



Buyer's Guide to Backup and Recovery

TABLE OF CONTENTS

- INTRODUCTION.....2**
- WHAT MATTERED YESTERDAY.....4**
 - Traditional Backup Requirements and Challenges.....4
 - Job and Scheduling Flexibility.....6
 - Replication.....8
 - Business Requirement Translation.....8
 - Tape Backup and Long-Term Retention.....11
 - Complex Installation and Configuration.....13
 - Cost.....13
- WHAT MATTERS TODAY.....14**
 - Modern Backup Requirements and Challenges.....14
 - Virtualization.....16
 - Simplicity and Automation.....16
 - Shorter Backup Windows Versus Larger Environments.....19
 - Cloud Usage and Application Agility.....19
 - Security and Access Controls.....20
 - Backup Security and Ransomware.....20
- WHAT SHOULD MATTER TO YOU.....22**
 - Selecting a Backup and Recovery Solution
for Today and Tomorrow.....22
 - Cloud Data Management.....24
 - Comprehensive Platform Support24
 - A Declarative Policy Engine and Automation.....26
 - Security and Compliance.....28
 - Easy Scalability.....29
 - Cost Versus Value.....30
 - Immutability and Ransomware Recovery.....31
 - Beyond Protection.....31
- CONCLUSION.....32**

INTRODUCTION

Backup and recovery need a radical rethink. When today's incumbent solutions were designed more than a decade ago, IT environments were exploding, heterogeneity was increasing, and backup and recovery systems were viewed as the protection scheme of last resort. They were intended only to provide a low-cost insurance policy for data, so companies patched together backup and recovery solutions under a common vendor management framework and tried to minimize costs by spreading data across different infrastructure and media.

What has changed? For one, IT departments have moved toward private cloud models with virtualization and are looking for converged architectures to replace multitier architectures. Second, because the amount of data under management has exploded, IT is challenged to do more with less. IT teams are now composed of fewer specialized roles, and more broad roles. Finally, public and hybrid clouds have opened up new data use cases such as analytics and test/dev, which create challenges for managing and securing that data.

Indeed, about the only thing that hasn't changed is the need for backups to be reliable and restores to be fast and dependable. But how do you accomplish this when so much about your infrastructure has changed? At the same time, how do you innovate quickly to align with critical business objectives and drive growth?

Any IT professional reviewing existing backup and recovery system investment—or considering a new investment to meet new IT and business needs—should reconsider whether the old assumptions are still relevant or whether a new approach is warranted. In this guide, we discuss the evolving demands of backup and recovery and the emergence of *Cloud Data Management*, which provides opportunities for protecting data, capturing new value, and making data available whenever, wherever it is needed.

We would like to extend our special thanks to those from the IT community who provided their perspective on these topics.



Businesses hinge on their data availability. Deciding what kind of workload infrastructure is best for your company will allow you to provide customers with top-of-the-line recovery.

Arminas Vrubliauskas


Data Protection Specialist, Core DataCloud

WHAT MATTERED YESTERDAY

TRADITIONAL BACKUP REQUIREMENTS AND CHALLENGES

The first batch of backup and recovery solutions were not built to address the challenges of application tiers powered by heterogeneous infrastructure. They were primarily built to be the platform of last resort, the platform from which one could rebuild after a catastrophic failure (or even after the inadvertent deletion of a critical file). Some backup systems needed to satisfy long-term data-retention requirements, generally using offsite tape archives, but most systems needed to focus only on moving large amounts of data across sprawling environments and managing it across multiple media types—disk and tape—to control costs.

Your Data Management Today




Legacy backup and recovery infrastructure—with its complex, multitiered architecture—cannot support forward-looking initiatives. It is time consuming to manage, stubborn to scale, and expensive to maintain.

JOB AND SCHEDULING FLEXIBILITY

Historically, the focus was often on backup windows and job schedules. These areas were necessary to achieve the desired recovery point objectives (RPO) and recovery time objectives (RTO) for a business. Unfortunately, this often turned backup engineers into glorified job schedulers—an unintended complication of complex architectures.

The RPO represents the point in time used for restoration and is determined by the frequency of backups. In the event of a primary system failure, a lower RPO means less data loss. Backup and recovery systems achieve low RPOs by taking more frequent backups at the expense of more traffic traveling across the network and more copies of data stored. In the case of mission-critical applications, RPOs need to be available as points in time measured in minutes, as opposed to hours or days.

