



DCIG

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

By Jerome M Wendt



The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Table of Contents

1	Executive Summary
2	Putting Some Context around the Looming Data Creation Explosion
2	On-Premises Storage Out of Sync with the Reality of Data Growth
3	The Overlooked Costs of On-Premises Storage Arrays
3	The Overlooked Hard Costs
4	The Overlooked Soft Costs
4	Exiting the On-premises Storage Business
3	Six Triggering Events that Drive Storage Array Acquisitions
4	Complex On-prem Data Migration Process
5	The Viability of Cloud Storage
5	The Tech Specs
5	OPEX is the Operative Word
6	Cloud Object Storage: A Fast Growth Storage Segment
6	Secondary Storage: A Logical Starting Point for Cloud Storage Adoption
6	Data Types Best Suited for Initially Storing on Cloud Storage
6	Capitalize on Innovations in Cloud Storage
7	The Golden Age of Cloud Storage Has Arrived

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Executive Summary

Data, data, and more data. Managing that data growth summarizes a challenge that almost every organization on the planet faces. Further, analysts forecast data growth into the petabyte, exabyte, and even zettabyte ranges in the years to come. To handle this continued data growth, organizations need to examine their current storage strategy.

By default, many look to on-premises storage arrays as their preferred means to store this data. They control them. They are familiar with them. They know how to manage them. These arrays have tremendously improved their performance features in recent years. Even their upfront acquisition costs continue to decrease.

However, organizations may fail to consider all the hidden hard and soft costs associated with acquiring and managing on-premises arrays. These hidden costs take many forms. They frequently acquire storage arrays that have more storage capacity than they need. This results in overpaying and overprovisioning. They must perform data migrations that take time and incur risk. They must install them, maintain them, cool them, secure them, power them, and then decommission them. Here again, all these processes incur more cost, risk, and time to perform.

Organizations need to ask themselves, *“What features do they really want on-premises storage arrays to provide?”* In short, they want the availability, durability, performance, reliability, scalability, and stability features they offer. They also want to avoid all the baggage that comes with those features. To date, organizations lacked any means to get the storage features they want without the accompanying baggage.

Cloud storage changes that. Using cloud storage, organizations get the storage features that on-premises cloud storage offers without its baggage. They also receive features that most on-premises storage arrays cannot natively deliver.

Organizations have the option to store data in multiple geographic regions. They can store billions of files. They gain access to global namespaces. They can store and access very large files quickly and easily. They only pay for the data they store. Cloud storage opens the door to these features and many more that on-premises storage arrays cannot easily provide.

The golden age of cloud storage has arrived. While no one predicts that on-premises storage will go away, organizations can no longer deny the benefits of cloud storage. Its OPEX model, fast on-ramp, and growing pool of features make it easier than ever for organizations to cost-effectively adopt. In so doing, it addresses their immediate storage challenges while positioning them to adeptly handle future data growth.

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Putting Some Context around the Looming Data Creation Explosion

Analyst forecasts of exploding data creation coupled with new terms that describe this explosion flood the media. These occur for good reason. These forecasts and terms pique everyone's interest. A November 2018 IDC report certainly helped to reinforce that perception.

This report predicted the world will create staggering amounts of new data over the next seven years. IDC expects that by 2025 the total amount of data created annually will reach 175 zettabytes (ZBs).¹ This represents an annual growth rate of approximately 27% based upon the estimated 33 ZBs of data generated in 2018.

The term zettabyte itself requires some context to fully grasp how much data this represents. Most organizations deal in megabytes (MB), gigabytes (GB), and terabytes (TB) with a few managing petabytes (PB) of data.

CHART 1

Storage Capacities

1,024 Gigabytes	1 Terabyte
1,024 Terabytes	1 Petabyte
1,024 Petabytes	1 Exabyte
1,024 Exabytes	1 Zettabyte
1,024 Zettabytes	1 Yottabyte

A zettabyte describes data at a scale beyond which any organization today currently manages. By way of example, the largest and densest storage arrays today only host about 10 PB of data. To host a zettabyte of data would currently require a company to acquire 100,000 of these arrays.

