

Higher Education: The Time Is Now For Cloud Storage

Brien Posey

INSIDE THE GUIDE:

- The Particular Challenges for Educational Institutions
- How To Squeeze More Out of Every Education Dollar
- Why Hot Cloud Storage Gets an A+

ENTERPRISE IT

Innovations

LEARNING SERIES

EXPRESS EDITION

Higher Education: The Time Is Now For Cloud Storage

Sponsored by Wasabi

AUTHOR

Brien Posey

Copyright © 2019 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America.

ACTUALTECH MEDIA

Okatie Village Ste 103-157
Bluffton, SC 29909
www.actualtechmedia.com

TABLE OF CONTENTS

Introduction: Why the Real Migration to the Cloud Will Begin with Storage.....	4
Chapter 1: ‘Hot Cloud’ Storage: The Birth of the Next Great Utility.....	8
What Happens to On-Premises Storage in the Wake of the Great Cloud Migration?.....	11
Pros and Cons of Cloud Storage.....	12
What About Hybrid Storage?.....	19
Multi-cloud Strategies.....	24
The Economic Case for Cloud Storage.....	27
Chapter 2: Education.....	29
The Unique Challenges of Storage in an Academic Institution.....	29
Wasabi’s Hot Cloud Storage as an Option for Educational Institutions.....	34
Conclusion: What Could You Do With Hot Cloud Storage?.....	39

INTRODUCTION

Why the Real Migration to the Cloud Will Begin with Storage

Over the last decade, cloud providers have had their marketing departments working overtime, trying to convince us of just how great the public cloud is, and how anyone who is still operating on-premises is completely out of touch. In spite of the marketing hype, only about 30% of the world has moved to the cloud. As such, it stands to reason that the long-anticipated mass migration to the cloud is still ahead of us.



In the “Cloud 2.0” era, organizations are finally able to store all of their data in the cloud, and leverage that data to derive business value.

The first generation of cloud adoption was all about compute and workloads. This “Cloud 1.0” concept – 13 years old as of 2019 – has led organizations to realize

that computers and workloads are really nothing more than the tools that enable data. Hence, the “great migration” is almost certain to be data driven.

In the “Cloud 2.0” era, organizations are finally able to store all of their data in the cloud, and leverage that data to derive business value.

It may be tempting to scoff at the notion of using the cloud to unleash the hidden potential of data. After all, “Big Data” has been discussed for a decade, but many companies haven’t experienced the promises.

The main reason for Big Data’s failure to live up to the hype is not that mining hidden business value from data is a myth.

Instead, it’s that the data hasn’t been BIG enough yet!

The new wave of Machine Learning (ML) and Artificial Intelligence (AI) advancements are strong evidence that the larger your dataset, the more powerful that data becomes. Larger data sets allow the data to be leveraged in more ways. This means that you can make more predictions, more projections, and gain greater insight from your data.

The real reason the “Big Data revolution” has been held back is not because the data has no hidden value; it’s there, for sure. The brake on the revolution has

been *the cost and the limited capabilities* of the underlying storage infrastructure.

The only way to overcome these challenges and fully realize the power of data is to leverage next generation, Cloud Storage 2.0. Although cloud storage has existed for many years, its unpredictable costs have caused many organizations to shy away from using it any more than is absolutely necessary.



The cloud is no longer just an archive or test/dev solution; it's becoming a platform for even the most active data applications, with performance that until recently was only possible in your data center.

But this is starting to change.

Solutions are emerging that combine the classic Cloud 1.0 advantages of scalability, reliability and resiliency with a predictable cost structure, high performance, and ease-of-use that makes it feasible for organizations of any size to consider moving to the cloud.

The cloud is no longer just an archive or test/dev solution; it's becoming a platform for even the most active

data applications, with performance that until recently was only possible in your data center.

This Guide is your roadmap to understanding this tectonic shift toward putting public cloud storage within the grasp of all.

