CDW PEOPLE WHO GET IT™

CLOUD STORAGE AND BACKUP

Whether to complement a storage strategy or replace it, cloud computing offers cost-effective models for optimizing today's complex IT infrastructures.

Executive Summary

The proliferation of cloud computing options has begun to change the way storage is thought about, procured and used. IT managers and departments need to think through how cloud options might fit into and complement their onsite data infrastructures.

This white paper explains cloud storage and backup, providing advice about the tools and best practices for its implementation and use. Whether in the IT office or on the management team, readers will find some useful takeaways about how to take advantage of cloud storage for high availability, backup and restore, archiving, business continuity, disaster recovery and other uses.

Table of Contents

2 Cloud Storage Benefits

2 The Big Picture

4 Converged Solution Bundles

5 Cloud Storage Modes: IaaS, SaaS and AaaS

5 Accessing Cloud Storage

6 Evaluating Cloud Storage Services and Solutions

SHARE THIS WHITE PAPER

Cloud Storage Benefits

There are many types of clouds to support storage, including public, private, hybrid and community. Cloud storage is available as a service, product or solution that in turn can be used for building or delivering other information services.

When considering cloud storage, the IT team should make sure the chosen option meets the often-wide variety of expectations and needs:

- ✓ Multitenancy, to protect the privacy of different users in a shared environment
- ✓ Measurability, for insight into how resources are used to deliver services
- ✓ A safe, secure, flexible, dynamic and resilient environment
- ✓ An opportunity to enhance service and availability using external resources
- ✓ The ability to scale and adapt in a dynamic fashion to changing conditions
- ✓ Stability in terms of performance, reliability, capacity and cost
- ✓ Rapid, self-service provisioning of resources or service features
- ✓ Efficient, cost-effective resource usage that meets quality of service and service-level agreement (SLA) policies
- ✓ Interface to existing technologies and applications
- ✓ Options for supporting business continuity, disaster recovery and archiving

The benefits and characteristics of the various flavors of cloud storage are many. Although the obvious upfront cost savings from reducing footprint and equipment maintenance requirements are commonly associated with cloud storage, there are other attributes that provide benefits as well (see the *Cloud Attributes* sidebar).

The Big Picture

The growth in cloud computing has led to the development of public, private and hybrid clouds — all of which can be used to deploy community clouds.

But how are these different? Here's a basic way to define them: A service provider manages and controls a public cloud for free or for a fee based on the offered capabilities and usage modes. An organization controls and manages a private cloud, on premises or offsite. A hybrid cloud combines public and private cloud computing capabilities.

Community clouds can be private to the member organizations with shared or private (separate) tenancy. For example, medical or healthcare organizations can create a community cloud to share resources and reduce management complexity (and costs) while meeting regulatory compliance and security demands.

Cloud storage can be used to complement or replace existing storage systems and services. One of the most likely scenarios involves using cloud services to supplement primary storage by providing a remote destination for backup, business continuity and disaster recovery.

Another complementary strategy would be to use a public-cloud storage service as an archive destination if the

Cloud Attributes

ATTRIBUTE	DESCRIPTION	BENEFIT
Geographically dispersed	Distance from production systems	Protects against local threats or disasters; provides enhanced operations survivability and availability of IT resources
Metering and metrics	Tracking who uses what and when to assign costs	Aligns spending to work being done or information being stored
Reduced capital expense (capex), power and cooling costs	For public cloud, a move from capex to operating expense (opex) along with lower power and cooling costs	Provides budget and spending flexibility by aligning resources to operations needs based on monthly use or subscription
Improved return on investment (ROI) for capex of private cloud	Reduced time to deployment and management	Improves ROI on capex for private cloud resources, agility and staff productivity
Dynamic and elastic	Fees based on what's used, on demand	Stretches budgets, and scales to support peak or workload spikes
Agility and mobility	Flexibility to use different resources when needed	Adjusts to changing business conditions or opportunities on the fly
Reduced management burden	Technology needs and services handled by provider	Lets IT staff focus on functions closer to organizational mission or operational needs
Service-level agreement (SLA) alignment	Service catalog and menus detail service options	Aligns to meet availability, security and performance requirements
Faster provisioning	User or self-provisioning of storage needs through automation integration	Increases productivity by allowing faster access to storage resources

organization has a data footprint reduction plan. By archiving active data from primary storage and moving it to a cloud service provider, an organization could gain more effective use of onsite storage capacity and also potentially reduce its storage footprint. In addition to general storage for file and data sharing, email, database, backup and archiving, cloud storage is also a good fit for server virtualization — to house virtual machines as well as virtual desktop infrastructures (VDIs).

