McKinsey Digital

Debunking seven common myths about cloud

Common misconceptions about cloud are holding companies back from capturing the full benefits available.

by Mark Gu, James Isenberg, Leandro Santos, and Isabelle Tamburro

October 2020

The need for superior speed and agility continues to push companies toward cloud adoption. But while historic predictions anticipated that upwards of 16 percent of enterprise workloads would be in cloud—infrastructure as a service (IaaS)—by 2019, there is a clear lag in 2019's actual figure, which is half as large, at less than 9 percent.[1]

For the most part, this delay in cloud adoption does not stem from a lack of ambition. Many company leaders have encountered major roadblocks along their path toward cloud or have gotten cold feet once they questioned its impact on costs, security, latency, and more.

Conversations with hundreds of CEOs and CIOs have revealed a consistent set of myths that lead to these roadblocks and questions, hampering progress and adoption. Companies that have effectively counteracted these myths are the ones that have derived the greatest rewards from their move to cloud.

Cost and value

Myth #1: The main value of cloud business cases is IT cost reductions.

The common industry introduction to cloud refers to the replacement of key IT activities, access to on-demand infrastructure, provisioned compute, storage, database services, and more. While all these descriptors are accurate, organization leaders often hear them and lose sight of the broader impact cloud can have on transforming the full IT operating model and, most importantly, on the business. Consequently, when they set out to write a business case, they spend months analyzing on-premises costs compared with cloud costs and focus much less time on the main value driver of cloud: the business benefits.

The reality is that the aggregation of business benefits can swamp IT cost efficiencies in cloud. For example, application-hosting spend at large companies often represents only a fraction of total revenues, perhaps 0.5 percent. Even if cloud could reduce this spend by 25 percent, this would be only "small potatoes" compared with the broader potential business impacts from cloud.

Any one of a number of cloud-enabled initiatives—improved analytics, faster time to market, stronger innovation—could generate a greater incremental contribution than IT cost reductions. Cloud can improve almost every aspect of an organization's products, services, or processes. Superior computing power can lead to a greater understanding of customer needs, for example, while extra processing capacity can be used to run more complex analytics or to create differentiated business insights. Innovation is quicker and less risky because experimentation and testing of new ideas cost less and take less time. All this drives revenue growth opportunities in a variety of ways, including acceleration of new-product lead time, entry into new markets, and response to competitive threats.

A health-insurance carrier moving to cloud, for example, drew out several billion dollars of additional revenue by accelerating multiyear projects into just months through superior agility and computing power. It found particular benefit in moving its provider-facing apps, accelerating the onboarding and revenue capture of new healthcare providers. A large financial-information provider found that by moving to cloud, it could enter and set up technology operations in new countries within weeks rather than months. This speed and flexibility gave the provider a first-mover advantage in markets at a fraction of the cost it had historically spent in new locations.[2] The common thread in these examples, and many more, is that the ultimate reason to move to cloud should be the business benefits rather than IT efficiencies.

Myth #2: Cloud computing costs more than in-house computing.

Cloud economics is one of the most contentious current questions in enterprise IT. The reality is complicated, as cost is highly dependent on

[1] See Kaitlin Buckley, "By 2019, 60% of IT workloads will run in the cloud," 451 Research, September 5, 2017, 451research.com; and William Fellows, *451 perspective: The cloud feast heralds the era of consumption, part 1*, 451 Research, August 12, 2019, 451research.com.

[2] Chhavi Arora, Tanguy Catlin, Will Forrest, James Kaplan, and Lars Vinter, "Three actions CEOs can take to get value from cloud computing," July 2020, McKinsey.com.

a company's starting point—and its ability to govern and optimize cloud consumption once there. For example, one financial institution runs on expensive proprietary UNIX systems at about $25,000 to $35,000 per operating-system instance (OSI). It anticipates up to 75 percent in savings from cloud adoption. In the next five years, by migrating 50 percent of its workloads to cloud, it expects to lower unit costs to $15,000 to $22,000 per OSI. On the other hand, a large insurance company found that through a combination of re-tiering and sourcing, it was able to improve unit-cost economics in its private environment, making a migration to cloud less attractive.

