



To the Cloud and Back

*Why some companies are repatriating data and workloads
and how their experiences can help you shape your rent-or-buy decisions.*

March 2021

Analyst: Mike Feibus

in association with

intel

Download your copy at:

<https://www.feibustech.com/blog/to-the-cloud-and-back/>

Contents

Introduction	4
Background and Context	6
Cost of Deployment	6
Size Matters	7
Deploy Smarter	8
Buy It Now	9
Performance-Minded Considerations	9
Fine-Grained Control	11
Ready for Anything	11
Beyond Security	11
Conclusions	13

Introduction

At first blush, the promise of the public cloud sounds almost too good to be true: shed the burden of owning and operating datacenters in return for access to seemingly boundless storage and compute capabilities you only pay for when you use. Which is why a steady stream of both public and private entities' workloads continue to flow to the cloud. And the pandemic, which effectively broke up face-to-face collaboration for a time, has only accelerated the trend.

But as most everyone discovers, something that sounds too good to be true often turns out to be just that. And the cloud is no different. Indeed, the ecosystem is full of stories about companies burned by their journey. Just a few of their complaints:

Cost: Contracts with public cloud vendors can result in sky-high bills, particularly if terms are poorly aligned with data storage and compute needs. Some CIOs say it can take years to become proficient enough at forecasting cloud demand to strike the right mix of unused capacity and overage surcharges.

Performance: Performance in general – and latency in particular – cannot match optimized private implementations.

Security: Privacy and security in a shared environment is inherently inferior to a customer-owned setup, which presents far more options to define and tune.

If the complaints sound familiar, it is for good reason: they are virtually identical to datacenter grievances that helped fuel the rise of public cloud in the first place.

So how could both the public and private options be simultaneously cheaper and more expensive than the other? Faster and slower? More secure and more vulnerable? The answer, quite simply, is that it depends.

One variable is size. Companies need the heft to buy server hardware in volumes large enough to tip the scales in favor of private datacenters and co-located facilities. Scale also makes it possible to hire and maintain adequate expertise to design, implement and manage the operations. As proof, on-premises advocates often point to exodus of Dropbox, which was born and bred in the cloud, from Amazon Web Services, or AWS. The first year in its own facilities, Dropbox reaped nearly \$40 million in savings.



But size is not the only consideration. Cloud native or not, Fortune 1000 companies from Apple to Zillow make extensive use of public cloud services. That said, few house their entire computing operations in the cloud.

Indeed, migrating to the cloud is rarely an all-or-nothing proposition. Companies increasingly are adopting a hybrid cloud approach, deciding on a case-by-case basis whether to park specific data stores, workloads or both in the cloud or across owned assets.

Netflix is a good example. The video streaming service's revenue is about 10 times Dropbox sales. Yet company executives show no sign of leaving AWS, where it hosts and manages myriad operations, including app hosting, subscription management and billing. At the same time, though, Netflix retains complete control over its core service: the content delivery network, or CDN, that actually delivers movies, TV series and comedy specials to its nearly 200 million subscribers.

And despite its very public divorce, [Dropbox continues to contract with Amazon](#), augmenting its own facilities with AWS' extensive global reach. Currently, the company stores something less than 10 percent of customer data on AWS.

There are also competitive considerations, particularly for cloud native companies. Amazon, Google and Microsoft, the three largest public cloud providers, all have their hands in many pots, some of which compete with cloud customers. Amazon Drive and Prime Video, for example, overlap services from Dropbox and Netflix, respectively.

So how should decision-makers weigh these and other variables for where to locate their data and workloads? This is, of course, a very complex decision colored by individual entities' existing assets, competitive advantages, needs and goals. To help add some structure to the decision process, FeibusTech has produced this research brief. The guideposts presented in this report are the result of extensive research and analysis, including interviews with executives from a wide range of organizations with broadly divergent experience in the cloud, including public and private companies, traditional and cloud-native businesses as well as government agencies across industries and functions.



Background and Context

It is impossible to assess the impact of the public cloud without acknowledging the myriad cloud-native services it helped enable. The cloud's flexibility and reach has breathed life into so many entrepreneurs' ideas, spawning a raft of billion-dollar internet services over the past dozen years, like Airbnb, Pinterest, Snapchat and Spotify.