Cloud storage, a variant of cloud computing, offers a different approach to deploying, managing, using and consuming IT resources than the traditional data center–focused storage environment. Instead of dedicated servers, storage, networking hardware and software management tools, an IT cloud shares these resources and makes them available for use.

An organization can procure a cloud service that is fully managed, partially managed or not managed by the provider. Additionally, the service can provide many options and tiers defining the types of networks, storage and server performance capabilities provided.

Depending on the service-level agreement that the organization sets with a cloud provider, optional fees may apply for upgrading from one type or tier of server to another (such as to faster processing or more cores), for adding more memory and storage capacity, or for boosting I/O and networking performance. Plus, the storage options through cloud services can include a variety of backup and continuity functions.

With cloud storage services, the organization can add capabilities as needed — even on a temporary basis — allowing for improved cost management.

Public Clouds

Shared cloud storage has several different meanings depending on the context. Most often it's viewed as using a service available to the public from a provider (such as Amazon, AT&T, Barracuda, CDW, HP, Microsoft or Rackspace, among others) to save and share files based on access privileges. Today, service providers are also offering semipublic, focused cloud storage services. These are essentially public clouds tailored to the needs of specific vertical industries, such as government.

For some organizations, a public-cloud service provides a way to let users save files or data while not authorizing them access to primary stored information.

When storing data in a public-cloud environment, a smart best practice is to maintain a copy of any critical data onsite or at another location. That alternate copy need not be current, depending on its value or importance, but unless the enterprise can live without it, it's prudent to have that copy.

Another consideration (particularly with services that the organization is using for storing data, such as backup or archive files or files to support business continuity and disaster recovery) is how quickly a restore of a large amount of data can be accomplished. It might be possible to restore individual files from an online service very quickly.

But hundreds of megabyte- or gigabyte-sized files, or terabytes of data, may take longer depending on the type of cloud service and the details of the SLAs. This could be a crucial factor in an enterprise's cloud choice for storage.

For this reason, the IT team should look into what mechanisms a service provider offers as part of its base package fee and what's available for an add-on cost to support bulk data import and export (and via what media).

It's also important to keep in mind that when sending data electronically to a cloud service, data footprint reduction (DFR) techniques such as compression and deduplication are usually employed on the organization's systems or via a gateway appliance. What this means is that data may flow faster to a cloud than it does when failing back over to enterprise systems.

That's often the case even when an organization's network supports faster download speeds. Why? Because during a restore, the download will typically involve a full file while the backup process might only require uploading changes in files or other portions of the data being protected.

Private Clouds

Private clouds are similar to public clouds except that they are intended for use or consumption on an internal organizational basis. A private cloud can be built and located entirely internally or it can leverage external and public resources.

The lines between a private cloud and a more traditional IT infrastructure can blur, depending on the definitions and solution offerings. In general, a private cloud has the same tenants or operating principles as a public cloud, including agility, elasticity, effective resource usage and metering for management insight.

In addition, private clouds can be extended to support chargeback or billing where applicable, as well as self-provisioning or procurement of resources by the enterprise's IT services consumers. Many of the same tools and technologies, from hardware to software and networking, are used for establishing both public and private clouds and traditional environments.

For example, EMC Atmos is one tool that can be used by organizations to create private clouds, or by service providers to create public clouds. Atmos is an object-based storage system that supports multiple interfaces and supplies built-in multitenant management and metering for accounting, chargeback and provisioning tools.

Server virtualization enables both public and private clouds by providing a mechanism for encapsulation of servers, their applications and their data. When encapsulated on a virtual machine (VM), apps can easily be moved to a private cloud or migrated to a public cloud as needed.

By decoupling data from the VM, similar to shared storage for physical servers, flexibility also exists in what storage capabilities are available. For example, a VM can be local yet point to remote cloud storage, or a VM can be moved to a cloud using storage provided by that service.

