ActualTech Media

rubrik

The Gorilla Guide to...®

Express Edition

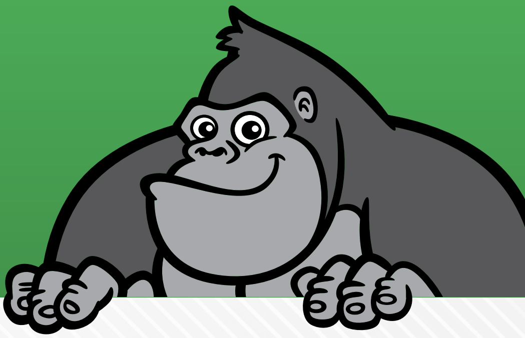

Database Protection in 2021

Joey D'Antoni

Inside the Guide

- How the Cloud Has Changed Database Backup and Recovery
- Top Challenges for Database Backup and Recovery Today
- What a Great Database Backup and Recovery Solution Looks Like

THE GORILLA GUIDE TO...

Database Protection in 2021

Express Edition

By Joey D’Antoni

ACTUALTECH MEDIA

6650 Rivers Ave Ste 105 #22489
North Charleston, SC 29406-4829
www.actualtechmedia.com

PUBLISHER'S ACKNOWLEDGEMENTS

EDITORIAL DIRECTOR

Keith Ward

DIRECTOR OF CONTENT DELIVERY

Wendy Hernandez

CREATIVE DIRECTOR

Olivia Thomson

SENIOR DIRECTOR OF CONTENT

Katie Mohr

PARTNER AND VP OF CONTENT

James Green

WITH SPECIAL CONTRIBUTIONS FROM RUBRIK

Kapil Chhabra, Director, Product Management

Michael Fal, Sr. Manager, Applications and Databases

Justin Ruiz, Sr. Product Marketing Manager

ABOUT THE AUTHOR

Joey D'Antoni is a Senior Architect and Data Platform MVP with over a decade of experience working in both Fortune 500 and smaller firms. He is currently Principal Consultant for Denny Cherry and Associates. He is frequent speaker at major tech events, and blogger about most technology topics. He believes that no single platform is the answer to all technology problems. He holds a BS in Computer Information Systems from Louisiana Tech University and an MBA from North Carolina State University.

ENTERING THE JUNGLE

CALLOUTS USED IN THIS BOOK

SCHOOL HOUSE

The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but are still important.

FOOD FOR THOUGHT

This is a special place where you can learn a bit more about ancillary topics presented in the book.

BRIGHT IDEA

When we have a great thought, we express them through a series of grunts in the Bright Idea section.

Takes you into the deep, dark depths of a particular topic.

EXECUTIVE CORNER

Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK

DEFINITION

Defines a word, phrase, or concept.

KNOWLEDGE CHECK

Tests your knowledge of what you've read.

PAY ATTENTION

We want to make sure you see this!

GPS

We'll help you navigate your knowledge to the right place.

WATCH OUT!

Make sure you read this so you don't make a critical error!

TIP

A helpful piece of advice based on what you've read.

INTRODUCTION

Welcome to The Gorilla Guide To...® (Express Edition) Database Protection in 2021. If you're struggling with how backup and recovery has changed along with the times, this book is for you.

Data is more important to your business than ever, and you have more of it, in more places than ever before. This data needs to be backed up and protected against the ever-growing threat of ransomware, in addition to the normal human or system faults that can cause data loss.

Cloud computing brings new risks, as you now have to protect different database platforms across on-premises and the cloud, which often leads to the need for different backup and restore options. Keeping track of hundreds of backup jobs across on-premises and potentially multiple clouds is a management and data challenge, to put it mildly.

This Guide is for those responsible for facing that challenge. That may include DBAs, decision makers in infrastructure and virtualization, cloud administrators, and more. Throughout this book, you'll get solid, actionable information to help you make informed decisions so that you can protect your organization's data from the myriad of threats and challenges facing it today.

CHAPTER 1

State of the Union in Cloud Data Management

As organizations move to the cloud, their data challenges only grow and get bigger. Managing and protecting data from a single platform becomes even more important as your data grows across on-premises and the public cloud. As organizations receive increasing amounts of data from Internet of Things (IoT) sensors, video, and telemetry data and other sources, they face enormous problems in managing that data.

