Creating value with the cloud

Digital McKinsey: Insights

December 2018

Digital/McKinsey

The articles in *Digital McKinsey: Insights* are written by consultants from Digital McKinsey together with colleagues across the firm.

The publication offers readers insights on digital transformations and the people, processes, and technologies that are critical to their success.

Articles appearing in *Digital McKinsey: Insights* also appear on McKinsey.com. If you would like to receive email alerts when new digital articles are posted, register at McKinsey.com.

To learn more about Digital McKinsey, please visit mckinsey.com/business-functions/digital-mckinsey/our-insights. To send comments or request copies, email us: digital_mckinsey_insights@mckinsey.com.

Director of Digital McKinsey Publishing: Barr Seitz

Editor: Josh Rosenfield

Managing Editors:
Michael T. Borruso, Venetia Simcock

Art Direction and Design:
Nicole Esquerre, Julie Schwade

Data Visualization:
Richard Johnson, Jonathon Rivait

Editorial Production:
Elizabeth Brown, Heather Byer, Roger Draper, Gwyn Herbein, Pamela Norton, Katya Petriwsky, Charmaine Rice, John C. Sanchez, Dana Sand, Katie Turner, Sneha Vats, Pooja Yadav, Belinda Yu

Cover Photo:
© Erik Isakson/Getty Images

McKinsey Practice Publications

Editor in Chief: Lucia Rahilly

Executive Editors:
Michael T. Borruso, Allan Gold, Bill Javetski, Mark Staples

Table of contents

Creating value with the cloud

Once a technological curiosity, the cloud has become integral to modernizing the IT environment and enabling the digital transformation of companies large and small. Cloud-based computing and storage platforms offer manifold advantages over conventional on-premise systems, from lower operating costs to better compatibility with the working styles of digital enterprises. But a large-scale move to the cloud isn't a matter of merely "lifting and shifting" applications and data from on-premises services to cloud platforms. It's a complex endeavor that requires companies to build new capabilities.

One often-overlooked capability is planning the cloud transition. IT leaders need to weigh the pros and cons of migrating each application or data asset. This often requires extensive dialogue with both cloud-services providers and software vendors so that companies can understand how their offerings are likely to evolve. Another key area of focus is managing cybersecurity during and after the transition. Companies should take stock of cloud-service providers' security resources and determine how to adapt their own cybersecurity practices to balance speed and protection.

Perhaps most important, companies will need to reorganize their operations so they can take full advantage of what the cloud can do. Some companies might choose to establish dedicated cloud-migration teams to set up cloud platforms and remediate applications or data assets so they function properly in the cloud. Others will entrust the migration work to existing teams. Either way, all IT specialists, from application developers to infrastructure teams, will have to learn the effective use of cloud-based services. Such a learning program should cover technical skills as well as agile methods, which enable teams to build and deploy cloud applications quickly.

Being smart about the use of cloud platforms and services can make the difference between gaining a competitive edge and falling behind rivals. With this volume, we hope to help you capture the value that the cloud can unlock.

Andrea Del Miglio
Partner, Milan

Will Forrest
Senior partner, Chicago

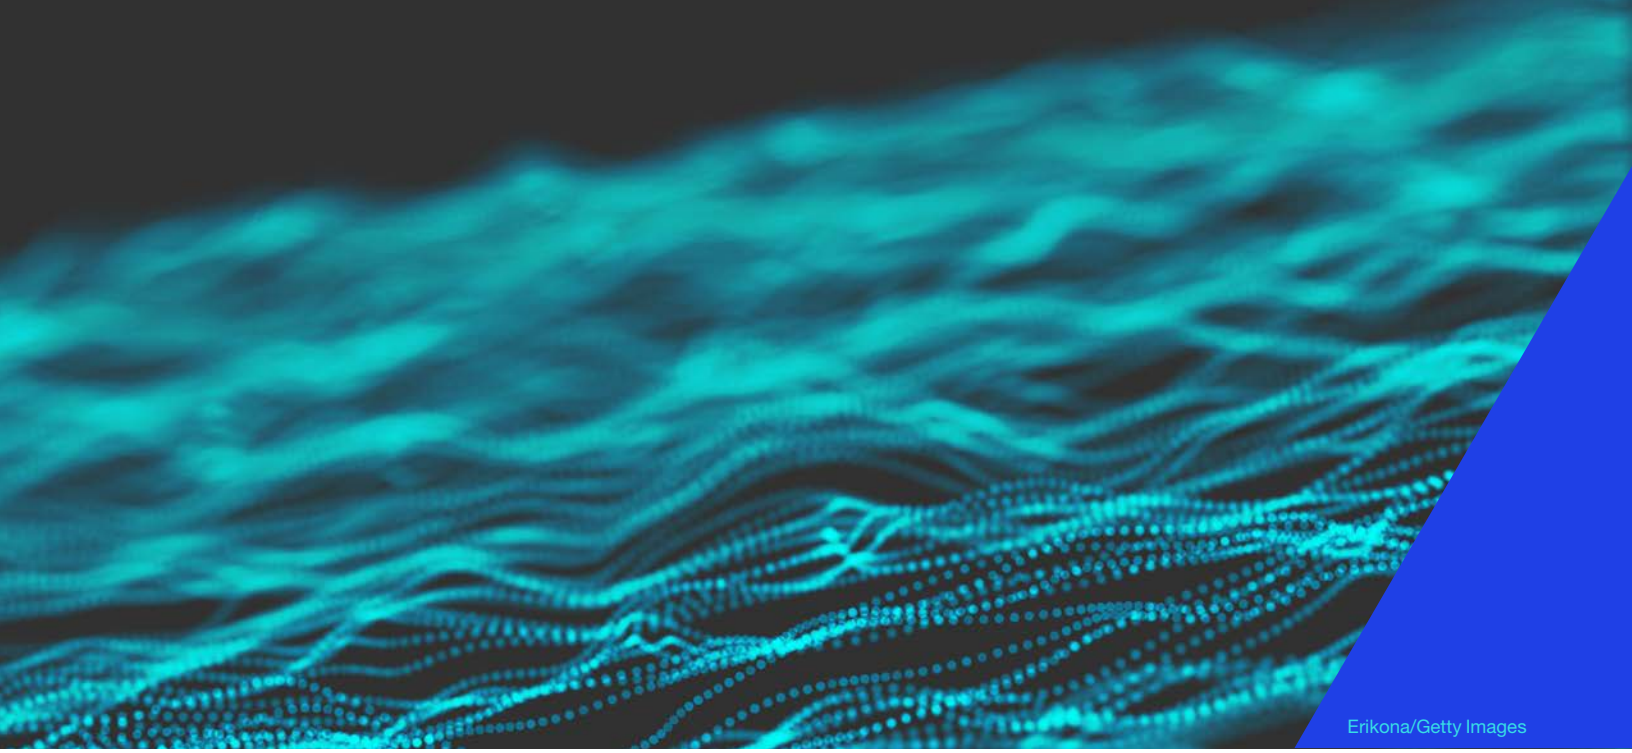

Erikona/Getty Images

The progressive cloud: A new approach to migration

Mark Gu, Krish Krishnakanthan, Anand Mohanrangan, and Brent Smolinski

Migrating applications and data to public-cloud platforms can be tricky. Companies can ease the transition with hybrid-cloud configurations that progressively combine private- and public-cloud features.

Moving processing workloads into the public cloud has helped leading companies lower their operating costs and build modern IT environments capable of rapid, integrated, and highly automated development and operations. But for large companies with complex IT architectures, moving applications and data to public-cloud platforms involves working through a formidable set of technology, security, operational, and financial issues. Those complications go a long way toward explaining the limited uptake of public-cloud platforms: some 60 percent of companies surveyed by McKinsey last year have migrated less than 10 percent of their workloads to the public cloud.

There are, however, ways to ease the transition to the public cloud. By progressively blending public-cloud and private-cloud solutions into hybrid-cloud configurations, companies can quickly take advantage of sophisticated cloud services and even move sensitive applications into the public cloud without disrupting their IT architectures and operations. Three practices are essential to implementing progressive cloud

models. Companies must first estimate the costs of operating a hybrid configuration. Next, they should devise a manageable sequence in which to migrate applications and storage to the cloud. With those priorities in mind, they should set up a dedicated unit to migrate applications and storage using agile practices and streamline operations with automated services. In this article, we provide a closer look at these three practices and how leading companies have used them to accelerate the movement of their workloads into the public cloud.

The best of two worlds: The progressive cloud

Cloud platforms come in two main varieties, public and private, both of which have pros and cons. Public-cloud platforms give companies easy access to a broad range of services, from basic storage and networking to innovative offerings like advanced analytics, machine learning, and virtual-reality development. And their menus of services expand all the time. Enterprises can easily take advantage of these cutting-edge services without having to develop their own or source them from other vendors. However, enterprises can be apprehensive about placing sensitive information and proprietary applications in the shared data centers that power public-cloud platforms.

Private-cloud platforms can be equipped with some of the same automation features as public-cloud platforms (for example, one-click provisioning of servers and automated scripting of architecture patterns), so companies can rapidly deploy new capabilities. Companies can also outfit private-cloud platforms with security controls of their choosing and thereby protect their critical applications and data. On the other hand, public-cloud platforms have more capabilities than private-cloud platforms: cloud-service providers (CSPs) invest heavily in developing new services, and third-party vendors tend to launch new services in the public cloud before introducing private-cloud versions.

To work around these trade-offs and bring public-cloud capabilities together with private-cloud security, companies can take a progressive approach to combining private-cloud and public-cloud services. Such hybrid-cloud systems come in three primary variants (Exhibit 1):

- *A private-front or backhauling topology* routes all traffic through private data centers and deploys applications partly or completely in the public cloud so that a company can apply internal cybersecurity controls and still take advantage of public-cloud services.

- *A public-front topology* also places applications in the public cloud but allows users to access them directly, with CSP-provided cybersecurity controls applied by default. Data are stored in a private cloud with additional security controls.

- *A public-cloud or cleansheet topology* places both applications and data in the public cloud. Enterprises apply cybersecurity controls from third-party services.

As companies develop more sophisticated cybersecurity controls and cloud capabilities, they can shift applications from a private cloud into a hybrid cloud with a private-front topology, then into a public-front topology, and eventually into a cleansheet topology. For example, an insurance company used a private-front topology to move some sensitive applications into the public cloud without having to overhaul its cybersecurity controls. Doing this allowed the company to migrate an additional 25 percent of its workloads into the public cloud, where it could use additional services while maintaining security controls.

Three essential practices for deploying progressive cloud systems

Since progressive cloud systems rely on some elements of public-cloud platforms, businesses

EXHIBIT 1

Progressive cloud systems come in three primary variants.

Private front or backhauling

A private-front topology routes all traffic through private data centers and deploys applications partly or completely in the public cloud.

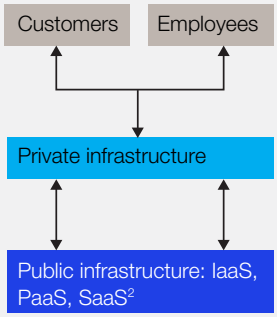

- Known and established security mechanisms
- Simplified monitoring and debugging
- Quick implementation
- Higher costs because of increased traffic

Public front or CSP default[1]

A public-front topology places applications in the public cloud but allows users to access them directly. Data are stored in a private cloud with additional security controls.

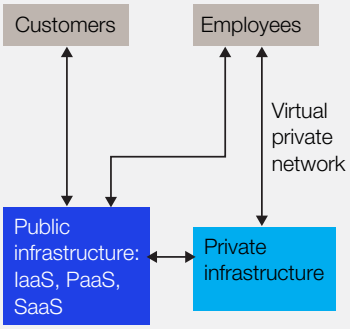

- Lowest-cost approach, but limited to offerings from CSPs
- Potential creation of gaps when limitations are not understood
- Greater scalability

Public cloud or cleansheeting

A public-cloud or cleansheet topology places both applications and data in the public cloud. Enterprises apply cyber-security controls from third-party services.

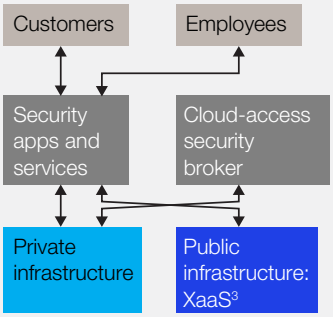

- Use of multiple solutions
- Enhanced user experience[4]
- Need for deep expertise in cybersecurity and cloud architecture; increased complexity and potentially IT costs
- High potential benefits (45–60% savings on data-center costs)

[1]Refers to the use of cloud-service-provider (CSP) security controls by default.
[2]IaaS = infrastructure as a service; PaaS = platform as a service; SaaS = software as a service.
[3]XaaS = everything as a service.
[4]For example, multiple device platforms with a single sign-on.

that opt for these hybrid setups will still need to manage some of the complexity that the public cloud presents. In our experience, three issues—finance, operations, and talent—typically warrant extra attention and can be managed effectively using the following practices.

Know the costs of progressive configurations

When devising a progressive hybrid-cloud configuration, most companies will want to consider more than innovative capabilities and security controls. Costs also matter. But comparing the costs of progressive options isn't always straightforward. Prices and pricing models for public-cloud platforms change over time. Companies have to keep an eye on those changes and assess their effects.

The characteristics of individual applications and data-storage systems affect technology costs too. A large financial-services company found that "input/output-intensive" (or "I/O-intensive") applications—those that read or write a lot of data—were costly to host in the public cloud because the CSP charged hefty "egress" fees whenever web applications made data calls to the company's private data center. Storage was cheaper in the public cloud, though,

so it was economical to run storage-intensive applications there. (In some cases, companies have to copy data into the cloud, rather than move them there, which expands their storage footprint and inflates their storage costs.) Hybrid-cloud systems may require investments in bandwidth and controls for the connection between private-cloud and public-cloud platforms.

Companies should also consider how migrating to the cloud will affect day-to-day expenses other than technology costs. For example, using an infrastructure-as-a-service capability in the public cloud still requires a company to perform many of the same maintenance activities that it would for a private infrastructure. But when enterprises use cloud solutions that sit higher in the stack, such as platform as a service and software as a service, they can pare down their IT operations and let CSPs handle operating responsibilities.