Moving Backup and Archiving Offsite

Cloud storage can provide a good alternative to supporting backup and archiving functions in house.

A cloud service provides a secure approach for an enterprise to offload management tasks, including media handling, to a service provider. Cloud storage services for backup, business continuity, disaster recovery and archiving enable the use of a geographically remote location for protecting data. Plus, backup functions tend to occur more frequently. The result is a shorter turnaround for recovery.

But what's the difference between backup and archiving?

In the case of backup, the provider maintains a copy of the users' production data for failover should a disaster bring down the users' production environment.

For archiving, the provider houses data (perhaps the only copy of the data) that is needed rarely or for recordkeeping purposes.

The SLA will define the recovery time objective (RTO — how soon) and the recovery point objective (RPO — how far back) for these storage services.

Archives tend to have a longer RTO and coarser RPO with a focus on restoring data over months, years or longer versus minutes, hours, days or weeks with backups.

Hybrid Clouds

When an IT team needs to control aspects of a private cloud for most services, yet also wants to leverage some application as a service (AaaS), software as a service (SaaS) or infrastructure as a service (IaaS) capabilities for other tasks, then a hybrid cloud is an optimal solution.

For some organizations, private clouds will be the initial starting point in their cloud journey. They provide an environment for conducting proof-of-concept research, and for gaining confidence in cloud services (and, ultimately, in the potential use of public clouds).

Following this logic, many organizations evolve from private-cloud environments to hybrid environments. Interestingly, for storage, some organizations have started with public-cloud use for backup and archive functions and are now determining how private or community clouds might fit into their environments — taking the reverse route toward hybrid clouds.

Community Clouds

Public, private and hybrid clouds get most of the attention. But organizations' concerns about security, privacy and compliance regulations have paved the way for yet another cloud model: the community cloud.

A community cloud may be déjà vu for some IT managers, who recall how compute and storage resources were made accessible to groups of users in the past through shared-services models.

The basic premise is that groups, organizations or consortia with the same or similar interests pool their resources and create a cloud for use by that community of users. Thus, the cloud is not public, but it's also not private, per se, to any one member of the community. Using multitenancy, each member essentially gains its own virtual private cloud.

Community clouds are a good option for local, state or federal agencies; schools; healthcare services; or other enterprises that have restrictions or concerns about public-cloud use, yet cannot afford to host their own private cloud.

Multitenancy: The Cloud Efficiency Model

One of the value propositions of cloud storage is the lower cost available from the sharing of resources in a safe and secure manner.

Service providers accomplish this through multitenancy. In essence, it's the data storage equivalent of shared occupancy in an apartment, townhouse or condominium complex. By leveraging a common infrastructure, yet providing separate private dwellings that offer residents security, environmental controls and utility metering, these multiuser housing complexes typically cost less per resident than separate physical houses.

With cloud multitenancy, different users can access the same shared-storage infrastructure, yet tap only their data and services. Cloud services leverage multitenancy, and it's a feature common among many storage solutions from leading storage vendors.

Converged Solution Bundles

Cloud and solution stacks come in many variations, from loose multivendor marketing alliances to integrated and tested interoperability technology reference architectures. They are essentially cloud building blocks that IT teams can acquire to speed the deployment of cloud services.

Organizations can choose to deploy solution bundles on premises or at a collocation or hosting site. And, as with other cloud tools, these bundled technologies can be managed

Storage Services Modes: A Snapshot

Mode	Characteristics	Functionality	Examples
AaaS and SaaS	Application or information services that eliminate the need to buy and install software and infrastructure	Archive, backup, email, office, payroll or expense, file or data storage, photo- or information-sharing services (focus is on consumption)	AT&T, CDW, EMC Google, HP Autonomy, Mozy by EMC, Quantum Q-Cloud, Rackspace and Seagate EVault
IaaS	Infrastructure components that provide cost, performance, availability and reliability options using different storage mediums and access	Back-end storage for AaaS and SaaS; remote storage tier for traditional apps such as general storage, file sharing, backup, continuity, disaster recovery and archiving	CDW, HP cloud services, Microsoft Windows Azure, Rackspace and Terremark

either by the IT team or a provider — depending on the human bandwidth of the IT department.