When you couple this massive data growth with the acceleration of digital transformation initiatives and cloud migrations, you can see how orchestrating data protection is no easy task. Moreover, you're likely to encounter a variety of different database engines that require different approaches to backup than are traditionally used by legacy systems. You are also likely to have a combination of both Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS) database services that you're responsible for managing. Having a single solution to back up all that data across a variety of systems across on-premises and cloud makes data management manageable. Otherwise, you're stuck managing multiple backup solutions and different types of systems individually.

The same applies to protecting databases natively in the cloud. Each public cloud has backup solutions at different layers—you can backup VMs and databases in some solutions. While Amazon's and Microsoft's native backup solutions can be a single integrated view of backup for your IaaS solutions, those solutions don't integrate with your Software-as-a-Service (SaaS) solutions, such as Office 365 and OneDrive. You need a way to ensure those services are secure. You also want to ensure that you have the most cost-effective solution to take advantage of the native storage tiering available in the cloud.

There are several different models for cloud computing that different companies implement. Let's review a couple of them and what it means for backup and recovery.

Hybrid Cloud

When organizations move into the public cloud, they typically first try a hybrid cloud model, which involves some combination of public cloud and private cloud, perhaps with some on-premises infrastructure as well. The most common initial patterns include burst capacity for periodic workloads (for example, retailers who add compute capacity for the holiday season) or new projects that quickly need capacity.

Some organizations might make their first moves into cloud computing by using a public cloud provider as a disaster recovery option—in this case the model is typically a virtual private network (VPN) connection to virtual networks in a public cloud region.

The organization may choose to use the cloud region as a full hot disaster recovery site with live data replication, or a cold site, shipping backups and/or data changes in real time using a disaster recovery service. An example hybrid solution for backups is shown in **Figure 1**.

In other cases, hybrid cloud means accessing a SaaS platform for email and productivity uses, while running the rest of their IT operations on-premises. In any case, having a cloud presence represents a new platform, and typically involves new sources of data the organization needs to protect.

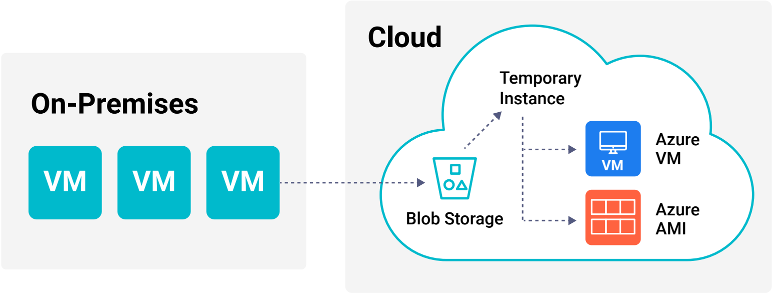

Figure 1: A hybrid solution for backups

While IT pros may joke about the importance of email, in reality it is one of the most critical data assets any organization has and needs to be protected as such. Critical to any data management strategy is understanding what data assets you have both on-premises and in the cloud, and how you're protecting those assets.

Over time organizations generally come to realize the benefits of the cloud platform, which often leads to plans to expand to a larger cloud environment than initially planned for. Thus, you want to ensure that whatever solutions you choose can grow with your deployments as needed.

Another challenge in a hybrid cloud environment is data movement between environments, which can be quite complicated. You need the ability to easily move backups and data between on-premises data centers and the cloud, and you may also want to have your backups replicated between your on-premises environment and your cloud to provide an additional layer of protection. It is essential that your backup solution can meet these data mobility requirements.

Multi-Cloud Deployments

Like hybrid clouds, the multi-cloud model has many variations, but it generally refers to having workloads in multiple public clouds. Multi-cloud has a couple of drivers: In some industries and countries, regulatory requirements mean workloads must be split across cloud providers. Sometimes the multi-cloud model is adopted with the idea of making applications more highly available, but, in fact, most cloud outages occur within a single region and a single provider, so similar levels of availability can be reached by using a simpler multi-region single cloud architecture.

Multi-cloud designs may be deliberate, such as a customer-facing application that has a front end in both Amazon Web Services (AWS) and Microsoft Azure. This is done to increase the availability of the application and protect against the failure of a cloud provider. In other cases, such designs can happen organically, as when a business unit decides to launch its own big data application using a SaaS solution that's not available on the IT organization's preferred cloud. In such implementations, it's important that business units work together to ensure awareness, monitoring, and data protection of all data.

