Data Protection Matters

Every data protection strategy should be cloud-powered

The real decision is which cloud(s) to use

September 2025

Like every other aspect of a modern IT strategy, data protection ought to be delivered through a hybrid architecture that leverages cloud-hosted infrastructure as its foundation. That said, there are a few key considerations that organizations need to be clear on before choosing which cloud to build their data protection strategy upon.

Cloud storage is inevitable for data protection solutions in organizations of all sizes. Data resilience, grounded in the foundations of business continuity and disaster recovery (BC/DR), has always and will always require offsite data. What has changed is the necessity for that offsite data to be agile, evolving from cold tapes to warm or hot cloud storage. Gartner forecasts that 90% of organizations will adopt a hybrid cloud approach by 2027[1], which includes using cloud storage combined with disk for backup and disaster recovery as part of broader infrastructure strategies. This near-universal embrace of local disk-plus-cloud is the result of ever-heightening recovery requirements combined with the economic pressures of reducing on-premises footprints without sacrificing retention, regulatory compliance, or agile recovery.

If 'Security' is part of the question, then cloud storage is almost assuredly part of the answer. A definitive element of cyber preparedness is a layered approach to protection storage. This helps prevent bad actors that breach one tier of backup repositories (e.g. on-premises disk within the backup server) from affecting the other tiers as well. Beyond just immutability (also critical), a cloud-tier can provide isolation via separate credentials and alternative protocols to those perusing the local network, while still offering the responsiveness of disk for frequent protection and agile recovery.

Good
- Cost effective retention
- Offsite survivability
- Immutability

Better
- Separate credentials
- Differing access protocols to cloud vs. local storage

Best
- Multi-factor authentication
- Multi-admin authorization for highly sensitive actions

This paper describes what to look for in cloud storage when designing solutions for data protection, data resilience, and cyber-preparedness.

[1] https://www.gartner.com/en/newsroom/press-releases/2024-11-19-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-total-723-billion-dollars-in-2025

It would be easy (and incorrect) to assume that all cloud storage is equal

The good news is that backup and other data protection tools can mask the intricacies of cloud storage as simple tiers of retention. The bad news is that the simplicity of usage often reduces the ability of consumers to understand the merits of purpose-built clouds versus generic or hyperscale storage, often with significant long-term operational or economic impacts. Three cloud storage architectures to consider:

Hyperscale storage
e.g., Amazon S3 or Azure Blob

Purpose-built clouds
offering protection storage

Backup as-a-Service
(BaaS)

There are a wide range of offerings across myriad hyperscale clouds, purpose-built solutions, and BaaS/DRaaS[2] providers, but here are a few architectural contrasts to frame your consideration.

	Hyperscale storage	Backup as-a-Service (BaaS)
What you get	You probably already have a subscription and maybe even 'free' credits if you consume other cloud services from that hyperscale vendor. But what you are getting is generic storage capacity (block, file, or object) that is completely self-maintained.	They can be as "turnkey" as turning on a household utility or subscription. Some offer concierge services, while others simply provide a curated repository. BaaS and DRaaS service providers are distinguished not necessarily by the storage they provide, but by the expertise they offer.
Cyber considerations	At best, you might get immutability for your storage repositories, which may or may not be supported by your various backup software solutions. Nearly all other cyber preparedness features are typically delivered by the backup software, not the cloud service, which is problematic since most orgs run more than one backup technology.	Most cyber capabilities within BaaS are based on each service provider's choices of underlying software choices – or within current backup software. Some service providers combine their preferred backup software, cyber toolsets, and storage platforms, which can work well, so long as you agree with their choice(s) of toolsets.
Economic considerations	Hyperscale (generic) storage is not meant for backup/recovery scenarios, which creates two significant budgetary dangers: 1) Surprise extra costs due to storage bloating when daily and weekly immutable backups are retained instead of overwritten and the increased API calls when immutability is enabled on generic cloud storage. 2) Egress fees incurred during recovery tests. A Cyber/DR plan that is not tested is not a 'plan' but a 'hope.' But egress fees can dissuade testing, thereby reducing your recoverability when you need it most.	BaaS is meant to be a combination of expert services and turnkey software-plus-storage that is outcome centric. That can work well, until you want a different backup/cyber stack that has different recovery features or protects different workloads – or you want a different level of expertise or service delivery. Then, you'll get a *Hotel California* lock-in where "*you can check-out, but you can never leave*" the original solution/provider without breaking your long-term retention and cyber-preparedness – unless you run two BaaS solutions until the first data set data expires.

2 Backup as-a-Service (BaaS) and Disaster Recovery as-a-Service (DRaaS) differ significantly in their outcomes and methods, as well as the amount of expertise and outsourced services. For the purposes of this paper, consider them as a managed service offering backup/replication to cloud repositories that are curated by the provider.