As the public cloud broadened and matured, it has extended its reach well beyond cloud-native companies, attracting many other types of public and private organizations in search of better reliability, global access and presence, higher performance and lower cost.

The pandemic greatly accelerated the trend toward the cloud, compressing timetables from years to months – seemingly overnight. According to [a recent enterprise IT survey](#), 87 percent of decision makers say that COVID-19 has accelerated cloud migration.

87 percent of decision makers say that covid-19 has accelerated cloud migration.

Of course, migrating to the cloud brings its own set of challenges – and they are strikingly similar to what many enterprises face with their on-premises assets. Security and cost, for example, rank as enterprise's top cloud challenges, according to [Flexera's 2020 State of the Cloud Report](#). Eighty one percent of survey respondents cite security as a top challenge, and 79 percent point to cloud spend.

81 percent of survey respondents cite security as a top challenge, and **79 percent** point to cloud spend.


Intertwined with security are a whole host of privacy-related directives, from protecting client data and meeting regulatory requirements to keeping competitive implementations and insights safe from prying eyes. And performance, of course, is top of mind for many companies that serve a broad, demanding customer base via browsers and applications.

Cost of Deployment

The public cloud isn't inherently more expensive than on-premise or co-located deployments, IT architecture decision makers say. But for all the vast storage and compute capability, cloud cost structures can be less flexible – and, consequently, more difficult to exploit efficiencies – than customer-owned assets. That said, taking full advantage of on-premise hardware efficiencies is no cakewalk. It takes effort, expertise and ingenuity.

IT architects stress that to realize savings in the public cloud, you must craft multi-year contracts at the outset that are aligned with what your actual demand turns out to be. That is a tall order, needless to say, that takes a savvy team with both technical knowledge and business acumen to successfully architect. But failing to do so invariably will result in higher overall costs, as you'll be paying for the capacity you reserved, whether you use it or not. As well, time-of-use and resource surcharges also can blow through budgets. The difference between actual and anticipated compute costs can be significant. On average, cloud spend is 23 percent over budget, according to the Flexera survey. Worse, decision makers expect their spend to inflate even faster next year, as the following chart shows:

On average, cloud spend is **23 percent** over budget.



Of course, it can be challenging to forecast storage and compute needs out several years when you are building out your own datacenters. But the public cloud, with its multi-year service contracts, is no panacea. Indeed, many IT decision makers who have migrated their organizations to the cloud say that they just traded one difficult forecasting exercise for another.

In addition, many IT executives say they have more options available on premises to help engineer workarounds for planned versus actual use mismatches. They also say that leasing someone else's equipment via the public cloud provides fewer opportunities to streamline architectures and fine-tune configurations to cut overall costs and improve performance.

Leasing someone else's equipment via the public cloud provides fewer opportunities to streamline architectures and finetune configurations.

Size Matters

When it comes to large-scale deployments, IT architectural decision makers say there are more opportunities to cut costs on premises than in the cloud. Of course, scale affords large enterprises leverage to negotiate better deals, both for leasing resources in the public cloud as well as purchasing hardware outright. But, many agree, there is more room to maneuver when buying your own hardware.

In addition to securing volume discounts, scale also means the organization likely has resources in place to devote to optimizing deployments. And there are far more opportunities to tune on premises, the IT architecture decision makers say. Which means they can do more with less.

But while decision makers agree that scale generally favors on-premises and co-located deployments over the public cloud, the crossover point – that is, the volume levels that tip the scales in favor of customer-owned hardware – is still very much up for debate.

Over the past few years, [Dropbox](#) has emerged as something of a high-water mark for assessing when companies have outgrown the cloud, and repatriating makes unequivocal sense. In late 2015, the cloud storage company moved 90 percent of its 500 million users' 500 petabytes of data from AWS' S3 service to its own network of datacenters.

In late 2015, Dropbox moved **90 percent** of its 500 million users' 500 petabytes of data from AWS' S3 service to its own network of datacenters.

The first year in its own facilities, Dropbox [netted \\$39.5 million in savings](#), according to SEC filings. Naturally, the company's capital-related expenditures rose markedly (\$53 million) that year – but nowhere near as much as its cloud bills fell (\$92.5 million). The following year, the cloud storage provider saved a further \$35.1 million.

Beyond sheer purchasing power, corporate girth affords companies like Dropbox the ability to deploy resources to ensure that owned assets are not only efficient, but secure. Furthermore, scale gives companies with outsized needs, the flexibility to streamline hardware so that there is not a lot of waste.