One of the challenges of a multi-cloud environment is that it's typically much harder to take advantage of features that are specific to each cloud, including native monitoring and backup services, as well as platform-as-a-service (PaaS) offerings, since each are implemented differently in each public cloud. As a result, having a backup and recovery solutions that allows organizations to orchestrate and manage their data protection across multiple cloud environments from a single platform is paramount.

Cloud-Native Computing

Another industry trend is cloud-native computing—a collection of open-source projects with Kubernetes at its center.

Solutions that utilize cloud-native components can scale more easily and gracefully and take advantage of modern distributed system design that interacts seamlessly with the cloud platform.

Keep in mind that your backup solution should be able to take advantage of the unique features of each cloud provider, including automatic storage tiering or cloud-specific backups. Such integrations are key to a good backup solution, and you'll also want tight integration with the cloud control plane access control so you can manage identity and access management from a single control plane. As your resources expand into the cloud and beyond having a single backup solution that allows you to have a single view on data protection across all of your environments is critical.

CLOUD-NATIVE COMPUTING

Here's how the Cloud Native Computing Foundation defines the issue:

"Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

The Cloud Native Computing Foundation seeks to drive adoption of this paradigm by fostering and sustaining an ecosystem of open source, vendor-neutral projects. We democratize state-of-the-art patterns to make these innovations accessible for everyone."

CHAPTER 2

The Challenges with Legacy Backup and Recovery for Databases

Beyond just backing up your VMs and SaaS data, you need to ensure your databases are backed up successfully in both your cloud and on-premises environments. Since databases are stateful applications, they require backup software that's "application aware," which simply means a consistent backup of the database. This means backing up a database is not as simple as taking a backup of the server—you need to interact with the database to ensure that you can restore the database to a specific point in time. Your database backup solution needs to be able to easily restore, archive, and replicate backups across both your cloud and on-premises environments, and allow you to have a single pane of glass.

Slow Backups and Recoveries

Data volumes continue to grow, and this trend will only continue as enterprises generate more data with the addition of devices like IoT sensors or applications like telemetry systems. These are just two examples—there are myriad examples of growing data volumes as organizations move forward with their digital transformations, whether it be log files, endpoint management, or video data. While some organizations manage petabytes and exabytes of data, even organizations dealing with mere terabytes have challenges. The sheer physics of moving large amounts of data around makes managing backups difficult, and standard backup solutions can't always back up many multi-terabyte databasea at once and can be slow to recover large amounts of data.

Moreover, storing multiple copies of those backups increases costs dramatically. Taking advantage of technologies like snapshots, deduplication, and compression allows you to improve the performance of your backup and restores while reducing your costs.

Lack of Self-Service Access

Developers and IT staff are expensive resources, and if they're idle it means you're wasting money. A common scenario is needing access to a production copy of a database where validation testing is performed. With a standard backup and restore process, you're likely dependent on an admin to process the request, and the request itself can take hours or even days to complete. Such a delay to restore production data to testing or QA environments keeps developers idle and makes it hard to meet your project goals. It is important that you be able to leverage your backup data to drive additional value for the business and the faster you can provide teams access to this data, the faster they can deliver value to the business. The faster you can provide teams with the data they need, the faster they can add value to your business.

Need for Added Storage

One of the most common data restoration tasks in organizations is bringing production data to a lower environment for validation and testing. Traditionally, this process is performed by a DBA, who will manually restore a database to a lower environment. In some cases, there may be some level of automation, which reduces the time to completion.

In addition to the time involved, storage costs increase for multiple copies of large production databases—you need the full amount of storage space for each copy of the database. For very large systems, this can represent a tremendous expense. Reducing storage costs by using virtualization technology is a key aspect of any backup solution. This also lets you decrease the time to restore and lower the storage requirements for software development.

Lack of Flexibility and Integration

The world, especially the cloud computing world, is driven by APIs, which can be automated using REST calls. API automation enables a wide variety of customization and integrations with automation tools like Puppet or Chef and IT service catalogues like ServiceNow, or VMware vRealize Automation. Such mechanisms allow backups to be easily integrated into automated software deployments to provide a quick rollback point in the event of error and make for more efficient DevOps workflows.