As it currently stands, the past decade of cloud storage solutions has been fraught with huge hurdles and technological challenges, many of which buyers don't find out until they're in the middle of a migration, or have implemented their public cloud platform and finally flipped the switch.

Then the first cloud storage bill comes due, and pow! They suddenly realize that Cloud Storage 1.0 is not going to be nearly as cheap or easy as they were led to believe.

CHAPTER 1

‘Hot Cloud’ Storage: The Birth of the Next Great Utility

Public cloud storage hasn’t grown as rapidly as some have predicted, although its usage is clearly on the rise.

One issue holding back more widespread adoption of cloud storage is that using it is anything *but* seamless. When you deploy a new application, (whether on-premises or in the cloud), that application probably isn’t going to automatically recognize and begin using your cloud storage. Instead, administrators typically have to set up cloud storage connectivity.

The process of attaching to and using cloud storage can vary widely from one provider to the next. There are several major cloud storage providers, and each one has its own way of doing things. Even if an organization opts to stick with a single cloud storage provider, using cloud storage is still going to have its ups and downs.

In an AWS environment, for example, as of early 2019, IT pros must choose between 6 different tiers of cloud storage, S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, S3 One Zone-IA, S3 Glacier and S3 Glacier Deep Archive.



DEEP DIVE

Hot Cloud Storage

“Hot Cloud Storage” is a single tier for active data, active archive and inactive archive.

Wasabi’s Hot Cloud Storage addresses three classes of data storage with a single, low-cost, high-performance, high-security tier:

- **Active Data:** Data that needs to be readily visible to the operating system and/or application software with which it was created, and must be immediately accessible to users.
- **Active Archive:** Archive or backup data available instantly online, rather than needing to be restored and rehydrated from an offline or remote source.
- **Inactive Archive:** Traditionally, data that is archived to tape, stored offsite, and used or referenced infrequently.

If you're using Cloud Storage 1.0 solutions, only after choosing the storage tier best suited to the task at hand can IT pros begin figuring out how to actually provision and connect to the storage.

With Cloud Storage 2.0 solutions, storage is treated as a utility – as common, reliable, and used every day as electricity. For example, AC power is delivered in a standard way. Every power outlet in America provides 110 volts at 60 Hz. (not including 220 volt outlets). Any electric-powered device can be plugged into any power outlet in America: it doesn't matter if the device is a fan, a computer, or a curling iron, because all these devices are designed to be powered in the same way.

Similarly, the same laptop that you can charge at home can also be plugged in at work, a coffee shop, or even a hotel room on the other side of the country. Electric power is delivered in a standardized way, and the devices that use that power also do so in a standardized way.

The end result of this is that electric power is a commodity. We use electricity on an as-needed basis, rarely considering where it comes from, whether it will work with our device, or how much it costs.

Cloud storage also needs to evolve to this point.

What Happens to On-Premises Storage in the Wake of the Great Cloud Migration?

With the tech media being so focused on all things cloud over the last several years, it's easy to lose sight of the fact that on-premises storage is still being used.

In spite of what some vendors might lead you to believe, on-premises storage is alive and well.

Cloud storage doesn't eliminate the need for on-premises storage, but rather changes some of the ways in which it is being implemented, maintained, and used. The truth is that some workloads are just better suited to use on-premises storage, while others may see significant benefits from the use of cloud storage.

There are any number of reasons why a workload may need to remain on-premises. The most common is that security or compliance requirements drive the need to keep a particular workload in-house. Other common reasons for continuing to run a particular workload locally might include limitations imposed by software licenses, or perhaps a desire to continue to capitalize on investments that the organization has already made in data center hardware.



THE 101

Private Clouds

One of the big trends in on-premises technology deployments right now is the rapid adoption of private clouds. What's a private cloud? It's a set of private and secure computing resources that's been provisioned and isolated to a single buying organization. It can be deployed in any secured setting that the buyer prefers, whether it's the customer's own data center or a preferred partner data center or hosting provider which has been vetted for security and performance.