One should note this report only foretells the amount of data created annually. It neglects to discuss the percentage of this newly generated data that organizations may retain. If organizations keep only 25 percent of this data, total data under management could swell to over 100 ZBs by 2025.

On-Premises Storage Out of Sync with the Reality of Data Growth

Admittedly, few organizations, if any, expect to create or manage a zettabyte of data, even by 2025. However, many

anticipate managing and storing petabytes or even exabytes of data by that date. To accommodate that data growth, they will need some place to affordably store and manage it.

"Affordable" becomes the operative word here. Organizations may still view an on-premises storage array as economical if they only consider the initial capital expenditure (CAPEX). However, its purchase price represents only a fraction of its total cost. Many, or even most, of its costs occur after organizations purchase it. These hidden costs can result in the inability of organizations to fully quantify their operating expenditures (OPEX).

Acquiring an on-premises storage array no longer represents the only or even the best option going forward. Rather, organizations should determine the total cost of ownership (TCO) associated with on-premises storage arrays. Then, once tabulated, they should compare that to the TCO associated with using cloud storage.

The Undocumented Risks of Data Migrations

All IT staff, regardless of their role in the organization, generally agree on one topic: they dread doing data migrations. The degree to which they dread them frequently depends upon the risks associated with completing them.

In a best-case scenario, IT staff find data migrations an annoyance. They have the expertise, software, and processes in place to successfully execute upon them. While they do not "love" them, they do not find them overly burdensome either.

In a worst-case scenario, IT staff loathe them. They may need to use multiple different software tools to migrate data. They lack the skills and expertise to manage these tools and complete the data migration quickly. They must manage a process that involves many people and touches multiple groups.

Love or hate them, data migrations create risk. Data migrations require that IT staff perform tasks that are often poorly documented or undocumented since they rarely perform them. (See *Complex On-prem Data Migration Process*.)

Aggravating the situation, everyone expects data migrations to occur without disrupting production applications. Expectations that applications experience minimal or no performance impact often accompanies these migrations.

To top it off, the responsible IT staff may have to work nights and weekends in addition to fulfilling their normal responsibilities. This combination of balancing competing priorities, working extra hours, and managing a complex, unfamiliar task introduces risks that companies may fail to fully grasp.

1. <https://www.forbes.com/sites/tomcoughlin/2018/11/27/175-zettabytes-by-2025/#1856fd8b5459>

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

The Overlooked Costs of On-Premises Storage Arrays

Calculating the true costs for deploying and managing on-premises storage arrays involves examining all their costs—acquisition, ongoing and disposal. To arrive at a TCO, they must identify the frequently overlooked costs associated with deploying a storage array on-premises. Here are just some of those costs.

The Overlooked Hard Costs

Hard costs consist of the additional monies that organizations will pay as a result of having an on-premises storage array. These costs may not readily present themselves when organizations purchase a storage array. Rather, organizations pay these costs at various times throughout the storage array's life on-premises.



Storage array overlap. Organizations can easily forget to include the costs of keeping multiple storage arrays on-premises at the same time. Rarely, if ever, does a new storage array arrive on-premises one day and the old one disappears the next. It may take weeks, months, and perhaps even years to migrate data from the old array to the new one.



Data center resources. Power, heating, and cooling for a storage array all incur real costs. These expenses can get buried in the overall operational costs of running a data center. Further aggravating the situation, multiple variables go into calculating these costs. Power costs will vary depending upon the region or city where the data center is located. The number and types of disks in the array will influence its power consumption. The array's heating and cooling costs will vary depending on the disk types (HDD or SSD) used.



Overprovisioned storage. No organization ever buys a storage array with just enough capacity to meet its current needs. Rather, it typically buys enough extra storage capacity to last until it acquires its next storage array. On the surface, that sounds like smart planning. Unfortunately, it also equates to an organization spending money before it needs to spend it. This capacity sits unused until an organization needs it, if the organization needs it at all. The entire time this storage capacity sits idle represents wasted dollars.



Storage networking. To connect servers to the storage array organizations will need host bus adapters and network cabling. They may also need FC or Ethernet switches depending on the number of servers they connect to the storage array.