REST APIS

Rest APIs use http calls to execute operations against a target resource. For example, a call of:

GET *http://example.api.com/users.json HTTP 1.1*

Would return:

```
{
  “user”: {
    “name”: “Jane Doe”,
    “siteRole”:  Owner”
  }
}
```

Other REST options include PUT, POST, and DELETE, which allow for updating and removing the state of the target resource.

In recent years, software tooling has vastly improved in a wide variety of products. Scripting options have gotten more robust, and vendors have become more supportive of both APIs and scripting interfaces to their products. Enterprise backup tools should support these capabilities to allow for easier management and integration using a variety of platforms and languages.

Complex Recoveries and Unreliable Backups

Modern backup solutions need to support rapid recovery that can scale across systems and provide varied recovery solutions. Whether you need to restore a single table in a database after accidental data deletion or perform a full system recovery after a ransomware attack, you need a rapid, repeatable process that allows for easy recovery in the event of failure. Organizations typically define their recovery mechanism using two metrics (see **Figure 2**):

- Recovery Point Objective (RPO)—how much data you can lose without impacting your business
- Recovery Time Objective (RTO)—how long your systems can be down without impacting your business

These two metrics, which should be agreed upon with your business leadership, define your recovery strategy. Note that while it's easy to ask for a 5-minute RTO and RPO for all systems, meeting metrics is expensive, and your executive leadership will suddenly be more

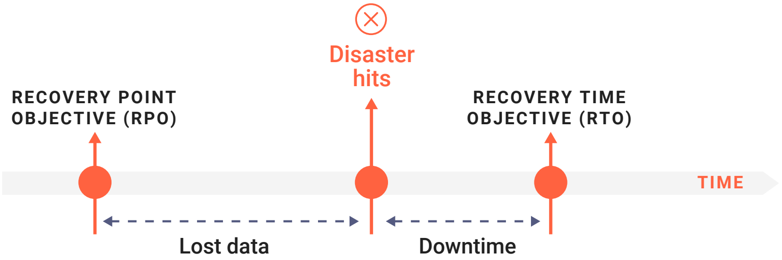

Figure 2: Recovery Point Objective (RPO) and Recovery Time Object (RTO)

understanding of potential downtime when you explain how much it might cost to meet that kind of objective.

Your backup and recovery tool should allow flexibility with your RTOs for different systems tiers. You also need to be able to easily test your data restorations, because if you don’t test recovery of your backups, you have no idea if you can actually restore those systems. A modern backup and recovery solution should automate backup validation as well as provide a means to do so through an API, provide visibility into your latest recovery points via easy to consume dashboards, and allow you to mount backups to quickly test.

Painful Scripting and Job Scheduling

Backup automation has always existed, from local cron jobs (Unix/ Linux job scheduler) to database job schedulers, but such solutions don’t scale. While the backup jobs may work on a given server, it’s hard to keep them in sync, and managing them can require accessing the individual server.

This may be viable when you only have a few servers, but it’s not workable as your system scales. And having to configure backups on individual VMs or services increases risk as implementation and job failure alerts can be misconfigured or simply missed in a sea of notifications.

All of these issues can be minimized by taking advantage of a unified solution that can manage your entire environment at once. Being able to easily deploy and monitor jobs from a single location can increase your productivity and ensure that all of your data is protected.

Security and Compliance

In addition to the challenges posed by cloud, growing data volumes, and painful scripting and job scheduling, organizations face threats from both hostile actors and even nation states via ransomware attacks. These threats are growing and becoming more aggressive. While approaches to protect against ransomware attacks are multi-faceted, having a robust backup solution that you can count on as a last line of defense is crucial.

You also need to be able to quickly identify malicious software as it enters your environment. And many industries, like the financial and medical sectors, have long-term data retention requirements—so being able to ensure compliance with business policies, audit user access and data protection policies, and the ability to recall data quickly is a necessity.

Staying Protected From Ransomware

IT organizations of all sizes are at risk for attacks—which are typically triggered through phishing attacks that target end users. Once an end user clicks on an infected attachment, the malware usually moves laterally within a network. In many cases, the malware will seek out administrator credentials, allowing the software to create new identities, or use existing administrator identities to perform malicious activities.