With all these costs in play, companies should take care to define the financial gains they want to achieve and the metrics they'll use to gauge performance. They can also benefit from assigning experts in cloud technology and pricing to model the costs of their technology stacks and operating models and to recommend adjustments as business needs and pricing schemes evolve (Exhibit 2). Organizations must not approach this as a one-time effort but rather as an ongoing business discipline like the procurement of other integral resources.

Develop a cloud-migration road map

It isn't practical to migrate all applications and data to the cloud at the same time. Companies need to sequence their migration efforts, ideally front-loading them with applications for which cloud migration can deliver big performance improvements or cost savings. One effective approach is to establish a rubric for assessing applications with respect to the performance and cost considerations described above and then engage colleagues from the business, from application development, and from IT infrastructure in conducting the assessments and scoring applications accordingly. The following issues are worthwhile to explore:

- dependencies on other applications
- security controls required by the application
- services consumed by the application
- data required by the application
- the underlying technology architecture
- the effort required to rewrite code and configurations and conduct testing
- the costs of cloud-deployment options
- the business risks of performing a migration

An insurance company used topics like these to evaluate the prospect of migrating its applications and data assets to the cloud. Then it developed a road map that called for migrating an additional 10 percent of applications in the first year, another 20 percent in the following year, and the remainder over the next few years (Exhibit 3).

Create an agile, automation-oriented cloud unit

Some companies think of the cloud as an infrastructure service, and so they ask their existing infrastructure teams to operate cloud services alongside legacy services. Such assignments often prove complicated. Established infrastructure teams can be slow to get familiar with new technologies, legacy systems seldom support the requirements of cloud solutions, and the additional work can exceed the infrastructure team's capacity. When one large enterprise put its infrastructure team in charge of cloud services, the team struggled

EXHIBIT 2

To get maximum benefit from cloud solutions, companies need to understand what cost savings are available with different cloud topologies.

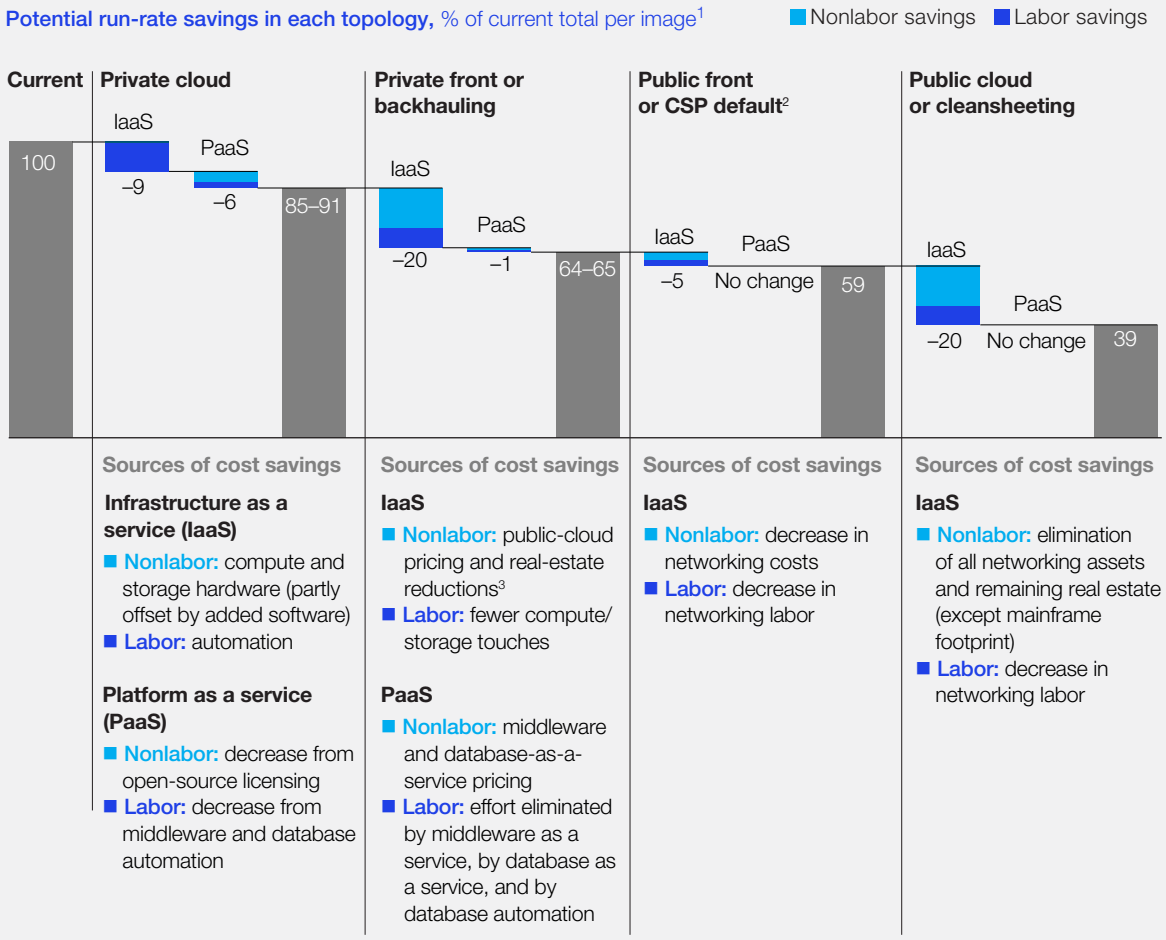

[1] 100% = $20,000 per environment ($246 million infrastructure baseline; 12,500 images). Saving scenarios range from low (only using IaaS solutions) to high (using IaaS and open-source PaaS solutions as well as optimized real estate).
[2] Refers to the use of cloud-service provider (CSP) security controls by default.
[3] Colocating data centers would make real-estate costs variable and allow the company to maintain a physical footprint for networking (~10%) and mainframes (~14%) only.

to learn new support processes, balance the added responsibilities with existing ones, and attract skilled cloud engineers. Its cloud program progressed slowly, migrating less than 10 percent of workloads to the cloud in three years.

A dedicated cloud-delivery team, on the other hand, can help ensure that the migration effort gets proper attention and expertise. This team is tasked with two main sets of responsibilities. One set covers designing, building, and maintaining the cloud

EXHIBIT 3

One company sought to double the share of workloads that it deploys on public-cloud and private-cloud platforms over a three-year period.

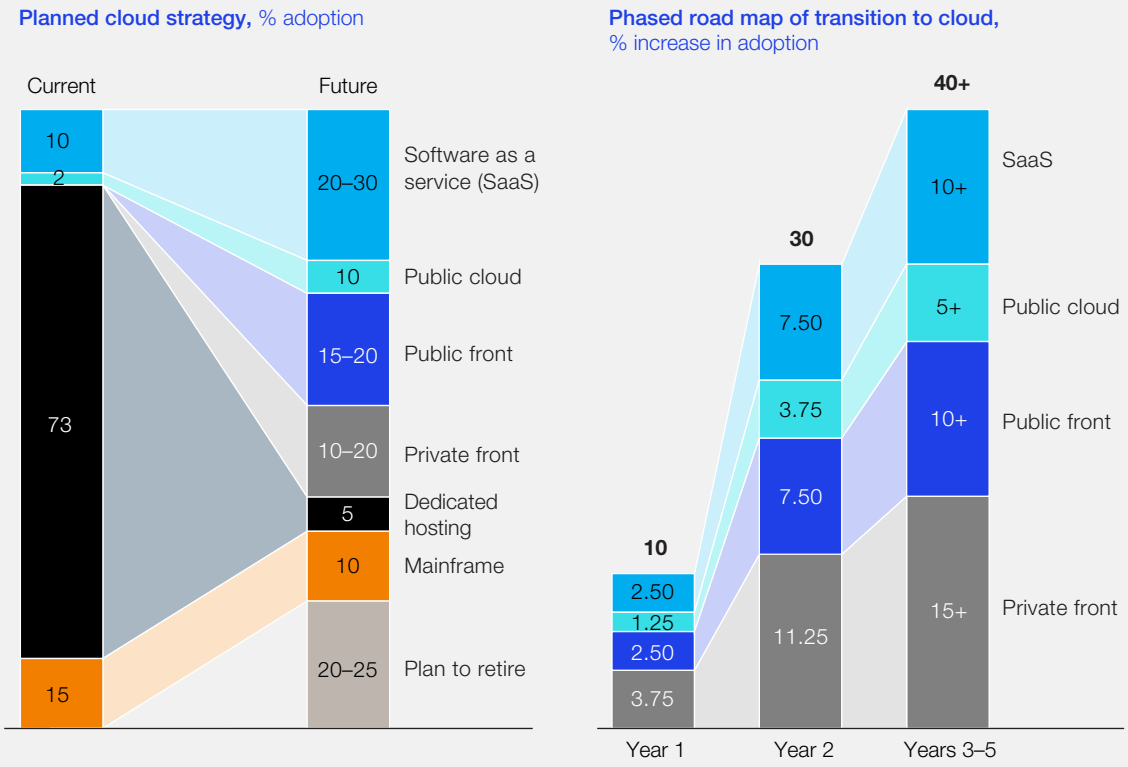

platform and training developers to use it. The other set of responsibilities covers the technical work of migrating applications, such as managing firewall and network settings, testing, writing code, and designing database structures.

A dedicated cloud team can be modestly sized to begin with: 30 to 40 people with a mix of skills in product management, system engineering, software development, user-interface or user-experience design, IT operations, and financial management. Most large companies will have ten to 15 people developing the cloud platform while the rest concentrate on migration work (typically for a period of about two years). In a recent survey, we found that companies with a dedicated cloud team migrated 52 percent of applications on average (from a minimum of 20 percent to a high of 95 percent), whereas companies without a dedicated cloud team migrated 29 percent of applications on average (between 8 percent and 55 percent).

Dedicated cloud teams tend to be most productive when they automate their work and adhere to agile development practices. Cloud teams can write scripts that perform virtually every task involved in operating cloud platforms (see sidebar, "How cloud teams can apply DevOps successfully"). They

can also build tools and application programming interfaces that let software developers deploy cloud services on their own. Cloud-delivery teams that follow these approaches write more code than conventional system-administration teams, so they find it advantageous to follow agile methods. They organize themselves into squads, prioritize service-development efforts by speaking with application developers and other cloud-service users, and roll out new offerings by developing them rapidly and making frequent improvements in response to users' feedback.

An automation-heavy approach typically results in higher productivity: one company found that its cloud team supported some 400 images per full-time employee, compared with the 80 images per employee in its traditional operations group. And as more applications get moved into the cloud, the cloud team's head count can increase, and the traditional infrastructure team can be scaled back. In this way, dedicated cloud cells can eventually replace traditional infrastructure functions.

One financial-services company's cloud team chose to let application developers and systems engineers contribute to the cloud platform's code base. Within two years, more than half of the application teams had voluntarily moved their applications to the cloud, and the remainder were eager to follow suit once essential capabilities were established.

◆◆◆

Cloud services can make IT organizations leaner and more nimble while giving companies access to innovative capabilities that will power their digital transformation. Migrating to public-cloud platforms

How cloud teams can apply DevOps successfully

Cloud teams can be tempted to move their development and testing work into the public cloud in order to save money by shutting down those activities for long periods while production takes place in private data centers. This works well in traditional IT delivery models with lengthy application-release cycles involving extensive manual effort.

But in a DevOps model, where virtually every activity in an application-release cycle is automated, moving development and testing into a public-cloud environment while production stays in the private cloud can cause trouble. Writing automation code that spans two or more environments can be more complex than writing automation code for a single environment, because the different environments might rely on different tools or protocols.

Applications can also perform differently in different environments, such that production exposes problems that could not be caught during testing in a separate environment.

To adhere to DevOps practices, cloud teams need to place their development, testing, and production environments onto the same platform. This lets them ensure that both functionality and hardware work as planned, and it lets them make needed adjustments quickly and cost effectively. After one large personal-insurance organization deployed an integrated DevOps platform on its private cloud, it was able to shift 12 of its most critical application-development teams into a DevOps delivery model without sacrificing application uptime or performance. In fact, the company achieved a 10 percent decrease in production errors for these applications.

poses real challenges, but these challenges can be overcome if companies progressively set up hybrid-cloud platforms according to the three practices described in this article. As leading companies have shown, the time and effort required by cloud-migration programs are more than offset by the resulting gains in the efficiency, quality, and speed to market of digital solutions. ◆

Mark Gu is a consultant in McKinsey's New York office, where **Krish Krishnakanthan** is a senior partner; **Anand Mohanrangan** is a senior expert in the Silicon Valley office; and **Brent Smolinski** is a partner in the Atlanta office.

The authors wish to thank Mishal Desai, Arul Elumalai, Marami Kar, Kevin Major, and Marc Sorel for their contributions to this article.

Cloud adoption to accelerate IT modernization

Nagendra Bommadevara, Andrea Del Miglio, and Steve Jansen

The cloud is a means, not an end. Success in modernizing IT through the cloud is driven by a complete standardization and automation strategy.

Cloud-computing adoption has been increasing rapidly, with cloud-specific spending expected to grow at more than six times the rate of general IT spending through 2020.[1] While large organizations have successfully implemented specific software-as-a-service (SaaS) solutions or adopted a cloud-first strategy for new systems, many are struggling to get the full value of moving the bulk of their enterprise systems to the cloud.

This is because companies tend to fall into the trap of confusing simply moving IT systems to the cloud with the transformational strategy needed to get the full value of the cloud.

Just taking legacy applications and moving them to the cloud—"lift and shift"—will not automatically yield the benefits that cloud infrastructure and systems can provide. In fact, in some cases, that approach can result in IT architectures that are more complex, cumbersome, and costly than before.

The full value of the cloud comes from approaching these options not as one-off tactical decisions

[1] John F. Gantz and Pam Miller, *The Salesforce economy: Enabling 1.9 million new jobs and $389 billion in new revenue over the next 5 years*, IDC, September 2016.

but as part of a holistic strategy to pursue digital transformation. Such a strategy is enabled by the standardization and automation of the IT environment through an open application-programming-interface (API) model, adopting a modern security posture, working in an automated agile operating model, and leveraging new capabilities to drive innovative business solutions. While the cloud is not a prerequisite for any of these features, it does act as a force multiplier. Companies that view cloud capabilities in this way can create next-generation IT capable of enabling business growth and innovation in the rapidly evolving digital era.

