EBOOK

Zero Trust Data Security™ for Databases

Keep Your Databases Running in the Face of Any Threat

rubrik

Table of Contents

Rubrik is a cybersecurity company, and our mission is to secure the world's data. We pioneered Zero Trust Data Security™ to help organizations achieve business resilience against cyberattacks, malicious insiders, and operational disruptions. Rubrik Security Cloud, powered by machine learning, secures data across enterprise, cloud, and SaaS applications. We help organizations uphold data integrity, deliver data availability that withstands adverse conditions, continuously monitor data risks and threats, and restore businesses with their data when infrastructure is attacked.

Database Protection Today

Organizations depend on the data in their databases for some of their most foundational operations. Budgeting, inventory tracking, order fulfillment, maintenance, and more all depend on databases to function.

Because this information is so essential, an organization can't afford to be without its database data for even brief periods of time.

Backups are the last—and best—line of defense aimed at keeping data safe and available to the people who need it. However, sophisticated cybercriminals have caught on to this fact and now target backups, so they can't be used to recover data.

Some of the greatest minds in cybersecurity have helped make significant strides in infrastructure and perimeter security, and those technologies do amazingly well at stopping the vast majority of cyberattacks. But they're not foolproof.

Evolving Risks

Cybercriminals have also made significant advances in their technologies and attack methods, often preying on individual users' curiosity and base desires to get access to an organization's systems. And all it takes is one mistake for an attack to succeed.

Recent data demonstrates just how vulnerable organizations still are despite heavy investments in infrastructure and perimeter security. A recent Rubrik Zero Labs report based on a survey of more than 1,600 IT and security leaders revealed that 92 percent of them are concerned they wouldn't be able to maintain business continuity in the event of a cyberattack. They also believe that one-third of boards have little to no confidence in their organization's ability to recover critical data and business applications in the event of a cyberattack.[1]

In addition, 82 percent of respondents to a 2020 survey from open source database consulting company Percona reported that their organization's database footprint had grown

1 "The State of Data Security by Rubrik Zero Labs," Rubrik, accessed December 13, 2022, https://www.rubrik.com/zero-labs.

by at least 5 percent during the last year. And 12 percent of respondents said it had grown by more than 50 percent.[2] As the number of databases grows in volume and variety, so, too, does the complexity of protecting all of those databases.

In order to be considered protected, databases need to be:

- Secure from threats
- Up, running, and always available
- Backed up according to predetermined SLAs
- Able to be recovered quickly

However, these protection tasks are often managed by multiple people across multiple teams. And because these teams all use separate tools to do their work and track their progress, they often struggle to work together to achieve their protection, backup, and recovery goals.

Couple the number and type of databases in an organization with the number of people and tools used to protect them, and something is bound to go wrong. And it often does. Databases get left unprotected or not backed up as often as they should be—all of which can put an organization out of compliance and result in fines. Or worse, in the event of a cyberattack, natural disaster, or operational failure, it can leave an organization without the data it needs to keep operations running.

Zero Trust Data Security™ for Databases

In the face of these new realities, database administrators and IT teams are turning to Zero Trust Data Security™ to protect critical databases against ransomware and other threats, streamline database discovery and protection, and, if needed, recover them quickly.

---

2 "Open Source Data Management Software," Percona, accessed December 13, 2022, https://www.percona.com/open-source-data-management-software-survey.

Zero Trust Data Security allows organizations to:

Keep databases secure and available with hyper-converged platforms that deliver air-gapped, immutable, access-controlled backups that can be easily replicated and archived to multiple locations.

Automatically manage database discovery and protection with a global management system that automatically discovers and dynamically secures all of the databases across an enterprise and in the public cloud.

Easily recover specific data or entire workloads as needed and make data rapidly available to people who need it.

This guide explores what Zero Trust Data Security for databases entails and examines the difference between Zero Trust Data Security and legacy approaches to database protection.

Using Zero Trust Data Security to Keep Database Data Secure

Cyberattacks that target backups have the potential to disrupt an organization's operations by denying it access to its data and blocking any attempt to recover it.

Legacy

Because cybercriminals have evolved their attack strategies to include backups, organizations' backup strategies and methods also need to evolve.

The legacy protection solutions that organizations generally use today consist of loosely coupled backup hardware, software, and secondary storage systems—offering a large attack surface for cybercriminals to exploit. Over the years, cybercriminals have gotten more sophisticated in their ability to find weaknesses, making it even more important for organizations to limit their attack surfaces.[3]

The volume and diversity of the databases these systems protect has also mushroomed, creating a corresponding increase in the amount of work teams need to put in to keep databases secure.

3 Katie Terrell Hanna, "What Is an Attack Surface and How to Protect It?," WhatIs.com (TechTarget, September 24, 2021), https://www.techtarget.com/whatis/definition/attack-surface.