The ransomware encrypts critical files and servers, and, in particular, backups. Some more rudimentary versions will simply look at file extensions, while more advanced variants will attempt to review file headers. After encrypting the file systems, the malware then asks for a ransom payment within a specific time frame in return for a decryption key.

Ransomware attacks can be debilitating at best, and in the worst cases can end a business. While there are several strategies for risk mitigation, such as network segmentation and anti-virus software, the only complete solution is a robust integrated backup solution that protects all of the data in your organization—ability to detect anomalies, analyze threat impact, and recover quickly from immutable backups that can't be compromised.

Keeping Up with Compliance Requirements

Beyond data protection, there may be regulatory requirements your organization needs to meet. Whether it's a bank that needs to keep a seven-year history of transactions, or a medical device company that needs to retain 20 years of device history data, these regulatory requirements can drive your backup strategy. This can add significant costs in terms of storage and management for an asset that doesn't provide much value to the organization.

Because these backups are rarely accessed, they can typically be optimized for cost over restore performance. This is ideally suited to cloud storage platforms, which allow easy tiering with dramatically reduced costs for archival storage tiers. While such archival tiers have high access costs and slower performance, they can be perfect for these regulatory requirements. Having a tight integration with these storage tiers is a key element for any backup solution.

CHAPTER 3

Capabilities to Consider for 2021

As organizations move to supporting databases across on-premises and cloud, they need backup solutions that match their infrastructure and services. Taking advantage of modern solutions allows a more robust and cost-effective data protection solution.

Simplified and Automated Backup

Backups were one of the first tasks that were automated in IT organizations. Because they need to happen on a regular basis and are of critical importance, handling those operations became a priority. As IT matured, this evolved beyond single-task shell scripts and moved into larger-scale automation solutions. While these solutions are frequently clever, they can be siloed and require constant management to maintain configuration as new services arrive.

Modern database backup and recovery solutions need to be able to automatically discover and protect database instances as they are created through a single SLA policy engine, eliminating the need for painful scripting and job scheduling. Not only does this give teams time back and peace of mind, but helps improve compliance with business SLAs. Aligning on a central SLA policy also helps bridge the divide between database administrators and backup admins, helping reduce data ownership ambiguity, costs, and compliance challenges.

API-Driven Extensibility

In addition to backups, your organization can automate mundane manual tasks. Typically, an administrator will identify a task—like testing a database restore in a VM in the cloud—and then work to develop code or templates to make it a repeatable, executable task. An example of this might be restoring a database to a lower environment for testing, as mentioned earlier.

Having a platform that supports automation technology like APIs and shell calls means that you can easily extend small automation solutions into complex and powerful workflows. These can be used to manage backup jobs—which were typically driven by RPOs and RTOs, leading to backup engineers becoming glorified job schedulers. In contrast, modern backup tools need API-driven solutions to support more complex and distributed data environments, to allow customization and configuration management. Having this API surface can enable richer solutions—you can provide rapid access to database clones, and you can even execute your backups over the API.

Beyond simply automating operations, API access enables easy integration with third-party monitoring tools like Splunk and Nagios, which allow you to stream metrics and analytics to a common web interface. This functionality also provides monitoring and alerting options in the event of any backup failures.

Extending Protection to the Cloud

Modern backup solutions need to be able to do more than simply perform backups from within a cloud VM. One of the major benefits of cloud storage is its cost effectiveness, especially for long-term cold storage. Having a backup solution that's not only cloud-aware, but also cloud-native can help you to reduce costs and maximize the efficiency of your data protection solution.

Seamless Access to Cost-Efficient Cloud Archival

Until quite recently, most organizations used tape for long-term backups. Tape has a lot of advantages—it doesn't need to be powered on; it can be easily moved to different physical locations; and it's very cost effective.

However, in terms of management, tapes are problematic as they need to be catalogued and it's challenging to execute a restore from tapes. Tape management and rotation was never easy and required specialty software or appliances that handled management. In most cases, tape backups were used to meet regulatory requirements only or as a final tier of defense for restoration. Tape faded from popularity when disk storage costs dropped and capacity increased, leading most organizations to move exclusively to disk for backup.

Cloud storage offers cost and functionality similar to tape (in some cases it actually uses tape behind the scenes), while providing nearline data storage that takes only minutes or hours to restore. Amazon started this revolution with its Glacier service, which is cheaper than anything you can accomplish on-premises, especially when factoring in power and cooling costs. Legacy Database Backup and Recovery solutions were not built for the cloud. You'll want a backup solution that can seamlessly take advantage of these capabilities.