Other starting-point differences we see are companies' maturity in on-premises life cycle, license commitments, and types of workloads. Companies facing large data-center upgrades, for example, will find cloud adoption attractive as a way of avoiding large capital expenditures on assets they may not fully utilize for years and that risk being deprecated faster than in the past. For companies that may have recently invested in a new data center, however, moving to cloud would duplicate some infrastructure costs. Another key difference is between companies with expensive license agreements that are hard to get out of and companies with limited penalties for transitioning. Finally, storage-intensive workloads are often less costly in cloud than those requiring lots of network bandwidth, as cloud service providers (CSPs) charge by the unit for network access.

Starting point aside, many companies moving to cloud have experienced cost benefits from cloud's shared-resource model and autoscaling. Rather than owning a cluster on-premises and paying for around-the-clock access, companies pay CSPs for CPU as they need it. Where the shared-resource model does not translate into total-cost-of-ownership (TCO) savings, it is often because companies lack correct resource governance, or they migrate applications designed to run internally without adjusting their resource consumption models. Such applications won't fully leverage the benefits of autoscaling and are more costly to administer and maintain than cloud-native applications. Therefore, to keep running costs low and maximize benefits, companies should assess their applications' architectures, remediate their portfolio where needed, and establish new transparency and governance processes.

The core question for cloud economics is whether the reduced run-rate cost on cloud justifies the up-front cost of remediation, assuming all configuration and governance are done correctly. Even in the cases where companies' starting point makes remediation too cost prohibitive, the business benefits explored in myth #1 often are a stronger reason for the transition to cloud and outweigh the short-term IT cost hurdles altogether.

Myth #3: The security I can set up and control in my own data centers is superior to the security on cloud.

Historically, executives have cited security of public cloud infrastructure as one of their top concerns and a barrier to cloud adoption.[3] In recent years, however, all major CSPs have made significant investments in their underlying security capabilities. A CSP's business model depends on best-in-class security, and they have each invested billions in cloud security and in hiring thousands of the top cyber experts. They have developed an array of new tools and methods to make cloud secure, in many cases requiring developers to take on the security responsibility, rather than relying on a traditional security team to carry the burden. This is particularly important because public cloud breaches have almost all been driven by enterprise customers' insecure configurations. Gartner, in fact, predicts that, through 2025, 99 percent of cloud security failures will be the customer's fault, not the security provider's.[4]

Developers, therefore, must be retrained to follow carefully defined governance and policies on how to configure the right security controls. For example, if it is policy that data must be encrypted, it is up

[3] Nagendra Bommadevara, James Kaplan, and Irina Starikova, "Leaders and laggards in enterprise cloud infrastructure adoption," October 2016, McKinsey.com.
[4] Kasey Panetta, "Is the cloud secure?" Gartner, October 10, 2019, gartner.com.

to the developers to invoke the correct application programming interface (API), telling the CSP they want data in a given storage bucket to be encrypted.

For these new policies to be successful, cloud requires companies to adopt a DevSecOps operating model, where security is a key element of every software project.[5] IT organizations should automate security services across the full development cycle and make them available using APIs or risk vulnerable configurations. More than one large financial institution has had to put its public cloud program on hold due to poor operating-model and configuration decisions. These institutions are now backtracking to invest in automated security controls for future applications, having discovered, like many other organizations, that they can no longer rely on manual security controls and traditional operating models if they want to transition successfully to cloud.

The key question for companies, therefore, is not whether cloud is more secure to begin with, but what measures they need to take themselves to enhance their cloud security. Companies that define the correct policies, adopt a secure DevSecOps operating model, and train or hire the right talent can actually achieve safer operations in their cloud environments than on-premises.