Six Triggering Events that Drive Storage Array Acquisitions

Organizations regularly acquire on-premises storage arrays when certain events transpire. These occasions provide a natural time for organizations to examine using alternative storage technologies. Here are six common business and technical events that prompt organizations to acquire an on-premises storage array.

1 Array no longer under warranty or at the end of its depreciation cycle.

Organizations often find it more economical to buy a new storage array than keep an existing one. The purchase or lease of a new array may cost less than putting maintenance in place for the existing array.

2 Reached its maximum capacity.

The storage array will not support adding any more disk drives to it. This creates a need to acquire another array.

3 Application performance demands exceed storage array's capabilities.

The storage array cannot meet the performance demands of specific applications. These applications may need less latency or faster throughput than the current array can provide.

4 Applications need specific storage array features.

An application can require any of multiple storage array specific features. It may need a specific type of storage networking connection such as Fibre Channel (FC) or Ethernet. It may require a storage array that supports encryption of data at-rest. It may need an array that supports multi-tenancy. These application-specific requirements may dictate organizations acquire another array.

5 Storage array at the end of life.

A storage array is a physical device and wears out. In these circumstances, organizations have little choice but to replace it.

6 Additional internal storage is no longer an option.

Applications that use internal disk drives (HDDs or SSDs) may be economical. Unfortunately, to add more drives to address application specific capacity or performance needs may require downtime. This application outage now costs more than the storage savings. These limitations of using internal storage media prompt the need for a more robust storage solution.

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Complex On-prem Data Migration Process

Pre-data Migration

- ▶ Create network routing between the servers and new storage array (New zones on FC switch)
- ▶ Assign storage capacity to the servers (LUN masking)
- ▶ Determine which software tool(s) will be used to migrate data
- ▶ Contact appropriate people and groups
- ▶ Schedule change control
- ▶ Find a date that will work for everyone
- ▶ Verify existing servers can see new storage

Data Migration

- ▶ Initiate the migration
- ▶ Verify all data is successfully moved
- ▶ Verify applications run as good or better on new storage

Post-data Migration

- ▶ Remove unneeded network routes (Remove unneeded zones)
- ▶ Remove unneeded storage assignments on storage arrays (Remove unneeded LUN masking)
- ▶ Delete data on the old array

The Overlooked Soft Costs

Hard costs consist of expenses that organizations pay out of pocket any time they acquire a new storage array. Soft costs consist of expenses that organizations incur regardless if they acquire a new storage array. Companies can spend so much time managing storage arrays these soft costs eventually turn into hard costs. Here are just some of the soft costs that an on-premises storage array can incur.

 **Data center resources.** Data center resources show up as a soft cost as well. An on-premises storage array consumes available data center floor space, though that cost may not show up on the bottom line. It also

requires organizations to have security measures in place (guards, check points) to control access to it.

 **Storage array acquisition time.** Storage array acquisitions differ significantly from buying a gallon of milk. One does not simply walk down to the corner store and buy a storage array. The IT staff must evaluate available storage arrays on the market and make its recommendation. The procurement team may get bids and negotiate prices. They then hand the contract over to the legal department for review. Finally, the accountants come into the picture. They put the storage array on the books and depreciate it. All this takes time and adds to the storage array's costs.

 **Storage array installation and maintenance.** Having selected and acquired a storage array, it goes to the IT staff who install it. The IT staff will oversee and perhaps perform its installation as well as any required storage networking. The IT staff will also assume responsibility for its ongoing maintenance, to include fixes, patches, and upgrades.

 **Data migrations await.** Once installed, organizations may need to migrate data from other arrays to the new array. This may incur further soft costs as IT staff plan for the migration, manage it, and then validate it completed successfully. The migration process may even incur some hard dollar costs in the form of software that migrates the data.

 **Storage array management.** Once deployed, more soft costs associated with the ongoing storage array management surface. These costs span the gamut. The most basic tasks consist of allocating capacity to servers, monitoring its performance, performing load balancing, and keeping its firmware up to date.

