

LEARNING MADE EASY

Veeam Special Edition

# Hybrid Cloud Backup

for  
**dummies**<sup>®</sup>  
A Wiley Brand



Ensure data everywhere is protected

Own and control your data across any cloud

Manage disparate data and workflows

Compliments  
of

**VEEAM**

Brett McLaughlin

# About Veeam

Veeam® is the leader in backup, recovery and data management solutions that deliver Modern Data Protection. We provide a single platform for Cloud, Virtual, Physical, Software as a Service (SaaS) and Kubernetes environments. Our customers are confident their apps and data are protected and always available with the most simple, flexible, reliable and powerful platform in the industry. Veeam protects over 400,000 customers worldwide, including 82% of the Fortune 500 and 69% of the Global 2,000. Veeam's global ecosystem includes 35,000+ technology partners, resellers and service providers, and alliance partners and has offices in more than 30 countries.

To learn more, visit [www.veeam.com](http://www.veeam.com) or follow Veeam on LinkedIn @veeam-software and Twitter @veeam.



# Hybrid Cloud Backup

Veeam Special Edition

**by Brett McLaughlin**

for  
**dummies**<sup>®</sup>  
A Wiley Brand

# Hybrid Cloud Backup For Dummies® , Veeam Special Edition

Published by

**John Wiley & Sons, Inc.**

111 River St.

Hoboken, NJ 07030-5774

[www.wiley.com](http://www.wiley.com)

Copyright © 2022 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact [info@dummies.biz](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN 978-1-119-91328-3 (pbk); ISBN 978-1-119-91329-0 (ebk)

## Publisher's Acknowledgments

### Development Editor:

Rachael Chilvers

### Project Editor:

Saikarthick Kumarasamy

### Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

### Business Development

Representative: Matt Cox

## IN THIS CHAPTER

- » Knowing the challenges to data that a hybrid cloud presents
- » Understanding the limits of typical data resilience
- » Unifying your approaches to data protection

# Chapter 1

# Delivering Modern Data Protection to a Hybrid Cloud

Every modern technology-driven organization understands data is one of its most valuable assets. It's no surprise, then, that the protection of that data has become an increasing concern. This concern is made much more complex, though, because organizations no longer store all their data in a single location.

With the advent of modern cloud — virtualized, on-premises, hybrid, full cloud, across multiple vendors — it's no longer enough to say “I back up my local data to keep it protected.” More sophisticated approaches are required to fit the evolving data and platform needs of the modern organization.



REMEMBER

If you've bought into or heard the idea of hybrid cloud merely being a precursor to a full cloud approach, you can move on from that thought. Thousands of organizations are using hybrid cloud as their permanent cloud solution.

# Realizing the Value of Hybrid Cloud

Hybrid cloud has become one of the most popular solutions for managing highly sensitive data. It provides many of the advantages of the modern cloud — elasticity, managed services, economies of scale — but allows a local component, which often is closely related to an organization's most sensitive data.

However, with hybrid cloud comes multiplicity:

- » **Multiplicity of data.** Organizations using any type of cloud tend to have more data. Access and transit logs, application data, and business data all tend to grow across all the components of a hybrid cloud.
- » **Multiplicity of server types.** You also see more servers: sometimes physical servers and virtual machines hosted locally, but increasingly virtual machines hosted in the cloud, as well as Platform as a Service (PaaS) and Software as a Service (SaaS) workloads.
- » **Multiplicity of data patterns.** With more data and more server hosting solutions, it shouldn't be a surprise that data begins to move through new pathways, be accessed in different ways, and, at times, even be stored in new ways.

To gain true value from hybrid cloud, you need to manage this multiplicity, and be able to scale with it. At some point, if you have an individualized solution for each server type, or data flow, or storage approach, you'll lose the ability to manage just one more flow or approach, as well as the ease of a unified approach to data management.



TIP

Don't see this multiplicity as negative. In fact, the ability to support a truly heterogeneous set of environments is an advantage of hybrid cloud, not a limitation. Instead, you can view this as a potential loss of value that you need to solve to ensure you get maximum value from your cloud.

