

9

CLOUD TRUTHS

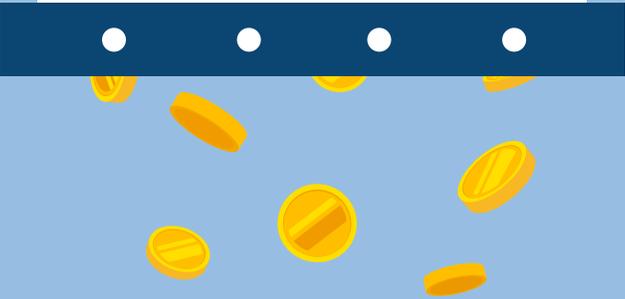
That No One Talks About

Cloud nine or cloud 911? After a pandemic-driven year of sometimes frantic cloud adoption, it's time to rediscover the essential cloud truths lost in all the bustle.

BY JOE MULLICH

2021 marks the cloud's 15th year of existence, making it a promising teenager on the cusp of maturity. But ready or not, the pandemic has supercharged the demand for cloud-based services: [IDC](#) reports that 80% of enterprises now intend to double the speed of their shift to a cloud-centric infrastructure. And the scramble to the cloud isn't going away when COVID-19 does: Gartner predicts [cloud spending will jump another 18%](#) in 2021.

In the midst of this gold rush, however, some fundamental truths about the cloud can easily get overlooked or papered over. That doesn't mean the cloud isn't delivering huge value, or that companies should hit pause on their cloud initiatives. But it can be difficult to focus on critical nuances when everyone is scrambling to transition to the cloud faster than ever before. So, to make it easier to understand just what you're getting into when you move critical applications and workloads to the cloud, here are nine important cloud truths that don't always get the attention they deserve:



1

The cloud may not be as cheap as you think—and that's OK

The idea that running your business in the cloud is always cheaper than using on-premises data centers has always been more myth than reality. “There is a lot of BS, to use a technical term, about saving money by using the cloud,” says consultant Joshua Greenbaum, Principal for Enterprise Application Consulting (EAC).

But many companies haven't gotten the message. In fact, according to a [2020 Flexera study](#), companies bust their cloud budgets by an average of 23%—a figure expected to rocket up to 47% in 2021. One problem is that companies waste some 30% of their cloud spend due to overprovisioning resources, and a lack of cloud infrastructure performance tuning can also lead to cost overruns.

These issues can be mitigated with the help of experienced cloud DevOps talent or better management tools, but not everyone is willing to make those investments. “People are biased toward saving money rather than spending money,” says Corey Quinn, Chief Cloud Economist with the Duckbill Group. “People misunderstand the nature of their business, and so they focus on shaving off pennies, rather than moving their business forward.”

Focusing on cost at the expense of strategic advantages falls under the category of penny wise, dollar foolish. Enterprises spend just a tiny fraction of their revenues on application hosting, so viewing cloud investments only through the lens of cost savings is limiting and not always strategic.

“A lot of the public cloud software forces companies to standardize on best practices, and that exercise can save a lot of time, complexity, and money,” says Greenbaum. “But it's ridiculous to move to the cloud just to try to get cheap software. That won't move your business forward. You want more from the cloud than cheap software. A lot more.”

The real cloud payoff is measured in speed, flexibility, and agility

2

“Treating the cloud just as an auxiliary to your existing data center can be a useful transitional step,” Quinn says. “But too many companies get stuck there, declare victory, and miss out on the real benefits of the cloud. Theoretically, the upper bounds of cloud savings is 100%, but you can make multiples of that by hitting the right market at the right time with the right product.”

The real payoff from cloud initiatives comes from bolder innovation, faster time-to-market, and better business agility. For instance, when you need to scale up workload processing power in the cloud, additional compute can be added in just a few clicks—or even automatically—making it vastly easier to experiment with new technologies and processes.

The cloud allows new applications to bypass the time-consuming provisioning and procurement processes that make on-premises technology so irksome. For example, [DBS Bank in Singapore](#) was able to leverage its cloud foundation to quickly spin up new, self-service applications, including one that allowed house-bound customers to submit their documents digitally. As a result, the bank landed [tens of thousands](#) of new online equity trading accounts while rivals floundered.

Growth, not savings, is the truly important part of the ROI equation. Disruptive technologies such as advanced analytics and artificial intelligence, for example, are underpinned by the cloud. In addition to speed, scale, and flexibility, the cloud's ability to manage and process huge amounts of data in real time and host apps that can unlock insights about customers and opportunities will be worth far more than reducing the number of servers on premises.

3

You won't optimize your cloud journey unless you optimize your data layer

The temptation is understandable: Just move your legacy databases with their existing data layers to the cloud and be done with it. What companies often don't realize is that seeming short-term gain can lead to long-term pain.

