


THE DEFINITIVE GUIDE TO

ENTERPRISE DATA BACKUP & RECOVERY ARCHITECTURES

Comparing On-Premises vs. Hybrid vs.
Hosted vs. Cloud-Native



From legacy to cloud data protection

The evolution of modern data protection strategies

The all-encompassing data center is a thing of the past. Modern data environments are distributed and include remote and branch offices, mobile devices, and the Internet of Things (IoT) as well as cloud solutions such as infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). There's more critical data to back up than ever before. Plus, data silos and fragmented management mean poor visibility, which can make it difficult to comply with regional data residency and security rules as well as service-level agreements (SLAs).

On-premises data protection has not kept pace. According to a survey by Fujitsu, 45% of IT managers said they had lost data or productivity related to data protection inefficiency within the last year.¹ Often, this is because on-premises backups are cumbersome and do not always happen on time.

A growing number of businesses are thinking about moving data protection to the cloud, because it promises to be:

- **Less expensive.** With cloud services, you reduce hardware, development, installation, and maintenance costs. You don't have to build and manage a data center. And, when you store all your data in the cloud, you only have to maintain one copy of that data, thereby minimizing storage costs and overall total cost of ownership (TCO).
- **More scalable.** You can add or subtract capacity to support large and rapidly changing workloads virtually instantaneously.
- **Easier to manage and maintain.** Data unification—storing all your backup data in the cloud—makes it easier to know exactly what data you have and manage it from a single control panel rather than multiple, incompatible tools. Plus, IT often doesn't have to manage routine software patches and updates for data protection infrastructure.
- **Highly flexible.** Cloud providers are increasingly adopting open standards, which means greater flexibility and more applications to choose from.
- **More reliable (and less risky).** Cloud providers are more likely to meet their SLAs for recovery time and recovery point objectives (RTO and RPO) when restoring business-critical data, and therefore ensure business continuity.
- **Today, there are many data protection architectures to choose from.** This paper will compare on-premises data protection with **hybrid, hosted, and native cloud** options, and help you determine which one is right for you.

¹ https://go.druva.com/rs/307-ANG-704/images/451_DisasterRecovery_DataBackup_RM.pdf

The problem with on-premises data protection

Although a growing number of businesses are adopting or considering data protection in the cloud, most still depend on on-premises backup and recovery. In a recent ESG survey, 49% of businesses said they still rely on tape as their primary form of direct backup.³ Why is this a problem? Downsides of on-premises backup and recovery include:

- **High costs.** Maintaining and upgrading traditional backup solutions can be incredibly expensive. Tape backups in particular can absorb significant IT resources and generate escalating maintenance costs.
- **Poor scalability.** On-premises data protection is difficult to scale as your business grows. When data growth outstrips the capabilities of your data protection system, backups may be late or simply not happen. An ESG survey of almost 400 IT and storage managers at midmarket companies revealed that only 70% of backup jobs are completed on time.²
- **Data loss.** When backups are late or incomplete—which is often the case with on-premises data protection during periods of rapid business growth—recovery from events such as outages or ransomware attacks can be compromised.³
- **Non-compliance.** Traditional backup solutions that store regional data in centralized hubs may not comply with evolving data privacy regulations and present significant legal risks.

² https://www.lto.org/wp-content/uploads/2014/06/ESG-WP-LTO-EVV-Feb_2016.pdf



Is the cloud secure enough for data protection?

Just a couple of years ago, public cloud was considered too dangerous and unreliable for storing sensitive company data. A lot has changed since then, including:

- **Improved physical security.** Cloud service providers have outfitted their physical premises with protective infrastructure and severely limited human access to critical servers.
- **Continuous monitoring.** Automated, 24/7 security monitoring allows cloud service providers to identify potential issues long before they impact customer data. Through 2020, Gartner predicts public cloud IaaS workloads will suffer at least 60% fewer security incidents than those in traditional data centers.³
- **Frequent security audits.** Cloud service providers conduct frequent security audits to ensure they're using the latest best practices.
- **New compliance measures.** Cloud service providers have introduced local archiving and thoughtful storage policies for full compliance with regional data privacy and security rules.

