



The Definitive Guide to Enterprise Data Backup and Recovery Architectures

Comparing on-premises, hybrid, hosted,
and cloud-native solutions



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 cloud solutions such as Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). There is 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 ESG, the 3 benefits of cloud-based data protection services were improved security (52%), improved recoverability and reliability of backups (46%), and reduced IT costs (41%).¹



¹ [The Evolution of Data Protection Cloud Strategies, ESG Research Report, May, 2021](#)

Businesses Are Moving Data Protection to the Cloud

Many organizations are already moving mission critical workloads to the cloud and a growing number of them are thinking about moving data protection to the cloud, because it promises to be:

- **Less expensive.** With cloud services, you eliminate hardware installation and maintenance costs. You don't have to build and run backup infrastructure in your data center. When you use the cloud to back up and store your data, you can eliminate off site infrastructure costs (leveraging storage redundancy in the cloud) and pay for capacity as you use it, thereby minimizing infrastructure costs and overall total cost of ownership (TCO).
- **More scalable.** You (or your solution) can add or subtract capacity (storage and compute) to support large and rapidly changing workloads virtually instantaneously.
- **Easier to manage and maintain.** Data unification—storing all of 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 teams don't have to manage patches and updates for data protection infrastructure.
- **Highly flexible.** Cloud gives you the flexibility to restore data directly to the cloud or to on-premises. Additionally, having data in the cloud enables you to take advantage of value added services such as cloud-based disaster recovery, e-discovery, and data governance.
- **More reliable (and less risky).** Cloud providers typically contract to meet SLAs for recovery time and recovery point objectives (RTO and RPO) when restoring business-critical data, therefore ensuring business continuity.

Today, there are many data protection architectures to choose from. This paper will compare on-premises data protection with hybrid, hosted, and cloud-native options and help you determine which one is right for you.



The Problem With On-Premises Data Protection

Although a growing number of businesses are adopting cloud or considering data protection in the cloud, most still depend on on-premises backup and recovery. An ESG report mentions that public cloud provides 38% more availability compared to on-premises resources.² Why then continue to rely on onsite data protection? Downsides of on-premises backup and recovery include:



High costs

Maintaining and upgrading traditional backup solutions can be very expensive. Tape backups in particular can absorb significant IT resources and generate escalating maintenance costs.

More vulnerable

Protecting the data stored on-premises is your responsibility. You must take care of the physical security and the cyber security of your entire data center. A small lapse is all it takes for hackers to gain access to your production or backup data. An ESG report mentions that cloud provides 41% better security compared to on-premises resources.³

Less reliable

On Premises backup and recovery solutions are not hardware agnostic. Any failure in devices, software, or environment compounds challenges with backup and restore. This increases the risk of data loss and the time taken to recover from events such as outages or ransomware attacks.

Non-compliance

Security is a greater challenge especially for distributed data centers, ROBO, and cloud workloads.

^{2,3} [The Evolution of Intelligent Data Management: ESG Research Report, January, 2022](#)

Is the Cloud Secure Enough for Data Protection?

Gartner says that by 2024, more than 45% of IT spending on system infrastructure, infrastructure software, application software and business process outsourcing will shift from traditional solutions to cloud.⁴ However, just a couple of years ago, the public cloud was considered too dangerous and unreliable for storing sensitive company data. Here's how cloud security has changed:



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.

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.