Lift-and-shift is not enough

Cloud services such as Amazon Web Services (AWS), Google Cloud, and Microsoft Azure appeal to many organizations because of their stated features, such as pay per use, ability to scale up or down based on usage, high resiliency, and self-service. All these benefits are expected to lead to much lower IT costs, faster time to market, and better service quality compared with traditional IT offerings.

However, traditional enterprises run into two major issues when moving to the cloud:

- The existing business applications were created using the traditional IT paradigm. As a result, these applications are typically monolithic and configured for fixed/static capacity in a few data centers. Simply moving them to the cloud will not magically endow them with all the dynamic features of the cloud.

- The typical technology workforce of an enterprise is well versed in developing business applications in the traditional IT framework. Most of it will need to be reskilled or upskilled for the cloud environment.

IT security is a good example. Most traditional IT environments adopt a perimeter-based "castles and moats" approach to security, whereas cloud environments are more like modern hotels, where a keycard allows access to certain floors and rooms. Unless the legacy applications that have been developed and deployed for a castles-and-moats security model are reconfigured for the new security model, migrating to the cloud may have an adverse impact on cybersecurity.[2]

Enterprises have been successful in adopting SaaS solutions mainly because they address these constraints in a simple fashion: the solutions replace the existing business applications and leave the development of new features to the SaaS provider. SaaS solutions have therefore become very popular for business functions such as marketing and sales, back office, and communication and collaboration. However, in most sectors, there are no mature SaaS solutions for core business functions such as billing for the utilities sector and core/online banking for financial services.

As a result, despite overall increased cloud investment, enterprise cloud adoption is maturing slowly. Many enterprises are stuck supporting both inefficient traditional data-center environments and inadequately planned cloud implementations that may not be as easy to manage or as affordable as they imagined. While some forward-thinking companies have been able to pursue advanced enterprise cloud implementations, the average enterprise has achieved less than 20 percent public- or private-cloud adoption (Exhibit 1).

Benefits of automating IT processes through the cloud

Historically, enterprise business applications have been designed to run on custom-configured IT systems, each application requiring its own

[2] "Benchmark your enterprise cloud adoption," Forrester Research, January 3, 2017, forrester.com.

EXHIBIT 1 **On average, enterprise cloud adoption remains low, at around 20 percent.**

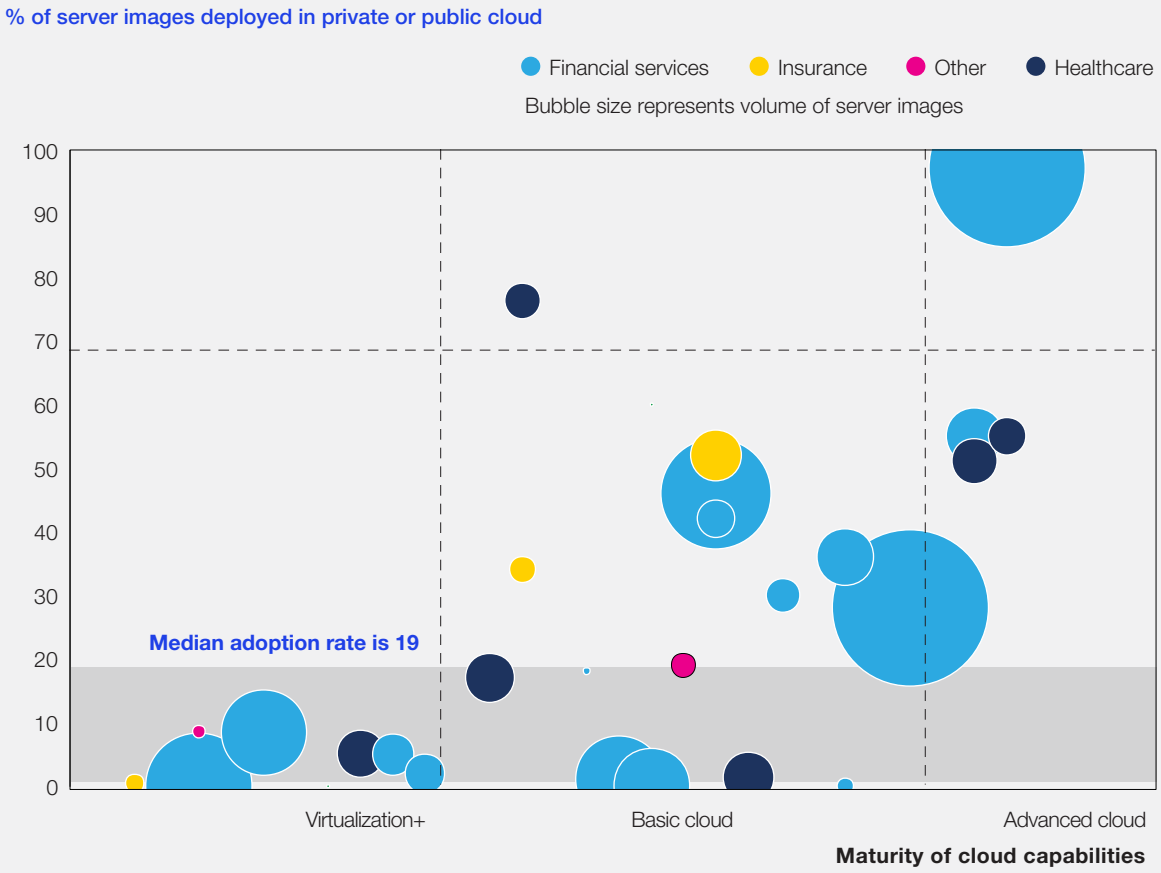

Source: McKinsey Enterprise Cloud Infrastructure Survey 2016

heavily customized configuration of computer storage and network resources. As a result, IT needed armies of administrators just to keep systems updated and running, to add new capacity manually when demand is high, or to apply quick fixes for issues such as low performance. As the number of IT solutions has increased, so has the overhead necessary for testing, integration, and maintenance. In a typical enterprise, just a fraction of IT personnel are focused on designing and developing the market-differentiating solutions the business cares about; the rest are working simply to "keep the lights on."

Standardizing system configurations and automating IT support processes can reverse that ratio. By enabling enterprises to manage their infrastructures better, companies can not only save on costs but also shorten times to market and improve service levels.

Adopting the cloud is a massive enabler of the necessary standardization and automation. With the cloud, companies can do the following:

- reduce IT overhead costs by 30 to 40 percent
- help scale IT processes up and down as needed, thereby optimizing IT asset usage
- improve the overall flexibility of IT in meeting business needs, such as more frequent releases of business features (cloud providers are increasingly offering much more sophisticated solutions than basic computing and storage, such as big data and machine-learning services)
- increase the quality of service through the "self-healing" nature of the standard solutions—for example, automatically allocating more storage to a database (we have seen enterprises reduce IT incidents by 70 percent by using cloud computing as an opportunity to rethink their IT operations)

Capturing these benefits from cloud adoption requires more than just a lift-and-shift approach when the business-application system configurations are heavily customized and IT processes are mostly manual. It requires a certain level of remediation to make IT systems more cloud oriented.

Netflix is one of the most public examples of this kind of commitment to and investment in cloud-enabled, next-generation infrastructure. It spent seven years on its transformation, adopting a cloud-native approach, rebuilding all its technology, and restructuring the way it operates. It employed APIs to reduce its monolithic legacy applications into smaller components, make them more flexible, and then move them to AWS. As a result, service availability has increased, nearing the company's stated goal of 99.99 percent of uptime. And Netflix has seen IT costs for streaming fall to a fraction of what they were in its own data center.

Recently, many established companies have made aggressive moves to adopt public-cloud solutions. Capital One is running the bank's mobile app on AWS, GE Oil and Gas is migrating most of its computing and storage capacity to the public cloud, and Maersk is migrating its legacy systems to reduce cost and operational risk while enabling advanced analytics to streamline operations.

Pioneer organizations are also actively seeking ways to leverage the new services on the cloud to create innovative business solutions. Progressive deployed its Flo chatbot on the public cloud; NASCAR is leveraging machine-learning solutions on the cloud to analyze real-time and historical race-car data to improve performance and simulate scenarios.

Even "born digital" companies that initially chose, for strategic reasons, to have their own IT infrastructure and systems are now opting to move to the cloud to leverage the scalability and the higher-order functionality it offers. Spotify is a prime example.

How to approach the cloud transformation

Fully embracing the cloud can have a significant upside but also requires substantial up-front investments in what is often a multiyear journey. For this reason, an all-in transformation approach needs active commitment and a clear mandate from the CEO and board over the long term (see sidebar, "A tale of an all-in transformation").

Specifically, there are four key topics companies should address for successful cloud adoption at scale:

1. *Decide on sourcing.* It is difficult for most companies to build their own cloud-technology stack and even harder to maintain it. Partnering with public-cloud providers to build and manage the cloud stack is the more typical approach. In most cases, the pragmatic way to start is with use

A tale of an all-in transformation

A Fortune 100 company with a $2.2 billion annual IT spend ($800 million on infrastructure costs alone) was struggling with the cost and complexity of its legacy IT environment. Its IT department was supporting 8,000 applications (including 150 instances of SAP) and 20,000 workloads. Not surprisingly, provisioning was slow. It took more than 45 days to set up a server, and the company knew this was not sustainable.

Consequently, the company invested more than $200 million in an aggressive digital transformation. It was a significant effort, but the company achieved a return on its investment in fewer than four years.

The company first defined its cloud-sourcing strategy, grounding it in an aggressive move to a hybrid model (both public and private cloud), as public-cloud options were still maturing in late 2013. It opted for a single strategic partner for each cloud and recently added a second public-cloud partner. It then created a cloud operating model, setting up a new 100-person team working within an agile operations framework.

Then, beginning in 2015, the company began its legacy-remediation work, moving all its applications to a private cloud, heavily incentivizing its application teams. It took an opportunistic approach to upskilling IT: every application team that wanted to use the cloud had to go through an in-house training program.

Within the first six months, the company had moved its complex SAP environment to a private cloud and adopted a cloud-first policy for all new applications. It replaced expensive colocated contracts and moved its systems to a software-defined data center.

Less than three years in, the company has moved more than 2,000 workloads and two petabytes of data to the public cloud. The company had reduced costs by $90 million at the two-year mark and is on track to cut another $60 million. Automation also significantly improved performance and agility. With the transformation on track to completion in 2018, the company is now one of the largest enterprises operating on the cloud.

of a single cloud-service provider while adopting the necessary guiding principles to avoid being locked into one provider.

After achieving a certain scale and level of maturity—in our experience, a good rule of thumb is to plan for an annual run rate of $30 million with the primary cloud-service provider—an enterprise can explore a second or third service provider for scaling up.

2. *Create a public-cloud operating model.* Unlike traditional operating models, the public cloud requires IT to manage infrastructure as code. This requires software engineers who understand the compute, storage, and security protocols of the public cloud (as opposed to network engineers or system administrators). For most enterprises, this translates to a massive upskilling of the infrastructure organization and the operating model in which they work. Specific

teams need to be assigned to configure and manage the production environment.

3. *Remediate legacy applications.* Existing applications will need to be refactored at the infrastructure and application layers to align with the security and capacity requirements of the public cloud. Security must be baked into these applications, and they must work in a more automated fashion. This requires significant attention from application teams, which can be hard to get.

   Companies can address this hurdle by creating a clear business case for legacy-application modernization, aligning the migration schedule with major application upgrades or replacements, and adopting foundational solutions (such as API frameworks) to make the remediation easier.

4. *Cultivate the right skills.* Professionals must be able to develop applications on the cloud (specifically on the vendor's system) securely and quickly. To do this, companies will need to hire and train cloud experts and then introduce them into development teams, retrain or upskill the existing workforce, and set up digital-innovation labs as needed, with an emphasis on cloud development.

This aggressive approach relies on true commitment from leadership in the form of money (one financial-services business is investing $300 million in a cloud transformation) and time (these programs can take two to three years). That is because, in executing a cloud transformation, multiple things need to happen at the same time. In many cases, for example, a core group of cloud engineers prepares for the cloud migration by setting up the cloud environment, hardening it, looking at applications to move, and creating tools for migration. Meanwhile, the main IT team is being trained in how to work in an agile way. This approach has significant management challenges, but with strong leadership, it is the fastest path to transformation.

Many enterprises, however, are not yet ready to take the full plunge into the cloud, perhaps because organizational buy-in is lacking, there is a reluctance to invest the required resources in a multiyear effort, or they face regulatory constraints. These organizations can achieve significant benefits in the short-to-medium term, albeit on a smaller scale, by adopting the cloud's agile and automated operating model within their traditional IT. This approach builds important organizational capabilities and prepares the business for a cloud transformation when it is ready.

Companies have eagerly adopted agile methods for application development and are actively pursuing automation or DevOps (such as continuous integration and continuous delivery), but the same approach can have an even greater impact on IT operations and infrastructure. By organizing the infrastructure function into tribes of small, cross-functional, self-directed squads with product owners to prioritize work and scrum masters responsible for removing barriers, IT departments can prioritize work in ways that increase productivity, quality, and speed. In addition, the continuous automation program, over time, can further infuse cloudlike capabilities into traditional IT, such as APIs for interactions between developers and infrastructure (Exhibit 2).

With the goals of improving service levels and reducing costs, one major life-insurance company adopted an agile approach within its 250-person IT operations groups. The company began by assessing the state of its current infrastructure—its core processes, organizational model, metrics, key performance indicators (KPIs), and historical demand—and developed a hypothesis about what it might achieve with a more agile approach. It created a leadership program appropriate to agile methods,

EXHIBIT 2

The agile/DevOps operating model is proving to be even more applicable in infrastructure than in application development.