Stacks can include products from the same or different vendors, purchased separately or using the same SKU/part number.

Stacks typically focus on one of four areas:

- Application or functionality (database, virtualization, Big Data, email and VDI);
- Platform or middleware, including hypervisors (VMware, Hyper-V, Citrix Xen);
- Infrastructure (server, storage, networking, hardware, software and management tools);
- Data protection, backup and restore, snapshot and replication, archiving, and business continuity and disaster recovery.

Whatever its focus, a bundle will combine servers, storage, networking, hypervisors and management tools into an integrated component that can be centrally managed.

Some examples include EMC VCE Vblock, HP Converged Cloud, IBM PureSystems, NetApp/Cisco Systems FlexPod and Oracle Exadata. These solutions are used to create public, private, hybrid and community cloud infrastructures.

The converged approach helps to take some complexity (and, thus, cost) out of the equation for the user. These solutions can help an IT department jump-start efforts to move to a dynamic and flexible abstracted virtual environment. Cost savings can be had through the ease of acquisition, ease of installation, ease of configuration and, if management tools are provided, how well those tools automate provisioning.

Cloud Storage Modes: IaaS, SaaS and AaaS

In addition to the different types of cloud services (private, public, hybrid and community), there are also multiple modes for cloud storage. These modes refer to the layer or type of functionality provided.

The modes include AaaS and SaaS, along with IaaS and IT as a service (ITaaS). Cloud storage services modes presented as an application include backup, archive, file or folder sharing, synchronization, social media, collaboration, and photo or audio sharing, and fall under the SaaS or AaaS mode. In a hybrid cloud, an enterprise might manage a private cloud that relies on an IaaS public-cloud provider for its remote storage.

Within these service modes, IT departments make use of many of the standard management tools that they would find in a more traditional storage environment, from basic file systems to object-based access with multiple protocol support. This includes common protocols such as Network File System (NFS), Common Internet File System (CIFS), Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), Representational State Transfer (REST), Simple Object Application Protocol (SOAP) and Torrent, among others.

But there are also newer technologies that have been developed for cloud environments, such as the Cloud Data Management Interface and Hadoop Distributed File System.

Accessing Cloud Storage

To tap their cloud storage services and solutions, organizations must depend on networks, either through the public Internet or a private connection.

The type of cloud service being accessed will determine the network infrastructure needed. Cloud services may provide an access program or utility that allows the IT team to configure when, where and how data will be protected, similar to other backup and archiving tools.

For example, some services can be accessed using a standard web browser, while others require plug-in or add-on modules. Some cloud services may require downloading an application, agent or other tool for accessing the cloud service and its resources, while others provide an onsite or on-premises appliance or gateway.

Some traditional backup and archiving tools have added direct or third-party support for accessing IaaS cloud storage services. Third-party access appliances or gateways enable existing tools to read and write data to a cloud environment by presenting a standard interface, such as NFS, that is mapped to the back-end cloud service format.

For example, if the organization subscribes to an IaaS storage service for backup, storage might be allocated as objects, and various tools would be used to access or utilize those objects. The cloud access software or appliance would understand how to communicate with the IaaS storage application programming interfaces and abstract those.

Access software tools or gateways, in addition to translating or mapping between cloud application programming interfaces (APIs), format many of an organization's other applications that apply to storage. This can include encryption, bandwidth optimization and DFR (compression and deduplication).

Cloud access software and gateways or appliances are used for making cloud storage accessible to local apps. The gateways, in addition to enabling cloud access, provide replication, snapshots and other storage functionality. Cloud access gateways or server-based software includes tools from CommVault, EMC, Quantum and Symantec, among others.

In addition, the makers of many management and data protection tools have added (or are developing) support for access to public-cloud IaaS services.

Lock It Down

It should be understood, although sometimes the obvious needs to be said, that any data being sent or stored elsewhere needs to be encrypted. This means the data must be encrypted before it leaves an organization's systems, while it's being transmitted to its destination and while at rest with the cloud service provider.

Most data protection software provides the ability to encrypt data, as do most backup and restore, network bandwidth optimization and cloud gateway or access appliances. Providers of AaaS, SaaS and IaaS also offer a variety of forms and levels of encryption.