# Addressing the Limitations of Cloud Resilience

Most cloud providers offer a variety of solutions for data resilience. The most common is some form of replication:

- » **Replication across data stores.** It's often trivial to turn on data replication, mirroring your data from one data store location to another (or several others).
- » **Replication across cloud regions.** Another common resilience strategy is to mirror or replicate data to entirely different geographical regions. This offers greater protection from a catastrophe at one location wiping out your data stores.
- » **Replication from local to cloud (or vice versa).** A slight twist on regional replication is using a cloud to replicate data from local stores, or from the cloud to local data stores. This often gives a feeling of greater protection because critical data exists in a local organization-owned facility.

However, it often takes more than these resilience approaches to truly secure your data. Consider that resilience won't solve any of these common problems:

- » **Data corruption.** Corrupted data can replicate just as easily as non-corrupted data, so you can often find an entire dataset corrupted *and* replicated across all your data sources.
- » **Data security.** Cybersecurity incidents like ransomware continue to increase, while also becoming more sophisticated. Successful attacks on production datasets will also be copied to replicated data.
- » **Application reconfiguration.** Your applications rarely automatically switch over to replicated resources. You'll more often have to reconfigure them or adjust replication to take a secondary source and replicate it back to the primary source.

All this leads to a critical understanding: data resilience isn't a complete solution for data retention or backup, and isn't by itself sufficient for protecting your data.

# Centralizing Data Backup and Management

There's a difference between data backup and data resilience. Data resilience attempts to make your data — current, as it stands, right or wrong — less likely to be deleted or vanish. Data backup approaches the problem of not just preserving your data, but ensuring that it isn't corrupt. If you have a database go bad and get replicated across multiple regions and platforms, it's your data backup that enables you to recover what was corrupted.

The best way to manage and protect data across a hybrid cloud is to select and utilize a singular solution for data backup. Ideally, that singularity comes in a number of forms:

- » **Singularity of product.** One product means one vendor, one set of configuration concerns, one set of training.
- » **Singularity of platform.** Look for a solution that can work across multiple platforms — local, cloud, in different varieties, and so on — but presents to you and your technology teams as a singular platform with one entry point common to all your data management services.
- » **Singularity of management.** You may need different tools to properly manage data sources across your hybrid cloud, but you also shouldn't need different dashboards or approaches to dealing with disparate data. When you backup, you should be able to backup all your data without special steps for each variation in storage.

Centralizing your data backups with a single vendor actually allows your data to become *more* complex, and *more* distributed, because that complexity and distribution doesn't get limited by your backup solution.



TIP

Data backup in a hybrid cloud world isn't easy. It's no longer as simple as just copying files to another medium or location. Because data backup is complex, your selection of a vendor and disaster and data recovery partner really determines if you can manage that backup complexity well.



## IN THIS CHAPTER

- » Avoiding provider-specific data strategies
- » Reducing the cost of operating data through uniformity
- » Unifying your approaches to data protection

# Chapter 2

## Adapting a Multi-Cloud Data Strategy

Your selection and configuration of your hybrid cloud often has a symbiotic relationship with your selection of data strategy — and the partner and tools that are attached to that strategy. In the typical case, the more complex your hybrid cloud, the more complex your data strategy.

However, you want to create a different kind of relationship: one in which your data strategy *creates* room for complexity and flexibility in your cloud. An effective data strategy must itself be flexible to support an evolving business, and should open up options to you, not close them off.



TIP

If a complex hybrid cloud sounds like a bad thing, consider that sometimes complexity comes with using multiple providers, different approaches to hosting (like virtual instances and containers), and different data sources and data flows. In these situations, complexity means “more options” and is something you want, not something to avoid.

# Maintaining Flexibility with Multiple Cloud Providers

Whether you use one cloud provider or many, you're going to quickly see that there are two basic approaches to data protection:

- » **Provider-specific options.** Often options within the cloud provider itself, or through tight partner relationships, your data backup is highly tuned to a specific cloud provider, and depends on provider-specific features or configuration.
- » **Provider-neutral options.** Provider-neutral options allow for an agnostic approach to data protection. These solutions typically work across a large set of possible cloud providers, as well as a multitude of on-premises datasets.