Zero Trust

Zero Trust Data Security addresses these challenges with hyper-converged platforms that provide a significantly smaller attack surface and make managing protection far less time consuming.

With a Zero Trust Data Security solution, instead of manually building out, configuring, and managing backups, replication, and archives, a user can simply determine the service level agreement for a database or a set of databases and assign it.

These platforms also provide increased protection to database data through creating air-gapped, immutable, access-controlled backups, so it's extremely difficult—if not impossible—for a cyberattack to affect database backups.

By using Zero Trust Data Security principles to back up their databases, organizations will have a clean backup of their data readily available in case of emergency.

Characteristics of a Zero Trust Data Security backup:

Air gapped backups are either physically isolated, meaning stored separately from any network-connected system, or logically isolated, meaning still connected to a network, but separated through logical processes, including encryption, hashing, and role-based access controls.[4]

Immutable simply means that once the data is saved, it cannot be changed, overwritten, or deleted.[5] So, an immutable backup, once written, cannot be altered in any way, ensuring that the owner always has access to a clean backup.

Role-based access-controlled backups are only accessible to the individuals within an organization that absolutely need access to do their jobs. These individuals are also only able to use those backups in a way that's pertinent to their role. Role-based access controls limit the amount of damage a single individual can cause either accidentally or purposefully.

---

4 "What Is an Air Gap and Why Is It Important?," Rubrik, accessed December 13, 2022, https://www.rubrik.com/insights/what-is-an-air-gap-and-why-is-it-important.

5 "What Is Immutable Data Backup?," Rubrik, accessed December 13, 2022, https://www.rubrik.com/insights/what-is-immutable-data-backup.

Managing Discovery and Data Protection

As the number of databases an organization manages grows and environments become more dynamic, it gets increasingly harder to keep up with growing data protection tasks, putting data at risk.

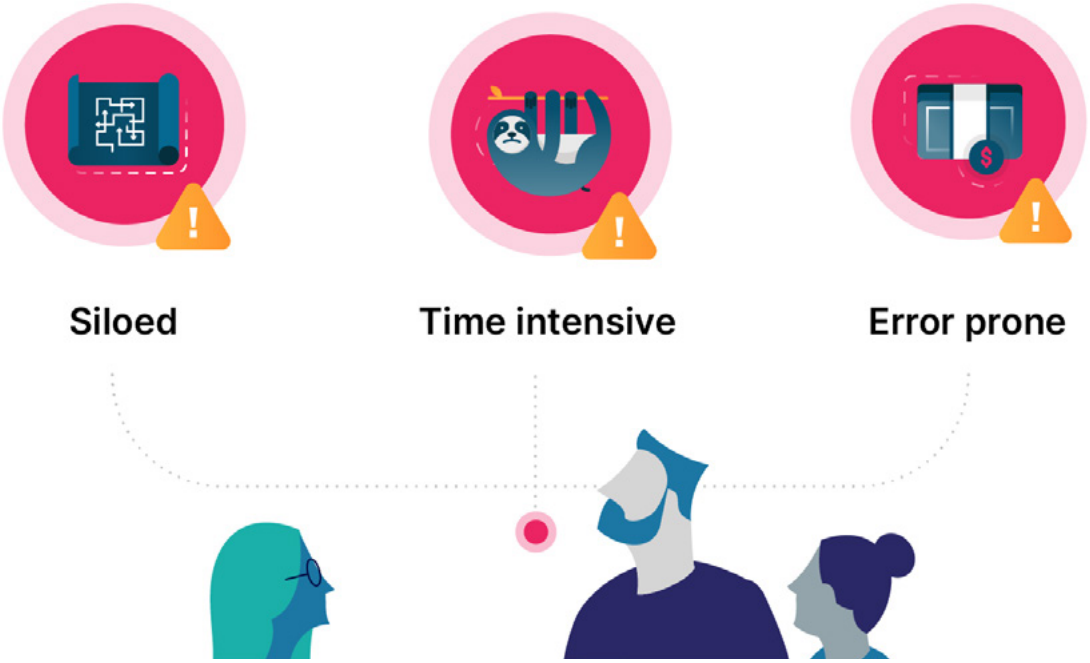

Legacy

Adding a new database to the backup schedule using legacy data protection is a multi-step, manual process. Each new backup job has to be created using native scripting, which not only requires writing the script (or modifying an existing one), but also installing it on the server and scheduling it to run—for instance, using cron on Linux or the SQL Server Agent job scheduler.

The person setting up the backup, often a backup administrator, also has to identify a storage target with enough capacity to satisfy the set retention schedule, which may require negotiating with the storage administrator to get the necessary space. Then, they need to coordinate backup schedules to make sure that too many backups aren't running at the same time, creating bottlenecks on servers, storage, or networks.

