The Storage Professional's Guide to

Surviving the BIG DATA Apocalypse

wasabi
hot cloud storage

She Who Has the Most Data Wins

Data is transforming the world. Advances in the way we capture, process, and analyze digital data are enabling greater efficiencies, better insights and decision making, even entire new business models.

Data is pushing the boundaries of what is possible in every industry. It's also taxing our current ability to store even a fraction of the amount of data that our new digital world is generating.

The New World is Data Driven

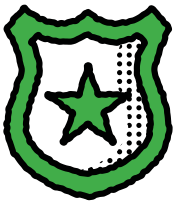

LAW ENFORCEMENT

Video surveillance and facial recognition software are helping police catch bad guys. The Boston Marathon bombers were identified by machine-analysis of faces that didn't look surprised when the bomb went off!

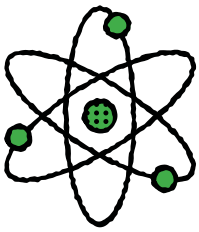

BIOTECH AND SCIENCE

Applying machine-learning algorithms to massive data sets are helping research scientists make incredible breakthroughs in every field—from quantum physics to astrophysics to human health issues here on Earth.

MEDIA AND EDUCATION

Virtual and augmented reality is shaping a new reality for the future of education as well as entertainment. Meanwhile, faster broadband and new streaming video services are creating hot new revenue streams for old content sitting in cold storage.

Stop. Children, What's that Sound?

There's a dark side to all this innovation. Hear that rumble? That's the sound of 163 trillion gigabytes of data that's being generated by millions of machines, mobile devices, sensors, high-resolution cameras, and fancy medical imaging equipment.

And it's heading our way.

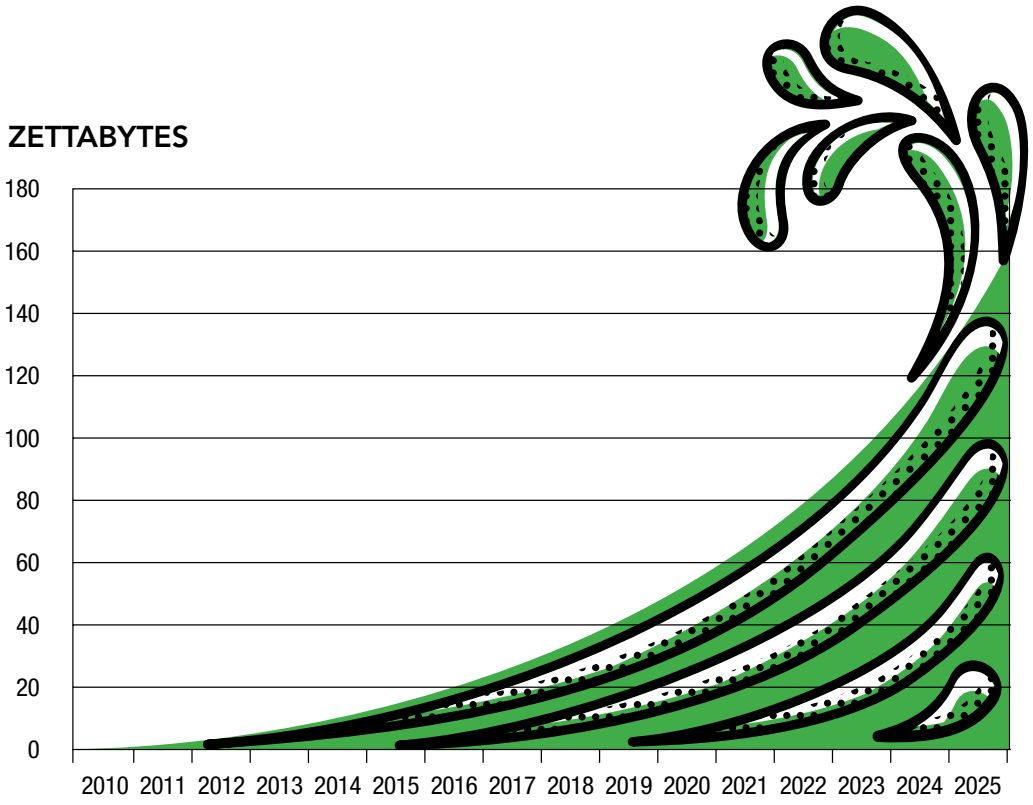

The world's digital data will double in size every two years.

EARTH 2016 / 16.1 ZB - - - - - - →

EARTH 2025 / 163 ZB - - - - - - →

It would take 16 billion 12TB enterprise HDDs to store this much data.