⁴ [Top Four Trends Are Shaping the Future of Public Cloud](#), Gartner Research

Popular Data Protection Architectures

A brief overview of on-premises, hybrid cloud, hosted cloud, 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. The following are the most popular solutions for enterprise data protection:

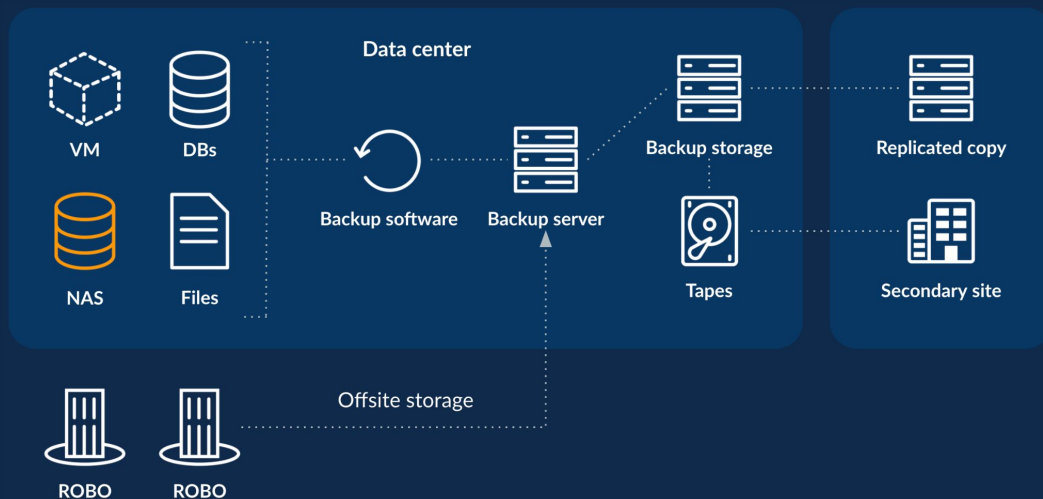


On-Premises

- **Where the data is stored.** Physical space owned or managed by the organization
- **Who manages these systems.** IT Team manages the backup platform and storage systems
- **Advantages.** Faster access to large amounts of data
- **Disadvantages.** Costs, overhead, and security of multiple systems often in multiple locations (silos of storage including offsite and archive, operational security)

Definition

Backup data is stored in a local data center. The IT team is responsible for the maintenance and upkeep of this infrastructure. As the servers are close by, the access speeds of data stored on-premises are thought to be very high. However, maintaining this footprint requires substantial manpower and costs. This strategy also requires organizations to move a copy of data offsite to a different media source for security and business continuity.

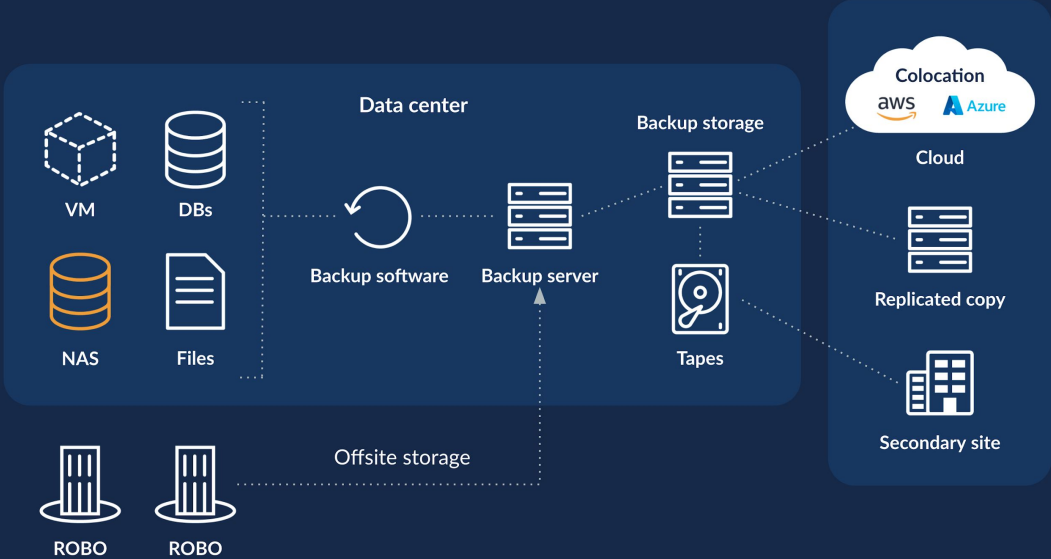


Hybrid Cloud

- **Where the data is stored.** On-premises for short-term and cloud for offsite or archive data
- **Who manages these systems.** IT team (on-premises) and/or third party (cloud)
- **Advantages.** Frequently used data can be stored on-premises, while cloud can be used for data archival and long-term retention
- **Disadvantages.** Complex management of storage, retrieval, and archival across multiple locations. Limited visibility into current and future costs. Cyber security and data loss risks managed internally.