An RTO represents the time it takes to recover an object such as a file, server, or datacenter. A lower RTO means less downtime in the event of a primary system failure, but at the expense of using more expensive, faster-access media (like a disk), as well as costly network switches to move data back to where it can be accessed.



Visual Representation of RPO and RTO



No one likes a disaster;
clearly, knowing your RPO and
RTO is a solid investment that
will yield positive returns.

Julio Lau-Cheng

Network Specialist,
Mount Saint Mary's University

REPLICATION

Replication is the capability to copy data from a primary location to a secondary location. This is often referred to as *disaster recovery* (DR) because it protects against a site-wide failure at the primary location. Replication only indirectly relates to RPO and RTO; the assumption is that most failures occur at the subsystem level rather than the site level.

Replication is a common requirement for critical applications, though. When it is required, it raises other questions that backup administrators and system architects have to answer. Should replication be synchronous or asynchronous? That's a question that RPO/RTO requirements can answer, but that's also a question that budgetary realities might answer. Synchronous replication is possible, but can be very expensive.

BUSINESS REQUIREMENT TRANSLATION

With all this in mind, the key requirement for any backup and recovery solution is to take business-level requirements for recovery time and data recovery, otherwise known as Service-Level Agreements (SLAs), and to translate them into a set of instructions for placing, retaining, and expiring data on different storage media.

The main problem with traditional backup and recovery systems is that the translation from business requirements to platform-executable instructions requires the engagement of professional services. In other words, traditional solutions have imperative versus declarative operating models. Moreover, after this translation is complete, these traditional solutions lack intelligence to optimize resources to avoid failed backups. This in turn leads to ongoing tuning and sometimes rearchitecting.

The best way to evaluate RPO and RTO in your current system is to ask an executive to pick some data (use different granular types) from random points in time. Quantify how closely you can achieve your RPO and RTO in the recovery. Compare this to the cost of downtime while recovery is in place. This has always been one of the most challenging areas of backup and recovery systems.



Successful backup architecture design relies on RPO and RTO. You need to understand how much data your company is willing to lose in the event of a DR, and how much time it will take to get back up and running.

Michael Alesi

Engineer, Moelis & Company



Designing and implementing a data-protection environment without understanding RPO and RTO requirements is like driving a car without a destination. Even if you're lucky enough to get there somehow, how will you know you've arrived?

Jon Heese

Sr. Systems Engineer, Flexential

TAPE BACKUP AND LONG-TERM RETENTION

Backup and recovery are typically designed for short-term data retention up to one month. Archive is used for long-term data retention with one to seven years being common timeframes. Long-term data retention is especially important in businesses that require regulatory compliance, such as health care or financial services.

Until recently, the only economically viable choice for archival has been tape. For all but the largest enterprises, tape involves manual handling, off-siting, logging, and rotation of tapes. Restoring from tape is time consuming, manual, expensive, and complicated because tapes are typically stored offsite, and a single file restore requires a broader system or volume restore. In addition, tapes degrade over time and must be refreshed.

Tape also lowers the value of data by sequestering it. Typically, tape-archived data is poorly indexed and limited in accessibility. It's inaccessible to your DevOps teams and your data analysts. By placing your most valuable strategic asset in a vault, you lock in your data and reduce its value to the business.

In some countries, tape archive was explicitly mandated to meet data-retention, legal, and regulatory requirements. Yet, an increasing number of agencies and jurisdictions are adapting data-retention policies to specify functional requirements rather than media.



Complexity in a solution
is hard to maintain and can
delay deployment.

Burhan Shakil

Systems Engineer, Harvard Law School

COMPLEX INSTALLATION AND CONFIGURATION

Configuration and installation of enterprise backup has always been a challenge. Almost all vendors require professional services to install and configure a backup system to the point at which all promised functionality is available. To use the system, administrators often must attend a week-long training.

There's an inherent risk in such complexity, though: If your enterprise backup system is too rigid, your backup solution might have difficulty accommodating an evolving operational infrastructure. Invariably, you will add new infrastructure components—some, perhaps, on-premises, some on the edge, some in the cloud—and you need to be sure that you're not creating unintended vulnerabilities or new infrastructure silos because your backup systems cannot support those new elements of your infrastructure.

One other consideration: Backup and recovery systems need to be easy to use, and complex customized configurations often are not. Operations should be so intuitive that the average administrator can use the systems without specialized training. In addition to being easy to setup and administer, backup should be automated as much as possible. Not only does automation facilitate consistency of backup execution, but it also enables your administrators to focus on other value-added tasks. Automation can also ensure that new assets—regardless of where they sit—are automatically protected when added.