The Four Horsemen of the Big Data Apocalypse

(The "Four Vs of Big Data" didn't sound quite as badass!)

1 VOLUME

The brainiacs that track this sort of thing[1] estimate the world's total digital data will exceed 163 zettabytes by 2025. That's 163 trillion gigabytes. Or 250 trillion James Brown CDs.

2 VELOCITY

The pace at which digital data is growing is increasing as well: 10x faster than traditional business data. Machine-generated data is growing at a whopping 50x faster rate. As the Internet of Things kicks into high gear, all those intelligent machines, smart cities, autonomous cars, cameras, and wearables will make the current data tsunami look like a drop in the pond.

3 VARIETY

Not long ago, data came in from a small number of predictable sources in a few different but highly structured formats. But in a world where everything generates information, not only has the number and variety of unstructured data sources expanded, so has the size of the individual files.

4 VALUE

We already touched upon the growing value of data, but it's worth repeating. As we continue to get better at collecting, connecting, manipulating, and analyzing data, more organizations will want (and need) to store more data more often for longer periods of time.

[1] Brainiacs include IDC's David Reinsel, John Gantz, and John Rydning

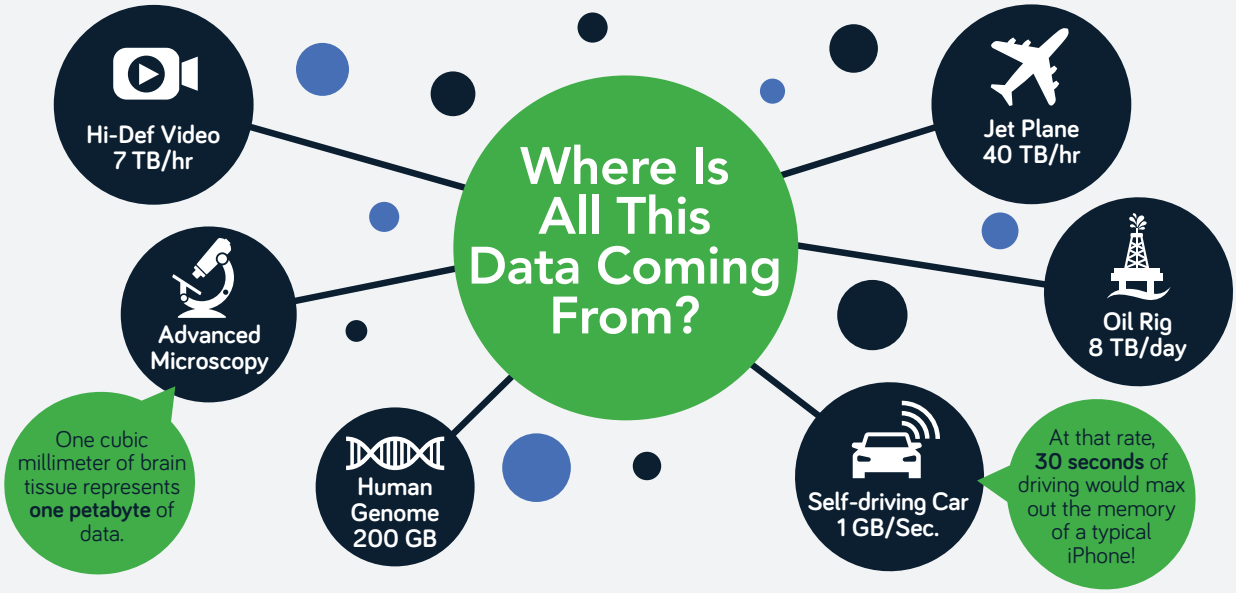

Why On-Premises Storage is a Flawed Premise

In the face of this data tsunami, your first instinct may be to throw more money at your existing on-premises storage infrastructure. But we all know that IT budgets, even when growing, are not infinite. The percentage of IT budget allocated to storage is already high—and there are many other high priority (and higher profile) initiatives competing for precious budget.

Before you sink another penny into beefing up your existing infrastructure, consider the following:

1 ON-PREMISES STORAGE IS EXPENSIVE

A decent on-premises storage server will run you about $344K in capital investment for 1 petabyte of storage. But a server does not a system make. On average, physical storage represents about 60% of total costs. Add in supporting gear like switches, routers, racks, cables, utility costs, and staffing. Oh, and the yearly maintenance costs (another 18%), and you're looking at approximately 2.3 cents per GB per month on storage over a five-year period.[2]

2 LOCKED-IN OBSOLESCENCE

That five-year period is an important data point to consider. Because that's how long you are typically bound to an on-premises service contract. That's an eternity in IT. Sure, you can benefit from incremental improvements with software upgrades, but slick new high-speed, high-capacity hardware won't come often, and won't come cheap. You are essentially locked into an investment that is losing value as it ages.