Always strive for the second approach — one where the solution you use is neutral to a specific provider, but still provides the tight integration across multiple platforms. This neutrality enables you to add additional providers as needed.



WARNING

You don't always need to be cloud-neutral, as you can lose a lot of the value of a particular cloud provider's technology. However, in the case of data backup, cloud neutrality has a *lot* of value. Applications can and often should take advantage of specific cloud features; data strategies generally should not.

## Choose a provider-neutral strategy from the start

If you build your data strategy from the beginning to account for multiple cloud providers, you won't have to adjust it when you want that flexibility. More specifically, your strategy should account for data in any number of locations: on-premises, hosted in this provider, hosted in that provider, hosted in a database or on an object store.

## Maintain provider-neutral data

You also need to ensure your data remains provider-neutral. This is less about format and more about ensuring your data can easily move in and out of a particular provider, and then easily move to

another provider. Data that starts in AWS should be able to flow to Microsoft Azure or Google without change, and your backup solution should be easily adjusting to the data location without needing to worry about format changes.



TIP

Think about your data protection as a platform in and of itself. In that light, you want that platform to be able to interact with all other platforms in your hybrid cloud in a singular, unified way.

## Building Uniformity of Data Management

It's a good start to adopt a provider-neutral data strategy and work hard to maintain data uniformity across your cloud. But that's only a first step. You also need to ensure uniformity of the *management* of all that data, wherever it is created and stored.

Your data solution, then, is your entry point to all your data, in any location. It should allow you to answer Yes to the following questions:

- » Can you manage data stored on-premises as well as data stored in a cloud through a singular interface?
- » Can you report on and aggregate all your data and treat it as a whole, not just individual silos in different locations?
- » Can you easily orchestrate data protection across all platforms at once?



REMEMBER

Data management isn't just about tooling, either. Over 60 percent of organizations in the 2021 Cloud Protection Trends Report indicated that the same individuals managing on-premises backups are managing cloud backups. That means that a uniform solution serves this singular group far better than forcing them to learn specific options for each cloud provider.



TIP

As uniformity of data management increases, you typically see a decrease in the cost of that data management. Your operations team will need less training and spend less time context-switching, which all means greater efficiency of operation.

# Mapping Multiple Backup Data Flows to Gain Flexibility

There's another important question you should ask of your data strategy, and the partners and tools you use:

*Can you restore data to the original location as well as alternate locations, irrespective of whether the location is in the cloud or on premises?*

This is a powerful question because it begins to move beyond a rigid backup solution. Rigidity in this sense is a solution that takes data from point A, backs it up to point B, and can then restore that data to point A — and nothing else. A more flexible solution can take that same backed up data at point B and restore it to point C, or D, or anywhere else.

This flexibility in data flow ties directly to your ability to orchestrate real-time disaster recovery at the application level as well.



TIP

Always try and pair advances in your disaster recovery and backups with advances in how your applications connect to the protected resources. As recovery becomes more sophisticated, ensure your applications gain the benefit from this increased sophistication and uptime.

## IN THIS CHAPTER

- » Distinguishing between backup data flows and data workflows
- » Automating repeated processes
- » Orchestrating automations to provide ultimate flexibility

# Chapter 3

## Managing Disparate Data and Workflows

Flexibility is the key when it comes to a long-lived and sustainable approach to backups. Flexibility is critical when it comes to operating in multiple cloud environments, and it's equally important when it comes to all the varying approaches to storing and categorizing data workflows.



WARNING

It's easy to confuse the flow of data as it is backed up from one location to another (discussed in Chapter 2) and the flow of data within your applications and cloud environments — the focus of this chapter. You can think about the flow of data from a source to a backup as a *backup data flow* and the flow of data within and through applications as simply *data flow* or sometimes *data workflow*.

## Backing Up Heterogeneous Data Models

The reality of a hybrid cloud world, and the modern technology organization, is that data is no longer confined to a singular database in a singular location.