Other tasks that administrators must perform sporadically include resolving firmware incompatibilities between the array and its attached applications and servers. They will need to plan and execute upon data migrations. They will need to perform recoveries should any data loss occur during a data migration. They will need to constantly evaluate new storage array models and plan for the next purchase.

Exiting the On-premises Storage Business

These tasks and others represent the factors that organizations must use to calculate the TCO for an on-premises storage array. They also represent the primary reasons that few organizations take time to determine the TCO for on-premises arrays. It takes too much time and is too difficult.

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Even assuming organizations can calculate the TCO, they should step back and ask an even more basic question, “Why are we in the on-premises storage business at all?” Managing storage arrays brings little intrinsic value to the business or the organization. They need the storage array’s capacity and performance – not the array itself.

Few organizations consider IT their core competency and an even smaller number would name storage as a specialty. Due to the specialized and critical nature of storage, they should explore options that allow them to exit the on-premises storage business. This assessment should include exploring cloud storage to determine if it meets their needs.

The Viability of Cloud Storage

A few years ago, many organizations dismissed cloud storage for various reasons. Its availability, immaturity, reliability, and security features were, among others, reasons they cited for not pursuing cloud storage more aggressively as an option. Fast forward to today and organizations can no longer point to them as reasons not to adopt cloud storage. If anything, cloud storage providers offer more robust features than what many organizations find in their own data centers.

The Tech Specs

The data centers that cloud storage providers use to host data almost universally meet or exceed any organization’s current technical requirements. Features they offer include:

 **Availability and redundancy.** Cloud storage providers deliver redundancy in two ways. Many deploy two or more data centers in a geographic region. They group these data centers into a region called an availability zone. Should one data center in an availability zone become inaccessible, the other data center(s) take over keeping data accessible. To do so, the cloud storage solution automatically replicates data across all data centers in an availability zone. To achieve even higher levels of data availability and redundancy, organizations may optionally replicate data to another region.

 **Certification.** Cloud storage provider data centers meet standard Tier 3 and even Tier 4 data center certifications that guarantee about four nines of uptime or greater annually. Many also provide certifications to meet stringent industry and government data security standards such as CJIS, FedRAMP, HIPAA, and PCI DSS.

 **Data durability and resiliency.** Cloud storage providers use storage platforms that guarantee the data’s integrity. Once they store data with a cloud storage

provider, organizations can have complete confidence they can retrieve it as well.

 **Fortification.** Cloud storage providers fortify their data centers to withstand natural disasters such as earthquakes, floods, hurricanes, and tornados, among others.

 **Scalability.** Cloud storage providers remove worries about running out of storage capacity and the work associated with obtaining it. If an application needs more storage capacity, organizations can immediately and seamlessly add more storage to it. The cloud storage provider assumes all responsibilities for managing the physical storage infrastructure.

 **Security.** The security cloud storage providers offer takes at least two forms. They offer physical security to control physical access to their building. This includes logging and tracking all activities in their data center. They also offer identity access management software that organization use to control access to the data they store in the cloud.

OPEX is the Operative Word

The technical specifications that cloud storage providers offer often pique the interest of organizations. Their pay-as-you-go model serves as the impetus for them to adopt it. Unlike on-premises arrays that require large upfront, capital investments, cloud storage providers only charge organizations for the storage capacity they use.

This pay-as-you-go approach closely aligns to the OPEX (operating expense) model that more organizations currently want to adopt. Using cloud storage, they avoid the need for making upfront CAPEX (capital expense) investments in storage.

They also avoid overbuying too much storage capacity that so commonly happens when storage array purchases occur. Organizations routinely must forecast how much storage capacity they require so they may appropriately size their new arrays. Cloud storage eliminates both the forecasting and the overbuying. Cloud storage allows them to only pay for what they use when they need it.

Cloud storage also encourages another corporate behavior that has fallen into disuse: good storage management practices. Many organizations either abandoned or only paid lip service to proactively managing the storage capacity on their on-premises storage arrays. The challenges of managing capacity in storage arrays often made it more practical to “just buy more.”