3 DOESN'T SCALE EASILY OR CHEAPLY

You will be generating and amassing more data. That's a fact. You know what else is a fact? Physics. There is only so much physical space inside a server – and only so much room in your facility. Scaling capacity means more boxes, more electricity for power and cooling, and less space in your office or data center for other important things. Like employees.

[2] Kevin L. Jackson, CEO and founder of GovCloud Network.

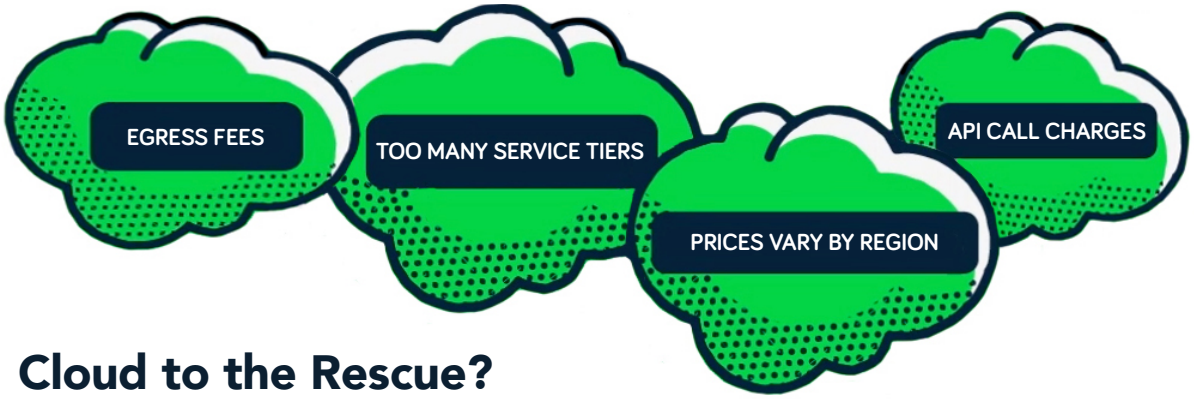

Cloud to the Rescue? NOT SO FAST

The cloud would seem to be the perfect solution. There are no upfront costs—the cloud is pay as you go. Cloud services are infinitely scalable. You can take advantage of their latest innovations immediately. And from a cost standpoint, you benefit from the cloud's economies of scale. **Or do you?**

As it turns out, there's a reason (or three) why most of the world's data is still stored in on-premises data centers.

What happened to that big migration to the cloud?

1 THERE'S NO REAL INCENTIVE TO MIGRATE

And, no, it's not because of security concerns. In 2018, most of us well understand how and why the major cloud service providers are even more secure than our own data centers. They have to be. Their businesses depend on it.

Surprisingly, there's no real cost benefit for switching. The baseline cost of storing data in Amazon's Simple Standard Service (S3) is 2.3 cents per GB per month—the same price as 1 PB stored on-premises when amortized over five years.

True, there are no upfront capital costs with the cloud, but there's also no incentive to risk such a momentous task of migrating all of your workloads to the cloud.

In fact, once you add in all the "hidden" costs for egress and various transactions and API calls, cloud storage service from the Amazons, Microsofts, and Googles of the world can end up costing more than traditional on-premises storage.

2 IT'S COMPLICATED

All the leading cloud storage providers offer a dizzying array of service tiers, each with different pricing models designed to lower your overall costs. On the face of it, it makes sense.

Don't need high levels of redundancy or resiliency? Pay less money. Don't need to access your data often—or at all? They have lower-cost options for you there, too.

Some companies hire expensive cloud storage professionals or consultants that do nothing but calculate which workloads to store where—in an attempt to control what can quickly become out-of-control cloud expenditures. It's that complicated.

New A.I. tools have entered the scene that can make these cost-saving calculations for you—even automatically transfer your data from one tier to the next. However, nothing can prevent the sticker shock you'll experience if you ever decide you need that data you've been storing in those "lower-cost" service tiers just to save money.

COLD STORAGE IS NOT THAT HOT

Cold storage solutions like Amazon's Glacier can significantly lower the cost of long-term archival storage. IF YOU NEVER, EVER, NEED TO ACCESS IT. If you do, the extra access fees can quickly eat up any savings you were counting on.