Skilled backup administrators can do these tasks with relative ease when they only have a few database instances to take care of. But any more than that, and things can quickly get out of hand—not only in terms of volume but also in the variety of databases, locations, and control planes in play.

To further complicate matters, backup administrators often do their work using separate processes and tools than DBAs. This division of labor helps each role make the best use of their skills during the backup process. But it closes off an important line of sight for DBAs, who need to work closely and quickly with backup administrators during recovery scenarios.

All these steps are time consuming, but the complexity also leaves the door open for errors that could result in:

- Unprotected or under-protected databases
- Prolonged backups that can hamper production
- Failed backups that put data and recoveries at risk
- Extended troubleshooting to identify and correct errors

At best, these mistakes can cost organizations untold hours of lost productivity. At worst, an overlooked yet critical database could be hit with ransomware and bring basic business operations to a halt.

Demonstrating Compliance and Recoverability

These processes also can put organizations in a precarious position with compliance requirements. Organizations need to be able to demonstrate they are in compliance with internal policies for data protection, government regulatory requirements, and cyber insurance requirements. This includes the ability to demonstrate that backups occur on schedule and that data is recoverable.

When organizations use legacy backup processes with a variety of scripts and schedules deployed to different servers, there's no central repository that shows all the backups and their statuses or throws up an alert when something goes wrong.

In traditional database environments, visibility into what's going on in the data protection domain is largely restricted to DBAs. Managers only know what DBAs tell them or what they can find out from spreadsheets and sporadic reports.

Demonstrating recoverability is a heavyweight task that can require DBAs to periodically set aside significant blocks of time (along with adequate compute and storage capacity) to validate recovery processes.

Zero Trust

Database protection solutions based on Zero Trust Data Security principles use SLA policy engines to automate the manual tasks associated with database protection. Once the software knows which servers run databases, it automatically discovers new instances, assesses their unique characteristics, and protects them using previously established custom policies, eliminating the need for error-prone and time-consuming manual scripting and scheduling.

Zero Trust Data Security-based solutions also create and execute optimized backup schedules based on an organization's operational constraints to prevent the risk of too many jobs running at once and causing problems.

Making It Easier to Verify Recoverability

In addition to saving time and avoiding mistakes during setup, a solution using Zero Trust Data Security principles can also help teams get better visibility into how their databases are being protected. Using a single interface, teams can quickly see if backups ran successfully and whether they're meeting the organization's SLAs for data retention, while monitoring compliance with government regulations and internal guidelines.

With centralized dashboards and reports, managers can see what's happening firsthand, organizations can better show how they're satisfying compliance requirements, and teams can easily identify and correct any problems before a crisis occurs.

A Zero Trust Data Security solution also makes verifying recoverability a much lighter weight activity. It can be as simple as quickly mounting an existing backup and verifying that it's readable—no additional storage required. Plus, with full APIs, important verification tasks can be automated to ensure they don't get overlooked.

With the visibility provided by a Zero Trust Data Security solution, everyone can sleep better knowing that critical databases are protected according to predetermined requirements. Teams not only save hours upon hours of time manually scripting, negotiating, and troubleshooting, they also sidestep possible errors that can lead to catastrophic outcomes and keep various teams updated on backup status and recoverability.

Managing Recoveries

Database downtime, from a cyberattack or otherwise, has a direct impact on an organization's bottom line. So, time is of the essence when it comes to recoveries.

Legacy

Anytime a database goes down, it's an emergency. DBAs and backup administrators need to work together swiftly to get the database working again as soon as possible. Unfortunately, restoring databases using legacy processes and solutions can be just as, if not more, time consuming than backing them up.

Traditionally, DBAs are dependent on backup administrators and storage teams to provide the right backup and the necessary storage to perform the restore. Coordinating across different teams using different tools takes time. If the files aren't online—maybe they're on tape—the process can take even longer.

Oftentimes legacy recovery methods force DBAs to recover entire workloads even if only specific datasets actually need to be recovered. Because legacy recovery tools can't recover just the data that's needed, the DBA, and the organization, loses valuable time copying over data that might not need to be recovered.

These factors combined can lead to missed recovery time objectives (RTOs) and more serious consequences, such as potential lost revenue for the organization.

Given the critical nature of database applications, DBAs often leverage other tools ahead of backups to ensure high availability. They also test their restore processes regularly to make sure everything goes according to plan. But should a DBA be faced with a scenario where they need to use a backup, these hurdles can exacerbate a crisis.

Creating Database Copies

Aside from recovery scenarios, DBAs are also being asked to make copies of databases for developers and business users. In recent years, companies have mined their data for better business insights and to offer new digital services. While finding new uses for data has been a boon for organizations at large, making copies of databases—whether for reporting purposes or for QA or development work— can put additional strain on DBAs.