Currently, the most commonly deployed encryption tools use 128-bit or 256-bit Advanced Encryption Standard (AES) keys. For example, an AaaS or SaaS tool can support AES encryption of data before it leaves a facility or location, and another layer of encryption can exist at the IaaS provider layer.

In addition to encryption, Secure Sockets Layer (SSL) security is also commonly supported, along with other authentication and authorization techniques, including user names, passwords and access keys.

Of course, the type and breadth of security tools available in a particular environment are of little value if the IT department fails to enable them and users don't take advantage of them. That's where a shared responsibility model can come into play. In such a model, the service provider or manufacturer of a product has certain responsibilities, as does the user of the service to make applicable and appropriate configuration decisions.

Evaluating Cloud Storage Services and Solutions

Which cloud storage service, product or solution is best for the organization? Answering that question will depend on the enterprise's and its users' particular needs and requirements. And that decision will have to take into account numerous factors, such as constraints on where data can be located (inside or outside the country) and whether storage resources are used on a continuous basis or on a varied schedule.

When evaluating cloud storage, an IT team can view the consideration process as similar to deciding between leasing a piece of technology or buying it. There are capital and operating expense trade-offs. There will also be long-term and short-term cost trade-offs. Plus, an organization's technology refresh opportunities may affect the decision and shape its particular needs.

The good news is that every IT department has options to choose from to meet its specific goals and requirements, not to mention the ability to mix and match scenarios to devise the ideal hybrid solution.

When evaluating cloud options, cost is clearly a driver. The ability to reduce footprint — and all the ancillary costs it engenders — make public-cloud services appealing to many organizations.

But when evaluating any cloud service, it's critically important that IT and management teams understand what a provider's fees include and the associated terms of service. For example, if the fee has an asterisk or qualifier, then that cost is based on some volume of use or level of service.

As with anything that is promoted at a low price, cloud storage must be purchased with care. The IT team will need to consider the fees incurred as more data is added (or removed) or as access increases (or decreases). Two other fee factors to ask about are whether services at a particular fee are limited to a certain number of users or sessions (total and concurrent) and how the amount of access using the service affects fees.

Additionally, the IT team will need to evaluate any network infrastructure investments that may be necessary because of network traffic increases generated by viewing information or sending data. Although service providers are striving to be successful businesses, they also aim to keep costs as low as possible, spreading resources across multiple users or subscribers, which can lead to performance or resource contention issues. Performance and latency or delays in accessing information from a cloud or network service provider also need to be kept in perspective.

Beyond fees and their drivers, there are also some other service factors to delve into. For instance, optional services typically include the ability to pick the availability level and

What Cloud or Mode of Service Is Best?

Cloud/ Service Type	Optimal Use
Public AaaS or SaaS	Offload storage functions (file or data sharing, general storage, backup or archiving) using a service that also provides tools. There are no restrictions on public-cloud services. Shift from capital hardware and operating software expenses to monthly payments for services used. If apps are in the cloud, that is where the data should be. If applications are local, then network latency needs to be analyzed before storing data in the cloud.
Public IaaS	Use if the storage target is general file or data sharing, for a destination for backups, replication, snapshots and archives. The service can replace on-premises backup, archiving or data protection storage, such as tape or disk, or complement it by adding an extra location and geographical separation.
Private	Gain control over resources, including sensitive data within the management domain (if security concerns prevent using public-cloud services). Applications can be kept local and allow for a large amount of changing data and performance needs. Gain cloud confidence by taking first steps using converged solutions while running proof-of-concept testing in public clouds for other purposes.
Hybrid	Mix and match public resources such as IaaS for archiving of not-so-sensitive data, with onsite storage for resources that have security, regulatory or other concerns.
Community	Pool resources with community members or peer organizations that must address concerns or restrictions around public-cloud use. This is often a good option for government and education organizations, allowing them to leverage common regulation requirements and cloud resources while maintaining separate and autonomous access to data.
Converged solution bundle	Use as a building block for private or hybrid services along with community cloud solutions to simplify acquisition, installation, integration and ongoing support of server, storage, network, hardware and management software tools – both proprietary and open-source.

the geographic region where data will reside. Specifying a geographic region may be important for data or applications that must meet regulatory compliance mandates specifying that the data not leave certain jurisdictional boundaries.