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Cloud storage should change this mindset as well for at least two reasons. First, it eliminates many of the ancillary storage array management and maintenance tasks they must perform now. This leaves them time to do it.

Second, and more importantly, cloud storage providers reward good storage management behavior. In the on-premises storage world, it rarely made sense to recover storage. The storage was already paid for, there was often excess storage capacity on the array, and it was risky to recover it. In the cloud storage world, when you stop using storage, you stop paying for it. Cloud storage incentivizes organizations to actively manage their stored data and only keep what they need.

Cloud Object Storage: A Fast Growth Storage Segment

Organizations may wonder about the feasibility of more broadly adopting cloud storage. They may not currently use cloud storage or only use it in a limited capacity. Further, cloud storage platforms use object storage as their means to ingest and store data, a solution that organizations may not routinely use. To those concerned about object storage and its interface, they can take comfort in the following two trends.

First, cloud object storage represents a rapidly growing segment of the storage market. One recent report puts cloud storage's annual growth at 22.49% from 2018 to 2023, growing to over \$100 billion by 2023.² This level of growth and adoption will ensure continued innovation and support from the vendors who offer this service.

Second, cloud storage providers increasingly make it easier for organizations to seamlessly adopt and use their platform. Many currently offer interfaces that present object storage in the more familiar and well-understood format of files and folders. This interface accelerates cloud storage's adoption since organizations often use file servers to store their archival, backup, and file data. On top of that, cloud storage providers offer manageability, resiliency, and scalability features rarely found in network attached storage (NAS).

Secondary Storage: A Logical Starting Point for Cloud Storage Adoption

The costs of on-premises storage do not, alone, always justify organizations immediately abandoning it for cloud storage.

Organizations recognize hosting data on-premises does provide certain benefits, to include the familiarity, performance, and proximity of using on-premises storage arrays.

Further, organizations often *"don't know what they don't know"* about cloud storage. This prompts them to cautiously approach adopting cloud storage by first figuring out a logical starting point.

Data that currently resides on on-premises secondary storage often serves as the best data to initially move to cloud storage. Organizations often store archival, backup, and capacity optimized data on secondary storage, though they may keep other data types on it.

Data Types Best Suited for Initially Storing on Cloud Storage

- ▶ Archival
- ▶ Backup
- ▶ Capacity optimized, low transaction work loads
- ▶ Data requiring long retention periods
- ▶ Latency tolerant that allows for millisecond response times

These data types make good candidates to move to the cloud. They have relatively low levels of visibility in the organization even though they can make up to 90% of their on-premises data. Their infrequent access rates and tolerances for millisecond response times when accessed also contribute to prioritizing them for placement in the cloud. These data types represent low levels of risk and significant savings by adopting cloud storage sooner rather than later.

Capitalize on Innovations in Cloud Storage

As organizations consider taking this low risk approach to cloud storage adoption, they simultaneously gain access to ongoing cloud storage innovations. Using on-premises storage arrays innovation may only occur when organizations acquire a new one. In cloud storage environments, innovation continually takes place. Best of all, organizations immediately have access to these innovations and may capitalize on them as they become available.

² https://www.researchandmarkets.com/research/jtccvk/global_101?w=4

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

Here are just a few innovations that organizations gain access to when they move to cloud storage:



Flexibility to store billions of objects. Organizations need more than lots of storage capacity in the cloud. As more organizations embrace artificial intelligence and machine learning, the potential exists for them to create billions of files. They will need a platform that can ingest that number of files, effectively manage them, and provide the performance they need. Most if not all cloud storage offerings already provide this level of scalability.



Data replication to ensure high availability. Behind the scenes, cloud storage providers keep data in physical data centers. This makes them susceptible to natural disasters and downtime. To mitigate against these occurrences, many cloud storage providers automatically replicate data stored with them between their multiple data centers. This approach gives organizations access high levels of data availability.



Geo-dispersed namespaces. Cloud storage providers make global namespaces available to enable organizations to uniquely identify and store objects anywhere in their cloud. In this way, organizations can store data closest to the individuals using it without losing track of it.



