

Cloud Buyer's Guide

Tips for Developing Effective Hybrid
Cloud Data Strategies.

Table of Contents

This buyer's guide provides advice to help you increase the effectiveness of your hybrid cloud with a focus on data strategies for running hybrid applications.

Introduction	03
Key Differences Between On-Premises Enterprise Environments and Public Cloud	04
How Can You Support Enterprise Applications in the Public Cloud?	07
How Can You Modernize Your Applications and Future-Proof Data Consumption?	11
How Can You Protect Data Across Your Hybrid Cloud?	13
Conclusion	16



Introduction

In the cloud era, enterprises are searching for the best cloud strategy to increase agility, decrease costs, and foster greater innovation. Some have experimented with moving everything to the public cloud but soon discovered that not every application fits there, requiring significant porting effort and increasing costs without delivering the expected benefits. Other companies have adopted a “cloud-first” approach where new projects, including application updates and substitutions, are targeted to the public cloud. This approach has helped IT teams become more cloud-aware, but the results have been mixed—for many of the same reasons.

The truth is that there's no easy substitute for thinking carefully about the requirements of your applications and your business. There are many considerations in hybrid cloud deployments – from migrating to cloud-native application architecture, to optimizing applications with partial

adjustments to take advantage of cloud-native features, to re-hosting by migrating applications to the cloud without code changes. That's why most enterprises now take a hybrid cloud approach, with each application workload running in the environment that best addresses both application needs and business requirements.

An effective hybrid cloud strategy should take into account the reality of your existing IT operations. You have to prepare for the future without ignoring the past. This may mean combining private cloud, SaaS and PaaS cloud services, and public cloud, supporting traditional enterprise applications and cloud-native applications, or providing the ability to run VMs and containers in the same environment.

How you handle data may have a big impact on your success. Cloud data storage is different from the enterprise storage you're used to. If you need application compatibility between environments, consistent APIs, and easier migration, shared data services that span on-premises and multicloud environments may help eliminate roadblocks.

This buyer's guide provides advice to help you increase the effectiveness of your hybrid cloud with a focus on data strategies for running hybrid applications. It includes suggestions for services that make the public cloud more friendly to enterprise applications, targeted updates you should consider for your data centers/private cloud, and data protection approaches suited to hybrid cloud environments and modern business needs. We hope it offers some insights that will help you make more informed cloud decisions for your organization.

93%

In 2019, 93% of organizations are committed to or interested in hybrid-cloud strategies.

Source ESG Research Report: Hybrid Cloud Trends: [Strategies for optimizing and managing on-premises and public cloud infrastructure, December 2019.](#)



Key Differences Between On-Premises Enterprise Environments and Public Cloud

A significant divide exists today between enterprise IT environments and the public cloud, with different management models, different consumption models, different application architectures, and different storage and data services.

- Traditional on-premises management is still centered around manual processes while the public cloud focuses on self-service and programmatic control of everything via well-defined APIs.
 - The public cloud replaces the traditional CAPEX model of the enterprise data center with on-demand consumption that matches costs more closely to usage.
 - The [cloud-native application](#) architecture relies on resources that scale out rather than scale-up. Enterprise applications have different requirements for compute and I/O performance and different expectations for availability and data protection that public cloud providers may not meet.
 - Cloud storage services differ in terms of storage types, resiliency, and native capabilities, as described above.
- If you want to create an effective [hybrid cloud](#), you'll need to identify which of these differences create the biggest challenges for your operations and take steps to close those gaps.

	Traditional On-Premises	Public Cloud
Management	- Manual management	- Comprehensive APIs - Everything controlled by software
Consumption	- Dedicated hardware (CAPEX)	- On-demand consumption (OPEX)
Application Architecture	- Packaged applications in VMs - Application silos	- Webscale applications in containers - PaaS services
Storage	- Dedicated arrays for block and file - Highly reliable - Rich feature set (snapshots, clones, replication, etc.)	- Object, file, and block storage services - Higher durability but lower availability - Simple but scalable





Enterprise IT Environments in the Real World

For the sake of simplicity, this buyer's guide primarily frames the discussion in terms of the enterprise data center. But in the real world, enterprise IT environments are complex and can encompass a mix of primary and secondary data centers, remote offices, hosted services, co-location facilities, and so on. While you may not need the capabilities described in this guide in every location, the advice still applies to any location where you do.

3 Things to Know About Cloud Storage

Public clouds offer a variety of data stores—with characteristics that are significantly different from the storage in your data center. Before diving into hybrid cloud, it's helpful to understand the differences.