For example, legacy monolithic applications—and their data—running in your data center may need to be broken up into microservices in order to take full advantage of moving to the cloud. Not taking the data layer into account can create data silos and unnecessary complexity. Legacy applications often rely on one relatively simple data model. But the cloud makes it easy to leverage a multiplicity of data models, and you won't get the full benefits if you're stuck with an obsolete, suboptimal data layer that doesn't address a variety of data models, such as graph, time series, JSON, and more. The [average company requires three different databases](#) to manage the data across their various

hybrid and public clouds, relying on APIs to share data among them.

Most importantly, traditional databases are often too slow and inflexible to keep up with web-scale demands in the cloud. At a time when making real-time prediction decisions has become a delineator of success, storing data at rest for later analysis is a non-starter. Your cloud applications need a fast, highly scalable, highly available database platform to deliver real-time application performance no matter the amount of data or the number of users.

A modern, cloud-native data layer must be able to scale with your application, and your organization, without causing your costs to soar as well. High availability of the data layer is also non-negotiable. So is responsiveness. If your cloud app responds to customers in the blink of an eye, you could be too slow! To satisfy users today, an app has to respond in about 100 milliseconds, including sending the data back and forth, which doesn't leave much time for the database to do its thing. Finally, in a global economy, the data layer must support applications and customers no matter where they're located, with no data loss, degradation of availability, or performance issues.

4

Most apps still aren't in the cloud, and some will never be

A few years ago, people were quick to tick off lots of reasons certain applications couldn't run in the cloud—security and compliance, serious data durability, and strict latency requirements were often considered deal breakers. Now, there are fewer “stories of what doesn't belong in the cloud,” Quinn says.

Yet, only about 20% of applications now run in public and private clouds or edge environments, according to [Accenture](#). Over the next five years, that number will rise to 80%. Given the cloud's ability to scale dynamically, workloads that are highly variable or unpredictable—or where self-provisioning is important—are obvious cloud use cases. But the migration should never be by rote. Deciding if an app belongs in the cloud requires a careful analysis of the application and factors

such as latency requirements, its value to the business, and more.

There are still plenty of examples of apps that shouldn't be in the cloud. High performance computing may be the next “it-can't-run-in-the-cloud” story to fade away, but for now HPC applications can be a stumbling block, at least economically. “Any application with massive data egress charges doesn't belong in the cloud,” Quinn says. “Even moving between two availability zones in the same region comes at a cost of \$.02 per gigabyte, which adds up to a lot of money when transferring multi-petabyte workloads.”

Latency can be a critical determinant of app placement. Data processing at the edge, often based on IoT inputs, is not yet commonplace, but interest is rising in industries as varied as manufacturing, healthcare, financial services, and communications. In those specialized circumstances, once the data is analyzed it is shuttled to the cloud for further analysis in data lakes or cloud-based data warehouses.

5 AI in the cloud is hot—but requires the right architecture

Companies are salivating about artificial intelligence in the cloud. From fraud detection to medical diagnosis to facial recognition, the speed, flexibility, and agility of the cloud combined with AI promises to transform all sorts of applications and industries, providing compelling insights, automating full lines of business, generating new revenue streams, and making people's lives easier.

Cloud-based AI databases must simultaneously ingest, explore, analyze, and visualize fast-moving, complex data within milliseconds. The goal is to lower costs, generate new revenue, and integrate AI models so that businesses can make more efficient, data-driven decisions, act with greater agility, and decrease risk while exploring new opportunities.

But as companies begin to shift from AI and machine-learning (ML) pilot programs to actual production systems, many don't see a potential bottleneck looming on the horizon: the need for an [enterprise feature store](#) and inferencing engine. In this context, "features" are highly curated sets of data that inform predictive models. For example, a fraud model might include a feature based on unusual credit card purchases. A feature store, meanwhile, catalogs and serves up all available features—ready for use by any machine-learning model. Centralizing these features in the cloud speeds development, improves model accuracy, and eases version tracking.

So, when you think about AI in the cloud, don't forget that creating "intelligence" and value takes more than just training models. You also need to consider inference and features. And because many AI use cases are real-time, you need to keep your features and inferencing close to your database in order to minimize application latency.



Moving to the cloud is changing IT jobs—and unleashing creativity

6

As workloads shift to the cloud, changes to the IT organizational chart are following suit. Although virtualization already ushered out many hardware-focused jobs, the cloud is making even more dramatic changes to rack-and-stack roles.

According to the [2020 IDG Cloud Computing Survey](#), 67% of organizations have added new cloud roles and functions. Data center specialists are being replaced with cloud administrators or cloud architects. (If companies can find them. It's not easy to hire "unicorns" with skills in traditional IT, public and private clouds, security, and governance.)