Ability to Lift and Shift

When mature organizations move to the public cloud, the pattern most follow is what's known as "lift and shift," in which existing applications are re-platformed onto cloud VMs (**Figure 3**). While this strategy sounds like a trucking model (and it sort of is), it's actually a fairly nuanced approach to cloud migration.

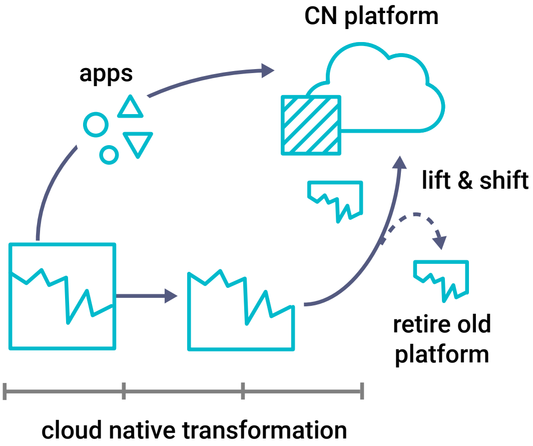

Figure 3: How “lift and shift” works

The first part of this strategy involves building networks and storage to replicate the source environment. This is typically the biggest challenge of any migration as the paradigms for networking and storage are quite different in the cloud compared to traditional on-premises environments.

Once your cloud storage and network resources are fully deployed, you can implement any cloud-necessary resources that can exist in parallel. For example, you can deploy domain controllers in both your cloud and your on-premises environments. These are among the key components needed to get up and running.

It’s important to understand the platform and the state of your applications as you migrate them to the cloud. Some applications, like web servers, are easy to move—you can simply redeploy your web servers from source control; there’s no need to physically migrate those VMs, unless they contain special software that needs to be installed manually. Having a backup and restore solution that allows you to execute testing to ensure your cloud resources function as intended. Being able to easily test greatly reduces the risks of your cloud migration without impacting your existing production resources.

However, stateful applications like databases require special handling to ensure you have no data loss. This requires the use of database-aware backup utilities that can be either native or third party. You might also use specialized database tools to manage near-zero-downtime migrations. Understanding your workloads and how you migrate them are key to successful cloud migrations.

Cloud-Native Protection

After migrating to the cloud, in addition to backing up IaaS solutions you need to ensure your SaaS solutions are backed up as well. A common misunderstanding is that fully managed SaaS solutions include point-in-time recovery. But that's not always the case. For example, while Microsoft provides a fully managed solution, Office 365 and OneDrive don't include backups as part of the service.

If you want to be able to restore, for example, an Exchange mailbox to a point in time, you need a backup solution in place for those services. And you need to ensure that all of your cloud IaaS workloads have data protection, especially when you're frequently spinning up new infrastructure.

One of the tenets of the cloud is that it's better to spin up new resources as needed rather than rely on physical entities. But this means VMs and containers may exist only temporarily, a potential problem if those resources are storing key data. Your backup solution should be able to identify newly created resources and ensure that they're backed up.

Beyond quickly spinning up resources, new solutions like NoSQL databases, PaaS databases require data protection, in the same way that more traditional enterprise solutions like Oracle and SAP HANA require backup protection.

Mitigating Data Risks

Backups are one of the prime targets of ransomware attacks. Since they're your only fail-safe against the attack, it's important to protect those backups, by network segmentation and other security and access controls.

Security is a multi-faceted problem that requires defense-in-depth approaches at each layer. It's crucial to quickly understand the impact and scope of any attack, so you can immediately identify what systems you need to isolate and recover. Beyond ransomware, protecting your data from accidental deletion, data exfiltration, and system failure are equally as important to maintaining your business operations that are dependent on data.

End-to-End Encryption

Backup tools need to be able to encrypt the backup at all stages. In many cases, default system backups are in plain text. Database backups are binary files, but they can be viewed in a hex editor and can potentially be restored to any computer running the database software.

Encrypting these backups means that only computers with the appropriate certificate can restore the backup, and the contents are not viewable as text. Moreover, the backups should be encrypted both in-transit (as they are being taken) and at rest. This reduces the risk of data exfiltration from external and internal threats. Backup solutions need to provide robust encryption options to meet all security and regulatory challenges.