“At the scale we’re at, we see tremendous economies from our on-premise datacenters,” Matt Singer, Senior Staff Hardware Engineer at Twitter, said. “Once we figure out what we need, we can build hardware that’s very, very close to that.”

Smaller companies, conversely, “might need to stay in a t-shirt size that isn’t really right for them,” Singer said. That might mean, for example, buying OEM server bundles with more expensive components – like RAID configurations or hot-swappable disks – than their application requires.

Deploy Smarter

For its part, Twitter repartitioned its datacenters into three classes of servers, each with no more cost or capability than is absolutely needed for the task:

- **Front-End Servers** include a whole lot of memory and processing power, and just enough storage to boot up,
- **Storage Servers**, which, as their name implies, are packed densely with drives and little else, and
- **Database Servers**, which are basically higher-performance servers with redundant, high-performance storage.

Twitter’s Three Classes of Servers

1 <i>Front-End Servers</i>	2 <i>Storage Servers</i>	3 “Database”
Servers with lots of memory and processing power, and just enough storage to boot up.	As the name implies, these servers are packed densely with drives, and little else.	Servers that feature high-performance disk sub-systems with built-in redundancy.
<i>CPU, memory dominate the design.</i>	<i>Evolved to provide huge storage capacity.</i>	<i>Disk capacity, performance, and RAS..</i>

Twitter’s server repartitioning is not unlike the concept deployed by modern-day dental offices, which boost revenue and slash costs by similarly splitting workloads into tasks of varying cost structures. Dentists used to see one patient at a time, cleaning the patients’ teeth and taking x-rays as well as doing an examination. With repartitioning, offices can boost throughput by hygienists to clean multiple patients’ teeth simultaneously. The one high-priced, highly certified dentist can drop in on each of them to do her examination.

Twitter saves quite a bit of money on deployments with this construct, Singer says.

As for where to locate data and workloads, Twitter is now adopting a more hybrid approach, on two fronts. Specifically:

- As a vehicle for handling rapid growth, traffic spikes and other sudden demand changes, Twitter announced late last year it is bringing up AWS to augment real-time serving functions. That will give the social media giant new load-balancing flexibility, enabling it to quickly respond to changes with AWS resources rather than relying purely on its own data centers, and
- Early this year, Twitter revealed it is expanding its relationship with Google Cloud Platform, or GCP, to handle offline, asynchronous machine learning workloads. With the expanded relationship, Twitter engineers can now save time by leveraging existing GCP services rather than writing one-off applications for individual jobs, Singer said.

Buy It Now

As with Twitter, testing, prototyping, AI training and other elastic workloads at eBay start in the cloud. "Once we've developed them, then we bring those workloads internal," said Lam Dong, Director of Platform Hardware Engineering at eBay.

Dong agreed that on-premise assets are much more cost-effective than the public cloud for repetitive, well-understood tasks.

eBay has migrated most of its critical corporate IT operations to the cloud, Dong said. But most anything associated with ebay.com, the ecommerce pioneer's core business, is on premise. The primary motivation for owning rather than renting? Cost, Dong said.

On-premise assets are much more cost-effective than the public cloud for repetitive, well-understood tasks.

"You have to reach a certain critical mass to be able to run your own more efficiently," Dong said. "Up to that point, the cloud can be more attractive to smaller companies."

Dong cautions that higher efficiencies and lower costs at scale are not a given. Even when the math adds up, he said, you still need the right talent in-house to realize the savings. That is why he believes some companies stay in the cloud long past when, on paper at least, it makes sense to repatriate.

Performance-Minded Considerations



As mentioned in the previous section, cloud-first companies like eBay and Twitter are comfortable relying on the cloud for applications where the difference between response times of milliseconds and seconds isn't consequential. That makes traditional IT workloads candidates for cloud migration.

As well, cloud-first companies often rely on the high-performance compute capability at the ready in the cloud for AI training workloads that will migrate to on-premises operations when they are mature and ready to live.

Public cloud providers have made great strides in improving responsiveness and reliability in recent years. Which makes the cloud a viable solution for some applications where time, you might say, is of the essence.

For those real-time applications, it comes down to three simple questions:

- Does the cloud offer acceptable performance for my application,
- Can I do better? And if so,
- Is it consequential to my business?