Even if people's job titles remain static, the daily

work they do is changing significantly. Offloading management of infrastructure to a cloud provider reduces the amount of control on-site personnel have over applications. Employees must develop new skills in contracts, vendor management, and helping end-users deal with new processes.

These workforce disruptions won't come without pain. Some jobs will be lost as companies close or de-emphasize their data centers. But there's a big upside, too. As noted earlier, the real opportunity for IT departments comes with offloading mundane chores so key personnel can focus on more strategic tasks. "There is a crying need for all enterprises to cut the umbilical cord of excessive operational maintenance cost," Greenbaum declares. For example, a [Harvard Business Review white paper](#) recounts how after one CIO moved his enterprise ERP system from an on-premises data center to the cloud, he was able to exchange five application administrators for five IT relationship managers who worked with the business units to identify their business requirements.

Another CIO says cloud capabilities changed the mindset of his team. The cloud allowed them to take snapshots of the environment and spin them up for tests or development projects, and then instantly revert back to a version in the cloud if needed. This gave them the "psychological safety" they needed to try new ideas.





Everyone says they have a “multicloud strategy,” but it’s mostly just talk

Clouds are like potato chips: Hardly anyone can stop at just one. Some 93% of enterprises say they have a multicloud strategy, according to the [2020 State of the Cloud Report by Flexera](#). On average, the study reports, enterprises use 2.2 public clouds and 2.2 private clouds.

Using multiple cloud providers can help companies avoid vendor lock-in, leverage best-of-breed features and capabilities, and boost leverage on pricing. Using a second cloud for data backup can be a sensible disaster-recovery strategy for worst-case scenarios. “However, few companies really approach multicloud from a strategic standpoint,” Greenbaum says. “They just end up there.”

This can happen for many reasons. A software vendor may nudge them to a specific cloud. Or a company might make a big acquisition and inherit three different ERP systems scattered in different clouds. Perhaps a rogue subsidiary cuts its own cloud deal. Or maybe it’s a shadow-IT situation where an ambitious development team spins up its own instances without clearing it with corporate.

So rather than implement an actual multicloud strategy, most companies tend to cluster their infrastructure on one cloud, with bits and pieces running on other clouds for a variety of reasons. But while it’s

most common to have different applications running on different clouds, companies are increasingly choosing to distribute a single app across multiple clouds to create globally distributed service that delivers low latency to end users no matter where they are located.

The nirvana of multicloud—when the same app running on multiple clouds can have its data seamlessly traverse clouds, or failing over one from cloud to another—is extremely complex and still fairly rare. To make it work, companies need to coordinate the connections among the various cloud services and platforms and any on-premises resources. Notably, moving stored data from one cloud to another for processing can wreak havoc on application latency rates and network bandwidth.

There are other multicloud issues as well. “You are at the whim of both the software vendors’ upgrade, update, and API-management strategy, as well as the maintenance and support strategy of the cloud vendors themselves,” Greenbaum warns. “There is a crying need for a single pane of glass to manage multiple cloud environments, and that doesn’t exist right now. You have to put in a lot of sweat equity to manage all those moving parts.”

If you have a cloud security issue, it's probably your own fault

8

Admit it: Putting data in the cloud is scary. A recent [Bitglass survey of IT professionals](#) found that almost three quarters (73%) were extremely or very concerned about security in the public cloud. Another 20% were moderately concerned. That sounds ominous—until you realize that three quarters of those same IT professionals also acknowledged that cloud-based apps and software were just as secure as or even more secure than on-premises systems.