Pros and Cons of Cloud Storage

The merits of both cloud storage and local data center storage have been hotly debated in an effort to promote one as a superior position and make it seem like it's the best choice in every situation. In reality however, there are positive and negative aspects to each. While by no means comprehensive, here are a few of the more significant cloud storage pros and cons.

BENEFITS

One of the biggest advantages to cloud storage is that, in at least some situations, it's far less expensive than

storage residing on-premises. Public clouds have long been marketed as “the cheap way of doing things.” While claims of bargain storage in the cloud are debatable, there are circumstances in which cloud storage is by far the less expensive option. Imagine, for example, that an organization needs to add a significant amount of storage, but its existing data center is already at its maximum capacity. There is no doubt that using cloud storage will be far less expensive than constructing a new data center, or the pain, cost and personnel time spent to do rolling hardware upgrades and migrations from old to new on-premises storage.

This concept holds equally true for organizations that need offsite disaster recovery capabilities. If an organization doesn’t already have a data center and hardware resources that can be used for disaster recovery purposes, cloud-based disaster recovery quickly becomes an attractive option.

Another benefit to cloud storage is that, although there are exceptions, it’s easier to scale than on-premises storage. Sure, there are the “big name” on-premises vendors that will be happy to sell you racks of local storage for a premium price, but is that really going to be an option, as storage needs skyrocket but your budget doesn’t? Can you afford to keep increasing your data center’s footprint like that – not only the huge

cost of the storage itself, but the associated costs of power, maintenance, personnel? Where does it end? When the cost of a petabyte of storage from a cloud provider is less than the annual maintenance fee from an on-premises vendor, how does that change the economics of your infrastructure? (A helpful comparison of the cost differences between on-premises storage and cloud storage can be found on the Wasabi site.¹)

This problem is compounded by the fact that data growth is projected to double every two years (or in the case of genomics, every 7 months), meaning that most organizations will meet or surpass the petabyte threshold in the very near future. This will only serve to drive up on-premises storage costs; because of this, smart organizations should consider switching to a more cost-effective alternative before getting locked into an expensive hardware expansion.

In contrast to that example, cloud storage tends to be far easier to scale. Again, each cloud storage provider has its own way of doing things, but most providers, including Wasabi, use object storage. Object storage uses a flat namespace (without a traditional file system) that supports massive scalability. A single object store can easily scale to multiple petabytes and can

¹ <https://wasabi.com/blog/on-premises-vs-cloud-storage/>

store billions of objects. Hence, additional storage capacity can usually be added any time that it's needed.

Cloud storage also reclaims valuable IT time by eliminating most storage maintenance tasks. This includes things like replacing aging hard drives and securely disposing of old drives. It also includes all the hardware-level storage monitoring and failure remediation. In a large data center, significant amounts of time can be spent on these types of tasks. When storage is shifted to the cloud, the cloud provider handles all of the hardware maintenance and monitoring. Your IT staff is suddenly freed up to focus on more important things that ultimately benefit the organization in a more bottom-line fashion.

DRAWBACKS

As you can see, there are tremendous benefits to cloud storage. But like every technology, cloud storage does have its disadvantages. Perhaps the greatest is that the costs of storing data in the cloud can be both hugely expensive and massively unpredictable. This is due to a combination of data egress fees, API fees, and frankly, the bewildering assortment of tiers to choose from to find “the right” tier (performance and cost) for very specific use cases.

Most cloud storage providers bill their subscribers a few cents per month for each gigabyte of data stored in their cloud. This is where you need to be extremely careful. A cloud storage provider may advertise a low cost per gigabyte in an effort to give potential subscribers the illusion that using cloud storage is inexpensive.