In order to create these copies, DBAs using legacy solutions have to go through all the steps and cross-team negotiations that they would do if they were doing a traditional restore. Because the DBA has to make a full copy of the database, storage often becomes the limiting factor in determining how many copies can be made, which affects the business if there is significant demand for various database copies.

Plus, once the user no longer needs the copy, someone has to clean up that storage, and it may sit idle until it's needed for the next project.

Zero Trust

Zero Trust Data Security solutions not only automate backup processes, they provide more flexible recovery options, help DBAs coordinate with backup administrators and other teams during intense recovery scenarios, and make it easier to provide users with copies of a database.

Database protection solutions based on Zero Trust Data Security principles use a global management system that acts as a central hub for all teams to reference while performing their various backup and recovery responsibilities. With a global management system, DBAs have rapid access to the backups.

What's more is that Zero Trust Data Security solutions offer flexible recovery options that both allow DBAs to restore only the data they need and take advantage of advanced capabilities that help make that data available in the shortest time possible.

By recovering only the data they need, DBAs avoid time-consuming and anxiety-inducing wholesale recoveries. DBAs can also make that data available near instantaneously by eliminating the need to copy files over.

Flexible recovery options for near-zero RTOs

- Copy database files directly to the original host in the event of a production recovery
- Easily copy your production database onto an alternate host
- Mount any backup to the original host for instant recovery during a production outage
- Mount any backup to an alternate host for immediate data access or to pull the needed data
- Access a files-only recovery to manage restores using pre-existing scripts

Instead, the DBA can quickly mount a backup directly from storage to the production host in the case of a recovery. Once the database is back in production, files can be migrated to the production host in the background to restore normal operation. The DBA could be recovering whole databases or any part thereof. From the outside, both scenarios would cause minimal disruption.

Solving the Storage Problem

With Zero Trust Data Security solutions, teams also no longer have to waste precious storage capacity by making full copies of their databases. Instead of making full copies of a database, Zero Trust Data Security solutions only back up the data that has changed. Since storage is only consumed by metadata and user changes, additional storage costs for copies come out to be a fraction of what they would have been otherwise.

Fulfilling Secondary User Requests

These same abilities help DBAs provide secondary users with copies of databases. Instead of having to manually create another copy of the database, DBAs can simply pick a point in time, and the software executes the restore, creating the necessary database copy quickly and reliably. As an added benefit, every time the DBA creates a copy of a database, they're essentially validating the recoverability of the data, increasing the organization's confidence should it have to restore production systems quickly during a crisis.

Organizations can also use APIs to automate the fulfillment of requests for database clones and refreshes. The user would simply request what they need from a service catalog and the API would help create it immediately instead of waiting for a DBA to personally respond to their request.

Key Takeaways

Databases are some of the most important assets in any IT environment and should be protected as such. However, cybercriminals are penetrating traditional security defenses and even targeting organizations' best and last line of defense: backup data.

Zero Trust Data Security backup and recovery capabilities protect an organization's entire database operation—potentially including thousands of database instances—from cyberattacks, while simplifying data management and compliance and giving precious time back to DBAs and backup administrators.

Database protection solutions based on Zero Trust Data Security principles:

Keep databases secure and available with hyper-converged platforms that deliver air-gapped, immutable, access-controlled backups that can be easily replicated and archived to multiple locations.

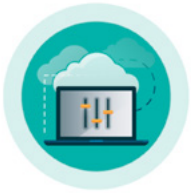

Automatically manage database discovery and protection with a global management system that automatically discovers and dynamically secures all of the databases across an enterprise and in the public cloud.

Help easily recover specific data or entire workloads as needed and make data rapidly available to people who need it.

About Rubrik for Databases

Rubrik for Databases delivers Zero Trust Data Security, ensures your critical databases—and all your data—are protected from cyberattacks, and gives you the ability to quickly and surgically recover data.

Rubrik for Databases enables you to protect mission-critical databases across on-premises and cloud from cyber threats, while unifying backup, replication, archival, and recovery into a single converged software platform.

With Rubrik for Databases, you'll be able to take immediate advantage of advanced features, including:

Automate Discovery and Protection

Eliminate painful scripting and job scheduling to free up DBA time

Full Control Over Recoveries

Choose from a range of automated recovery options or continue to leverage your own scripts

Near-Zero Recovery Time Objectives

Access databases in minutes regardless of size with Rubrik Live Mount and Instant Recovery

Self-Service Clones

Create self-service clones without added storage or impacts to production

To find out how Rubrik can help you enhance data protection for your database environment and increase the productivity of your entire team, visit rubrik.com


rubrik