³ http://www.gartner.com/imagesrv/books/cloud/cloud_strategy_leadership.pdf



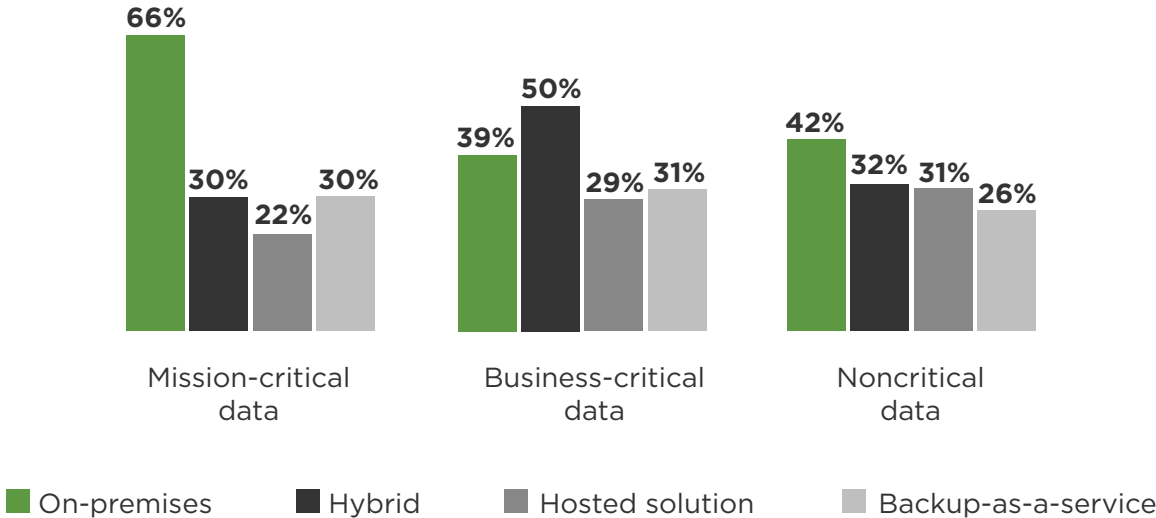
Popular data protection architectures

A brief overview of on-premises, hybrid, hosted, and cloud-native solutions

While enterprises are moving some of their data to the cloud and actively considering the cloud for data protection, most still rely on some form of on-premises architecture, especially for mission-critical data. According to a study commissioned by Forrester, most businesses use multiple approaches to data protection, including on-premises architectures as well as hybrid, hosted, and native cloud.

On-premises architectures are still commonly used for data protection

What approach(es) are you using to back up your data?



Base: variable; IT professionals at organizations in the US and Canada
Source: a commissioned study conducted by Forrester Consulting on behalf of Druva, November 2017

The following are the most popular solutions for enterprise data protection:

On-premises. On-premises data protection architectures are deeply entrenched in many organizations. Phasing them out may represent a financial loss. In addition, some IT teams may believe on-premises data protection gives them more control over critical data and minimizes security risks.

Hybrid cloud. A hybrid cloud data protection strategy relies on on-premises infrastructure and software for recent backups, and uses cloud for long-term archiving. Hybrid solutions may include data replication and tape backups. TCO may be better than traditional, on-premises data protection, but it is hard to scale, may take months to build, and can be complex in a way that makes achieving visibility into global data more difficult. This kind of cloud offers IT more control, but also requires active IT involvement at all phases of development.

Hosted cloud. Hosted cloud is traditional on-premises backup software running in the cloud and using warm storage. It offers greater availability than hybrid cloud, but it isn't necessarily optimized for a cloud environment; it doesn't fully deliver the economies of scale the cloud has to offer. TCO is similar to or slightly better than hybrid cloud, and it still requires considerable IT oversight.

Cloud-native. Cloud-native data protection is optimized for performance and scalability over the public cloud. It offers centralized management of backup and recovery processes, consistent performance even with petabytes of data, and lower TCO than hybrid and hosted solutions. It also scales fast. IT has much less day-to-day involvement than with hybrid and hosted models, which can mean greater vendor risk.

How hybrid, hosted, & cloud-native data protection frameworks stack up

This chart compares the most popular options for data protection in the cloud and how they compare to on-premises solutions

	On-premises	Hybrid	Hosted	Cloud-Native
Resiliency/Durability*	99%	99.5%	99.99%	99.99999%
Availability	99%	99%	99.5%	99.95%
Time to Market	Months	Months	Days	Days
Total Cost of Ownership**	High	Mid	Mid	Low
Offsite Options	Replication/Tapes	Replication/Tapes	Limited Replication	3-Way Replication
Retention Options	Disk/Tapes	Disk/Cold Storage	Disk/Cold Storage	Object/Cold Storage
Network Throughput	LAN	LAN + WAN	WAN (Restricted bandwidth)	WAN (Unrestricted bandwidth)
Scalability	On-Request	On-Request	On-Request	Scale-Out (unlimited)
RTO/RPO	High	High	Low	High (cache/disaster recovery)
Global Reach (new data centers)	On-Request	On-Request	On-Request	On-Demand

How fast is fast enough?