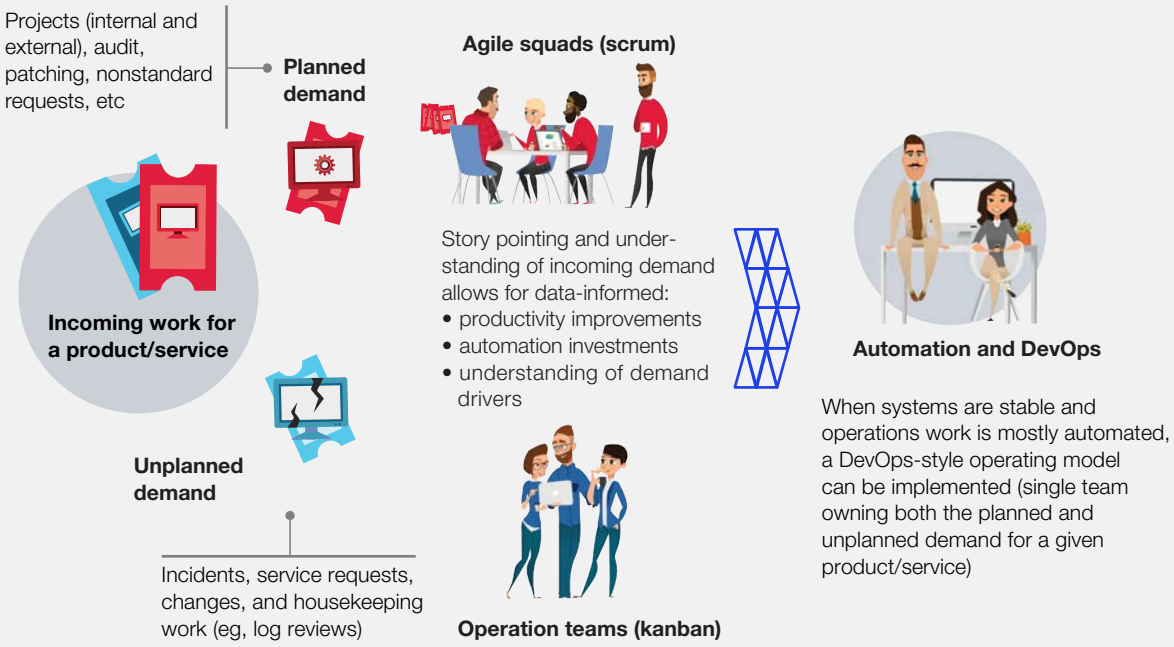

Source: McKinsey analysis

adopted the necessary tools, and conducted an agile-for-infrastructure boot camp for stakeholders.

Within six weeks, the IT infrastructure group started planning for ongoing projects, conducted training sessions for senior leaders and infrastructure teams, and set a goal for what ongoing operations should look like. It fully leveraged the scrum methodology for planned work such as projects and kanban—a methodology for managing the creation of products emphasizing continual delivery—for unplanned work such as incidents and service requests. By the end of the second month, the company had achieved the operational model it envisioned and was able to begin designing service-management processes and launching automation initiatives.

It completed the initial transformation in six months, cutting IT costs by more than 35 percent and doubling productivity. The insurer plans to automate up to 80 percent of its operations work, driving costs down even further and significantly improving its service levels. Today, it is well positioned to move more aggressively to the cloud in the future.

The rules of the cloud game

There are many actions enterprises can take that have proved valuable to early adopters of cloud-enabled next-generation infrastructure. These include but are not limited to the following:

- *Evaluating the current IT portfolio.* Before beginning any cloud development or migration,

take a dispassionate look at the existing IT portfolio to determine what is suited for public-cloud platforms or SaaS alternatives.

- *Choosing your transformation approach.* Involve all key stakeholders in determining whether your enterprise will be an aggressive or opportunistic transformer.
- *Articulating IT and business goals.* Create a well-defined set of outcome-oriented aspirations for both the short and long terms in line with your approach.
- *Securing buy-in.* Ensure commitment and investment from senior management, particularly finance leaders, who must support the transfer from capital to operations and maintenance investment/accounting.
- *Addressing change management.* A heavily automated agile operating model will require significant shifts in IT behaviors and mind-sets. Invest in both change management and the development of cross-functional skills across infrastructure, security, and application environments.
- *Adopting new KPIs.* Measure and reward your technology team for standardization and automation rather than, say, for availability.

◆◆◆

By viewing cloud computing as a starting point for IT automation, companies may be able to have it all: scalability, agility, flexibility, efficiency, and cost savings. But that is only possible by building up both automation and cloud capabilities. ◆

Nagendra Bommadevara is a partner in McKinsey's New York office, **Andrea Del Miglio** is a partner in the Milan office, and **Steve Jansen** is an associate partner in the Charlotte office.

The authors wish to thank Thomas Delaet, James Kaplan, Pankaj Sachdeva, and Anand Swaminathan for their contributions to this article.

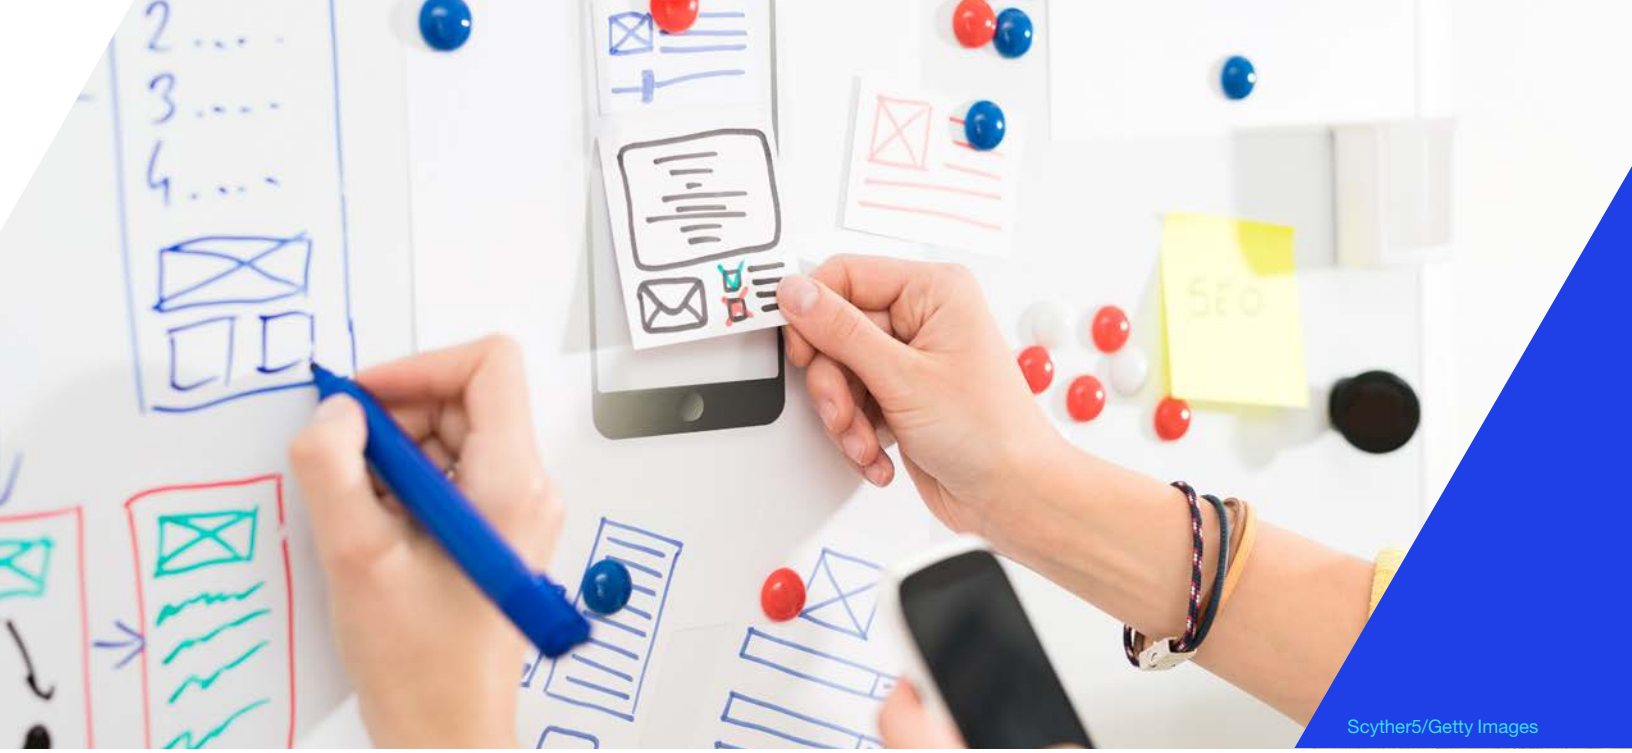

Scyther5/Getty Images

Reimagining software services for the cloud and the digital world

Chandra Gnanasambandam, Rahul Mangla, and Jigar Shah

Customers expect software firms to do more to help deliver outcomes. Software vendors must therefore evolve their professional-services capabilities to meet the new needs.

The growing prevalence of subscription business models and next-generation technologies is fueling large-scale digital transformations to make companies more productive, smarter, and faster. These trends portend a significant change in the way B2B software vendors support newly digital companies.

In the past, the professional-services arms of software companies focused on installing, customizing, and deploying applications for customers. Today, they must help customers design, implement, and adopt new technologies (for example, machine-learning-based applications and blockchain) and migrate workloads to the cloud. In short, software companies are now called on to be partners, not just vendors. And this means that the software industry is being challenged to reassess its entire approach to professional services.

We find that many software vendors encounter challenges navigating these shifts. Until now, their primary areas of focus have been R&D, sales, and marketing. For some companies, the professional-services unit was viewed as a cost center or, at most, a low-margin revenue generator. Many professional-services businesses therefore haven't invested in the

new tools and capabilities they need to propel their operations. That's a mistake. Software vendors must strengthen their professional-services offerings to meet their customers' new demands and to maintain or increase their market share.

To transform the services business and position it for the future, software companies must act along five dimensions: defining the strategic vision for services, reimagining the services portfolio, investing in skills, adapting the services-sales model, and delivering services more efficiently.

Define a strategic vision for the services business

The first step in such a transformation is to define the vision and strategy. Specifically, software vendors need to consider the service business's role (market making or value delivery), economic purpose (growth or profit maximization), and size (the share of the services ecosystem in the company's revenues). That effort should include a thoughtful evaluation of the company's product capabilities and market landscape as well as the maturity of the partner ecosystem. Services should evolve to fit a product's evolution—as customers move their applications to the cloud, for example, the services organization must move away from serving on-premises products.

The role of the services organization must also match the company's goals. A vendor with new products may need the organization to play a greater market-making role by helping to increase their rate of adoption. But a software vendor with a mature product line may instead need a services organization that helps the vendor's partners provide third-party services to the vendor's customers. Leaders must determine whether the services organization should be a growth engine to drive the adoption of products or an efficient operation to maximize profits.

Finally, it's crucial to establish the desired size of the services business and how much work should be left to third parties (Exhibit 1). We find that top vendors seek to provide 10 to 15 percent of the professional services their products require, with the balance provided by partners. For new, unestablished software products, however, a vendor might provide 40 to 50 percent of the services for the first two years and then taper off as third parties take over.

Reimagine the services portfolio

To help customers succeed throughout a digital transformation, B2B software vendors must typically provide a mix of advisory, implementation, and customer-success services—on top of basic installation. Advisory services make it possible for vendors to help conceptualize and design large, complex digital-transformation projects. A big multinational company moving its finance functions to the cloud, for example, would probably need its software-as-a-service vendor to help it design

EXHIBIT 1

The role and size of the professional-services organization must be defined for each product area and customer segment.

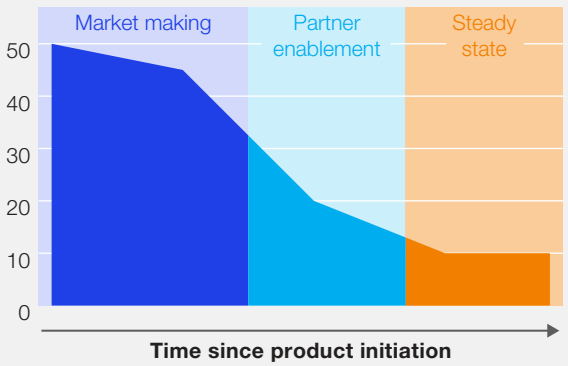

cloud-based processes, the on-premises and cloud architecture, data models, and more.

Implementation services can help vendors give customers the speed they require by helping them to deploy transformative products rapidly and to use advanced technologies, such as the Internet of Things (IoT) and machine learning. A manufacturing company looking to implement an IoT-based digital supply-chain solution, for instance, might use a vendor's implementation services to establish proof of concept and rapidly integrate it with the company's existing supply-chain-management system.

Customer-success services, another integral component of the new professional-services offerings, help customers maximize the value of their software purchases—for example, by using analytics to increase business value. For the vendor, these services not only promote adoption and usage but also reduce churn and therefore boost subscription revenues, which capital markets value disproportionately.

Invest in skills

We often find that to accomplish this substantial portfolio shift and offer these new capabilities, software leaders must fundamentally rethink their people and partner strategy. Training and hiring for the new roles requires a wholly different approach.

Instead of providing standardized training for all members of the services organization, a company should help its employees to learn through the lens of the specific services they will provide. Training ought to focus on this type of role-based learning; an architect and a salesperson, for example, would benefit from very different kinds of training, and account managers transformed into customer-success managers would need broader training. Learning journeys are an effective tool to manage the different training programs required to accommodate all the new services customers need (Exhibit 2).

Even with the most thorough training, the nuances of professional-services roles are learned on the job, especially when customers go through their own digital transformations. Supporting the

EXHIBIT 2

Learning journeys and maps are effective tools to guide retraining of the professional-services workforce.

professional-services team with a resource library that its members can access in any situation will be critical. The customer-success organization, for example, would be able to draw on a tactical tool kit and resource repository if a customer struggled to start a new installation.

Professional-services organizations now require many different skill sets to support their customers' digital transformations, so they must often hire external talent—a process that should start with data. The ability to mine profile and skills data on LinkedIn can be a key differentiator in hiring; cluster analyses on LinkedIn data, for instance, can help sort skills into categories used to find candidates and make decisions. Companies that mine these data with machine-learning tools can hire more effectively.

It's not only the software vendors' professional-services organizations that need to adapt—so must systems integrators and the relationships in this ecosystem. For example, when a customer adopts a bleeding-edge software solution, the vendor's professional-services organization should provide implementation services to establish the new product in the marketplace. As the product matures, however, the implementation tasks can be transferred to systems integrators. To make this shift possible, vendors should invest in training the partner community. Like similar efforts in the vendor's own services organization, these are most effective when conducted through the lens of role-based learning.