01: Cloud Storage Has Different Resiliency Characteristics

Cloud data services can be highly durable, but may not offer the availability that enterprises are used to. This means the cloud won't lose your data, but you may not be able to access it. Native cloud storage (of all types) may offer up to five-nines (99.999%) availability; enterprise applications often expect six nines (99.9999%).

That corresponds to almost an hour of unplanned downtime per year versus less than a minute of unplanned downtime. This difference is significant. Recent well-publicized outages such as the weather-related [Azure outage in September 2018](#) underscore the risks.



99.9999% = less than one minute of
unplanned downtime per year



02: Object Storage Predominates

Enterprise IT teams—and traditional enterprise applications—rely on block and file storage. However, while block and file are available in the public cloud, object storage is more widely used. Rather than the blocks, sectors, and tracks of block storage or the directory hierarchy of file storage, each data item is treated as a discrete object and has a globally unique ID rather than a name. This enables the namespace to be vast and to span devices, making object storage simpler to use and eliminating some of the complexities of scale during application development.

Enterprise applications have traditionally been built with block storage where data is stored in fixed-sized chunks. Block storage has historically been used for applications that require high performance because the data is quickly addressed piecemeal and there is no overhead for metadata. While block

storage is available in the cloud it is primarily just basic storage and lacks enterprise data services.

03: Native Cloud Storage May Lack Features You Depend On

The storage in enterprise data centers has rich features such as efficient snapshots, cloning, and replication. Your team may have incorporated these as part of your [data-protection strategy](#) and in common automations and workflows. For example, fast, efficient snapshots can be indispensable for data protection, Dev/Test, and other functions. Cloud data services focus on simplicity. While they may offer snapshots and replication, these capabilities lack the efficiency that results from deduplication, thin provisioning compression, and careful metadata management. In practice, this means cloud services may not take the place of the functionality you're used to.

Common Public Cloud Services

	AWS	Azure	Google Cloud
Object	Amazon S3	Blob Storage	Google Cloud Storage
Block	Amazon EBS	Disk Storage	Cloud Datastore
File	Amazon EFS	Azure Files	Cloud Filestore (beta)



How Can You Support Enterprise Applications in the Public Cloud?


Most enterprises have dozens, even hundreds, of traditional enterprise applications running on-premises. There are many reasons that you may want to move some of these into the public cloud. You may just want to get less-critical applications out of your data center to free resources for more important ones.

There are several possible options for moving applications to the cloud:

- Popular enterprise applications are increasingly available via a [software-as-a-service \(SaaS\)](#) model. Switching to SaaS not only gets the application out of your data center, it eliminates most of the management overhead (but not your responsibility for company data).
- You may consider re-platforming some applications, so they take on more cloud characteristics, for instance, replacing a traditional relational database with cloud database services.
- Some important applications may need to be completely refactored (re-architected) to run in the cloud.

Re-platforming and refactoring are generally only considered for the most important applications. In many cases, the ideal option may be to simply to lift and shift an application into the cloud with the minimum of effort, but most enterprise applications have requirements that public clouds don't satisfy very well.

Fortunately, there are third-party solutions available that [address storage and performance challenges](#) and provide greater API and management consistency, making it easier to run enterprise applications in the cloud without significant changes.



How Do You Deliver Enterprise Storage Characteristics in the Cloud?

Enterprise applications usually require block or file storage and expect data services to be highly reliable, to deliver high performance, and to provide enterprise features for data efficiency, data management, and data protection.

This is a marked contrast to the characteristics of the native block and file data services that are available in the public cloud, which are optimized for cloud-native applications that have resiliency built-in and grow performance by scaling out.

Re-architecting enterprise applications is risky, time-consuming, and expensive. An easier alternative is to choose data services that have been designed to meet the needs of enterprise applications. Fortunately, several block-and-file data services from well-known vendors are now available for enterprise workloads in the cloud. These services layer enterprise capabilities on top of existing cloud data stores, offering pay-as-you-go consumption while delivering the high availability and efficiency your applications expect so that applications don't have to be redesigned. They also offer snapshots, cloning, and [replication for data protection](#), disaster recovery, and other management functions.

These solutions may include mobility services between your data centers and the cloud.

Consider using these services in place of native cloud data services to support any enterprise application you move to the cloud. Some services may also help you address other important use cases including backup to the cloud, DR to the cloud, migration to/from the cloud, and cloud-to-cloud high availability.

What About Enterprise Application Performance?

Native cloud data services range in performance and cost from slow and cheap to fast and expensive. To make sure the cloud meets the performance needs of enterprise applications, IT teams tend to provision the fastest and most expensive data service by default. This means you often overpay for storage. To overcome this, there are now enterprise data services available in the public cloud that aggregate performance across multiple data stores to deliver predictable performance while reducing your total storage costs when it's not.



How Can You Deliver API and Management Consistency?