Can data protection in the cloud meet your RTOs?

When an event happens that takes your business offline, time is money. If the problem can't be fixed and data must be restored from backup, you want to recover your data as quickly as possible. A recovery time objective (RTO) is how long you have to restore data from backup without incurring business consequences. Your ideal RTOs may be different for different applications and lines of business.

The ABCs of RTOs in the cloud

RTO is a common metric in data protection SLAs, both internal and external. Research shows that only 68% of recovery jobs are achieving their RTO SLAs.⁴ In many cases, cloud data protection—especially hosted and cloud-native solutions—can offer greater reliability as well as significantly better RTOs than on-premises solutions for:

- **Single system recovery.** Most backup and recovery architectures can support single-system recovery, but hybrid, hosted and cloud-native solutions offer coverage for both on-premises and cloud applications. Hosted and cloud-native models also offer extremely fast RTO for endpoints and remote offices.
- **Archive recovery (cold data > 1 year).** On-premises and hybrid cloud solutions that rely on tape backups are significantly slower than hosted and cloud-native solutions that use cold storage.

RTO COMPARISON FOR SINGLE-SYSTEM RECOVERY

	Application Restore/Recovery			
	On-Premises	Hybrid	Hosted	Cloud-Native
Endpoints	20 minutes	20 minutes	5 minutes	5 minutes
Remote Offices	30 minutes	30 minutes	20 minutes	20 minutes
Small DC (100TB)	30 minutes	30 minutes	1 hour	1 hour
Large DC (100TB+)	30 minutes	30 minutes	1 hour	1 hour
Cloud Applications	NA	NA	1 hour	1 hour

RTO COMPARISON FOR ARCHIVE RECOVERY

	Archive Recovery (Cold Data Older Than 1 year)			
	On-Premises	Hybrid	Hosted	Cloud-Native
Remote Offices	Days	Days	4-8 hours	4-8 hours
Small DC (100TB)	Days	Days	4-8 hours	4-8 hours
Large DC (1PB)	Days	Days	4-8 hours	4-8 hours
Cloud Applications	Days	Days	4-8 hours	4-8 hours

⁴ https://www.lto.org/wp-content/uploads/2014/06/ESG-WP-LTO-EVV-Feb_2016.pdf

Hosted cloud vs. cloud-native

What's the difference?

Both hosted and cloud-native data protection services are typically delivered as SaaS.

Unfortunately, it can be difficult to tell the difference between the two because buzzwords like “cloud enabled” are often used to describe a broad spectrum of cloud frameworks.

Generally speaking, hosted cloud solutions repurpose traditional data protection software by running it in the cloud. Cloud-native solutions rely on software designed specifically for the cloud.

A quick comparison of hosted and cloud-native data protection



Hosted

A hosted solution is limited by the on-premises software it uses.

- Uses traditional software in the cloud
- Doesn't completely leverage cloud native technologies for scale and efficiency
- Expensive to build and manage
- Limited search or disaster recovery capabilities
- Single-tenant, manual updates



Cloud-Native

A cloud-native solution utilizes the full capability of the cloud.

- < > • Uses software designed to take full advantage of the cloud
- < > • Leverages cloud-native technologies for scale and efficiency
- < > • Lower cost and effort
- < > • Higher value of data with integrated disaster recovery, archival, search, and analytics
- < > • Multi-tenant, automatic updates



Finding the right cloud data protection strategy

Ultimately, the best data protection for your business is the one that meets your requirements.

The first step to moving some or all of your data protection to the cloud is understanding your RTO and RPO needs for mission-critical, business-critical, and noncritical data. Once you've defined your requirements, it's time to research cloud solutions and data protection solution providers.