COST

The cost of backup and recovery has always been a significant part of the IT budget. As the volume of data to be backed up has grown exponentially, so has the cost of backing up and storing that data. Sometimes, data protection costs even more than primary storage—particularly if you need to deploy a forklift upgrade to scale up the backup and recovery infrastructure and stay in sync with business expansion, or to accommodate new infrastructure (such as the cloud). Experienced IT organizations often budget for two to three times the primary data cost to cover data protection and backup costs.

WHAT MATTERS TODAY

MODERN BACKUP REQUIREMENTS AND CHALLENGES

The fundamental need to protect, recover, and archive data has not changed. But can the same be said about the character of those needs? Downtime today means much more than it did in the past, affecting not only business processes but customer satisfaction and business reputation. At the same time, there is more data than ever, in more places than ever, and all of it needs to be managed and protected. The need to be able to quickly and accurately recover in the event of a lost machine—or a breach—has grown more acute.

If you built a backup and recovery solution to meet present-day needs, how would it resemble a traditional backup solution? How would it be distinct? The requirement to support customizable SLAs based upon RPO and RTO, disaster recovery, and archival capability would stay the same, but you must also factor in new requirements. IT departments are increasingly adopting hybrid cloud models; they need hyperconverged infrastructures with modular scalability and increasing levels of virtualization. A modern solution would provide support for easy data migration from on-premises to cloud, APIs and policy-based cloud data management (CDM) services, and cost-effective approaches to DR. It would support dynamic technologies like SQL and NoSQL, service delivery systems such as software-as-a-service (SaaS), and IT automation tools such as ServiceNow and others. It must encompass the Internet of Things (IoT), Big Data, and the dynamic needs of DevOps teams that are eager to draw value from existing deep pools of data. And it would do all this with an unblinking eye on security, as ransomware attacks and data leakage remain ongoing threats.



Backup, archive, and replication form the backbone of today's data-protection and data-security solutions. CDP is a real game changer; with Rubrik's Radar product, now you can automate recovery from potentially massive incidents.

George Pickersgill

IT Infrastructure Engineer,
Shakespeare Martineau

VIRTUALIZATION

Most backup systems were originally designed to support physical hosts. Virtualization was the last major wave of computing innovation. Before virtualization, systems' RAM and CPU were underutilized, which provided resources during off hours for backup processes. Through virtualization, overall RAM and CPU usage were driven much higher, and storage moved onto a central array. Without careful planning or newer technology, backups can push virtualized systems past their resource maximums.

SIMPLICITY AND AUTOMATION

Virtualization initially reduced the number of servers that needed to be maintained in a datacenter, which initiated a thinning of the ranks within IT organizations. Today, with an increasing amount of infrastructure residing in the cloud, the ranks of IT organizations are growing even leaner. In many organizations, the role of backup administrator is falling to IT personnel with less specialized and more generalized skill sets.

That shift has an impact on backup and recovery strategies. Today's solutions need to be simple to learn and simple to use, with low-level instructions handled by heuristics and intelligence. They need to incorporate sophisticated data-management tools to help the IT generalists manage what is, in fact, a more distributed and complex environment than the datacenter of old. Such tools need well-designed user interfaces as well as APIs that can facilitate extensibility, integration, and automation through configuration management and infrastructure-as-code (IaC) tools such as Chef, Puppet, Ansible, and others.



With virtualization, our recovery time went from days to minutes, just as quickly as we can provision them.

Brandon Morris

Systems Administrator, City of Sioux Falls



As a CIO, I am always looking for products that will streamline our business and deliver speed to provide the best customer experience. With Rubrik's platform, we can meet these objectives while freeing up our time to use the solution for use cases beyond backup.

Leonard De Botton

CIO, Berkeley College

SHORTER BACKUP WINDOWS VERSUS LARGER ENVIRONMENTS

Modern IT departments are capturing and managing more data than ever, even while the windows of time in which to perform backups continue to shrink. How can you protect more data, more reliably, in a shorter period of time? Newer approaches such as snapshot-based backups remove the need to stop applications for backup. They also reduce the time required for backups and eliminate resource loads (placed by backup agents) on hosts.