Performance acceleration for large files. Everyone hates the spinning wheel of death as they wait for large files to upload or download. Cloud storage providers have not overcome physics. However, they do provide technologies to accelerate the uploads and downloads of large files. By breaking large files into smaller segments, they can concurrently send multiple chunks of data that they reassemble once downloaded or uploaded.

Due to their size and reach, cloud storage providers have relationships and work with all the large media providers who all store data in the cloud. These organizations store everything from video to images to sensor data to PDFs with cloud storage providers. Supporting these providers requires that cloud storage providers consistently bring to market ever better networking and storage technologies to meet these demands.

Best of all, all customers of these cloud storage providers benefit from their innovations. Further, they can gain access to them with no capital expenditures. These new features are just like all the other features that cloud storage providers offer. Organizations only pay for what they use.

The easy access to these features gives organizations the flexibility to trial them by starting and stopping them whenever they want. In this way, organization can experiment with these latest innovations at any time. This mitigates the limitations of on-premises storage arrays whose features can become dated as soon as they deploy them.



The Golden Age of Cloud Storage Has Arrived

The tables have turned for the use of on-premises storage arrays. As organizations retain increasing amounts of data, they must examine the practicality of using on-premises storage as their preferred target. While no one expects on-premises storage will go away, the arguments for using cloud storage in its stead continue to mount.

Cloud storage delivers on the new metrics that organizations both better understand and want. Using cloud storage, they may treat storage as a utility and pay for it as they go. In so doing, they can apply the same best practices for managing cloud storage as they do to managing utilities to achieve similar results.

Cloud storage eliminates the parts of storage management that add little or no value to the organizations. Storage array installs, firmware patches and updates, hardware refreshes, data migrations ... the list of undesirables goes on and on. Organizations derive no intrinsic value from performing these maintenance and management tasks. If anything, they distract them from why acquired storage in the first place.

Cloud storage takes the stress out of storage. It offers organizations the availability, capacity, durability, performance, reliability, and scalability that they want storage to deliver. It simultaneously removes all the hidden and overlooked time and costs associated with managing it.

The Growing Business Case for Using Cloud Storage to Replace On-premises Storage

The golden age of cloud storage is upon us.

In so doing, cloud storage frees companies to put their focus back on their core business. It absolves them of having to focus on IT management, data storage, and constantly monitoring and managing on-premises storage systems. This translates into more time spent leveraging data to innovate in an increasingly competitive world while spending less time managing data.

Best of all, organizations may begin to adopt and implement cloud storage in a low-risk manner. Archival, backup or any data that resides on secondary storage represents the best data to initially move to cloud storage. Taking this approach, they gain exposure and experience with cloud storage while simultaneously lowering their overall costs and overhead.

The golden age of cloud storage has arrived with cloud storage being mature and ready for immediate adoption. Organizations obtain all the technical benefits that on-premises storage offers without all its accompanying backend management headaches. In so doing, cloud storage accomplishes three key objectives for them. It simplifies their storage environment, lowers their costs, and positions them to more aggressively adopt new features going forward. ■

About DCIG

DCIG empowers the IT industry with actionable analysis that equips individuals within organizations to conduct technology assessments. DCIG delivers informed, insightful, third party analysis and commentary on IT technology. DCIG independently develops and licenses access to DCIG Buyer's Guides. It also develops sponsored content in the form of blog entries, executive white papers, podcasts, Technology reports, webinars, white papers, and videos. More information is available at www.d cig.com.



DCIG, LLC // 7511 MADISON STREET // OMAHA NE 68127 // 844.324.4552
dcig.com

© 2019 DCIG LLC. All rights reserved. The DCIG Technology Report Executive Edition is a product of DCIG, LLC. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. Product information was compiled from both publicly-available and vendor-provided resources. While DCIG has attempted to verify that product information is correct and complete, feature support can change and is subject to interpretation. All feature support represents the opinion of DCIG. No negative inferences should be drawn against any product or vendor not included in this report. DCIG cannot be held responsible for any errors that may appear. Licensed to Wasabi with unrestricted and unlimited distribution rights. This report was commissioned by Wasabi.