Making the case for purpose-built protection storage – in a cloud

Decades ago, backup software used generic disk storage, often simply a cheaper per-terabyte (or older) variation of their production storage and then came purpose-built protection storage appliances. While most originally chose these appliances for features like deduplication, two key benefits were:

- Their performance was architected for backup scenarios, including IO optimizations that allowed for faster ingest and more efficient retention.
- Their benefits were typically consistent regardless of which backup software you initially chose; so, if you later switched or added a different backup software, you could effectively retain the data for as long as your regulatory mandates required it without sustained usage of the software. When you were ready, you could repurpose the storage capacity for use by your new backup software while preserving your investments. Thus, the same protection storage appliance could store the repositories of multiple backup software stacks, with secure, siloed access regardless of what was writing to it.

Even in modern data centers, most organizations still run multiple backup software stacks, now integrated with multiple cyber toolsets. Sometimes, the diversity of tools is forced by the diversity of production workloads (on-premises and cloud-based) to be protected. In addition, mergers, regional preferences, and 'shadow IT' often introduce diverse tools into even the most standardized environments.

The benefits of purpose-built protection storage enjoyed in on-premises appliances are similarly attainable with can apply to purpose-built cloud storage – including performance optimizations for key usage scenarios (e.g., immutability and cyber/DR testing) and the ubiquity of capabilities when multiple backup software tools are in use.

Why purpose-built cloud storage makes so much sense (and cents) for data resilience and cyber preparedness

The right cloud storage for data protection unlocks significant agility and flexibility for recovery, which is why your data protection solution(s) should utilize cloud storage for most of your retention and resilience goals. The key role for on-premises storage is inarguably fast recovery, especially for operational restorations and potentially for cyber resilience. That said, cloud storage that is designed with cyber/DR preparedness in mind has some notable characteristics to look for:

- When cloud storage is designed for immutable backups and cyber restorations, it does not unpredictably bloat, nor do the API calls increase exponentially. This means that a cloud designed with cyber-preparedness in mind won't surprise you when you see the monthly bill.
- Immutability (and other elements of cyber preparedness) shouldn't have to radically vary by the backup software, especially since you are likely running more than one and will eventually want to swap tools while retaining data.
- If your recovery tests or BC/DR exercises are incurring egress fees (which will invariably reduce your willingness to test), then you need a different cloud.

Key topics to cover with your teams:	
Technology choices • **Which backup solution(s)** are protecting our production data? • **Which cloud(s)** are each of those tools using as protection storage? • **If using BaaS/DRaaS**, does the provider allow us to specify the cloud storage that fits our preparedness strategy?	*Understanding which backup solutions are in use, throughout your business units/regions across your data centers & cloud-hosted production environments, will reveal your first heterogeneity challenge – and potentially an opportunity to consolidate.* *Consider your isolated silos of protection storage, including backup appliances, on-premises arrays, and clouds (hyperscale and BaaS/DRaaS) – and whether each of those silos supports a modern and layered approach to data protection and cyber preparedness.* *One very compelling differentiator for BaaS providers is whether they are only using their own storage, back ending to hyperscale storage, or allow you to choose the storage clouds that fit your needs.*
Impacts to our organization • **What costs do we incur** when we conduct large-scale DR or ransomware test recoveries pulling data from the cloud back to our data center as part of our preparedness testing or audit? • **If we migrate to a new backup platform**, how will we preserve the 5+ years of long-term retention mandated for regulatory compliance? • **What if a threat-actor gains administrative credentials** and wants to delete our backups?	*Each of these questions has significant operational implications:* • *If your cloud incurs egress or other storage fees when you are testing your recoverability from disasters or ransomware, you will inevitably have tension between rigorous testing and budgetary caution. Remember, if you are not regularly testing, you do not have a 'plan', you have a 'hope'.* • *If your retention is tied to a complete software/service stack and you change software tools, then you have two choices: (1) run both stacks for potentially years, or (2) lose your historical copies.* • *In the face of ransomware, using separate credentials (with MFA and potentially multi-admin authorizations) is likely the difference between having data to recover from or paying the ransom.*
Moving forward • **When was the last time we assessed purpose-built cloud storage** as a potential part of our data protection, BC/DR, and cyber-preparedness strategies?	*Most data center backup solutions started with generic storage that grew organically as the amount of data being backed up continued to expand. Candidly, there is rarely a convenient time to reimagine your IT architecture; yet the benefits of embracing purpose-built storage are measurable in dollars saved as well as capabilities gained.* *The same is true for cloud storage for data protection, where generic hyperscale capacity was easy to initially add and BaaS can seem like an 'easy button'. For many organizations, the benefits of purpose-built clouds for data protection unlock operational capabilities and economic savings while maintaining the independence to choose software stacks that make sense (and cents) for your organization.*