CLOUD USAGE AND APPLICATION AGILITY

Increasing numbers of organizations are using the cloud to run Agile workloads and cloud-native apps. They are also embracing cloud storage and cloud archives due to rapidly declining costs. Although these options might make perfect business sense for the enterprise, they create management challenges for IT. Which assets are on-premises now, and which are in the cloud? Are they in a private cloud or public cloud? Are assets migrating among *multiple* clouds or multiple physical and virtual locations? Other organizations are attracted by the possibilities presented by the cloud, but their questions are more fundamental: How do we even embark on a cloud journey?

More to the point, though, how—and where—are you backing up and securing all of this data? How quickly and effectively can you restore data to the right location if you need to? Finally, on the subject of agility, in a scenario in which you use the cloud itself for backup and disaster recovery, do you still need to use tape to ensure that you have multiple backups in separate locations? Some organizations might be subject to regulations that require such tapes, but others might find themselves in a position to review their backup strategies and the tools they use to implement such strategies.

SECURITY AND ACCESS CONTROLS

Data-targeted cyber threats have grown more frequent and more sophisticated, and when they succeed, their impact can be monumental. Some attacks do emerge from the shadows of the dark web, but the attacks that more frequently do damage often originate from within the organization's own firewall.

Whether your data is on-premises, in the cloud, or both, it needs to be protected—at rest as well as in transit. That calls for data-encryption strategies, access-control strategies, and mechanisms designed to ensure the inviolability of your data even when stored in backup form.

BACKUP SECURITY AND RANSOMWARE

Backups are a favorite target of hackers and identity thieves, so special measures are required to prevent bad actors from stealing or tampering with your data. Can you easily examine your backup system to see if it is vulnerable to underlying OS security issues? Are your backups immutable? Are they protected against a ransomware attack even if the system itself is misconfigured? The ability to manage backups is critical, particularly if they contain sensitive data that is subject to regulatory compliance.

These days you must also consider your vulnerability to *ransomware*, a type of cyberattack that blocks access to an infected system until money is paid to acquire a code that will unlock the device. Although estimates vary, more than 40% of organizations will likely experience a ransomware attack each year, and these attacks can have major financial and reputational impacts on a company. The key to thwarting a ransomware attack—without paying the ransom and with minimal downtime for your organization—lies in the ability to restore an infected system quickly and effectively from a recent snapshot that is uninfected. Ideally, an organization would not even need to restore an entire system, and would be able to restore only those files that were compromised during the attack. If they can recover from a ransomware attack by quickly restoring only the files that were compromised in the attack, businesses could recover faster, minimize costly data loss, and employees could get back to work quickly and with maximum confidence.



Cloud is best for disaster recovery because you can instantly spin up servers without maintaining any infrastructure.

Burhan Shakil

Systems Engineer, Harvard Law School

WHAT SHOULD MATTER TO YOU

SELECTING A BACKUP AND RECOVERY SOLUTION FOR TODAY AND TOMORROW

Changing backup vendors always requires some level of technical or organizational effort. So, if you're considering a change, here are some questions to ask: Is the backup solution something that has been around for years with little change? Is it something created from multiple company acquisitions that have been cobbled together? Does it require multiple systems and interfaces for managing backup and recovery, replication, archival, and compliance? Does it require dedicated employees with specialized skillsets to manage it? What kinds of professional services will this change require? What kind of personnel training?

You will want to know not only how well a new solution will accommodate the complexities of the environment you currently have, but even more importantly, how well it will support the environment you will have in the future. If in that future environment you see a widening array of applications and technologies—from locally attached network-attached storage (NAS) appliances to cloud archives, from SaaS offerings to NoSQL databases—you need to be sure that the vendor you select can provide the support you will need to achieve your vision.

Backup innovation begins with recognizing that a new approach needs to align with the dramatic technology changes and data growth in centers over the past decade. Today, there are new approaches from industry visionaries who understand how to meet business needs in a period of rapid change.



Our team is small, but efficient and effective. We can accomplish great things with the relevant technology, which is why we partner with best-of-breed solutions like Rubrik. Not only does the solution fit nicely into our company initiatives, but also there's immense potential for the future.

Shadrach Kisten

CTO, Sesame Workshop

CLOUD DATA MANAGEMENT

Cloud Data Management services should be designed to orchestrate the management of mission-critical application data regardless of where it resides—whether in on-premises datacenters or within private and public clouds—while unifying backup, instant recovery, replication, search, analytics, archival, compliance, and copy data management in one infinitely scalable, cloud vendor-agnostic software fabric. Any solution built for the cloud generation should be able to take advantage of policy-based automation to manage data programmatically throughout its life cycle. It should be able to use APIs to facilitate data management functions across multiple applications, clouds, and protocols—ensuring the flexibility you need without locking you in to any particular cloud.