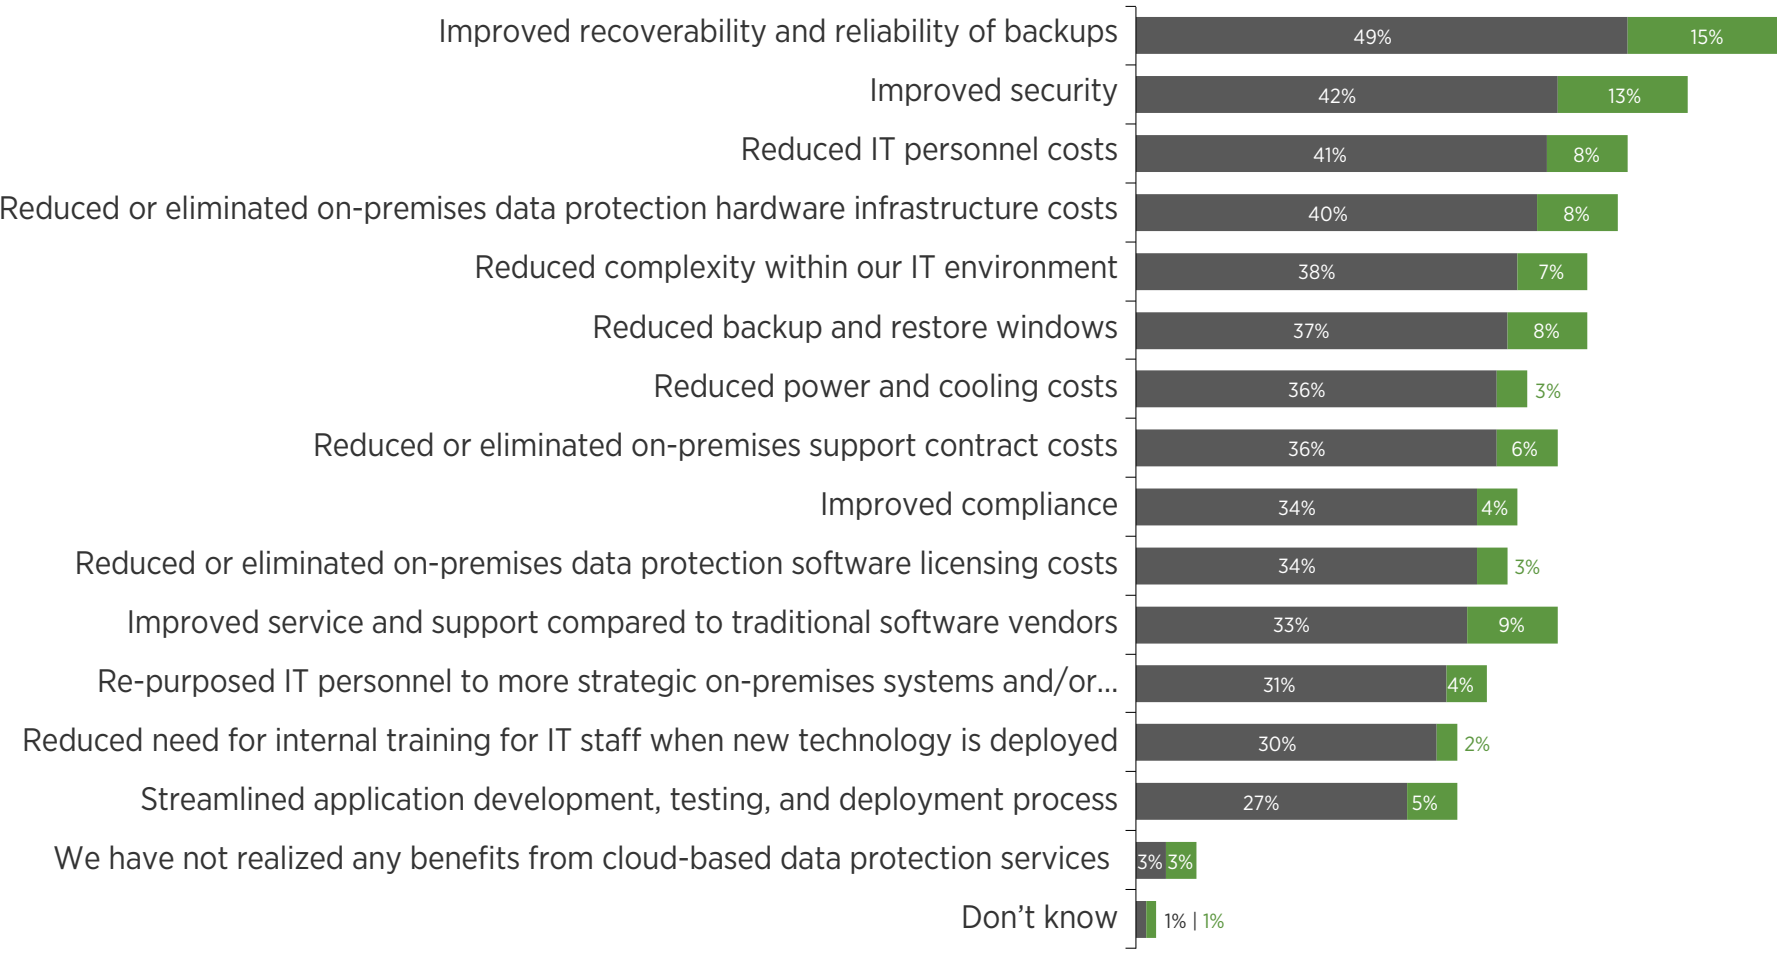
When you're talking to solution providers, these questions can help you see beyond the "cloud" buzzwords and identify which solutions will work best for you:

