is met.

The Big Takeaways

With a market share estimated at around 21%¹ (plus the fact that you've read this book), it's probably safe to say that your organization has some portion of its' operations in Azure and needs to protect it with backups.

The use of Azure, while likely maintaining some form of on-premises environment, can look like it will complicate the issue of backups. But, by getting back to the basics and looking past the buzzword-worthy data protection hype, there are some very concrete and, in some cases, rudimentary backup truths that can serve as useful best practices.

By putting the 12 backup best practices we've outlined into use, you'll be able to develop an effectual, cost-effective, and secure means of protecting your Azure investment.

¹ <https://www.statista.com/statistics/967365/worldwide-cloud-infrastructure-services-market-share-vendor/>

Sponsor Chapter – Veeam Backup for Microsoft Azure



Organizations dependent on Microsoft Azure for their operational needs must take steps to protect the data, applications, and virtual systems hosted there. But, doing this should align with all of the best practices outlined in this book to achieve expected levels of efficiency, cost-effectiveness, and

productivity you've come to expect from your on-premises backups.

So, the solution used needs to be designed with the cloud specifically in mind so that it can take advantage of all that Microsoft Azure has to offer. It also must simultaneously augment features in innovative ways to ensure the highest levels of data protection of everything important you maintain in Azure. You've no doubt seen the following grid outlining the shared responsibility model every cloud provider promotes today.



Source: Microsoft

Microsoft Azure fits the model in the SaaS, PaaS, or IaaS categories, depending on the service being used. But when thinking about this from a data protection perspective, regardless of the service Azure provides, the “information and data” are *always* your responsibility.

This is where Veeam Data Platform comes into play.

Veeam Data Platform

Veeam Data Platform includes Azure-native data protection, providing organizations with the ability to use Azure VM snapshots for frequent recovery points and reliable recovery of everything from a single file to an entire VM.



Veeam supports the backups of over 450K+ customers globally today.

This product touts some impressive Azure-centric features to deliver fast recovery, no matter the data loss.

Azure-Native

It makes sense that since your organization is operating in the cloud, your backups of that environment should equally be both hosted and designed for the cloud.

Veeam Data Platform meets this need by offering:

- **Rapid cloud-based deployment** – Veeam Data Platform can be subscribed to and launched directly from within the Azure Marketplace or traditional methods. Data protection of your Azure IaaS instances like Azure VMs, as well as PaaS instances like Azure SQL, Azure Files and more can begin in, literally, minutes.
- **Multiple subscription support** – Backups can be centrally managed across subscriptions for improved ease of use and greater security.
- **Agentless** – We start with native Azure VM snapshots to allow for fast and frequent restore points and even faster recoveries. Image-based backups can then be

created from these snapshots, subsequently tiered off to Azure Blob for cost effective longer-term retention.

- **Set and forget simplicity** – Snapshot, backup and retention policies can be automatically configured and managed to ensure more reliable backups.
- **Enterprise Scalability** – Designed to protect even the largest environments with a single appliance, Veeam Data Platform easily grows with your organization without sacrificing performance and cost-reduction.
- **Fast and flexible recovery** – Sometimes you need an entire VM and sometimes you just need certain files. Veeam Data Platform offers flexible full- and file-level restore options, allowing the organization to get back to work quickly.

Cost-Effective

Storage in any cloud – even Azure – can get expensive if your use isn't managed. Microsoft Azure offers a number of storage tiers, each with lowered cost/GB tied with reduced response times and recovery speed. But, because some backups are needed for long-term retention while others are needed for instant recovery, cost-effective backups are only achieved when you can efficiently manage your storage use.

To assist, Veeam Data Platform offers:

- **Cost calculator** – Azure costs can be controlled while still optimizing your data protection with an industry-first built-in backup cost estimation tool. This helps avoid unexpected costs before they're incurred and ultimately lowers your bill.

- **Low-cost retention** – easily tier snapshots and backups across Azure Blob object storage classes, including archive, while hitting retention and compliance objectives.
- **Community Edition** – smaller customers can take advantage of free Azure and hybrid cloud backup and recovery in Veeam Data Platform.

Secure

Backup data is still at risk of data theft, deletion, and ransomware encryption, so it's imperative that your backups in Azure are still secure.

Veeam Data Platform ensures your backups are secure with the following features:

- **Delegated role-based access** – access to backups can be granularly controlled, based on three specific roles. Azure Key Vault support allows for additional security with backups encryption.
- **Isolated backups** – Logically air-gap and secure backup data from production with support for cross-subscription and cross-region configurations.
- **Layered defense** – Protect backup data from security breaches and cyberattacks with support for multifactor authentication.
- **Immutable data** – backup integrity is maintained by keeping data stored in a *write once, read many* (WORM) state through immutable storage for Azure Blob.

Hybrid-Ready

- **Cloud mobility** – Backup, recover or migrate any on-premises or private cloud workload to, from and across Azure and other clouds with no Veeam charges.
- **One platform** – Veeam Data Platform unifies Azure backup with other cloud, virtual, physical, SaaS, and Kubernetes environments under a single, powerful interface.

Protect Your Investment in Azure with Veeam

Organizations today face the need to ensure their operations in Azure remain available and secure. Having an ability to quickly, flexibly, and cost-efficiently backup and recover Azure workloads and their data is an absolute necessity.

Veeam Data Platform takes the native snapshot toolset offered by Microsoft and layers automation and multiple backup options to ensure data is reliably protected. Added intelligence like cost management and recovery flexibility enables organizations to achieve the highest levels of recoverability at the lowest total cost of ownership.

veeam

Veeam Data Platform

Cyber resilience and recovery
for Microsoft Azure and
hybrid-/multi-cloud environments



Layered Security



Powerful Recovery



Cloud Efficiency

Free Trial



Microsoft
Azure

Quickly become conversational about Microsoft Azure backups.

It's probably safe to say that your organization's investment in Azure is backup-worthy, right? Microsoft offers native functionality to backup Azure VMs, but what's the best way to ensure you truly are protecting your Azure environment? We'll discuss twelve backup best practices to help point you in the right direction.



About Brien Posey

Brien Posey is a 22-time Microsoft MVP, a published author and conference speaker with 20+ years of IT experience, and a Commercial Scientist Astronaut candidate.



ConversationalGeek®

For more content on topics geeks love visit

conversationalgeek.com