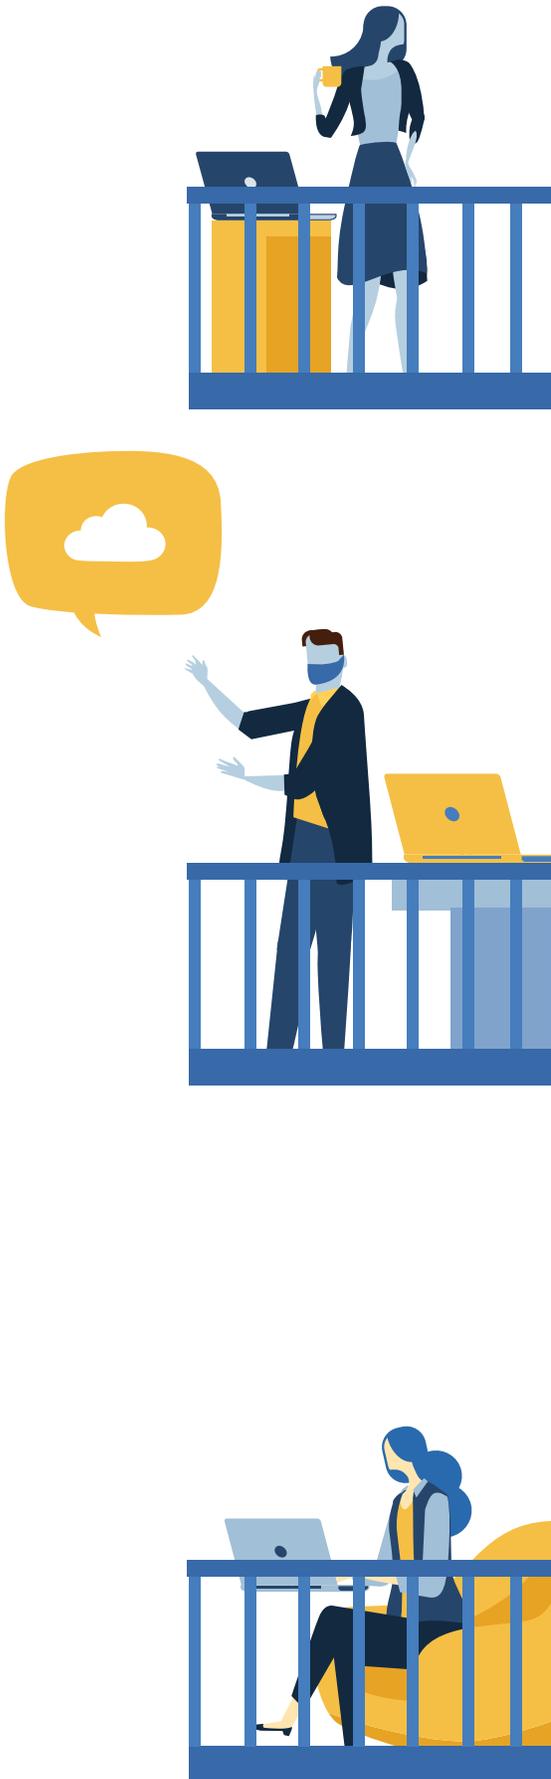
The real point? Putting data *anywhere* is scary for IT professionals.

The truth is the cloud is more secure than on-premises data centers. Unlike many enterprises, cloud providers have no choice but to spend heavily on the latest versions of security apps and services, and are usually two-to-three generations ahead of most enterprises. For example, few on-premises systems encrypt their data in transit and at rest as all the major cloud providers now do.

Enterprises are slowly but surely catching on to this complex reality. (In the Bitglass survey, 28% said clouds were more secure, while 27% said they were less secure.) As the CIO of a large consumer goods company put it in a [Harvard Business Review white paper](#), “This is not our grandfather’s cloud where it was completely vulnerable from a security standpoint.”

When cloud security issues do occur they are almost always due to customer error—misconfigured implementations, poor user management, and other user mistakes that the cloud provider does not control. The real question, then, is not whether the cloud is secure, but whether your company is using the cloud in a secure manner by implementing proper procedures and controls. Because the cloud is only as safe as you make it.





9

The COVID cloud rush will have big but unexpected repercussions

As noted at the beginning of this story, the COVID-19 pandemic caused lots of companies to expand their use of the cloud dramatically as they sought to support and secure a newly remote workforce, ensure resiliency, and optimize IT costs. In a telling example, Deloitte Consulting's Chief Cloud Analyst David Linthicum recalls clients whose servers went down during quarantine, and no one was allowed into their buildings to restore them. "People realized that having data in the cloud was much safer than having it on premises," he says.

A hallmark of the pandemic spending spree is the cloud's surging share of the enterprise pie. Driven by greater needs for cost efficiency and business continuity, Sid Nag, Gartner Vice President, predicts [cloud spending will hit 14.2% of the total global enterprise budget in 2024](#), up from 9.1% in 2020. With remote work likely to remain widespread even after the pandemic eases, many companies are looking for modern cloud-native applications to help them do their jobs.

However, there might be an unexpected hitch—in the rush to cope with COVID-19, companies didn't always think things through strategically. The emergency expansion of cloud usage also resulted in lots of unanticipated complexity.

With little time for coordination between cloud migration and cloud development teams during the pandemic, some companies now face a lot of work fixing mistakes they may not even realize they made. So even when the pandemic eventually eases, enterprises will likely continue to boost cloud spending to address issues stemming from rushed cloud adoption as well as to drive ongoing business and operational transformation.

THE

If you enjoyed reading this story, explore more perspectives for designing a future-proof cloud strategy in the full magazine.

READ

Rediscover Magazine: The Cloud Edition

NOW AT

redislabs.com/rediscover-magazine/



redislabs
HOME OF REDIS

9 Powerful Cloud Truths

From AI to security to multicloud, the cloud secrets no one talks about

Q&A with Kelsey Hightower

Google's cloud guru on Kubernetes as a design principle, open source as a service, and much more

Searching for Greatness in the Cloud

How the cloud helps us rediscover winning products, companies, and technologies