This cost per gigabyte, however, will often be a relatively small portion of the overall cost of cloud storage. In reality, most cloud storage providers charge a hefty data egress fee each time that a subscriber accesses their own data. Thus it's not the stored data, but *getting to* that data that they sock it to you. And moving data into and out of international locations is often more expensive than moving data domestically, adding another layer of expense onto that.

Data retrieval fees can be terribly expensive; so costly, in fact, that it locks the company into keeping their data in the cloud, afraid to pry it out because of the charges they'll incur.

This is referred to as “cloud lock-in,” because the data egress fees make transferring data cost prohibitive, whether you are transferring data to another location, such as back to your own data center for normal usage, or another cloud provider.



Another potential disadvantage to using cloud storage is the limitations imposed by bandwidth. If an application running in the cloud is creating data, it makes sense to keep that data in the cloud. If, however, you have large amounts of data on-premises, it may not be feasible to transfer all that data to the cloud over the Internet. Typically, you are going to need to use an offline data transfer method to accommodate large-scale data migrations; typically, it would involve having hardware shipped to you from the cloud storage provider. You load it with data, then ship it back to be copied to your cloud storage location.

One more thing to consider is the fact that so many organizations are now undergoing their “digital transformation.” Bill Schmarzo of CIO Magazine offers one of the better definitions of digital transformation out

there. Schmarzo says “Digital Transformation is the application of digital capabilities to processes, products, and assets to improve efficiency, enhance customer value, manage risk, and uncover new monetization opportunities.”²

Essentially, a digital transformation represents a new way of doing things.

Because digital transformation changes both business and IT processes, it also creates new IT infrastructure requirements.

This is especially true of storage. In many cases, on-premises storage is ill-equipped to handle these new challenges; while cloud storage is.



Because digital transformation changes both business and IT processes, it also creates new IT infrastructure requirements.

The problem with cloud storage as it exists today, however, is that the pricing model most providers use make the cost unpredictable. Unpredictable costs are something every organization tries to avoid – while cost underruns MIGHT be a pleasant surprise, cost

² <https://bit.ly/2N1tWlW>

overruns that are vastly higher than expected are not good for anyone.

This is where Wasabi comes in. They've created a new model that builds in predictability with no egress or API fees. That's right – it doesn't cost you anything to pull your data out of their cloud. Consider the advantages: with no data egress or API fees, the boundaries of the data center can finally disappear, at least from a storage standpoint. Organizations can finally move any or all of their data to the cloud, while still maintaining control over the data and being able to accurately predict related storage costs.

Wasabi sees a great opportunity here to save you money, since most cloud migrations are still on the horizon. This suggests that companies are still hesitant, for whatever reason, to make the leap into the cloud. But it's clear that they also know that it's the future, and are starting to plan for that eventuality.

What About Hybrid Storage?

Today, most organizations that haven't gone 100% cloud are basing their cloud strategy around a hybrid storage and on-premises approach. The hybrid storage approach can be any combination of public cloud or private cloud alongside on-premises choices, or any

combination of the three. Vocal cloud evangelists may tell you should be “all in the cloud,” but that may not be the best strategy for you.

While an all public cloud stack can be a great option for startups with little capital to invest in IT resources, in many cases it wouldn't be prudent for an established organization to completely abandon its existing on-premises infrastructure just because the press and industry analysts insist you “have to.”

This is why hybrid storage implementations have become a popular choice: they offer the best of both worlds. Organizations can continue to leverage their on-premises IT investments, while also taking advantage of public cloud capabilities.

Perhaps even more importantly, hybrid storage strategies give organizations the flexibility to run workloads in the location that makes the most sense based on each workload's own unique requirements. If an organization has data that it has deemed must remain on-premises for compliance reasons, then a hybrid storage strategy will certainly allow the organization to continue keeping that data in its own data center. Conversely, if an organization has other data that it has deemed to be less sensitive, it can reduce its costs by shifting the data to the public cloud.

DIGITAL TRANSFORMATION

“Digital Transformation is the application of digital capabilities to processes, products, and assets to improve efficiency, enhance customer value, manage risk, and uncover new monetization opportunities.”