Role-Based Access Control

Keep in mind that your assets won't be truly secure unless you control who has access to backups and backup infrastructure. You want to be sure the principles of least privilege are applied, and that you meet separation of duties rules for audits. Many modern systems use an easy-to-understand approach involving role-based access control, often containing the following basic roles:

- Reader: Can see all resources and review metadata, but can't perform any actions
- Contributor: Can create and modify resources, but can't grant access to others
- Owner: Can create and modify resources and assign permissions

This is a basic example but having a simple access pattern allows you to easily audit permissions and privileges. Of course, you'll want to provide the ability to create custom roles that assign specific privileges to meet specific organizational requirements. In a backup utility, this means defining the scope of what objects a specific administrator has access to, and what they can restore.

Immutability

Storing backups on standard media is risky because files are subject to manipulation or, in the case of ransomware, to being encrypted. The way to protect against these sorts of attacks is ensure no process can modify any backup once the backup operation has completed. This function of data protection is known as immutability, and it means that even if a ransomware attack hits your backups, you're still protected since the data can't be modified (see **Figure 4**).

This kind of protection needs to be implemented at the file system level of the backup storage target. This is a fundamental pattern that must be a part of any backup solution as it protects against unauthorized access or accidental deletion.

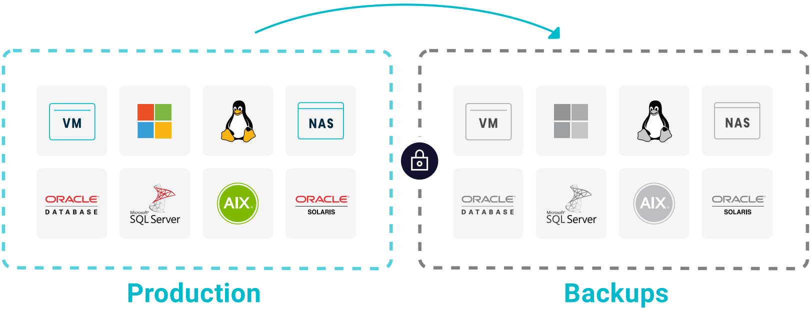

Immutable filesystems prevent attacks from accessing or encrypting data.

Figure 4: Immutable Protection

Impact Analysis

Understanding the real impact of system downtime on your business is crucial because it can help you drive relevant technology decisions that allow you to maximize your technology investments. Unfortunately, maintaining high levels of availability for systems requires a hefty level of investment.

Many organizations look at these as IT-only problems, and not the business-impacting events that they are. Involving your business's leadership and performing a business impact analysis ties into your RPO and RTO numbers by quantifying the specific impact—both financial and non-financial—of downtime or data loss to your business.

This can help you refine your RPO and RTO times, and help you understand where to make investments in redundant infrastructure and which systems to prioritize when recovering from a major failure. It can also give the business leadership a sense of ownership and involvement in your IT strategy.

It's Complicated, but Help Is Available

As you can see, it's an increasingly disparate, complicated world of data out there. This short Gorilla Guide has given an overview of those changes as they relate specifically to databases, and what's necessary to protect these most precious assets.

When you start searching for products to help, consider what Rubrik has to offer. It provides modern protection for a wide range of databases and can save you time and money in the process.

You can visit the company website to look at the full range of offerings, contact them for more information, or sign up for an on-demand webinar.

ABOUT RUBRIK

Rubrik helps enterprises achieve data control to drive business resiliency, cloud mobility, and regulatory compliance. Rubrik bridges the gap between owned, on-premises infrastructure and the cloud by decoupling data from the data center through a software-defined fabric and offering a single management plane for all data, whether on-prem or in the cloud. Comprehensive data management is delivered through instant access, automated orchestration, and enterprise-class data protection and resiliency.

ABOUT ACTUALTECH MEDIA

ActualTech Media

ActualTech Media is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.

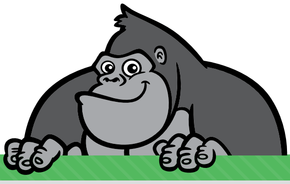

If you're an IT marketer and you'd like your own custom Gorilla Guide® title for your company, please visit https://www.gorilla.guide/custom-solutions/