Transform services sales

As the services portfolio shifts, so should the go-to-market model that sells it. In the past, generalist account managers sold services. But today, when these services involve much more than just implementation, the savviest vendors recognize that they must rethink their approach to service sales (Exhibit 3).

EXHIBIT 3

A changing product portfolio requires a transformation in the sales of software services.

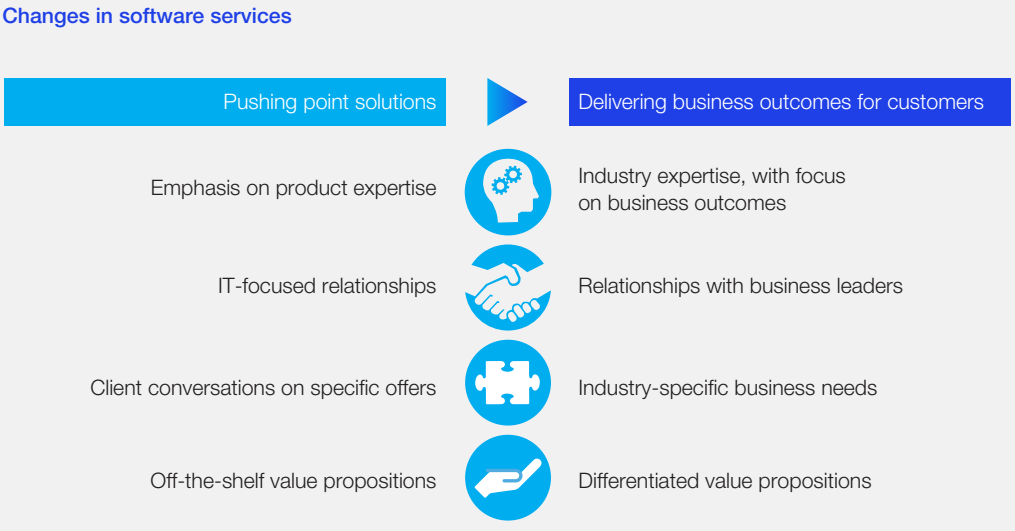

As the need for more types of services grows, roles in the sales process must adapt, along with the orchestration among them. A salesperson courting a bank that's looking to digitize more of its operations, for example, might need to work with services-organization specialists who can guide both the prospect and the salesperson through specific regulatory requirements.

In the same vein, the coverage model also needs to evolve. In the past, a sales organization might have been staffed mostly by generalists, with a small subset of specialists. Now that customers demand so many highly specific services, however, the proportions of specialists and generalists have nearly flipped.

Other aspects of sales that must evolve include the way success is measured (something vendors should consider early on) and the organizational model, which involves weighing trade-offs, such as revenue accountability versus speed of innovation. Institutional capabilities should also be reexamined, especially because services organizations often lag behind their product counterparts in developing a granular understanding of customer needs at the account level.

Focus on efficient delivery

As customers demand that more services be bundled with—or even be enabled by—the platform itself, software vendors must adapt the way they think about managing the cost of services. This isn't cost cutting; it's investing intelligently in the right areas. The key is establishing a balance among services resources and maintaining that balance vigilantly.

Talent is the most prominent driver of costs, which is why investing in skills is so important. Finding the right balance between people hired from inside and outside the organization, and the training involved for each, is also critical. So is deciding how much to use offshore talent and contractors.

A perception has increasingly taken root, for example, that the offshoring or nearshoring of talent is less effective or impossible when software vendors shift focus away from account management and toward customer success. We've found that while customer-success managers certainly need to spend time at customer sites, the services engine can continue to employ nearshore and offshore resources. Newer offshoring locations, such as Eastern Europe, offer access to excellent talent.

The use of contractors must also be managed carefully. Contractors are alluring because they bring high-quality skills without overhead and can be deployed on short notice. But they are expensive, have built-in incentives to become indispensable, and may form important customer relationships that really ought to be held by full-time employees.

Through a structured, concerted effort targeting these and other cost drivers, enterprises can often improve the run-rate cost of services operations by 10 to 25 percent. In addition to improved margins, there are other benefits to improving efficiency in this way: it boosts customer satisfaction and creates headroom for investment in new capabilities and skills.

◆◆◆

Professional services have historically been a required—but often uninspired—offering to help software companies win more enterprise customers. But the world has shifted as those customers adopt subscription products and pursue digital transformation. Today, they are moving much faster, so they need software companies that move with them as partners, not just vendors.

This is a great opportunity for software companies to cement long-term relationships and loyalty. But it also challenges them to rethink their services-business models, to develop new capabilities, and to find the right balance among advisory, implementation, and customer-success services. ◆

Chandra Gnanasambandam is a senior partner in McKinsey's Silicon Valley office, where **Rahul Mangla** and **Jigar Shah** are associate partners.

John Lund/Getty Images

Making a secure transition to the public cloud

Arul Elumalai, James Kaplan, Mike Newborn, and Roger Roberts

As enterprises scale up their use of the public cloud, they must rethink how they protect data and applications—and put in place four critical practices.

After a long period of experimentation, leading enterprises are getting serious about adopting the public cloud at scale. Over the past several years, many companies have altered their IT strategies to shift an increasing share of their applications and data to public-cloud infrastructure and platforms.[1] However, using the public cloud disrupts traditional cybersecurity[2] models that many companies have built up over years. As a result, as companies make use of the public cloud, they need to evolve their cybersecurity practices dramatically in order to consume public-cloud services in a way that enables them both to protect critical data and to exploit fully the speed and agility that these services provide.

[1] For more, see Nagendra Bommadevara, James Kaplan, and Irina Starikova, "Leaders and laggards in enterprise cloud infrastructure adoption," October 2016, McKinsey.com. Also see Arul Elumalai, Kara Sprague, Sid Tandon, and Lareina Yee, "Ten trends redefining enterprise IT infrastructure," November 2017, McKinsey.com, which primarily addresses the impact of infrastructure as a service (IaaS) and platform as a service (PaaS), rather than software as a service (SaaS).

[2] By cybersecurity, this article means the full set of business and technology actions required to manage the risks associated with threats to the confidentiality, integrity, and availability of systems and information. Some organizations may refer to this function as information security or IT security.

While adoption of the public cloud has been limited to date, the outlook for the future is markedly different. Just 40 percent of the companies we studied have more than 10 percent of their workloads on public-cloud platforms; in contrast, 80 percent plan to have more than 10 percent of their workloads in public-cloud platforms in three years or plan to double their cloud penetration. We refer to these companies as "cloud aspirants" (Exhibit 1).[3] They have concluded that the public cloud offers more technical flexibility and simpler scaling for many workloads and implementation scenarios. In some cases, using the public cloud also reduces IT operating costs. As a result, companies are both building new applications and analytics capabilities in the cloud and starting to migrate existing workloads and technology stacks onto public-cloud platforms.

Despite the benefits of public-cloud platforms, persistent concerns about cybersecurity for the public cloud have deterred companies from accelerating the migration of their workloads to the cloud. In our research on cloud adoption from 2016, executives cited security as one of the top barriers

EXHIBIT 1

Cloud aspirants: Nearly 80 percent of companies plan to have 10 percent or more of their workloads in the public cloud or double their public-cloud use within three years.

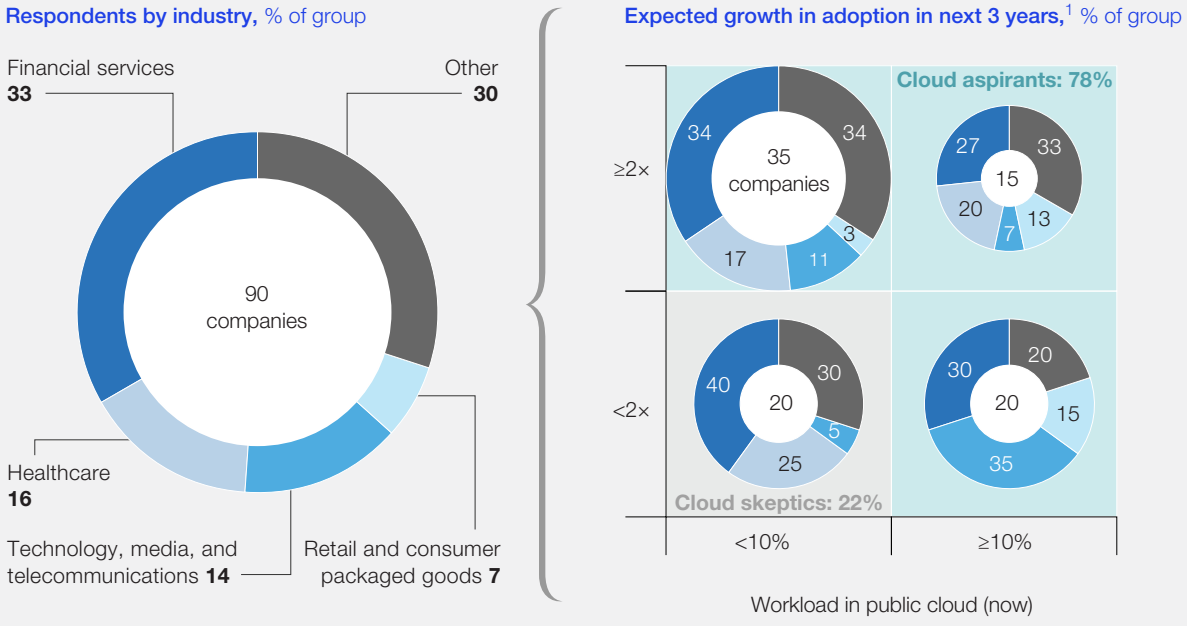

[1]Figures may not sum to 100%, because of rounding.

Source: McKinsey analysis

[3] McKinsey conducted a global survey and in-depth discussions with IT security executives at 97 companies between August 2017 and November 2017, receiving 90 complete survey responses. Forty-one percent of these 97 companies generate annual revenues of less than $3 billion, 22 percent generate $4 billion to $10 billion, 20 percent generate $11 billion to $22 billion, and 17 percent generate more than $22 billion. Thirty-five percent of the 97 companies are in the financial-services industry; 15 percent are in the healthcare industry; 13 percent are in the technology, media, and telecommunications industry; 6 percent are in the retail or consumer-packaged-goods industries; and 31 percent are in other industries.

to cloud migration, along with the complexity of managing change and the difficulty of making a compelling business case for cloud adoption.[4]

Interestingly, our research with chief information-security officers (CISOs) highlights that they have moved beyond the question, "Is the cloud secure?" In many cases they acknowledge that cloud-service providers' (CSPs') security resources dwarf their own and are now asking how they can consume cloud services in a secure way, given that many of their existing security practices and architectures may be less effective in the cloud. Some on-premises controls (such as security logging) are unlikely to work for public-cloud platforms unless they are reconfigured. Adopting the public cloud can also magnify some types of risks. The speed and flexibility that cloud services provide to developers can also be used, without appropriate configuration governance, to create unprotected environments, as a number of companies have already found out to their embarrassment.

In short, companies need a proactive, systematic approach to adapting their cybersecurity capabilities for the public cloud. After years of working with large organizations on cloud-cybersecurity programs and speaking with cybersecurity leaders, we believe the following four practices can help companies develop a consistent, effective approach to public-cloud cybersecurity:

- *Developing a cloud-centric cybersecurity model.* Companies need to make choices about how to manage their perimeter in the cloud and how much they will rearchitect applications in a way that aligns with their risk tolerance, existing application architecture, resources available, and overall cloud strategy.
- *Redesigning the full set of cybersecurity controls for the public cloud.* For each individual control, companies need to determine who should provide it and how rigorous they need to be.
- *Clarifying internal responsibilities for cyber-security, compared with what providers will do.* The public cloud requires a shared security model, with providers and their customers each responsible for specific functions. Companies need to understand this split of responsibilities—it will look very different from a traditional outsourcing arrangement—and redesign internal processes accordingly.
- *Applying DevOps to cybersecurity.* If a developer can spin up a server in seconds but has to wait two weeks for the security team to sign off on the configuration, that attenuates the value of the public cloud's agility. Companies need to make highly automated security services available to developers via application programming interfaces (APIs), just as they are doing for infrastructure services.

Developing a cloud-centric cybersecurity model

For a company that has only begun to use the public cloud, it can be tempting to build a public-cloud-cybersecurity model using the controls it already has for on-premises systems. But this can lead to problems, because on-premises controls seldom work for public-cloud platforms without being reconfigured. And even after being reconfigured, these controls won't provide visibility and protection across all workloads and cloud platforms. Recognizing these limitations, cloud aspirants are experimenting with a range of security strategies and architectures, and a few archetypes are emerging.

[4] For more, see Nagendra Bommadevara, James Kaplan, and Irina Starikova, "Leaders and laggards in enterprise cloud infrastructure adoption," October 2016, McKinsey.com.

The most effective approach is to reassess the company's cybersecurity model with respect to two considerations: how the network perimeter is defined and whether application architectures need to be altered for the public cloud. The definition of the perimeter determines the topology and the boundary for the cloud-cybersecurity model. And choices regarding application architecture can guide the incorporation of security controls within the applications. These two key choices also inform one another. A company might opt, for example, to make its applications highly secure by adding security features that minimize the exposure of sensitive data while the data are being processed and making no assumptions about the security controls that are applied to a given environment.

Choosing a model for perimeter security

Among cloud aspirants, the following three models for perimeter design stand out (Exhibit 2):

- *Backhauling.* Backhauling, or routing traffic through on-premises networks, is how half of cloud aspirants manage perimeter security. This model appeals to companies that require internal access to the majority of their cloud workloads and wish to tailor their choices about migrating workloads to fit the architecture they have. Companies with limited cloud-security experience also benefit from backhauling because it allows them to continue using the on-premises security tools that they already know well. But backhauling might not remain popular for long: only 11 percent of cloud aspirants said they are likely to use this model three years from now.

- *Adopting CSP-provided controls by default.* This model is the choice of 36 percent of cloud-aspirant companies we studied. Using a CSP's security controls can cost less than either of

EXHIBIT 2

Architecture options: Three models for perimeter architecture stand out among cloud-aspirant companies.