And don't expect to get your data right away. It's in COLD storage. That means you will need to submit a request for your data. Then wait, possibly up to 24 hrs. before your data is even ready for you to begin to download.

Not exactly the user experience we've all come to expect—and require—in the fast-paced, always-on, information-at-your-fingertips world of the digital era.

So, if cloud isn't the answer. And on-prem storage isn't the answer. How are we supposed to keep our heads above water and make it through this data tsunami alive?

FUN FACT:

With Cloud 1.0, you never know what your total storage bill is going to be each month.

Cloud Storage 2.0: A New Hope

Cloud storage 2.0 is cloud storage built for the big data era. Incredibly fast, but ridiculously inexpensive. So, there's no need for complicated storage tiers or figuring when to move from hot storage to cold. It's all hot. Red hot.

Wasabi Hot Cloud Storage

Hot cloud storage is 2.0 technology invented by Wasabi co-founders David Friend and Jeff Flowers.

They were out to dramatically lower the cost of storage. By laying data on a disk in sequence instead of in blocks they were able to maximize every micron of disk space to deliver a storage solution that is 80% cheaper than Amazon S3.

They suspected that this approach would also boost performance since less mechanical movement would be required to read or write sequential data. But faster than the competition?

The test results blew them away. They had just invented the world's highest performing cloud storage…and they could sell it for less than the cheapest cold storage tier.

The hot cloud storage movement was born.

wasabi™
hot cloud storage

- Price – 80% cheaper than AWS S3
- Performance- faster than the competition
- Protection – 11 nines of durability, plus **data immutability** (a huge deal.)

PLUS FREE UNLIMITED EGRESS!

Reliability, Immutability, and the Real Reasons for Data Loss.

The top tier cloud storage services, including Wasabi, all offer 99.999999999% (that's 11 nines) of durability. Statistically, that means for every 1 million files you store, you might lose one every 659,000 years or so! Let's agree to call that "pretty darned reliable."

But . . . two thirds of all data loss has nothing to do with hardware failure. The true villains are a combination of human error, viruses, software bugs, and malicious employees or intruders. Which is why Wasabi introduced the "immutable bucket."

Data stored in an immutable bucket cannot be erased, overwritten, or modified by anyone. Not by your storage admins. Not even by Wasabi. Once you write it, it remains until the time you designate expires.

Immutability is ideal protection against accidental deletions and for ensuring compliance with strict government and industry regulations.