Put another way: if the operation is crucial to how customers perceive your business, then lean toward the cloud for workloads that any cloud provider could perform as well as you can on-premise, and hold close, what you believe to be a competitive edge.

This helps explain why Netflix houses the back end for its video streaming application on AWS, but operates its own content delivery network: Netflix views front-end app capabilities as a commodity, but sees real differentiation in retaining end-to-end control over the content delivery experience.

Netflix views front-end app capabilities as a commodity but sees real differentiation in retaining end-to-end control over the content delivery experience.

At a time when new entrants like Apple, Disney, HBO and NBC are flooding the market with choice, Netflix can ill-afford any subscriber dissatisfaction. Nearly two-thirds of consumers who are considering adding one or more of the new services intend to downgrade or cancel at least one of their existing subscriptions, according to [PwC's recent video streaming survey](#). And [a new Deloitte survey](#) found that the pandemic was accelerating those decisions.

As part of its CDN, which is called Netflix Open Connect, Netflix installs servers at internet service providers' own datacenters for free, provided the ISP handles enough Netflix traffic to qualify. Netflix has optimized the appliances, which are Intel-based servers, for the application. And ISPs welcome the installations because they ease congestion on their network and slash costs. And because Open Connect improves streaming quality, their subscribers are happy. Which is good for their business as well as for Netflix'.

The care and feeding of the proprietary Netflix Open Connect also acts as a barrier to entry for newcomers in the streaming business. Thus far, Netflix has deployed Intel-based Open Connect appliances in nearly 1,000 locations globally.

The investment appears to be paying off. In late April, about a month after the pandemic shut down much of the economy, Netflix reported the fewest technical issues of any streaming service, according to [a JD Power study](#). It was also the service more subscribers said they would keep if they could only have one.

Many others follow a similar model. Zoom, for example, hosts back-end processes on AWS, but operates real-time traffic on its own servers.

Of course, latency is essential for more than just ensuring quality streaming. It is also critical for producing content in the first place. And the proliferation of virtual workstation technology in video production has only accelerated since the pandemic.

Those who engineer production deployments for studios say that latency is a primary consideration. To help ensure a quality remote production experience, architects say they need to cap latency to just a few seconds. One way to achieve this is to locate private facilities as close to the production as possible. Some also leverage the reach of the public cloud to incorporate filmmaking talent located around the globe.

As well, architects at some very large cloud-native operations say they can minimize latency across far greater distances by employing dedicating leased lines. Of course, leased lines means skirting the internet, which effectively eliminates the public cloud as an option for those operations.

For video editing and other demanding applications, savvy on-premise optimization of the tradeoffs between performance and energy consumption also helps to ensure companies get the most out of each rack. Of course, public cloud operators are also intent on maximizing performance per watt – but for a more general class of applications.

Fine-Grained Control

Optimized hardware, leased lines and pinpoint server location are just a few tools that architects rely on to achieve superior latency over what the cloud can provide. Many cloud-native companies, in fact, take it much further than that.



Modern, distributed applications have many performance, reliability and security benefits. But for some large, cloud-native companies, it sometimes means that the required microservices are scattered across hundreds of far-flung servers. Architects say that requires a level of coordination that's simply not possible in the public cloud.

Twitter is a great example. Its system processes more than 1 trillion events each day across about 100,000 servers globally.

"When you load your timeline, your phone is going to interact with maybe 10 or 20 top-level Twitter services that are going to fan out into requests across hundreds of services," Twitter's Singer said. "If you want quick response times, then you can't afford to add even a millisecond to all of those calls. You have to have control over everything."

Ready for Anything

Perhaps the best reason for maintaining a hybrid cloud approach is that the technology is constantly changing. This means what looks like the leading candidate for a particular workload today may be lagging tomorrow.

"A data platform is not now, never has been, and never will be a single choice," said Charles Boicey, Chief Innovation Officer at ClearSense, a data-as-a-service platform provider for healthcare. "A technology is going to come out that is far superior to what you are using. So you had better be able to migrate to it without having to rearchitect everything. There is just no time for that."

*"A data platform is not now, never has been, and never will be a single choice,"
said **Charles Boicey**
Chief Innovation Officer at ClearSense*

ClearSense counts four of the 10 largest US healthcare systems as customers. And like Twitter, some of ClearSense's biggest customers have a penchant for Google BigQuery. And ClearSense is set up for that.