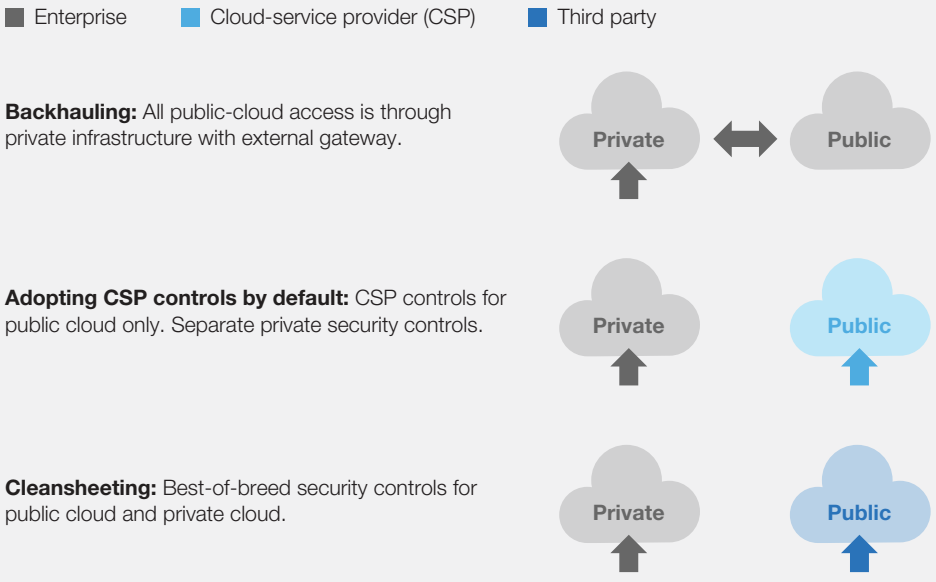

Source: McKinsey analysis

the other perimeter models but makes it more complex to secure a multicloud environment. For larger and more sophisticated organizations, using CSP-provided controls appears to be a temporary measure: 27 percent of cloud aspirants say they will use this model in three years (down from 36 percent today).

- *Cleansheeting.* Cleansheeting involves designing a "virtual perimeter" and developing cloud-specific controls from solutions offered by various external providers. Used by around 15 percent of cloud aspirants, this approach enables companies to apply the best perimeter-security solutions they can find, switching them in and out as needed. Since changing solutions creates technical demands, companies typically practice cleansheeting when they have enough in-house cybersecurity expertise to select vendors and integrate their solutions. Although those efforts can slow the migration of workloads into the cloud, cleansheeting appears to be on the rise, with 47 percent of cloud aspirants saying they will use cloud-specific controls in three years. Despite the high cost and complexity of cleansheeting, organizations choose this approach so they can support multicloud environments and replace point solutions more easily as their needs evolve.

Backhauling is now the most popular model for perimeter security among the cloud aspirants we researched. However, enterprises are moving toward a virtual-perimeter model, which they develop through cleansheeting (see sidebar "A progressive outlook on perimeter-security design"). Cleansheeting is the least popular practice for managing perimeter security today, but more executives say they will use cleansheeting over the next three years than any other model.

Deciding whether to rearchitect applications for the cloud

The second choice that defines a company's cloud-cybersecurity posture is whether to rearchitect applications in the public cloud, by rewriting code or altering application architectures (or both). Just 27 percent of the executives we interviewed said their companies do this. The benefits are compatibility with all CSPs (with container architectures, for example), stronger security (with changes like tamper detection using hash, memory deallocation, and encryption of data flows between calls), superior performance (for example, by allowing horizontal scaling in the public cloud), and lower operating costs (because app-level security protections reduce the need for a company to choose best-of-breed security solutions). However, rearchitecting applications for the cloud can slow a company's migration rate. Because of this, a large majority of enterprises in our survey, 78 percent, migrate applications without rearchitecting them for the public cloud.

The choice of perimeter-security design, along with the choice about whether to adapt applications to the public cloud, create six archetypes for cloud cybersecurity. In our experience, five primary criteria inform enterprises' decisions about their overall cloud-cybersecurity model: public-cloud-security effectiveness, their desired cloud-migration rate, their willingness to pay additional security costs, their expertise implementing new security programs, and the flexibility they desire from their security architectures (Exhibit 3).

Rearchitecting applications for the public cloud improves security effectiveness but can slow down migration. Backhauling extends existing controls that companies are already familiar with to public-cloud implementations. Using default CSP controls is the simplest and most cost-effective approach.

A progressive outlook on perimeter-security design

A cybersecurity executive we interviewed at a large pharmaceutical company described a forward-looking view of perimeter-security design that is fairly typical of cloud aspirants. As the company increases its use of the public cloud, it is backhauling as a stepping stone but intends to move to a flexible architecture that leverages cloud-service provider controls where available and third-party controls for areas that CSPs do not support. Said the executive: "We lift and shift applications to the public cloud, and backhauling is an intermediate step. However, we see that CSPs and third-party tools provide more secure technology. We appreciate the shared responsibility with our CSP, but we require additional third-party tools to go beyond default CSP capabilities."

EXHIBIT 3

Assessing architectures: Cloud-cybersecurity models generally follow six archetypes, which are defined by their designs for perimeter and application architectures.

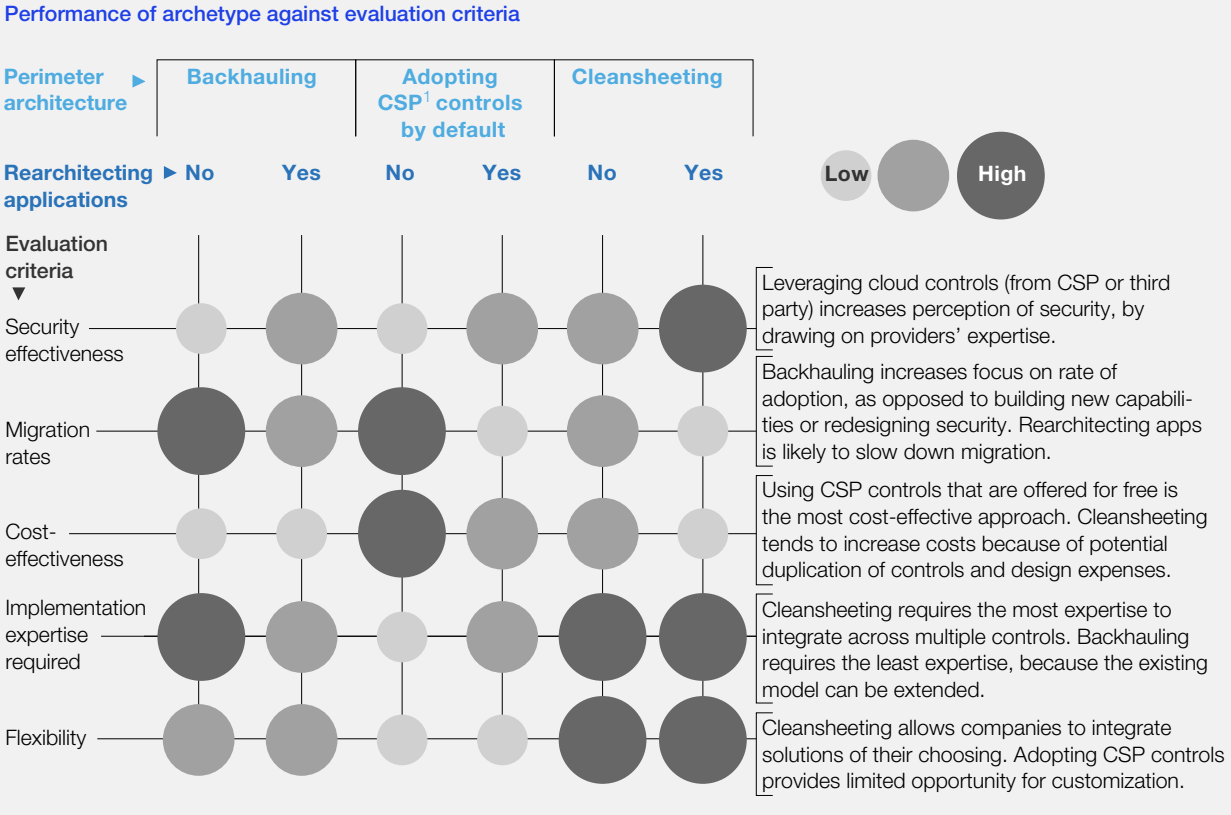

[1]Cloud-service provider.

Source: McKinsey analysis

Cleansheeting controls calls for substantial security expertise but provides flexibility and support for multiple clouds. Organizations can use these criteria to choose the best methods. That said, a company need not apply the same archetype to its entire public-cloud profile. It's possible, even advantageous, to use different archetypes for applications with different requirements: for example, backhauling with a single CSP for a core transaction system to enable faster migration and familiar controls while using CSP-provided security controls for low-cost, accelerated deployment of new customer-facing applications.

Redesigning a full set of cybersecurity controls for the public cloud

Once enterprises have decided on a security archetype (or a mix of archetypes, with each archetype matched to a group of workloads with similar security requirements), they can design and implement cybersecurity controls. Understandably, companies are experimenting with a variety of designs for controls, and, given the pace of progress, cybersecurity executives anticipate considerable change to these controls over the next three years. Cybersecurity controls can be categorized into eight areas, which organizations need to think about in combination:

- *Identity and access management.* IAM solutions for cloud-based applications and data are gradually shifting into the cloud (see sidebar "Moving into the next generation of IAM"). Sixty percent of interviewees reported that they employ on-premises IAM solutions today, but only half as many expect to be using on-premises IAM solutions in three years. By that time, 60 percent of interviewees anticipate that their enterprises will rely on a third-party IAM service that supports multiple public-cloud environments and unifies IAM controls across on-premises and public-cloud resources.

- *Data.* Encryption of cloud data in motion and at rest should soon be standard practice. Eighty-four percent of cloud aspirants expect that within three years they will encrypt the data they store in the cloud. Over time CISOs would like to have more practical mechanisms for encrypting data in memory as well. However, interviewees have different approaches to managing encryption keys for cloud workloads: 33 percent prefer to have CSPs manage keys, 28 percent keep them on premises, and 11 percent prefer to have third parties manage keys (see sidebar "Why companies manage keys differently").[5]

Moving into the next generation of IAM

A Fortune 500 healthcare company we spoke with has redesigned its identity- and access-management (IAM) controls for the public cloud by using the automation and analytics features of its public-cloud platforms. Specifically, it has created automated authorization schemes, based on identity services provided by a cloud-service provider (CSP), to eliminate human factors from provisioning and deprovisioning. The company has also developed a risk model that predicts each user's behavior based on monitoring data from the CSP and compares that behavior with what is observed to determine whether the user should gain access. As a company executive told us in an interview, "Passwords are obsolete. Even MFA [multifactor authentication] is a step backward. Behavioral authentication is the next generation. With the training data from CSPs, we are taking a risk-based approach and building continuous authentication."

[5] Twenty-eight percent of interviewees declined to discuss key management.

- *Perimeter.* Enterprises are moving toward a virtual-perimeter model. Around 40 percent of enterprises are routing traffic via on-premises data centers today, using on-premises security controls with some form of virtual private network or direct connectivity between on-premises and public-cloud workloads as the only way to access applications or data on public-cloud platforms. But 49 percent of interviewees say they expect their companies to use third-party perimeter controls over the next three years. The transition to these perimeter-control models will typically involve developing cleansheet designs that draw on a combination of services, such as security secure web gateways, web-application firewalls, and network monitoring from different third parties that support multiple clouds.

- *Applications.* Most interviewees (84 percent) define security-configuration standards for cloud-based applications and depend on CSPs to implement them. But 85 percent said their companies are likely to drive more developer governance as workloads move to the cloud. This is likely to be soft governance, with only 20 percent of enterprises using application-security tools or templates.

- *Operations monitoring.* Sixty-five percent of enterprises rely on their current security information and event management (SIEM) tools for monitoring cloud apps. This allows them to maintain a single view of their on-premises and cloud workloads. Another 30 percent use other native monitoring tools provided by their CSPs or request logs from CSPs to generate insights using proprietary data-analytics solutions. Since CSPs can provide a wealth of monitoring data, it is critical for organizations to collaborate with them on selecting solutions that provide a unified view of on-premises and public-cloud workloads.

- *Server-side end points.* Interviewees are mostly confident in the server-side security offered by CSPs: 51 percent indicate that they have a "high" level of comfort with CSP-provided security

Why companies manage keys differently

Companies determine their key-management practices based on various factors, such as regulatory compliance and security benefits. Two examples from our interviews show why approaches differ. An IT services company has opted to generate and manage keys using a localized private system so it can use key ownership as a mechanism to stay "in the loop" if cloud-service providers (CSPs) are forced to hand over data. The executive explained, "We are holding the key ourselves because it gives us and our compliance people confidence that only local employees have access to keys, and data cannot be accessed without our knowledge. That control gives peace of mind."

A global pharmaceuticals and medical-products company takes a different approach, drawing on its CSP's key-management capabilities to improve cost-effectiveness and performance. The executive we interviewed said, "Our public-cloud application functionality is improved when keys are stored in the public cloud. Public-cloud applications need the keys to decrypt public-cloud data, and so we see less security benefit to storing keys privately. We get better performance having keys closer to apps, and encryption and decryption cost less with publicly stored keys."

for server-side end points. Many companies, especially ones that have less sophisticated security programs, believe that CSPs have more insight into and control over their server fleets than they could ever achieve internally.

- *User end points.* Moving workloads to the cloud ordinarily necessitates changes to controls for user devices, mainly for data-loss prevention and for protections against viruses and malware. Seventy percent of interviewees said using a public-cloud infrastructure requires their enterprises to change users' end-point controls.

- *Regulatory governance.* Most cybersecurity programs are governed by regulations on data protection (such as the European Union's General Data Protection Regulation), data location and sovereignty, and personally identifiable information. Financial institutions and healthcare organizations are also subject to industry-specific regulations. More than 50 percent of the executives we spoke with indicated that they would like their CSPs to be jointly responsible for compliance with regulatory mandates.

In selecting controls, organizations should consider all eight areas in conjunction and build a comprehensive cybersecurity architecture rather than following a piecemeal approach. Companies can start to design controls based on threat scenarios and levels of security required, and then they can apply an appropriate security-model archetype (such as backhauling or cleansheeting) to determine the best security controls and their scopes. Companies can also work with CSPs to determine which of their controls to use and which ones to procure from third parties. Finally, companies should short-list and prioritize controls that can be standardized and automated and then implement them in agile iterations.