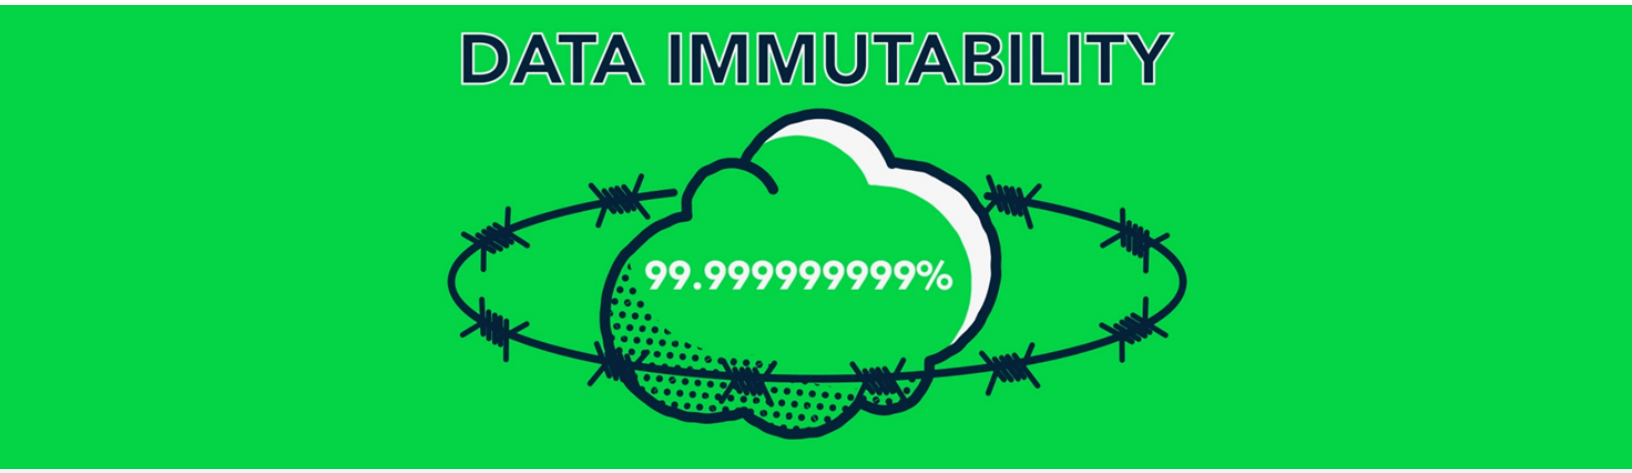

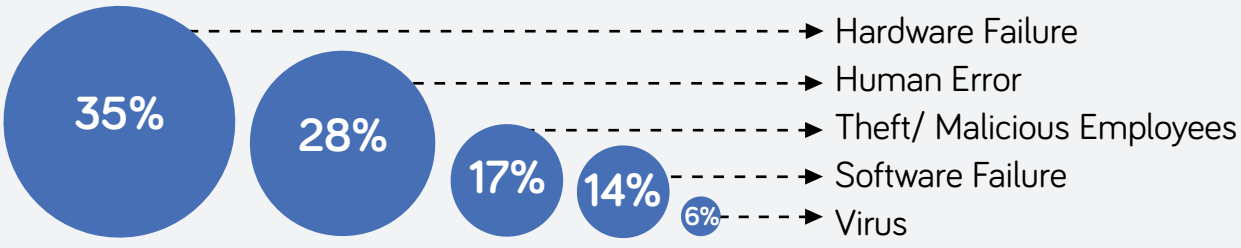

Source: http://tweakyourbiz.com/technology/2015/06/19/5-greatest-causes-online-data-loss/

Let the Great Migration to the Cloud Begin

Thanks to near limitless, inexpensive and lightning-fast cloud 2.0 storage like Wasabi hot cloud storage, there is finally an affordable alternative to expensive on-premises infrastructure and 1.0 cloud services. Hot cloud storage not only enables you to cost-effectively hold on to more data longer (price!) and keep it safer (protection!) – by eliminating storage tiers (performance!) and all those pesky transaction and egress fees, you'll never be surprised by cloud storage cost overruns again.

Hot cloud storage costs are not only the lowest on the planet, they're the only ones you can predict with absolute surety.

COST COMPARISON FOR 1 PB OF STORAGE			
COSTS	ON PREMISES	CLOUD 1	CLOUD 2
Capital costs depreciated over 5 years	$69K	$0	$0
Maintenance fees	$62K	$0	$0
Equipment and power (networking, computers, racks, cable, power)	$87K	$0	$0
Data Center and staffing	$52K	$0	$0
Total	$270K	$276K	$49K

Want to see the numbers crunched? Read The Future of On Premises Storage.

Cold Storage Is Dead. Say Hello to Hot Archiving.

Forget storage tiers. Forget complicated pricing tables or storage lifecycle calculations. Forget having to figure out—in advance—what you might need to access in the foreseeable future.

With hot cloud storage, ALL of your data is readily accessible whenever you want it . . . and faster than the competition.

Blazing fast active cloud storage at cold storage prices = hot cloud storage

There's No Stopping the Data Tsunami.

But You Can Avoid the Apocalypse.

So long as technology continues to advance, and data is viewed as a valuable asset, there will be no stemming the tide. You can delete data and hope you didn't just throw away the future of your business…or the cure for cancer. You can throw an increasingly higher percentage of your IT budget into current generation storage solutions.

But wouldn't you rather use that money for high-performance computing, A.I., and all the other analytics tools required to milk more value from all that stored data in the first place?

Then the answer, my friend, is **Wasabi hot cloud storage**.

Check out our price comparison calculator. Test it out for yourself.

Try it for Free for 30 Days

You'll quickly realize that Wasabi will not only help you survive the big data apocalypse, it can become a key component in your storage strategy for thriving in this rapidly evolving digital world we all find ourselves in.

ABOUT WASABI

Wasabi is the hot cloud storage company delivering disruptive storage technology that is 1/5th the price of Amazon S3 and faster than the competition with no fees for egress or API requests. Unlike first generation cloud vendors, Wasabi focuses solely on providing the world's best cloud storage platform. Created by Carbonite co-founders and cloud storage pioneers David Friend and Jeff Flowers, Wasabi is on a mission to commoditize the storage industry. Wasabi is a privately held company based in Boston, MA.

Tel 1-844-WASABI-1
Email info@wasabi.com

wasabi hot cloud storage
www.wasabi.com