COMPREHENSIVE PLATFORM SUPPORT

A modern backup and recovery solution should support, optimize for, and integrate with *all* of the elements in your IT environment. Those might be physical elements, from sandboxed NAS systems to server clusters running SQL Server, MySQL, PostgreSQL, Cassandra, MongoDB, or some other production database. They could be virtual elements built on VMware, Hyper-V, Nutanix, or something else. The elements might not even belong to you but to AWS, Google Cloud Platform, Microsoft Azure, or some other cloud provider.

Your infrastructure needs are ultimately driven by your vision, your customers, and the opportunities that lie before you. Your backup and recovery systems should help you achieve your goals, not constrain or leave vulnerable the infrastructure that will enable you to reach your goals. That calls for solutions built on openness, flexibility, extensibility, and comprehensive platform manageability.



We're growing 33% year
over year in server workloads;
automation is the only way to
maintain that growth.

Brandon Morris


Systems Administrator, City of Sioux Falls

A DECLARATIVE POLICY ENGINE AND AUTOMATION

As IT departments come to rely more heavily on IT generalists, older backup management strategies that relied on imperative scripting techniques have become too complicated to sustain. Backup systems that anyone can manage have become non-negotiable for many organizations.

Modern backup and recovery systems can be managed by generalists when they take advantage of a declarative management model. In a declarative management model, an admin enters the desired state for a workload into a policy engine. After a policy is set, the system automatically and intelligently executes the jobs that need to be performed to achieve that state.

Let Your Policy Engine Do the Thinking



SLA policies allow you to collapse multiple manually implemented settings into a single easy-to-configure and zero-maintenance policy.

A strong policy engine can facilitate other aspects of service automation as well, reducing the number of manual steps that a generalist IT admin might otherwise be required to undertake to accomplish a task. If the backup and recovery solution has an API-first architecture, the organization gains even greater benefits. An admin could use these capabilities to integrate backup and recovery into an IT service catalog (e.g., ServiceNow, VMware vRealize Automation, or vCloud Director), simplify management of large, distributed environments via configuration management or IaC tools (e.g., Puppet, Chef, SaltStack, and Ansible), automate lifecycle data-management workflows, and centralize monitoring and reporting (e.g., Splunk or custom monitoring dashboard).




Automation is the wave of the future. If you're not working on it now, you're already behind.

Brad Craig

Senior Engineer,
Toyota Industries North America

Additionally, automation enables regular backup validation—a requirement to mitigate the “testing gap” risk, as shown in the illustration that follows. If backups aren’t tested regularly, IT cannot guarantee their validity to the business.



Without regular testing, guaranteeing reliable restores is impossible.

SECURITY AND COMPLIANCE

In the area of security and compliance, today’s IT teams need to do much more than use OS tools to monitor system events and manage access control lists (ACLs). Data security today involves encryption both at-rest and in-transit, key management, and the ability to instantly recover from events ranging from system failure to data breaches and ransomware attacks. Depending on the industry in which you are operating (as well as the countries in which you are operating), you might be subject to regulations ranging from the Health Insurance Portability and Accountability Act (HIPAA) to the EU’s General Data Protection Regulation (GDPR). You might be building payment systems that rely on the Payment Card Industry Data Security Standard (PCI DSS) or conducting financial transactions that are subject to scrutiny under the terms of the Sarbanes-Oxley Act of 2002.

Each regulation places distinct demands on an organization—for data protection and privacy, for data retention, and more. Some require distinct levels of encryption; others stipulate what data you may and may not retain as well as where you must and must not retain it. Your backup and recovery systems are as subject to regulatory compliance as your production systems, which means that you need to be sure that your backup and recovery system can provide the security and compliance support that your business requires.

EASY SCALABILITY


Like the modern operational environments they support, backup and recovery solutions need to scale—quickly and easily—or they risk being the bottleneck to growth. Thus, a modern solution should be built to run on any hardware with scale-out software and simple cluster management. This solution should be able to scale from terabytes of data to petabytes of data with consistent performance and usability.

When selecting a backup vendor, it's important to know how easily the solution scales and the maximum size to which it can scale. Metadata and data should be distributed across all nodes within the cluster and support global deduplication. No single management node should be a point of failure, and the system should have self-healing capabilities. When the system has a node failure, find out if system restores are as efficient as when the system is fully functional.