Technical implications

Myth #4: There is greater latency among applications running on cloud providers' networks than there is on in-house networks. Some organizational leaders fear that when they transition to cloud, they will experience higher latency on a CSP's network than on their own. Latency, however, is often the result of the IT department attempting to backhaul its data through in-house data centers. Backhauling, or routing traffic through internal networks, creates higher latency, extra complexity, and poor user experience. IT departments that choose to backhaul usually either lack experience or trust with cloud security (believing they will have greater control by backhauling) or need to access critical data or apps that are in on-premises data centers.

It is important for IT departments that are backhauling for increased security to realize that CSPs now offer strong perimeter options and that there is no need to tolerate latency for security. While backhauling was the most popular model for perimeter security in 2018, companies are now adopting alternative methods, most popularly cleansheeting, or designing a "virtual perimeter" with cloud-specific controls. Indeed, in a McKinsey IT security survey, only 11 percent of cloud users said they are likely to be using a backhauling model by the end of 2021.[6] IT departments that are backhauling for critical data or apps should prioritize creating a data lake with their CSP and move the bulk of their data and analytics processing to cloud and use data replication only where absolutely needed. This will allow them to unleash the power of cloud-enabled analytics while simultaneously solving any latency issues.

Once companies stop backhauling their data, they are unlikely to experience greater latency on cloud, as there is no inherent difference between a CSP's IP circuits, pipes, and cables and their own data center's. In fact, companies may even experience lower latency in cloud, due to CSPs' advantages in content delivery. With their diverse, global footprint of data centers and their heavy investment in content-delivery-network services, CSPs can provide content at optimal speed, depending on location, content request, and server availability, on a scale that companies would be hard-pressed to achieve in-house. Given both the advantage CSPs have in content delivery and the shift away from backhauling, companies should not fear high latency during their move to cloud.

Myth #5: Moving to cloud eliminates the need for an infrastructure organization. The idea of infrastructure as a service (IaaS)—that an external provider will manage your underlying network, hardware, and resources—is an exciting

[5] Santiago Comella-Dorda, James Kaplan, Ling Lau, and Nick McNamara, "Agile, reliable, secure, compliant IT: Fulfilling the promise of DevSecOps," May 2020, McKinsey.com.
[6] Arul Elumalai, James Kaplan, Mike Newborn, and Roger Roberts, "Making a secure transition to the public cloud," January 2018, McKinsey.com.

proposition for many organizational leaders. The misconception arises, however, when leaders interpret IaaS as a full replacement for their infrastructure organization. While cloud radically changes the activities, talent, and operating model required in an internal infrastructure group, and beyond it, it does not altogether replace the need for infrastructure management.

When companies transition to cloud, they will encounter hundreds of services that can be combined and configured to affect performance, security, resiliency, and more. They need an infrastructure team that can build and manage standard templates, architectures, and services for use by their development teams. As infrastructure in cloud is managed through code, this infrastructure team will require different skill sets (for example, committing code) so they can operate much like an app-development team. Without this infrastructure team creating standardized services and platforms, many enterprises will simply replicate the fragmentation and chaos they experienced on-premises.

To accommodate this shift in function, infrastructure organizations must transition to a proactive (rather than reactive) operating model. Instead of responding to bespoke requests from development teams, which take months and can quickly become costly, cloud infrastructure teams should proactively consider organizational needs and turn this into a reliable, automated platform on cloud. In doing so, the ownership lands more squarely on development teams themselves, who have more flexibility in quickly configuring the resources they need. Not only will application teams gain more direct responsibility over costs, but this increased flexibility will lead to greater productivity and faster speed as well.

Shifts in infrastructure are not only helpful in managing cloud but also necessary in order to see the full range of cloud benefits. A large entertainment company saw that when it shifted to a cloud-compatible operating model, its infrastructure team could deploy to production on demand, support a larger infrastructure footprint with leaner teams, and improve time to market, going live in six new locations in record time.

In general, traditional infrastructure teams running cloud would be too large and too costly and would miss the benefits of having app teams own shared responsibility for the run costs they incur. On the other hand, having no infrastructure team at all would wreak havoc on an organization's ability to manage and benefit from cloud. Instead, a leaner, more specialized infrastructure organization is required to achieve the full range of agility, innovation, and performance benefits of cloud.

