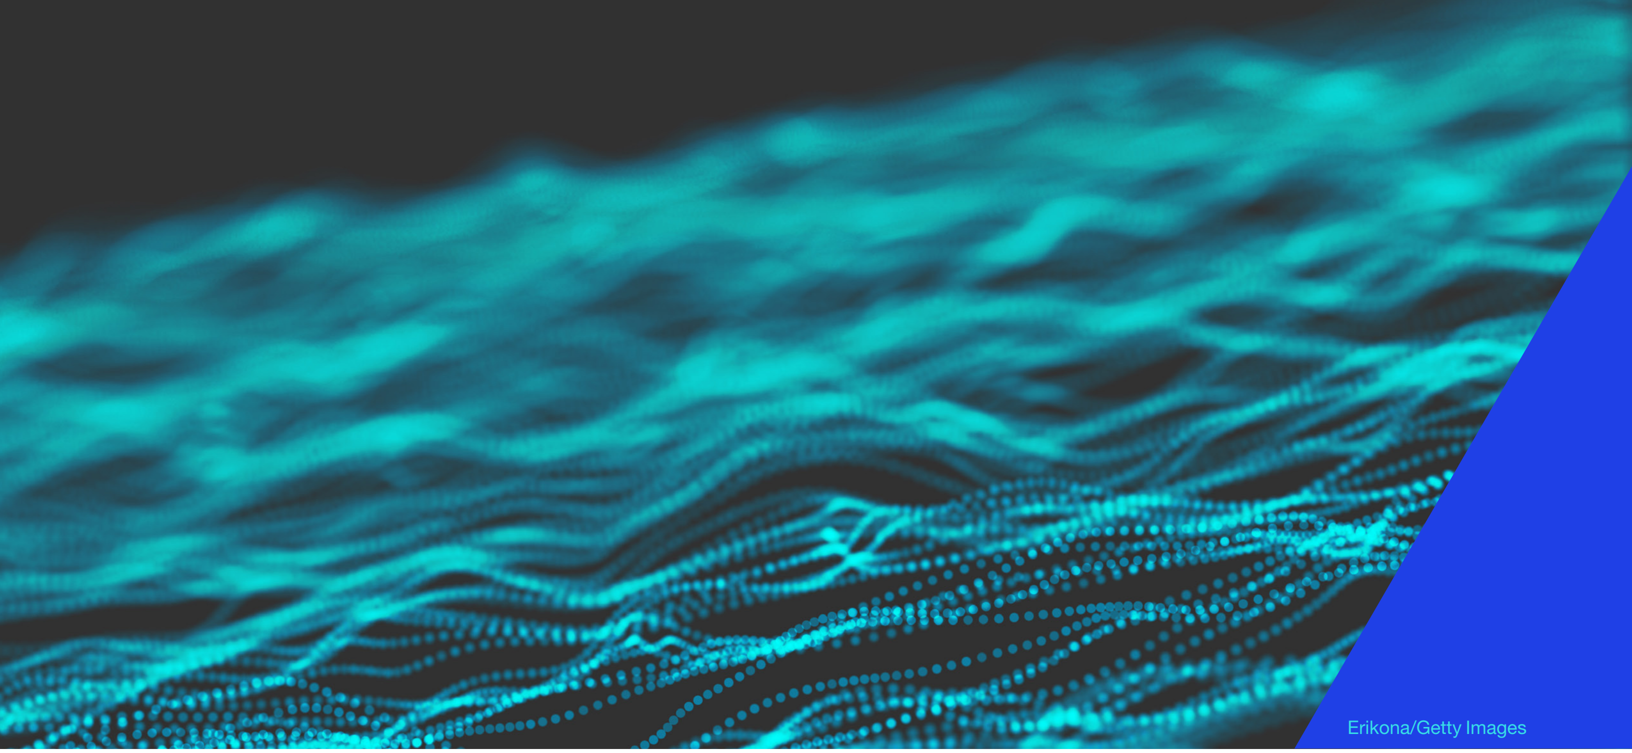
Erikona/Getty Images

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 to 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 into 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 back-hauling 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 clean sheet 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 clean-sheet topology. For example, an insurance company used a private-front topology to move some sensitive applications into the public cloud without having to overhaul their 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 back hauling

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

Public front or CSP default[1]

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

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.