Typically, backup solutions also have a specified amount of data that can be backed up. When this limit is reached, a separate system is required. A truly scalable solution should allow you to add backup nodes that take advantage of current technology and scale to your entire environment. This allows you to find data from a single source and to take advantage of global deduplication. Adding nodes should be an easy process that does not require days of data rebalancing or professional services to manage.

Storage efficiency is another important component of a scalable solution. Modern backup solutions use techniques such as *erasure coding* to make optimal use of storage while simultaneously ensuring fault tolerance and high performance.

Data Efficiency via Erasure Coding



Modern protection methods enable faster rebuilds
with lower storage space overhead.

COST VERSUS VALUE

Understanding the true cost of backups is extremely difficult. You need to know how much data you have, the type of data (structured or unstructured), the amount of granularity required for RPOs, and how long the backups will be stored. From an OPEX perspective, the calculable cost of backups includes the software and hardware, the cost of the WAN for replication and backup, colocation costs, and so on.

But there are other costs associated with backup and recovery that are far more difficult to calculate. What is the cost of business revenue loss and productivity if recovery is delayed—or worse, unavailable because the backup was compromised by ransomware? What is the cost to your organization's reputation if you are unable to recover in a manner commensurate with your brand? What is the cost of productivity lost by using a solution with a steep learning curve and hours of management time recovered? The real value of a backup solution lies in its ability to ensure that you are back up and running without a costly interruption or permanent loss of data.

IMMUTABILITY AND RANSOMWARE RECOVERY

As noted earlier, ransomware attacks have become increasingly common, and one component of a ransomware defense strategy involves the use of a backup and recovery system capable of creating immutable backups—that is, backups that cannot be encrypted by ransomware. A second part of a protective strategy would involve machine learning tools built in to the backup and recovery system that monitor application metadata to detect and alert you to signs of anomalous activity that might be indicative of a ransomware attack. Ideally, these tools would provide insight at a very granular level so you could quickly identify specific data blocks that had been compromised—and then quickly restore just those blocks rather than entire files, with a single click. This combination of machine learning tools operating in real time, coupled with the ability to recover infected data quickly from immutable backups, should be part of your protection strategy going forward.

BEYOND PROTECTION

For many organizations, the rationale for robust data protection is no longer just a matter of insurance and compliance. Many businesses are discovering that they can take advantage of their backup and recovery platform to explore new use cases—to create archives in the cloud, for example, or spin up test and DevOps environments with “real” data.

Your backup and recovery system needs to deliver on the fundamental tasks associated with rapid, reliable, and comprehensive backup and recovery—first and foremost. But many forward-looking organizations are also contemplating other uses for the data those systems are managing. It is worthwhile to consider how you might derive other value from your backup data and whether your backup and recovery solution provides the APIs and automation features that would enable you to capitalize on your existing data to explore new opportunities. Make sure the additional capabilities are not simply “checkbox” features but can provide true value to your business.


CONCLUSION

When considering a backup and recovery solution, you must look in two directions. There's the backward look: The system needs to be able to capture and secure key moments in your organization's digital life, and it needs to be able to quickly, reliably, and securely perform that function so that you can rapidly restore that digital life in the wake of whatever data disaster confronts you later. At the same time, there's a forward look: The system needs to be scalable enough and flexible enough to accommodate future growth and business needs.

Your organization's infrastructure will evolve in response to new business opportunities and customer demands, and the backup and recovery solution in which you invest needs to be able to evolve in step. Are you entirely on-premises? Partially in the cloud? Dynamically moving between multiple clouds from multiple providers? Whatever the case, it shouldn't matter to your backup and recovery system.

Your IT teams will evolve, too, so you need a solution that a generalist can master and manage effectively. That calls for simplicity, but it also calls for APIs and automation capabilities, so that a small team can manage your infrastructure effectively.

Finally, you must look at cost. Your data might be of incalculable value, but your budget is not unlimited. The total cost of ownership of a modern backup and recovery solution, one that can provide you with both the forward and backward-looking capabilities you need, should be lower than what you have been paying for your legacy system, even though it offers the innovative features that can enable you to do more with your data than ever before.





Rubrik delivers a single platform to manage and protect data in the cloud, at the edge, and on-premises. Enterprises choose Rubrik's Cloud Data Management software to simplify backup and recovery, accelerate cloud adoption, and enable automation at scale. As organizations of all sizes adopt cloud-first policies, they rely on Rubrik's Polaris SaaS platform to unify data for security, governance, and compliance.

For more information, visit www.rubrik.com and follow @rubrikInc on Twitter.