"Data governance is at the center of everything we do," Boicey said. "So, if there is something special they want to run on another environment, they can go ahead and do that."

Beyond Security

Beyond the obvious – protecting data and processes from hackers, thieves and nation states – the term security means many different things to different people.

Often, privacy figures in the calculus for what information to protect, and how best to lock it down. Organizations do their best to keep prying eyes from myriad customer data – everything from addresses, birthdays and social security numbers to bank and credit card coordinates. And in this era, meta data like household income, political party affiliation, purchase history and sentiment assessments can be sought-after assets as well.



For some highly-regulated industries like banking and healthcare, companies are required to take steps toward safeguarding customer and patient data and could find themselves on the hook for breaches.

That is influencing cloud migration plans at Manulife-Sinochem Life Insurance, a joint venture between two Fortune Global 500 companies, ManuLife Financial Corp. and Sinochem of China. The venture has been offering life insurance in the China market for more than 20 years.

As an insurer, Manulife-Sinochem does not have the real-time front-facing requirements of an eBay or a Netflix. As such, it is executing on a broad migration of its core systems to the cloud. The one exception is regulated data, which will remain on-premise.

"Cloud will be the trend for the next few years," said Ellen Wu, CIO of ManuLife's China business. "The only systems to remain on-prem are those required by regulatory policy."

And then there is the matter of protecting information for competitive reasons. Wearables makers like Garmin or Samsung, for example, might be reticent to store customers' step counts, heartrate and sleep quality data in Google Cloud, now that the search juggernaut plans to buy Fitbit. Companies might also want to protect platform breakthroughs that improve the platform ease of use, for example. Or, like Netflix and Zoom, to shield proprietary delivery innovations.

"Cloud will be the trend for the next few years. . . the only systems to remain on-prem are those required by regulatory policy."

*said **Ellen Wu**
CIO of ManuLife in China.*

In retail, for example, [Target](#) has been scaling back its AWS usage. And Walmart, Amazon's largest retail customer, is [supplementing its own datacenters with cloud services from Microsoft Azure](#), not AWS.

Conclusions

With the touch of a button, the public cloud can provide instant flexibility, scale and global reach to anyone willing to pay. But as with any service, there are costs that come with the convenience the public cloud offers.

For some, as with cloud-native startups that require global reach from day one in order to realize the vision, the public cloud is effectively the only option. As well, companies that are expanding faster than their IT capabilities can handle will find the cloud can be a great way to get a handle on things. And even for at-scale companies well-grounded in their datacenters, the cloud is a great option for highly elastic workloads.

If there's one ground truth to cloud migration, however, it's that it will cost more than you expect. It can be difficult to tightly align usage agreements with actual demand, so waste and surcharges add up. Indeed, it's interesting to note that while cost is rarely a primary motivator for migrating data and workloads to the cloud, it often factors heavily into repatriation.

For those who do decide to repatriate data and workloads from the public cloud, here are a few things you should ask yourself:

- **Size:** Has my operation reached the scale needed to justify repatriation from a cost perspective?
- **Privacy/Security:** Are there competitive reasons to keep my data close? Regulatory reasons? Other considerations?
- **Performance:** Do I have workloads that differentiate my business? Is there an opportunity to tune on-premise assets for those workloads?
- **Expertise:** Do I have – or can I access – the expertise required to realize those benefits?

At scale, regular, predictable workloads are almost always less expensive to run in company-owned assets. That's because there are more dials to turn in your own facilities. As well, those dials can translate into better performance as well as lower cost. And many companies find the security advantages of the datacenter to be compelling draws.

Of course, there's no denying the convenience of the cloud. For many businesses, the mix of cost, performance and security of the public cloud is just fine for their needs. Even organizations with exacting demands for some data and workloads find the cloud to be acceptable – even preferable – for other applications.

One thing is for sure, and that's change. Technology is advancing quickly, which means that what's optimal today could hold you back tomorrow. Even so, the public cloud may never be able to match IT architects' ability to tune hardware configurations in their own datacenter.

FEIBUSTECH

·
· clear · critical · independent
·

FeibusTech

P.O. Box 25685

Scottsdale, AZ 85255

www.feibustech.com

+1-480-922-3244

Copyright © 2021

All Rights Reserved