Definition

The hybrid cloud uses a mix of on-premises servers and public cloud to store data. The IT team manages the on-premises infrastructure and the cloud provider manages the public cloud. The primary benefit of this storage system is getting the best of both worlds. The primary benefit of this approach is that it helps eliminate the use of tapes for offsite storing of replicated data.

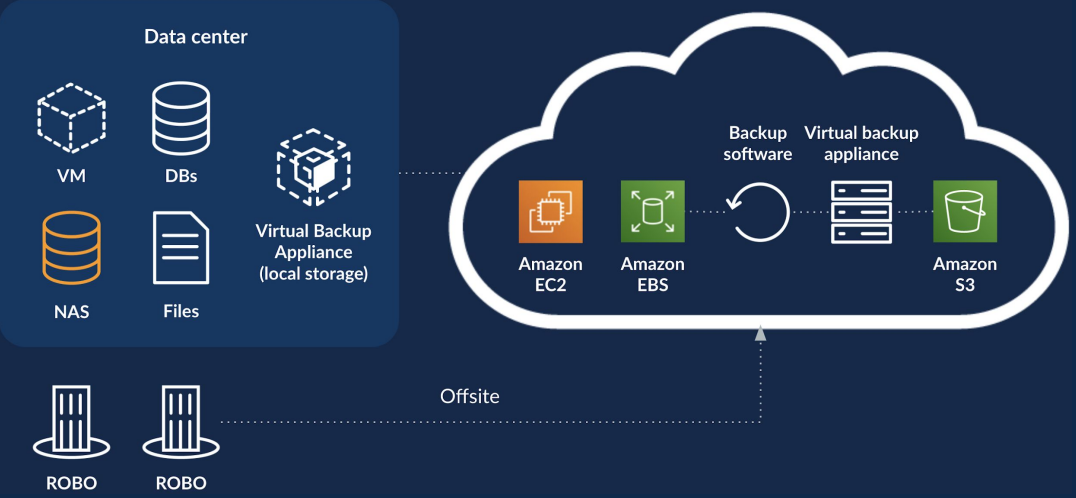


Hosted Cloud

- **Where the data is stored.**
On-premises and/or in the cloud
- **Who manages these systems.**
IT team and/or third party (vendor)
- **Advantages.** Gaining some of the benefits of the public cloud without having to share the resources or data, guaranteed single-tenant architecture
- **Disadvantages.** Difficult and expensive to scale. Requires dedicated CPU and storage resources.

Definition

The data is stored in the cloud, offsite and managed by a third party on behalf of the organization. This is similar to remotely accessing data stored in on-premises servers. This setup allows an organization to reap some benefits of public cloud without having to share the hardware resources or data. However, this approach is difficult to scale. You need to inform (and also pay) the third party if you want to upgrade or add hardware resources. This could add significant time to deploying new backup resources.