Clarifying internal responsibilities for cybersecurity, compared with what providers will do

When enterprises migrate applications and data to the public cloud, they must depend on CSPs and third-party providers for some security controls—but they should not depend on these parties to provide all of the necessary controls. Unless companies and CSPs clearly divide all the responsibilities for cybersecurity in public-cloud environments, some responsibilities could fall through the cracks. This makes it essential for companies to develop and maintain a clear understanding of what controls their CSPs provide by having CSPs provide a comprehensive view of their security operating models, along with timely updates as those models change. (CSPs organize their cybersecurity responsibility models differently, and take various approaches to sharing them, so each situation needs to be handled carefully.) That way, companies can design and configure controls that work well in multiple cloud environments and integrate well with various tools, processing models, and operating models.

Based on our experience and research, we find that enterprises can benefit greatly from collaborating with CSPs across the full cybersecurity life cycle, from design to implementation and ongoing operations. However, four main areas emerged as top priorities for collaboration between companies and their CSPs:

- *Transparency on controls and procedures.* Companies should get CSPs to provide full visibility into their security controls and procedures, as well as any exposure incidents. Companies will also need to understand each CSP's ability to conduct security audits and penetration testing.

- *Regulatory-compliance support.* Companies should ask their CSPs to provide detailed descriptions of the assurances they provide with

regard to regulatory compliance, inquire about how they stay abreast of regulatory changes for each industry, and update their compliance mechanisms accordingly.

- *Integrated operations monitoring and response.* Companies will likely have to collaborate with CSPs when it comes to integrating their SIEM tools in a way that supports centralized security administration. Companies should request that their CSPs provide them with comprehensive reporting, insights, and threat alerts on an ongoing basis. They can pass on insights to help CSPs develop new capabilities for all their tenants. They must also ensure that CSPs make logs readily available in formats that companies can process using on-premises analytics tools.

- *Multicloud IAM capabilities.* Companies should insist that CSPs provide native multifactor authentication. Those that use identity as a service (IDaaS) or on-premises IAM solutions will need to work with CSPs to integrate them properly, so they have adequate support for multiple public-cloud environments. Companies should also have their CSPs share their IAM road maps so the companies can plan to take advantage of features such as behavioral authentication and role-based access.

Applying DevOps to cybersecurity

DevOps is an increasingly prevalent approach to integrating development and IT operations that supports continuous delivery of new software features, in part by providing developers with APIs to access operational services. Secure DevOps (sometimes called "SecDevOps" or "continuous security") integrates security reviews, implementation of security controls, and deployment of security technology with the DevOps approach that many teams have already adopted for movement into the cloud. Integration is achieved by automating security services across the full development cycle and making them available via APIs (Exhibit 4).

Secure DevOps enhances all categories of security controls for the cloud by shortening deployment timelines and reducing risk. For example, some companies have policies requiring the classification of all data. But when data can only be classified manually, the necessary effort adds time to deployment schedules. With secure DevOps, mandatory data classification becomes much more practical, because all data receive a default classification based on preset rules. As a result of that improvement, and others provided by secure DevOps, organizations can decrease their risk of breaches in public-cloud environments while reducing or removing delays that would have been caused by manually classifying data before they are stored.

Adopting secure DevOps requires companies to foster cultures in which security is a key element of every software project and a feature of every developer's work. Many developers will need additional security training to provide effective support during and after the public-cloud migration. Training will also help developers understand the security features of the tools they are using, so they can make better use of existing security APIs and orchestration technologies and build new ones.

Companies should streamline their security-governance procedures to make sure they do not cause delays for developers. As companies automate their security controls, they can make controls fully visible to developers. That way, developers can independently check whether controls are working properly in the background, rather than delaying work to consult with security specialists. Automating the processes of auditing security mechanisms is also helpful. For example, companies can require

EXHIBIT 4

Traditional security models make it harder to take advantage of cloud's speed and agility.

Cloud-deployment process with secure DevOps

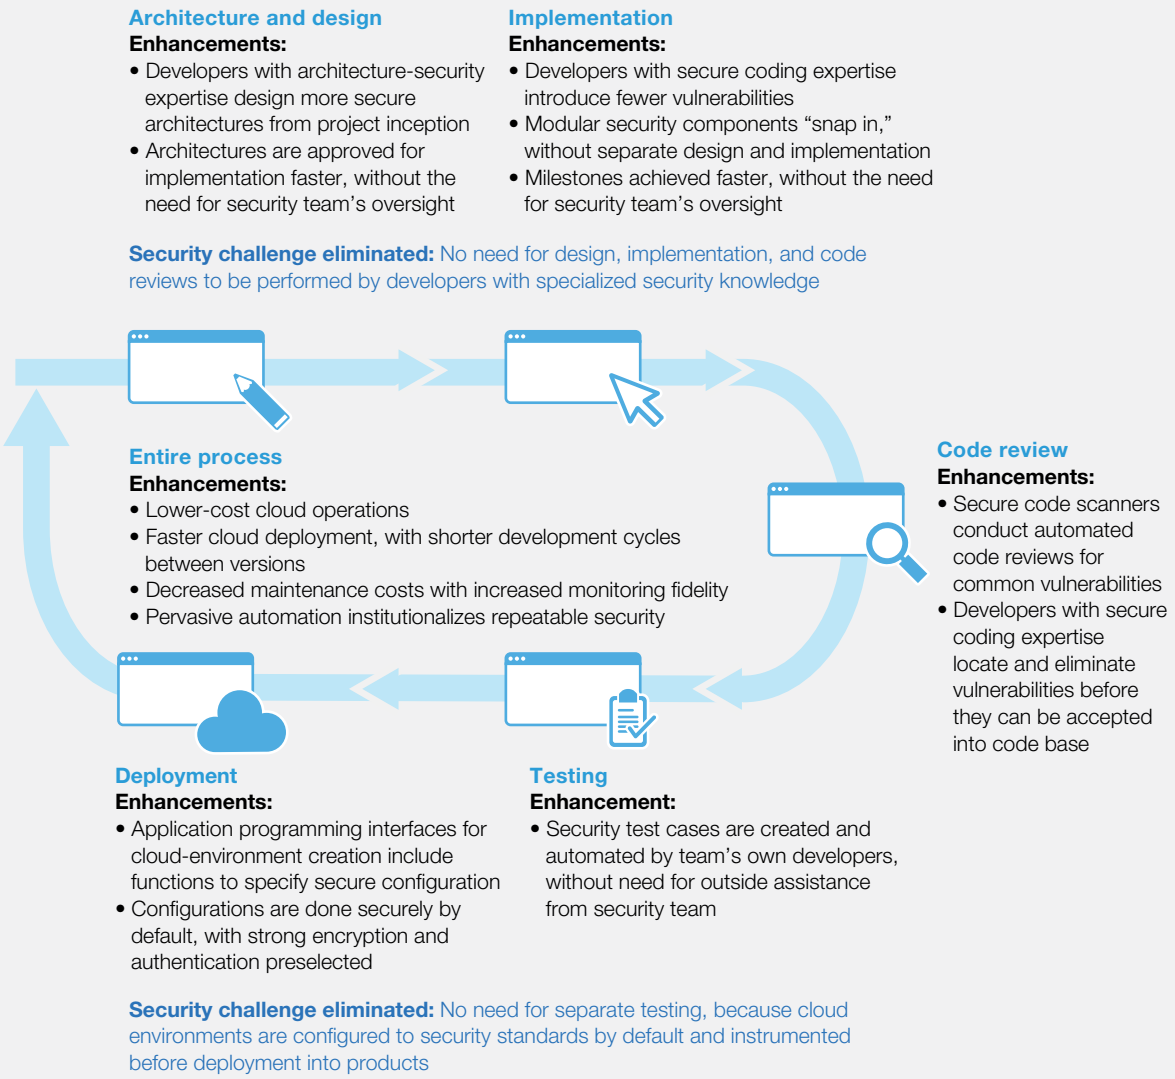

Source: McKinsey analysis

that code is automatically scanned every night for compliance with policy and integrate build-time checks of security components into applications.

To implement secure DevOps, companies also change their IT operating model so security implementation becomes a part of the cloud-development and -deployment processes. In such an operating model, a properly trained development team is the security team; no outside engagement is needed to obtain the right security expertise. Embedding security expertise in the development

team eliminates delays in the cloud-deployment process and permits the development team to iterate much faster than traditional security models allow.

How companies can begin strengthening cybersecurity in the cloud

The four practices we have described for structuring a public-cloud-cybersecurity program should enable companies to take greater advantage of public-cloud platforms. Nevertheless, setting up the program can be a complicated task, because companies have multiple cloud workloads, CSPs, on-premises and private-cloud capabilities, locations, regulatory mandates, and security requirements to account for. This ten-step workplan will help companies stay coordinated as they move through design, development, and implementation of their public-cloud cybersecurity programs:

1. *Decide which workloads to move to the public cloud.* For example, many organizations choose to move customer-facing applications or analytical workloads to the public cloud initially, while keeping core transaction systems on premises. Then they can determine security requirements for workloads that are migrated.

2. *Identify at least one CSP that is capable of meeting security requirements for the workloads.* Companies may choose multiple providers for different workloads, but these selections should be consistent with the objectives of the company's overall cloud strategy.

3. *Assign a security archetype to each workload based on the ease of migration, security posture, cost considerations, and internal expertise.* For example, companies can rearchitect applications and use default CSP controls for customer-facing workloads, and they can lift and shift internal core transaction apps without rearchitecting while backhauling for data access.

4. *For each workload, determine the level of security to enforce for each of the eight controls.* For example, companies should determine whether IAM needs only single-factor authentication, requires multifactor authentication, or calls for a more advanced approach such as behavioral authentication.

5. *Decide which solutions to use for each workload's eight controls.* Given the capabilities of the CSP (or CSPs) identified for each workload, the company can determine whether to use existing on-premises security solutions, CSP-provided solutions, or third-party solutions.

6. *Implement the necessary controls and integrate them with other existing solutions.* This requires the company to gain a full understanding of CSPs' security capabilities and security-enforcement processes. CSPs need to be transparent about these aspects of their offerings.

7. *Develop a view on whether each control can be standardized and automated.* This involves analyzing the full set of controls and making decisions on which controls to standardize across the organization and which ones to automate for implementation.

8. *Prioritize the first set of controls to implement.* Controls can be prioritized according to which applications a company migrates and which security model it chooses to apply.

9. *Implement the controls and governance model.* For controls that can be standardized but not automated, companies can develop checklists and train developers on how to follow

them. For controls that can be standardized and automated, companies can create automated routines to implement the controls and to enforce standardization, using a secure DevOps approach.

10. *Use the experience gained during the first wave of implementation to pick the next group of controls to implement.* Drawing on this experience will also help improve the implementation process for subsequent sets of controls.

◆◆◆

Companies are steadily moving more of their applications from on-premises data centers and private-cloud platforms to public-cloud platforms, which provide superior levels of cost-effectiveness, flexibility, and speed in many situations. But public-cloud migrations will only succeed if companies maintain the security of their applications and data—a task that some have struggled with.

Our experience and research suggest that public-cloud cybersecurity is achievable with the right approach. By developing cloud-centric cybersecurity models, designing strong controls in eight security areas, clarifying responsibilities with CSPs, and using secure DevOps, companies can shift workloads to the public cloud with greater certainty that their most critical information assets will be protected. ◆

Arul Elumalai and **Roger Roberts** are partners in McKinsey's Silicon Valley office, **James Kaplan** is a partner in the New York office, and **Mike Newborn** is a senior expert in the Washington, DC, office.

The authors wish to thank Yash Agrawal, Rich Cracknell, Srikanth Dola, Lisa Donchak, Dan Guo, James Manyika, Brent Smolinski, and Adam Tyra for their contributions to this article. They also wish to express their thanks to the security team members at Google Cloud for their input and insights and to the more than 100 security executives who shared their practices and plans, without which this article would not have been possible.

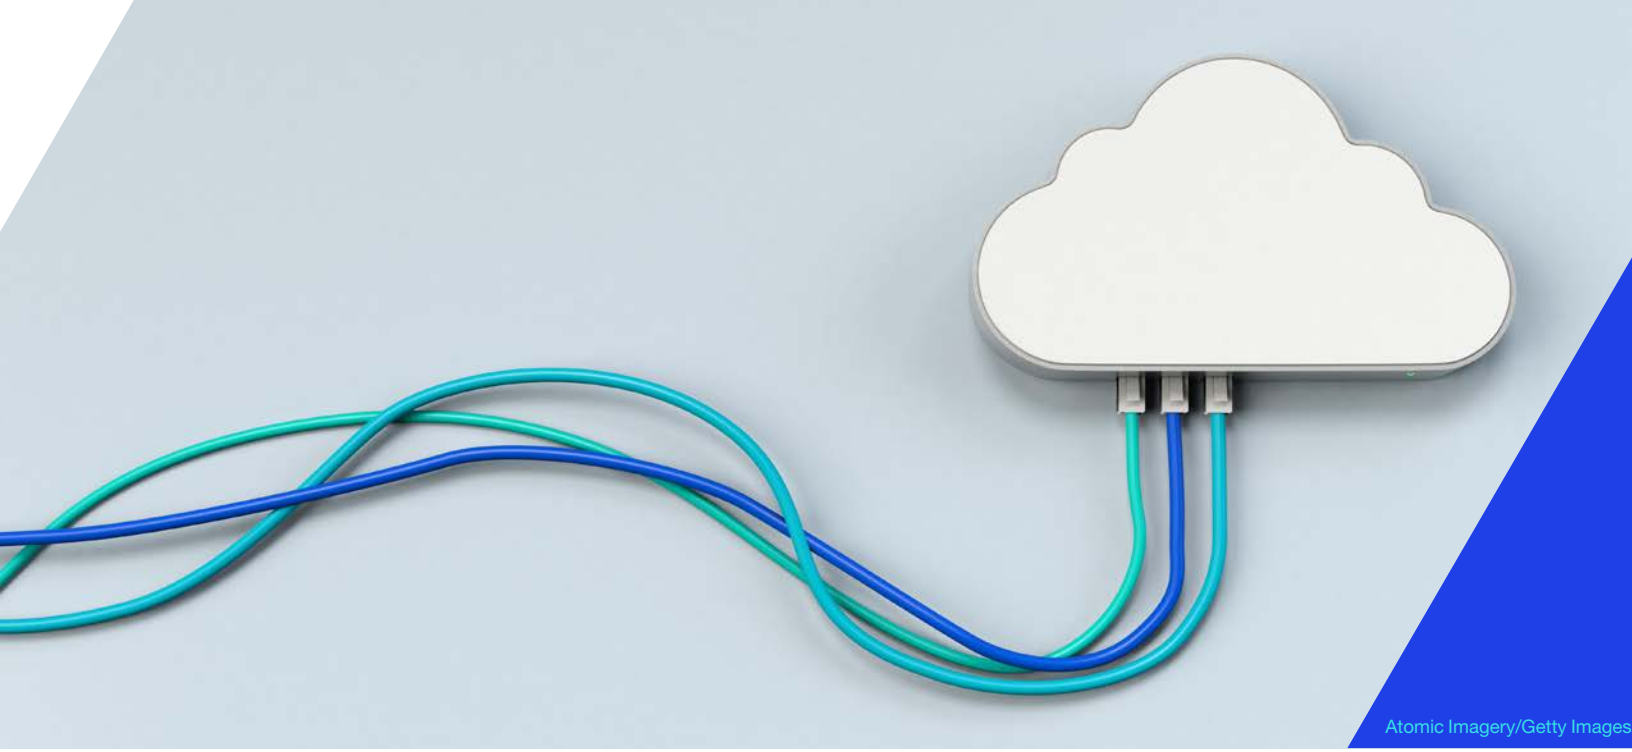

Atomic Imagery/Getty Images

Learning from leaders in cloud-infrastructure adoption

A crucial benefit of cloud adoption is a decrease in time to market for new applications, which in turn can drive down costs and quickly improve product quality.