- Bill Schmarzo



Even though it's perfectly acceptable to use a hybrid storage approach as a long-term solution, hybrid storage is often used as a part of a bridge strategy to systematically move from on-premises to the cloud in manageable, focused projects.

Suppose for a moment that your organization has extensive on-premises resources, and has made the decision to migrate everything to the public cloud. Even if you truly wanted to, the migration could not be performed instantaneously.

Migrating all of an organization's IT resources to the cloud is a major undertaking that requires extensive planning.

Rather than attempting to move everything all at once, it would be better to follow a hybrid storage model,

and gradually migrate data to the cloud as it makes sense to do so.

For example, if a particular workload's storage is running on aging hardware that's nearing its refresh date, it may make sense to go ahead and migrate that storage to the cloud, while saving workloads that are running on newer hardware for another time.

When used as a bridge strategy component, hybrid storage strategies accomplish three main things.

First, they allow both data and workloads to be migrated at a comfortable pace. Because hybrid storage can be used as a long-term solution, there's no reason an organization has to rush its cloud migration. Each workload can be migrated as the organization feels comfortable with doing so.

The second advantage of hybrid storage as a bridge strategy is that it helps out with workloads that are difficult to migrate. In spite of what some of the cloud vendors may tell you, not every workload is ideally suited to running in the cloud.

If a workload was developed assuming on-premises resources, and there was no prior intention to run that workload in the cloud, it may be constructed in such

a way that makes it extraordinarily difficult – or even impossible – to migrate it entirely.

In these situations, as an example, the workload's compute portion of the application can be run indefinitely on-premises, leveraging a hybrid storage approach to store data in a public cloud storage instance. Over time, the organization may decide to replace the entire workload with something better suited to the cloud (a web-native application), find a way to migrate or convert the compute portion of the workload to a cloud instance, or maybe even decide to keep the workload running locally.

Regardless of the decision made, hybrid storage strategies offer the freedom to deal with difficult-to-migrate workloads in whatever way makes the most sense.

The third benefit of hybrid storage in a bridge strategy is to treat the cloud as a mechanism for business continuity and disaster recovery. This means, for example, that the storage workloads that continue to run on-premises can be replicated to cloud storage as a way of providing extra protection against data loss, or catastrophic events that might otherwise result in an outage.

If a disaster occurs, workloads running on-premises can be configured to fail over to the cloud, where they can continue to function. Once the disaster has been

mitigated, those workloads can be brought back in house, or left to continue running in the cloud.

Multi-cloud Strategies

Another trend that's taken hold in the last couple of years is multi-cloud.

Multi-cloud is based on the idea that because it's possible to establish communications between clouds, there's no need for an organization to be limited to using a single public cloud. Instead, an organization can scatter workloads and services across several different clouds, and have them operate as seamlessly as if they were all running inside the same cloud.

From a technical standpoint, it's admittedly "easier" to host all of an organization's resources in a single cloud.

Creating a multi-cloud architecture involves configuring a variety of gateways, authentication mechanisms, and firewalls.



Creating a multi-cloud architecture involves configuring a variety of gateways, authentication mechanisms, and firewalls.

Even so, the benefits of creating a multi-cloud architecture can make all the work worthwhile.

The first is that it allows for best-of-breed service usage.

Even though the major cloud service providers tend to offer similar services, it doesn't necessarily mean that those services perform equally well from one cloud to the next.

This isn't to say that some public clouds are inferior to others, but rather that each provider has things that it does really well, and other things that it may not do quite as well. You may find that one provider has really strong directory service offerings, while another has the best platform for hosting virtual machines.

Using a multi-cloud architecture gives organizations the ability to take advantage of what each provider does best.



If a major cloud provider suffers a massive, cloud-level outage, every region could be impacted. Replicating data to storage located within a different cloud vendor provides protection against this type of scenario.