## Recognizing data categories

Consider all the ways in which data is typically categorized:

- » **Application data.** This is the common case where data (often in a relational database) serves dynamic applications.
- » **Log data.** Log data is now a first-class data citizen and is often required for compliance as well as operations.
- » **Internal user data.** Documents created and managed by your internal employees and contractors, from spreadsheets to presentations, are just as important to consider as traditional application data.
- » **Derived data.** With the explosion in importance of AI (artificial intelligence) and ML (machine learning), there's now a wealth of data generated from other data hosted in various locations.

## Understanding predominant data models

In addition to data categories, each of the above categories can be stored in a variety of data models:

- » **Relational data.** Again, the most common model remains data stored in relational, SQL-based databases.
- » **Structured, non-relational data.** The rise of NoSQL and technologies like Firebase have made data in key-value pairs and other non-relational stores extremely common.
- » **File-based data.** Logs are a great example here, as they're often stored in files in a directory-like structure.
- » **Object-based data.** Most commonly associated with Amazon Simple Storage Service (Amazon S3), object storage goes beyond basic file storage and adds organization without going to a relational model.



REMEMBER

The term *data model* is used a bit loosely here. Rather than the more formal “data schema” usage, it applies to the different approaches to organizing and storing data at a high level, each of which brings different considerations to management and backup.

Your approach to data backup must account for *all* these data flows.

# Building a Plan for Disparate Data Flows

In many ways, the best plan to handle different categories of data in different models is to *avoid* a plan for each of those categories and models. Put another way, unless you intend to be a data backup company, you'll rarely have the time, resources, and expertise to accommodate each category and each model in a long-term resilient manner.

Your plan should be less about the mechanics of a particular model, and more about selecting a provider and set of capabilities that address your current and future data needs. Inherent to a good selection here is support for disparity. Here are just a few valuable questions that you'll want to ask:

- » **Are there are data model limitations?** You want a “No” answer here. Assume that the *one* unsupported data model will always be the one you'll need most on your next project.
- » **Can heterogeneous data be treated as homogenous in backup?** That's a mouthful, but it boils down to this: you want your backup solution to allow you to treat data as data, not as “structured data” or “relational data” or “object data.”
- » **Can data be backed up from anywhere in its flow?** Obviously, you should be able to back up data from where it ultimately rests (such as a database). But can you also take snapshot backups from transient data or as data moves from one instance to another?



WARNING

Part of the reason you must work to avoid complicated model-specific strategies is that there's always another model around the corner. Imagine having spent hundreds of hours on a viable relational database backup strategy ten years ago to find that non-relational data is now critical, and little of your work translates!

Your plan, then, is less about solving for each individual model and workflow and more about finding a solution that handles heterogeneous workflows in general. Good solutions here are data model agnostic, as well as cloud and infrastructure agnostic. And, ultimately, the best solutions give you the most powerful tool you'll ever have: freedom to make choices that are right for your organization.

## DATA IS ALWAYS MOVING

If you're not already convinced that building a robust backup strategy is complex, consider that most data is never at rest for long. Data in one model may flow into an application and then be stored in another model, potentially with a new format or structure. This is just another reason that a comprehensive strategy that covers all types of data models is essential.

## Automating and Orchestrating Everything Possible

Flexibility is a key ingredient in a solid backup strategy. But make no mistake: flexibility without automation is extremely limiting. Automation here simply means that the processes which pull together data across your hybrid cloud and manage backups and restores are basically “button pushes.” You push a button and a complex process kicks off, manages data, and completes without intervention.

Once you have automation — long sequences of commands and actions wrapped into easy-to-use workflows — you can look to add orchestration through your backup solution. Orchestration takes all the automated processes, each doing single things (like backing up a specific relational database or restoring yesterday's snapshot to a production workload) and combines and manages them across a complex organization.

Take this approach with every bit of backup functionality you need and add to your data strategy: Automate it and then ensure that it's folded into your overall data orchestration. You'll find your life gets easier and your organization will run smoother than ever before.