There are additional considerations when you move an enterprise application to the cloud:



Automation

Existing scripts that you've developed for use with an application in your data center may not run as expected—or not at all. This is usually because the automation depends on tools and services supported in different on-premises and cloud environments. Plus, differences in data services and capabilities like snapshots, de-duplication, thin-provisioning, and availability may require changes to application design for portability.



APIs

Applications written to make calls to specific APIs present similar challenges since APIs are different on-premises and in the cloud, or between any two storage backends. Scripts for storage provisioning must be rewritten to help with application portability in these cases.



Management

Because the management interfaces and tools in the cloud are different than those on-premises, there can be a significant learning curve. Anyone who switches back and forth between environments is likely to be less efficient. This adds to the burden on your IT team and increases the risk of errors. Administrators may need retraining to manage and adapt to the tools and interfaces of diverse infrastructure environments.




Challenges

Because these challenges usually involve data storage in some way, choosing data services that provide enterprise capabilities is often the fastest and easiest solution. You can take this a step further by choosing solutions that offer consistent automation tools, APIs and management across your on-premises and multiple public cloud environments.

This may enable applications to run in the cloud without changes; at a minimum it will reduce the porting effort necessary to get scripts and applications working. A consistent management interface across environments can also unify management tasks and reduce the risk of operator error.

Key Success Factors

Choose cloud data services that:




Satisfy

Satisfy the resiliency, performance, data protection, and security needs of enterprise applications.




Facilitate

Facilitate efficient and easy application/data mobility between environments.



Aggregate

Aggregate performance and capacity across multiple block stores for greater economy.



Integrate

Integrate with your on-premises storage and seamlessly extend to multicloud environments.



How Can You Modernize Your Applications and Future-Proof Data Consumption?

What Are the Alternatives for Increasing Data Center Automation?

Traditional enterprise data centers rely on ticket-based systems. Administrators satisfy user requests taken from a queue with little reliance on automated processes. The limitation of this approach is that it doesn't scale. If your goal is to become more agile and more responsive to changing business demands, you'll need to look for ways to increase automation. Developers who are used to working in cloud environments want to consume resources on demand. Providing automated self-service for development teams can decrease time to market for applications and digital services.

A variety of cloud automation platforms are available to automate common tasks. For example, the [VMware vRealize Suite](#) provides a hybrid cloud management platform that is now significantly more automated and more developer-friendly. OpenStack and Kubernetes- based Platform-as-a-Service technologies such as Red Hat

OpenShift are also popular options. Make sure you understand your goals and needs before selecting automation solutions and consider standardizing on one platform or as few as possible. If you are adding or refreshing infrastructure, look for deep integration with your chosen platform.

VMs or Containers?

Containers are becoming mainstream as the adoption of modern, microservices-based application architectures grows. You'll need the ability to run containers to support these applications in your data center. Stateful applications running in containers require the ability to quickly provision, scale, and recover persistent storage volumes. Unless you are certain that your hybrid cloud will be strictly VMs—or strictly containers—for the foreseeable future, the smartest course of action may be to make sure your data center choices allow you to support both VMs and containers, including integration with Docker and Kubernetes.

To future-proof your data center, consider making all your storage resources available “as-a-service” and programmable through REST APIs.

Can You Adopt On-demand Consumption in Your Data Center?

Another characteristic that distinguishes enterprise data centers from public cloud is the consumption model. Today's application workloads consume resources elastically, but resource procurement for the data center can be both capital intensive and inelastic. You may find yourself overprovisioning and paying for resources months or years before you need them, but this situation is changing. Forward-looking storage vendors offer on-demand storage services for data center needs, enabling you to adopt an OPEX model for storage consumption across a hybrid cloud that more closely matches your costs to growth in usage.



Key Success Factors



Choose an OPEX Consumption Model

Choose an OPEX consumption model that allows you to flexibly assign resources across on-premises and public cloud environments. Modernize and automate your data centers.



Add Object Storage and Storage-As-A-Service

Add object storage and storage-as-a-service to support cloud-native applications.



Choose Storage That...

Choose storage that...

- Supports both containers and VMs
- Can be programmed using REST APIs
- Delivers flash performance and reliability
- Includes features for integration with public cloud



Moving to a hybrid cloud model provides the perfect opportunity to re-examine your data protection strategy and objectives.

Some questions to consider include:

- How will you protect enterprise applications that move to the cloud?
- Will you use the public cloud as a data backup and/or disaster recovery target?
- What additional protections do cloud-native applications need?

However, before you settle on answers to these questions there's an even more important question to ask:

- Do your current data protection and disaster recovery methods still meet your business needs? If not, what should your objectives be?