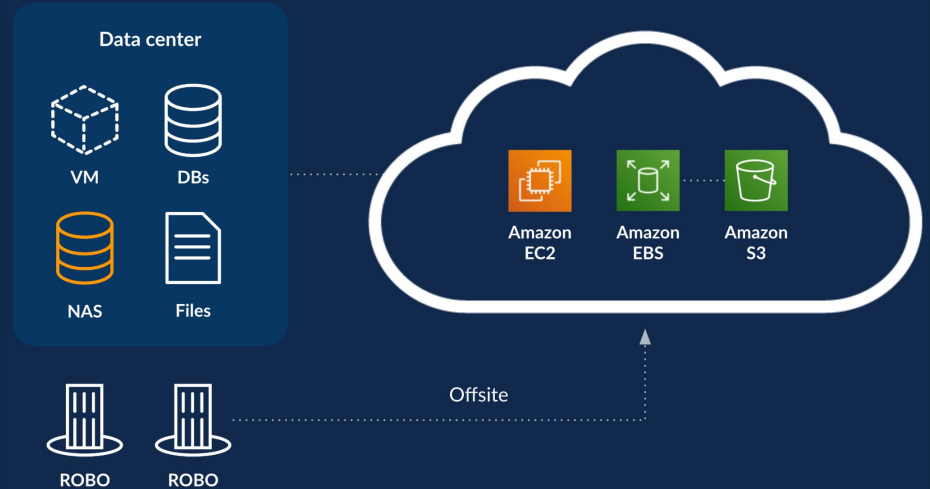


Cloud-Native

- **Where the data is stored.** Public cloud providers such as Amazon or Azure
- **Who manages these systems.**
Shared responsibility model – Vendor / Customer
- **Advantages.** You don't need to buy or maintain any backup infrastructure. You can scale up or down automatically without having to source and deploy new hardware. Cloud-native solutions allow for automatic patching, upgrades, and features to be done by the backup provider without impact to the customer. A cloud-native backup platform can be more secure than traditional approaches because it provides a logical air-gap and can take advantage of cloud security models.
- **Disadvantages.** You need internet to backup or restore data. Initial seeding for large environments may be time consuming

Definition

Backup data is stored in data centers owned and operated by the public cloud provider. To protect data from site-level failures, errors, and threats, the stored data is auto-replicated across three or more centers within a region. This ensures that your data is always safe, available, and accessible from anywhere anytime. As the cloud eliminates the need of buying and maintaining backup infrastructure, it is the most cost-effective, agile, and easy-to-manage data protection method.



Can Cloud Data Protection Meet Your RTOs?

When an event happens that takes your business offline, time is money. You want to recover your data as quickly as possible and if the problem system can't be recovered, data must be restored from backup. 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

Meeting RTOs from the cloud is often far easier than most organizations think. Of course WAN bandwidth and latency are important considerations, but direct recovery from the cloud is often undervalued. Why? The existing infrastructure for recovery may also be impacted or unavailable adding additional challenges to recovery. By contrast, cloud-native systems scale on-demand for multiple recovery requests and can support a local recovery cache. When combined with greater reliability, cloud-native (and hosted) often enable better RTOs than on-premises solutions for:

- **Active system recovery.** Most backup and recovery architectures can support active-system recovery, but hybrid, hosted and cloud-native solutions offer coverage for both on-premises and cloud applications. Hosted and cloud-native models also support extremely fast RTOs for endpoints and remote offices.
- **Archive recovery (cold data > 1 year).** On-premises and hybrid cloud solutions that rely on “shelf stored” tape backups add significant time to find the physical media and then recover, so are significantly slower than hosted and cloud-native solutions that use cold cloud storage.



Hosted Cloud vs. Cloud-Native

What's the difference?

Both hosted and cloud-native data protection services are consumed with a browser-like interface and may even offer consumption-based pricing. 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 cloud

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 (i.e., step-based consumption)
- Customer responsible for patches, updates, and security operations
- 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
- Pay as you grow and consume
- Vendor responsible for patches, updates, and security operations
- Multi-tenant, automatic updates

Finding the Right Cloud Data Protection Strategy

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

Not all data needs the same level of protection. Define your RTO and RPO needs across key categories like mission-critical, business-critical, and non-critical data. Identify your top-level goals for change and improvements and then pick a data set or location to start your evaluation.

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

Security Considerations

- Is the backup storage air gapped and truly immutable?
- What happens in the event of accidental or malicious deletion of backup data?
- What levels of resiliency are included in the backups or the ecosystem?