- » Ensuring you own and control all your data
- » Choosing a flexible backup solution

# Chapter 4

## Five Hybrid Cloud Backup Challenges

Here are five challenges you'll face and overcome to plan and execute a successful hybrid cloud backup strategy.

### Redefining IT in a Cloud World

You can't simply take your existing IT group, throw a few training courses on cloud and data at them, and expect great results. You need to invest in redefining what IT in your organization means through training, upskilling staff, adding new staff, and setting new requirements and expectations. IT is no longer about watching lights and changing out hard drives, and it's up to you to effect this change in your own organization.

### Avoiding an "Only Local Data is Safe Data" Mindset

Gone are the days when the cloud was a risky venture, even for data hosting. While there's value in having a copy of all or part of your data in a local facility, that's not the ultimate goal of a balanced approach to data backup and integrity. Backup your data

well and often, and store data where it makes sense — including in the cloud as a primary source.

## Maintaining Ownership of Data

While you shouldn't be afraid to store data in the cloud, you should still be cautious to ensure you *own* all of your data. That means you control where it's stored now and in the future, how it's replicated, and where your backup solution stores artifacts. If you can't at any point ensure ownership of all of these steps in a data flow, change the flow (or the backup solution!). Better yet, make sure that every time your data is moved (even when backed up), you're in total control of the locations of that data. Don't settle for a vendor-controlled offsite backup. Back your data up to where you want, on your terms.

## Balancing Cloud Neutrality with Value

It's easy to think that the best approach to flexibility in the cloud is to never use any cloud vendor-specific services. That does make your code and data more portable, but it drastically reduces the value of each cloud vendor's offerings. Seek balance, not absolute neutrality.

Your backup solution should ideally support common vendor-specific paradigms, working *with* the cloud providers' strengths. If one cloud provider has a strong replication and mirroring model, use it — and then ensure your backup solution supports that functionality. The result is a high-value use of cloud-specific features *and* the choice of a flexible backup solution.

## Avoiding a Single Backup Model

Always strive for a single backup *solution*, but that isn't the same as a single backup *model*. As your applications grow more complex, you'll typically use data in and from more formats and sources. Allow the differences in backup model through choosing a highly flexible backup solution. Ensure your solution can accommodate a variety of approaches to backup, rather than forcing all your disparate data into one specific (and therefore inflexible) mold.

# Own your Data. Any Cloud.

Veeam Platform is the #1 Hybrid Cloud Backup solution, giving you complete ownership of your data across any environment so you can:

- ✓ Protect, control and manage cloud data
- ✓ Leverage the cloud for data protection and security



## Own your Data. Any Cloud.

Veeam Platform. #1 Hybrid Cloud Backup



Protect, control and manage cloud data

Cloud-native backup

Cloud Mobility

SaaS Backup

Kubernetes Backup



Leverage cloud for data protection and security

Backup and archive

Disaster recovery

Ransomware protection

Migrate and modernize



AWS



Azure



Google



M365



Salesforce



K8s



On-prem



Request a cloud consultation today  
<https://vee.am/hybridcloudinquiry>

# Reap the rewards of a hybrid cloud backup solution

The benefits of the public cloud are significant, but the cloud comes with challenges, especially related to data protection, management, and security. Your organization needs a modern data protection platform to match your evolving data protection needs, but where to start? Step forward, *Hybrid Cloud Backup For Dummies*. Read on to explore how to ease management complexities, lower costs, and avoid lock-in to specific platforms. Discover the importance of employing purpose-built backup and recovery for each of your environments, the simplicity of centralized management, and retaining ownership of your data.

## Inside...

- Find a modern solution for data protection
- Lower costs with centralized management
- Manage the multiplicities of hybrid cloud
- Successfully scale your solution

# VEEAM

Go to **Dummies.com**<sup>™</sup>  
for videos, step-by-step photos,  
how-to articles, or to shop!

ISBN: 978-1-119-91328-3

Not For Resale



for  
**dummies**<sup>®</sup>  
A Wiley Brand

# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.