Companies that have taken the initiative to adopt cloud infrastructure rather than rely on server technologies have found that the advantages are well worth the investment of resources. In this transcript of a *McKinsey Podcast*, McKinsey partner Irina Starikova speaks with McKinsey Publishing's Roberta Fusaro about what laggards in the enterprise cloud-infrastructure space can learn from leaders finding business uses for cloud technologies.

Roberta Fusaro: Let's start this discussion on the ground. What is the cloud, and what are some examples that we might run across in our day-to-day lives?

Irina Starikova: Put very simply, the cloud is a network of distributed servers that are hosted on the internet, and those servers are managed in a highly automated way. They're also shared by many applications at the same time, and that results in three kinds of outcomes.

First, you have much lower cost of hosting applications and data. Second, you have much faster speed of putting new applications on that infrastructure. Lastly, you have much better reliability and security for your applications.

Those servers can be either internal for your enterprise—and we call them private cloud—or they can be owned or managed by a third party. In that case, you would call them public cloud or managed private cloud. We use applications and data that are hosted on cloud technology every single day. In our

personal lives, there are very few things that you do when you're turning on an application on your phone or you're sharing data with someone that would work without cloud technology in the back end.

The examples run the gamut of everything you do in your daily life. You can be shopping on Amazon. You could be watching Netflix, sharing pictures with your family, getting an Uber, ordering food on DoorDash. Or you could be booking your SoulCycle session.

That all involves some sort of cloud technology in the back end to make it work. Similarly, when you think about our clients, most large companies today use cloud technology quite extensively. That could be a private cloud that they're managing in their own data center, or they could be using services by public-cloud providers such as Amazon Web Services, Google Compute Platform, Azure, or IBM.

Roberta Fusaro: How have cloud technologies and the market for cloud solutions evolved over the past three to five years?

Irina Starikova: The overall market for those services has really taken off. If you look at the latest reports by all leading market analysts, everyone is putting it well above $200 billion.

There's hardly any debate about this being a huge thing happening. Secondly, when you look at enterprise adoption of cloud, that also started to change dramatically, and it's shifted a lot from private cloud to public cloud.

To give you some numbers, through our surveys, we found out that more than half of all enterprises of any size plan to shift at least some applications completely to the public cloud in the next two to three years. That's the change that we started to see happening in the last two years.

Those things have a huge impact on the overall enterprise-technology ecosystem. If you think about several years back, enterprises were direct buyers of 35 to 40 percent of all server and storage technology. Now some analysts expect that the share will shrink to less than 20 percent, and that will happen as soon as the next two years. That has huge implications, obviously, on all providers of server-storage networking technology as well as service providers that exist in the ecosystem around that.

Roberta Fusaro: How have companies' discussions about the cloud changed over the past three to five years?

Irina Starikova: In addition to this shift of enterprises to use public-cloud services a lot more, we also see that there's a shift in conversation to the scale of adoption. People are talking about what it's like to be using the cloud for a majority of applications in their portfolios. Another big set of conversations that has changed significantly is related to the security and compliance requirements of the public cloud. Let me take those one by one.

On scale of adoption, companies are no longer happy to be using the cloud for just a small share of their overall data-center footprint or a small share of their application portfolio. There's a lot of focus on what it would take to really adopt the cloud at scale and what it would take to adopt public-cloud services at scale.

On the security and compliance side, we've gone away from talking about how that is the hugest barrier to using public-cloud services. Now you have a lot more advanced conversation on what the right controls are and what the right standards are to protect information in the public cloud.

Security is still very important, and compliance is still a nonnegotiable thing for many of our clients. But what is happening now is that instead of saying, "OK, we're just not even going to discuss cloud because of those constraints," people are saying, "OK, well, those constraints are there. Let's talk about specifically how they're going to

be addressed when we use public-cloud services." And, frankly, even for clients that are coming from highly regulated industries that have to worry about highly sensitive patient information or customer information that is considered highly personal, we already see many examples of those companies moving to adopt public-cloud services at scale for a pretty large variety of different applications.

Roberta Fusaro: McKinsey's Enterprise Cloud Infrastructure Survey sheds light on what's really going on with cloud adoption. When was it conducted? And who participated?

Irina Starikova: We started the survey in 2014. Over time, we've collected information from more than 50 large enterprises that are based either in North America or in Europe. We wanted to understand what cloud technology they were adopting, how they were adopting it, and at what pace.

For a good majority of those enterprises, we have multiple observations across this time period, so we can see how they have evolved over time. We were able to include companies here from a variety of different industries. So we have just as many companies from nonregulated as well as regulated spaces as well as company sizes and different levels of cloud adoption and sophistication.

Companies are still investing in pretty complex private-cloud platforms. And those companies, we believe, first went down this path because they thought that the public cloud was not secure enough or not meeting compliance requirements they have. Some of them chose more sophisticated platforms to build something that can meet the needs of many different applications in their portfolio. They did that over choosing a more practical and simpler approach that is going more aggressively after broader adoption—and, frankly, better impact—from using simpler solutions, while some companies are continuing to build those complex private-cloud platforms. We sometimes talk about that as a big, hairy science project. There are clearly companies that are emerging as leaders in cloud adoption, and we are calling them "cloud savvy." They have achieved a lot higher adoption of cloud.

We measure that as a share of their overall hosting environments that are based on cloud technology. The difference between leaders and laggards here is pretty stark. We're talking in some cases about a gap of 40 to 50 percent. Some leaders in the same market and in the same industry would have over 40 or 50 percent share of their environments on cloud, whereas the laggards would have single-digit percentage share. What leaders have done differently in those cases is that they focused a lot more on building organizational capabilities rather than overinvesting on technology engineering.

They were not striving to create a perfect technology solution but were, first of all, focused on getting meaningful results. So they tested and learned and adjusted their strategies along so that they focused a lot more on getting results rather than science projects.

Roberta Fusaro: Clearly your research found leaders and laggards—a lot of companies that have a way to go with their cloud programs. What lessons can the laggards take from the leaders?

Irina Starikova: The benefits are quite significant, and there were multiple types. The number-one benefit that many leaders saw from adopting cloud was in time to market. What that means is that they were able to deploy new applications using cloud services a lot faster than they were able before. Sometimes we were talking about the difference between weeks cut down to a few hours and sometimes less than one hour.

The importance of that time to market is that the business of those organizations were able to deploy changes to their products a lot faster than they were ever able before or they could change some of their internal processes that they were transforming a lot faster.

What comes clearly in the second and third place in terms of benefits are cost reductions and quality improvements. What that means simply is that the total cost of operating your hosting infrastructure has gone down quite significantly because of the cloud. Similarly, the quality, the reliability, of that service has improved a lot in the same time.

Roberta Fusaro: I noticed that one of the major themes that emerged from the research was this notion around openness to the public cloud. This point has been cited in a lot of external media. Can you talk a little bit more about this point?

Irina Starikova: In part, this has been happening because some of the cloud-service vendors have become a lot more aggressive. They have invested a lot in their enterprise sales forces and have been beating on the doors of a lot of them.

In parallel, the economics of public-cloud services have changed a lot in the last three years and have become comparable to what some of the most efficient private-cloud environments were able to achieve.

So it has become a lot easier for our enterprise clients to be able to see that they can save quite a bit by moving to the public cloud. Of course, it also happened because the security standards started to emerge for the public cloud. As we already said, the conversation around security and compliance has shifted from that being the major barrier to it no longer being a major barrier but instead being something that needs careful understanding and analysis and engineering before any applications can be shifted to the public cloud.

Roberta Fusaro: There've been wide reports of a number of security breaches in government agencies and companies and so forth. I'm wondering if any of that has had any impact or could have any impact on the data points that you cited.

Irina Starikova: Absolutely. There will always be concerns. All of the cybersecurity questions and unfortunate incidents recently have brought it back to the top of mind for everyone. There's a much better understanding of how security in the public cloud works, how it is different from what companies have been able to build internally in their own data centers within their own walls, and understanding where the public cloud could be better, stronger, than what folks are able to do today. You start to understand a lot better what the weaknesses are and what the available tools are for you to address those weaknesses.

At the same time, what's been interesting to see is what other concerns have become the top barriers on the top of mind of enterprises for adopting public cloud, much more practical questions, such as, what is the cost? What is the complexity to move away from what the enterprises have accumulated in their own data centers?

Another one that often comes up in conversation is related to vendor lock in. Many enterprises are concerned about the concentration that is happening in the provider space. Increasingly, the top four players are gaining bigger and bigger market share away from all of the other players.

Roberta Fusaro: Looking at those two particular concerns—these notions of moving away from legacy systems and avoiding vendor lock in—did your research turn up any best practices or any advice for avoiding those traps? Or mitigating those traps?

Irina Starikova: A number of companies are starting to ask for better standards or interoperability commitments from the biggest vendors, so that

it becomes easier for enterprises to shift between those players and avoid the vendor lock in, avoid being attached to one single one.

Roberta Fusaro: Notwithstanding the very legitimate issues that were surfaced in the survey, do you think everything is going to end up in the cloud? Storage, computing, everything?

Irina Starikova: I love this question. Let me explain what I mean by that. By the year 2020, which is not that far away, I can see that up to 80 percent of enterprise applications can be in the public cloud. Whereas the remaining 20 percent would be in their own data centers in the private cloud because of legacy, cost, or security reasons. What I also believe is that the 20 percent might be even a smaller figure for some companies in nonregulated industries.

What I am also fascinated by is learning stories about digital-born companies, so those companies that have existed for ten years or less. When you ask about how they're doing their infrastructure and what they're doing with cloud, you almost never hear that they're building their data centers. They have all embraced the public cloud as just the right thing to do.

They, frankly, are saying, "This is not our competency. Why would we build our own electrical power station? No one does that anymore." Similarly, we see those companies completely move away from the concept of building infrastructure by themselves. They have clearly stated that they will not own their own data centers.

Roberta Fusaro: For the companies that do own their own data centers, what lessons can they take from digital-born companies and other leaders that have kind of gone in another direction?

Irina Starikova: The four big lessons that we've learned from the leaders in cloud adoption from our survey are all about building organizational capabilities rather than technology.

The first one is, focus on the migration road map and focus on getting meaningful migration results, basically executing on your plan. The second one is to look for ways to improve the experience for application-development teams, iterating on that as you go, because you will never get it right the first time. The third lesson is around being very clear on the business case and understanding, as you go with the migration, how that business case is realized and what kind of incremental decisions are changing that business case or helping you to realize the benefits you went after from the get-go. The final lesson learned is around understanding the operating-model implications of using the cloud services at scale. There are really huge implications on what kind of skill sets are required, how different teams within your IT department would operate with each other and with the business units.

The cloud leaders in our research have embraced and have done a lot against all of those four areas.

Roberta Fusaro: I had one last question about supporting a cloud-operating model. I'm just wondering, how hard or how easy is it for companies to make that wholesale change? And what are some key questions that executives need to ask themselves if they're thinking about making this journey?

Irina Starikova: That's a great question, Roberta. This is, frankly, one area where we've heard from a lot of companies we've been working with—that operating model is the hardest thing to get done right when migrating to the cloud at scale.

Even companies that anticipated that it would be hard were surprised by how much harder it was than they initially thought. What we are talking about here is that you not only change the skill sets quite fundamentally, you are rescaling a big portion of your infrastructure teams. You're also changing some of the processes: what those folks are working on day to day and how they interact as well as how they are working with other teams inside IT.

Roberta Fusaro: That's interesting, because you think of the term "cloud" as being very ethereal, right? But the actual work on the ground, there's a lot of nuts-and-bolts tactics that executives need to be involved with in order to adopt enterprise cloud and be successful with it.

Irina Starikova: Yes. None of those changes happen in a short period of time, either. ◆

Irina Starikova is a partner in McKinsey's Silicon Valley office. **Roberta Fusaro** is a member of McKinsey Publishing and is based in the North American Knowledge Center.

About Digital McKinsey

We help imagine and deliver digital reinvention by bringing together the best of McKinsey's digital capabilities. We work with clients to first uncover where meaningful value exists and then create and implement the right solution—from building a new business to developing an IT architecture to delivering a customer experience.

Digital McKinsey brings together more than 2,000 experts from across our global firm—including more than 1,500 developers, designers, IT architects, data engineers, agile coaches, and advanced-analytics experts.

For more information, visit DigitalMcKinsey.com.

Digital/McKinsey

December 2018
Designed by Global Editorial Services

McKinsey.com
@digitalmckinsey
linkedin.com/showcase/digital-mckinsey/
facebook.com/DigitalMcKinsey/