In addition, check to see, in the case of a regional disaster, if the service provider will automatically move data to a different region, and whether for free or for a fee. Additional fees may apply for different levels of service, including faster storage and network access, improved data resiliency, and availability or long-term retention.

Another factor? A service provider's financial and business stability. The enterprise will want to know how its data can be retrieved should the cloud service cease to exist. As is the case with any emerging technology or solution provider, look for organizations that are flexible and that have a growing list of active customers. In some situations, a lower à la carte pricing model may be a good option. For other situations, a package that has a higher fee but is more inclusive may be a better value.

Clouds do not have to be an all-or-nothing proposition. Cloud storage and computing can be complementary to what the data center or IT department already provides. Therefore, the IT team should view cloud storage options as providing another tier of IT resources or as a tool to deliver information services.

Perhaps most critical is the need for situational awareness. The IT team needs both an enterprise understanding of its data and storage components, as well as the skill to navigate in and around clouds. And it is important to determine whether moving applications or data to a cloud is simply moving a problem or whether it's an opportunity to improve overall operations.

Even if the organization is just beginning to test cloud use, it need not be a scary proposition. Determine IT team, senior management and user concerns and then figure out how to address them. Rethink when and why the enterprise protects its data, and analyze how cloud storage can be an enabler. Establish best practices and management policies for the cloud storage program. Finally, leverage data footprint reduction to efficiently move data to the cloud.

Ready to take the next step toward cloud storage? You can learn more about the various topics, technologies and best practices discussed in this white paper at cdw.com/cloud and by engaging your CDW account manager for a consultation.

Checklist for Evaluating Cloud Storage Services

- ✓ What management tools are included or supported?
- ✓ What capabilities are needed to move data, applications and VMs into a cloud?
- ✓ What are the information privacy and security requirements (logical and physical)?
- ✓ Is there compliance and audit reporting?
- ✓ Does the service provide interoperability through application programming interfaces or other software tools?
- ✓ Are there any geographic location considerations for compliance or regulatory purposes?
- ✓ How are metering, measurement, reporting and chargeback handled?
- ✓ Are the organization's networks able to move the required amount of data to and from the cloud?
- ✓ What are the service management options?
- ✓ Are there any extra access fees?

salesforce chatter

Salesforce Chatter is the engine of the social enterprise, helping organizations get more done through instant collaboration via enterprise social networking. One single, secure environment connects you to everything you need for your social enterprise, and eliminates islands of collaboration by enabling employees to work together on sales deals, service cases, marketing campaigns, files, dashboards – any business process – on one, trusted platform.

CDW.com

Microsoft

Microsoft® Office 365 for large and small organizations is a subscription service that combines the familiar Microsoft Office Apps with a set of web-enabled tools that are easy to learn and use, that work with your existing hardware, and that come backed by the robust security, reliability and control you need to run your organization.

CDW.com/microsoft

hp

HP Cloud Collaboration as a Service enables organizations to manage data and enable file/information sharing with extremely low latency. With its economics and agility, HP Cloud Object Storage enables flexible and cost-effective capacity, tight security and control, and anywhere, any-device collaboration at enterprise scale.

CDW.com/hp

Symantec.

Even as individuals and organizations realize the potential agility and cost savings benefits of cloud computing, concerns about security and availability of clouds persist. Gain confidence in your cloud with protection from Symantec™. Whether you want to consume services directly, build your own cloud for internal operations or external reach, or extend into third-party clouds safely and efficiently, Symantec delivers the path to a protected cloud.

CDW.com/symantec

SHARE THIS WHITE PAPER

The information is provided for informational purposes. It is believed to be accurate but could contain errors. CDW does not intend to make any warranties, express or implied, about the products, services, or information that is discussed. CDW®, CDW·G® and The Right Technology. Right Away® are registered trademarks of CDW LLC. PEOPLE WHO GET IT™ is a trademark of CDW LLC.
All other trademarks and registered trademarks are the sole property of their respective owners.
Together we strive for perfection. ISO 9001:2000 certified
121663 – 130701 – ©2013 CDW LLC

CDW® PEOPLE WHO GET IT™