Building a multi-cloud architecture can also be far more cost-effective than using a single cloud provider's services.

Just as there are differences between the services that cloud providers offer, there are also differences in the cost of those services. Using a multi-cloud architecture gives an organization a way of getting the best price for each service it uses.

It's entirely possible for instance, that one provider offers the best price on hosted compute power, while another offers the best price on storage.

Finally, creating a multi-cloud architecture can provide you with a degree of redundancy, especially on the storage front.

Public cloud providers in general do a great job of providing redundancy for their services.

A provider may, for example, offer its customers the ability to host a virtual machine in one region, with the option of failing over to a different region in the event that the primary region suffers from problems.

But while this approach will shield an organization using cloud services against regional-level disasters, it does nothing to provide protection against a cloud-level outage. If a major cloud provider suffers a massive,

cloud-level outage, every region could be impacted. Replicating data to storage located within a different cloud vendor provides protection against this type of scenario.

The Economic Case for Cloud Storage

From the time cloud computing first became practical, cloud vendors have presented economic models and case studies to prospective customers in an effort to convince them that running workloads in the cloud is the “most cost-effective way of doing things.”

While some of the cost-related claims vendors have made in the past by Cloud Storage 1.0 vendors are dubious (often costing the same as on-premises implementations³), there are good reasons it makes financial sense to use Cloud Storage 2.0 rather than on-premises or Cloud Storage 1.0 options. The economics have changed more than you may realize.

One fundamental issue is that the technology innovation curve is steeper than the financial depreciation curve. This means that an organization’s storage hardware is likely to be outdated before its value has been fully depreciated.

³ <https://info.wasabi.com/wasabi-vs-on-prem>

Because hardware refresh cycles are often tied to depreciation schedules, your business may have to continue using legacy storage hardware for a considerable amount of time, giving up the benefits of innovations in storage technologies.

Compared to on-premises implementations, cloud storage providers are able to take advantage of the latest and greatest storage hardware, as well as their own innovations in design and engineering to squeeze out both performance increases and cost decreases on an ongoing basis.

This is especially helpful for those working in the field of education, as you will see.

CHAPTER 2

Education

Educational institutions might not be the first organizations that come to mind when you think of Big Data and heavy storage needs. After all, student records have a relatively small data footprint. However, educational institutions often create video courseware or research data. And like nearly every other industry today, they are consuming far more storage space than they did just a few years ago.

The Unique Challenges of Storage in an Academic Institution

Schools are widely known for being strapped for cash, so it would be easy to assume that budgetary concerns are the greatest challenge facing academic institutions when it comes to data storage. Although funding plays a major role in the storage options they consider, the volume of data being stored is also a driving factor.

In addition to student records, digital coursework and course books, it is increasingly common for schools to store video lectures and supplementary materials for

classes. Some schools have even gone so far as to create entire digital libraries that are accessible to students and staff.



Perhaps the largest contributor to storage consumption is the research being conducted by universities. Scientific studies, by their very nature, often yield huge amounts of data that must be retained for the life of the study, and possibly retained indefinitely.

In addition, some schools (particularly at the university level) have made it a matter of policy for instructors to scan all research papers submitted by students. The scans are indexed, then added to a searchable coursework database. By doing so, the school is able to prevent student plagiarism, and prevent research papers from being reused.

Perhaps the largest contributor to storage consumption is the research being conducted by universities. Scientific studies, by their very nature, often yield huge amounts of data that must be retained for the life of the study, and possibly retained indefinitely.

One thing that makes the storage needs of academic institutions so unique is that, in many cases, schools



EXECUTIVE CORNER

Education

Education is being radically changed by technology. At both the primary and secondary school levels, all the way up through college, then graduate and post-graduate education, students are gaining the benefits of data that's available at the speed of Google.