How Can You Protect Data Across Your Hybrid Cloud?

When it comes to important applications, downtime has a direct impact on revenue—and your brand. If you've been using the same [data protection strategy](#) for a while and haven't examined your SLAs for recovery lately, you could be under-protected. Recovery from tape—or even from disk—may no longer be fast enough. You may need to consider a data protection strategy based on flash for rapid recovery and object storage as a long-term archive.

Why flash for data protection?

An ideal data protection strategy is to combine flash primary storage with a secondary flash copy and long-term retention in a cloud object store. Modern deduplication technology can make a strategy based on flash and object storage efficient and cost-effective, offering cost parity with disk and tape methods. Having a second data copy on flash ensures the fastest possible recovery in the event of an unforeseen failure or data corruption.

Why cloud for long-term retention?

Tapes typically accumulate in an offsite facility and you can't get at the data easily to use it for other purposes. Backup data copies in a cloud object store can provide fast and simple offsite storage that's both more durable and more convenient than tape. You can retain copies of the most critical data in multiple cloud zones/regions to achieve extremely high durability. Data copies in a cloud object store remain accessible for [recovery in the cloud](#) or to your data center, and the data can be easily repurposed for dev/test, analytics, or other needs that may arise in the future.

How do you implement flash data protection?

The simplest way to implement data protection based on flash and object storage is to choose flash storage that incorporates cloud data protection any time you refresh storage systems. The right flash systems make it simple to combine local snapshots with snapshots stored in the cloud. Choose snapshot technology that is efficient and self-describing to enable both fast local recovery and easy restore in the cloud.




In any data center, there's always a certain amount of heterogeneous hardware that resists all attempts to modernize. Solutions are now available that will allow you to back up this data to a secondary flash array on-site and then replicate to the cloud, providing the same fast recovery and easy reuse benefits.

Do enterprise applications in the cloud need protection?

At a minimum, enterprise applications that you move to the cloud need the same level of backup and disaster recovery they received in your data center. Backups are necessary to enable you to recover to a previous point in time should data corruption occur and as a protection against malware and other malfeasance. If you choose cloud data services with the right enterprise features, as described earlier, you can implement a flash and object strategy for applications in the public cloud.

What protection do cloud-native applications need?

Because cloud-native applications are often architected to be resilient, the importance of data protection can get overlooked. These applications still require backup and may need additional levels of disaster recovery as well. Solutions are now available to enable [fast and efficient data protection](#) for applications of all types running on object stores in the cloud. These let you implement flexible data protection and DR while minimizing your cloud footprint and cost.



Key Success Factors

To modernize your approach to data protection for all application types:



Adopt a Strategy

Adopt a strategy based on flash primary and secondary storage for fast recovery plus object storage for long-term retention.



Update On-Premises Storage

Update on-premises storage to increase performance and enable cloud connectivity.



Choose Cloud Data Services

Choose cloud data services that enable these same protections for applications running in the public cloud.



Conclusion

Your infrastructure choices can make or break your hybrid cloud strategy. By asking the right questions as you design your [hybrid cloud](#), you maximize your chances for success. What do you want to achieve? Identifying your goals will help you zero in on the biggest pain points and attack those first.

Your primary goal may be to find better ways to run traditional applications in the public cloud. You may need to enable cloud-native applications and cloud-native development in your data centers. Or your initial goal may simply be to figure out how to leverage the public cloud for data protection or Dev/Test.

How can you better support enterprise applications in the public cloud?

One option is to choose data services that meet the feature and performance needs of enterprise applications. Taking that a step further, selecting data services that are compatible with your data

center storage ensures that APIs and management are consistent across environments. This enables applications to move more easily between your data centers and the cloud—or even span environments.

How can you prepare your data centers to support cloud-native applications and integrate more smoothly into your hybrid cloud?

A few specific data center improvements will enhance your overall hybrid cloud results. Adding [container support](#) and object storage can enable cloud-native applications to run on-premises. Automation tools for persistent storage delivery across an elastically scalable storage pool can increase your organization's agility and facilitate cloud-native application development.

How do you provide the right protection for all your applications and data?

As your hybrid cloud takes shape, it may be wise to re-evaluate data protection and [disaster recovery](#) to ensure that the protection you are delivering meets your business needs. If your recovery objectives have changed such that recovery is no longer fast enough, you may need to consider backing up to flash storage for accelerated recovery. If you need to repurpose backup data for other uses, using object storage for long-term backup retention keeps data more accessible than tape.

Next Steps

Pure Storage® offers cloud data services that bridge on-premises and public cloud, and enable true hybrid operations. Learn how a Modern Data Experience can power your hybrid cloud: www.purestorage.com/cloud.

purestorage.com

800.379.PURE