- How much data will be stored given your current data protection footprint, and how much will it cost?
- How much more will you pay if you need to support more data? How do costs go down when data is purged?
- What is the archiving model? How is data moved from warm to cold storage? What are the associated costs?
- How quickly does the system scale when you need more capacity?
- What about management across multiple regions?
- Does your solution use block or object-based storage? If block, how does it provide replication and resiliency, and how does it scale as capacity increases?
- How does your solution handle bandwidth constraints?



It's better in the cloud

How cloud-native data protection provides value



Source: ESG Research Report, 2016 IT Spending Intentions Survey, February 2016. (Percentage of respondents, N=212)

■ All Benefits
■ Primary Benefits

The major benefits of cloud-native data protection are:

- **Improved reliability and recoverability of backups.** This is no surprise, as most cloud data protection models—and cloud-native in particular—are built to meet aggressive SLAs.
- **Increased security.** A growing number of businesses are recognizing that data protection in the cloud is actually more secure than on-premises data protection.
- **Reduced IT personnel costs.** Cloud data protection—especially hosted and cloud-native—require much less IT involvement than on-premises.
- **Reduced or eliminated on-site data protection hardware infrastructure costs.** For midmarket and larger companies, this savings can be substantial, freeing up resources for innovation.
- **Reduced complexity within our IT environment.** Legacy, on-premises data protection systems can be complex and time-consuming to maintain. Cloud data protection can help streamline IT infrastructure overall.
- **Greater manageability and compliance.** If all backup data is stored in the cloud (e.g., data unification), businesses can manage that data from a single dashboard and more easily comply with data residency and other rules as well as SLAs.

Druva Phoenix: cloud-native data protection & management

Enterprise cloud backup with integrated disaster recovery and archival

Druva Phoenix, built natively on AWS, delivers data availability and governance for enterprise infrastructure with a unique cloud-first approach. It offers high-performance, scalable backup, DR, archival, and analytics to simplify data protection, improve visibility, and dramatically reduce the risk, cost, and effort of managing today's complex information environments. Key features include:

- **Offsite infrastructure.** It automatically protects data without the need for expensive on-site hardware and administrative overhead.
- **Improved business agility.** Response times for failover are nearly instantaneous with RTOs measured in minutes. Likewise, it can seamlessly replicate and move VMs across regions.
- **Workload mobility.** It supports workload mobility, which brings data protection to local applications and addresses their testing and development requirements.
- **Simplified management.** It provides a single dashboard for managing all your server backup and DR policies, so you don't have to rely on multiple storage, compute, and network management tools.
- **Low TCO.** Its pay-per-use model means you pay only for what you consume. Because it performs backups, DR, archiving, and analytics on a single, unified data set, you also avoid the costs associated with managing multiple data silos.

Learn more at www.druva.com/phoenix

About Druva

Druva is the global leader in cloud data protection and management, delivering the industry's first data-management-as-a-service solution that aggregates data from endpoints, servers, and cloud applications and leverages the public cloud to offer a single pane of glass that enables data protection, governance, and intelligence, dramatically increasing the availability and visibility of business-critical information while reducing the risk, cost, and complexity of managing and protecting it.

Druva's award-winning solutions intelligently collect data and unify backup, disaster recovery, archival, and governance capabilities onto a single, optimized data set. As the industry's fastest-growing data protection provider, Druva is trusted by over 4,000 global organizations and protects over 40 petabytes of data.

Druva, Inc.

Americas: +1 888 248 4976

Europe: +44 (0) 203-7509440

APJ: +919886120215

sales@druva.com

www.druva.com

© Copyright 2018 Druva, Inc. All rights reserved