Here's how one university CIO put it:¹

Higher education institutions are under increasing pressure from public and private stakeholders to increase efficiency, effectiveness and output. These pressures stem from motivations to expand access while decreasing costs and increasing the number of graduates to fill positions that fuel economic growth and community prosperity. Higher education governing boards are calling for increased accountability and transparency, and regulatory bodies seek evidence of compliance in an environment of reduced public funding and increased competition for students.



These new requirements cry out for a solution like Wasabi's hot cloud storage.

¹ <https://evollution.com/technology/metrics/the-past-present-and-future-of-big-data-in-higher-ed/>

don't heavily benefit from the use of data reduction techniques such as compression or deduplication (the elimination of multiple copies of the same file, or redundant information within a file).

In order for deduplication to be effective, there must be redundancy (usually at the block level). This means that deduplication works really well for reducing the storage footprint of virtual machines and data backups, but there are other data types for which deduplication doesn't work very well. Data that cannot benefit from deduplication commonly includes structured databases, multimedia files, and scientific data. These are some of the most popular data types in schools.

Compression is another commonly-used data reduction technology. Unlike deduplication, which usually works at the block level, compression works at the file level. Compression works by finding redundancy within files, and replacing the redundant bits with a token. Once again, keep in mind that compression can only reduce the size of a file if redundancy exists within the file. Scientific data tends not to benefit much from compression, and compression can actually *increase* the size of multimedia files.

The end result is that schools commonly end up with large amounts of data that can't be significantly

shrunk, and yet must still be stored in spite of a very limited budget.

Like other organizations, schools have two basic options for data storage.

These options include storing the data on- premises, or storing the data in the cloud. On-premises data storage isn't always practical because it's expensive, has scalability limitations, and can be difficult to maintain.

Using cloud storage allows a school to avoid costs associated with storage hardware acquisition and maintenance, but can be just as expensive (if not more so) in the long run due to data egress fees and other charges imposed by the cloud provider.

Of course, costs are far from being the only consideration when schools make a storage purchasing decision. An equally important consideration is security.

Schools face security threats that far exceed those of many other organizations, due to the "open" nature of academic institutions – open from the inside and out – as compared to a Financial Services organization, as one example.

Some of these threats may include computer science students anxious to test their skills by hacking the school. Or inadvertent ransomware infections brought

on by students or teachers who opened a message, file or visited a site they shouldn't have. In any case, any IT resources in use by a school must be hardened against these and countless other threats.

Wasabi's Hot Cloud Storage as an Option for Educational Institutions

As previously noted, schools commonly suffer from low budgets and a high level of security threats. As such, educational institutions need low-cost, high-volume storage that is both scalable and secure. Wasabi provides a cloud storage platform that meets these needs.



Having the ability to free data from the archives is especially beneficial to research universities. By maintaining research data online, a university can make the data available to other researchers who are working on related projects.

Wasabi has structured its cloud storage pricing in a way that provides high-volume storage at a very low cost per gigabyte.

As mentioned previously, but bears mentioning again – particularly in environments in which every penny

counts – Wasabi doesn't charge the egress or API request fees that have become so pervasive among cloud storage providers. Wasabi bills its customers based on the storage that they consume, without the hidden fees.

Because Wasabi provides storage that is both scalable and inexpensive, schools are able to move data off of archive/cold storage, and onto cloud storage where it can be made readily accessible. This might include anything from old course materials to lecture recordings.

Having the ability to free data from the archives is especially beneficial to research universities. By maintaining research data online, a university can make the data available to other researchers who are working on related projects. Accessing this previously inaccessible data can potentially improve the quality of research, and may reveal trends or correlations that would previously have gone unnoticed.

For research and educational organizations who are part of the over 60,000 members of the Internet2 community,⁴ they can choose Wasabi from the Internet2 Cloud Exchange as one of the layers that powers their research initiatives. The benefits are both storage at far lower costs than they are used to, and the ability to take advantage of the high-capacity, low-latency

⁴ <https://www.internet2.edu>

network that connects Internet2 members and providers including Wasabi.