The transition

Myth #6: The most effective way to transition to cloud is to focus either on applications or on entire data centers.
It is a common misconception that an organization must opt for one of these two extremes to transition successfully to cloud.

In the application-by-application approach, organizations face unattractive scale dynamics. They will continue to pay for on-premises data centers and IT support, while simultaneously paying CSPs for hosting a subset of applications. Moving a subset of applications also does not lead to business benefits if those applications constitute only part of a business domain's portfolio. For example, if a business moves a set of applications within the customer-onboarding domain to cloud, but leaves behind the application that generates and stores user profiles, the time-to-market benefits of cloud cannot be fully realized. On the other hand, organizations that move an entire data center to cloud may face substantial up-front investment and risk. Many of the hundreds of applications in a data center probably were not designed to run in cloud. Companies will need to invest in various forms of remediation, which can become expensive and risky when executed all at once.

Instead, organizations should look to move *business domains* to cloud (such as customer onboarding, early-stage drug discovery, consumer payments). By transitioning the business domains, companies

will experience the full range of potential cloud benefits: faster time to market, greater agility, stronger reliability, and more. In addition to the business benefits, moving a business domain is a much smaller lift than moving an entire data center, meaning that cost and risk will be more manageable. Once one business domain begins to experience these improvements in time to market, agility, and reliability, it will be easier to make the business cases for the remaining domains.

Myth #7: To move to cloud, you must either lift and shift applications as they are today or refactor them entirely.

When companies make the commitment to move to cloud, they often face pressure to move fast, minimize costs, and maximize business benefits. As a result, leaders feel they must choose between a quicker and cheaper "lift and shift" transition strategy (to move fast and minimize costs) and a time-intensive and costly refactoring strategy (to capture business benefits).

While lift and shift—virtualizing the application and dropping it into cloud as is—can be a faster and more cost-effective way to move many applications into cloud at once, it fails to harness the majority of cloud's benefits. That's because there is no change to the application's architecture, which is often not optimized for cloud and so won't benefit from features like autoscaling, automated performance management, and more. Furthermore, the non-native application will likely face higher latency or other performance issues, and its preexisting problems will now simply sit in a CSP's data center rather than the company's.

On the other hand, a complete refactoring of the application and its architecture to optimize for cloud takes a lot of time, skill, and money. It achieves the benefits that lift and shift ignores, but so slowly and at such great cost that breakeven is often impossible. It also puts the transition at greater risk of error during complex recoding, configuration, and integration.

Many companies find they are better off using a "best of both worlds" strategy that takes advantage of specific techniques such as automation, abstraction, and containerization. These techniques are less costly and time-consuming than full refactorization but still allow companies to achieve the business benefits of greater agility, faster time to market, and enhanced resiliency. One pharma company, for example, is blueprinting the deployment of its applications and leveraging a CSP's continuous integration and delivery (CI/CD) pipeline. This will allow developers to run their development and testing environments only as needed and will largely automate the otherwise lengthy release management process. A global electronics OEM migrated its e-commerce application to cloud as is but rearchitected the infrastructure allocation algorithm to scale up during peak seasons, taking advantage of cloud's dynamic allocation. Both of these approaches were less costly than refactorization but still allowed the companies to benefit from agility and additional techniques that lift and shift would have ignored.

Many of today's beliefs about cloud are based on misconceptions fed by stories of adoptions gone wrong or fears of significant change. These beliefs get in the way of deeply understanding the positive business, operational, and economic impacts of cloud and must be addressed to enable organizations to capture cloud's full value.

Mark Gu is an associate partner in McKinsey's New York office, **Rich Isenberg** and **Leandro Santos** are partners in the Atlanta office, and **Isabelle Tamburro** is a consultant in the Chicago office.

The authors wish to thank Chhavi Arora, Thomas Carr, Will Forrest, Steve Jansen, James Kaplan, and Arvind Ramachandran for their contributions to this article.