Management and scale

- How quickly does the system scale when you need more capacity and are there any surprise costs (for example, ingress or egress fees)?
- Across how many regions can you store data today? How can you add new regions?
- How does your solution handle bandwidth constraints?

Efficiency and costs

- Will your current storage footprint and costs decrease when switching to a new model?
- How easily can you understand and project future costs for both backup and recovery (e.g. ingress or egress fees)?
- What is the archiving model? How is data moved from warm to cold storage? What are the associated costs? Who manages the cloud archival storage?
- 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?

Reach for the Cloud

The value of cloud-native data protection

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.
- **Minimal IT personnel costs.** Cloud data protection—especially hosted and cloud-native—requires much less IT involvement than on-premises protection.
- **Reduced or eliminated on-site data protection hardware infrastructure costs.** No legacy environments to maintain, secure, or manage frees up IT resources for business innovation.
- **Reduced complexity within an IT environment.** Legacy, on-premises data protection systems can be complex and time-consuming to maintain. SaaS data protection solutions eliminate the need for you to manage on-premises and cloud infrastructure.
- **Reduced risk of business downtime.** Cloud-native solutions can go beyond backups to enhance ransomware protection, detection and any type of recovery without the need for infrastructure at an alternate site.
- **Greater manageability and compliance.** If all backup data is stored in the cloud, businesses can manage it from a single dashboard, better understand trends and predict costs, and more easily comply with data residency and business SLAs.



Druva: Cloud-Native Data Protection and Management

The industry's first and only at scale SaaS solution for data resilience

The Druva Data Resiliency Cloud, built on AWS, enables cyber, data and operational resilience for every organization and securely protects data from mobile devices and endpoints to data centers and multi-cloud environments. Customers can radically simplify data protection, streamline data governance, and gain data visibility and insights as they accelerate cloud adoption.

Key benefits include:

- **100% SaaS.** No hardware or software required.
- **Improved security.** The Druva Data Resiliency Cloud offers a logical air-gap for data, built-in user security like multi-factor authentication (MFA), continuous operational monitoring, 3rd party SIEM/SOAR integrations, and a range of government and privacy certifications (i.e., FedRAMP, FIPS, SOC2).
- **Improved business agility.** Free up IT resources and get more value from your data including eDiscovery, sensitive data governance, accelerated ransomware recovery, and cloud disaster recovery.
- **Simplified management.** A centralized, web-based console is easy to use and always updated, putting the most current protection for on-premises, SaaS, and cloud environments at your fingertips.
- **Low TCO.** Eliminate underutilized HW/SW and pay only for what you consume versus forecasting, buying, and managing storage you never fully consume. Avoid the operational costs and storage inefficiencies of managing multiple products and data silos for business continuity and compliance.





About Druva

As the industry's first and only at-scale 100% SaaS platform for data protection and cyber resilience.

[Watch the video](#) to see how Druva eliminates hardware, software, and operational complexity, and [visit the platform page](#) for a deep dive into how customers protect data while gaining new insights across endpoint, data center, and cloud workloads.



Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976	Japan: +81-3-6890-8667
Europe: +44 (0) 20-3750-9440	Singapore: +65 3158-4985
India: +91 (0) 20 6726-3300	Australia: +61 1300-312-729

Druva enables cyber, data and operational resilience for every organization with the Data Resiliency Cloud, the industry's first and only at-scale SaaS solution. Customers can radically simplify data protection, streamline data governance, and gain data visibility and insights as they accelerate cloud adoption. Druva pioneered a SaaS-based approach to eliminate complex infrastructure and related management costs, and deliver data resilience via a single platform spanning multiple geographies and clouds. Druva is trusted by thousands of enterprises, including 60 of the Fortune 500 to make data more resilient and accelerate their journey to cloud. Visit druva.com and follow us on [LinkedIn](#), [Twitter](#) and [Facebook](#).