Because so much of the research being done at the university level is in some way related to medicine, it's worth noting that Wasabi's cloud storage platform is fully HIPAA compliant.

Of course, schools perform many other types of research as well, ranging in scope from mechanical engineering to materials science. Although such fields are not related to medicine and are not subject to the HIPAA regulations, using HIPAA compliant storage can still be beneficial. Even if the research data itself isn't subject to HIPAA, information about the researchers may be. Even documenting something as simple as a special-needs student's condition and requirements can, under the right circumstances, mandate HIPAA compliance.

Of course, HIPAA isn't the only set of regulations that schools need to be concerned about. The Family Educational Rights and Privacy Act (FERPA) was signed into law in 1974. FERPA sets forth standards for the protection of student records, and FERPA compliance is mandatory for any school, school district, college, or university that receives federal funding.

An educational institution that fails to fully comply with FERPA can lose federal funding, and may be fined by the state, or be subjected to other sanctions. As such, it's critically important for educational institutions to maintain FERPA compliance. Even if a school isn't technically required to comply with FERPA, it's conceivable that the school could one day receive federal funding and be required to comply, so it would be prudent for such a school to work toward compliance before it's actually required. A good first step is for the school to store student data in the Wasabi cloud, which is certified to be FERPA compliant.

As previously mentioned, data security is of the utmost importance for schools. Not only is good data security necessary for maintaining regulatory compliance, but security also helps to protect data against unauthorized access, ransomware, and other forms of cyber-vandalism. Wasabi's immutability feature allows educational institutions to guarantee that high-value data cannot be changed, erased, or encrypted by a ransomware attack.

Another advantage to using Wasabi's cloud storage platform is that doing so can help to prevent vendor lock-in.

Wasabi storage is designed to be fully compatible with Amazon S3, which means that data can be seamlessly transferred to or from the AWS cloud on an as-needed basis.

Schools no longer have to worry about the consequences of becoming overly dependent on a single vendor, and that vendor's offerings.

CONCLUSION

What Could You Do With Hot Cloud Storage?

One of the things that originally made Cloud 1.0 solutions so popular was the promise of unleashing potential that simply couldn't be realized by using on-premises resources. As the cloud has matured, and the Cloud 2.0 era has been ushered in, the cloud continues to push the envelope of what's possible.

Much of this new potential centers around improvements to cloud storage. Although cloud storage was at one time primarily suited for archival purposes, hot cloud storage has changed the game, enabling cloud data to remain online and readily accessible – and at a price that puts it within the reach of all.

This means that it may finally be possible to unleash the value in all of your old data by moving it to the cloud. Just think of what you may be able to achieve when you apply Artificial Intelligence or Machine Learning to all that data.



Wasabi's Hot Cloud Storage delivers a low cost, high performance, secure and S3-compatible storage offering that is truly next generation cloud storage.

Remember, it's no longer about things like compute and workload management – they're simply tools to help extract the gold that is the real motherlode of Cloud 2.0: data mining and analysis.

Given the factors we've looked at, it's clear that Wasabi is a company worth considering strongly.

Wasabi's Hot Cloud Storage delivers a low cost, high performance, secure and S3-compatible storage offering that is truly next generation cloud storage.

Their lower cost allows organizations of all kinds to get more mileage from their cloud storage budget, thereby allowing a larger volume of data to be stored in the cloud than would otherwise be possible.

And most importantly, that cost savings can be put back into the truly valuable work of your organization, rather than consuming an increasingly large portion of your budget.

After all, the purpose of your organization isn't really about storage, is it?

Whether you've already begun using Cloud Storage 1.0 solutions, or were waiting until the price came down and performance improved, it's time to start forming your Cloud Storage 2.0 strategy. There are both significant cost savings and new opportunities waiting for you. You can climb aboard, or risk being left at the station.

Wasabi's Hot Cloud Storage is the promise of Cloud Storage 2.0, realized.

The choice is yours.