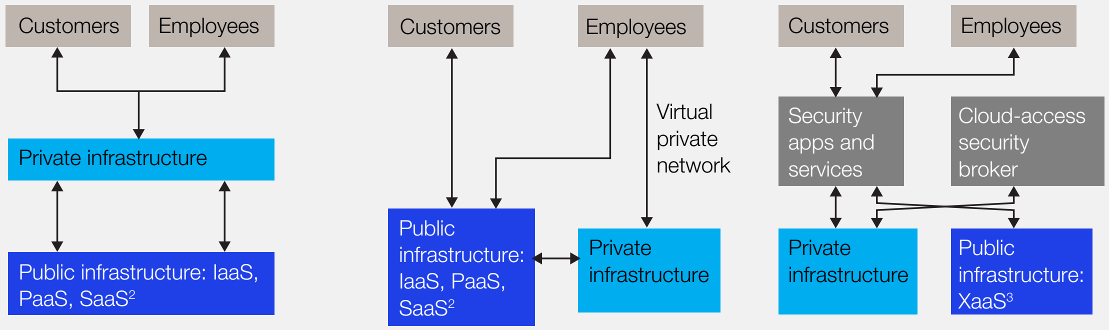
Customers
Employees
Private infrastructure
Public infrastructure: IaaS, PaaS, SaaS[2]
Customers
Employees
Virtual private network
Public infrastructure: IaaS, PaaS, SaaS[2]
Private infrastructure
Customers
Employees
Security apps and services
Cloud-access security broker
Private infrastructure
Public infrastructure: XaaS[3]

- Known and established security mechanisms
- Simplified monitoring and debugging
- Quick implementation
- Higher costs due to increased traffic

- Lowest-cost approach, but limited to offerings from cloud-service providers
- Potentially creates gaps when limitations are not understood
- Greater scalability

- Uses multiple solutions
- Enhanced user experience[4]
- Requires deep expertise in cybersecurity and cloud architecture, increases 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 it 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. 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 also benefit from assigning experts in cloud technology and pricing to model the costs of their technology stacks and operating models, and 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 questions 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.

Potential run-rate savings in each topology, % of current total per image[1]

🟦 Nonlabor savings ■ Labor savings

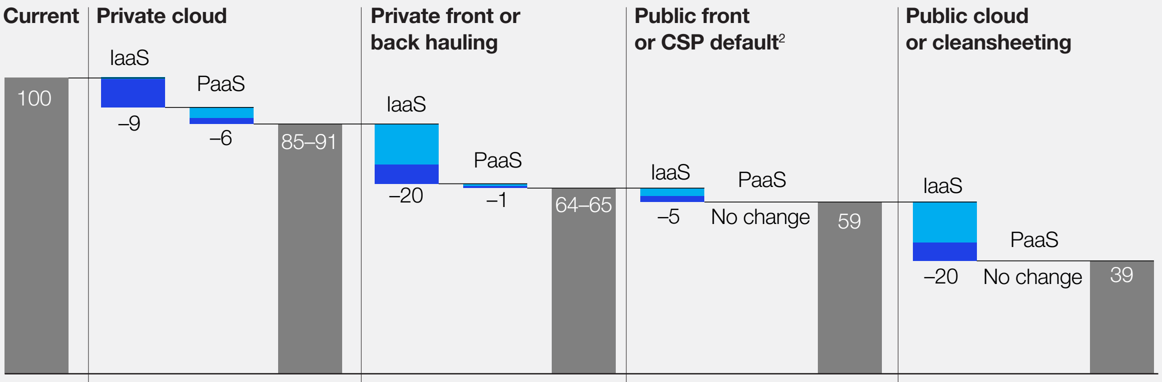

Private cloud

Sources of cost savings

Infrastructure as a service (IaaS)

- Nonlabor: compute and storage hardware (partly offset by added software)
- Labor: automation

Platform as a service (PaaS)

- Nonlabor: decrease from open-source licensing
- Labor: decrease from middleware and database automation

Private front or back hauling

Sources of cost savings

IaaS

- Nonlabor:public-cloud pricing and real-estate reductions[3]
- Labor: fewer compute/ storage touches

PaaS

- Nonlabor: middleware and database-as-a-service pricing
- Labor: effort eliminated by middleware as a service, by database as a service, and by database automation

Public front or CSP default

Sources of cost savings

IaaS

- Nonlabor: decrease in networking costs
- Labor: decrease in networking labor

Public cloud or cleansheeting

Sources of cost savings

IaaS

- Nonlabor: eliminate all networking assets and remaining real estate (except mainframe footprint)
- Labor: decrease in networking labor

[1] 100% = $20,000 per environment ($246 million infrastructure baseline; 12,500 images). Saving scenarios range from low (only using infrastructure-as-a-service solutions) to high (using infrastructure-as-a-service and open-source platform-as-a-service 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.

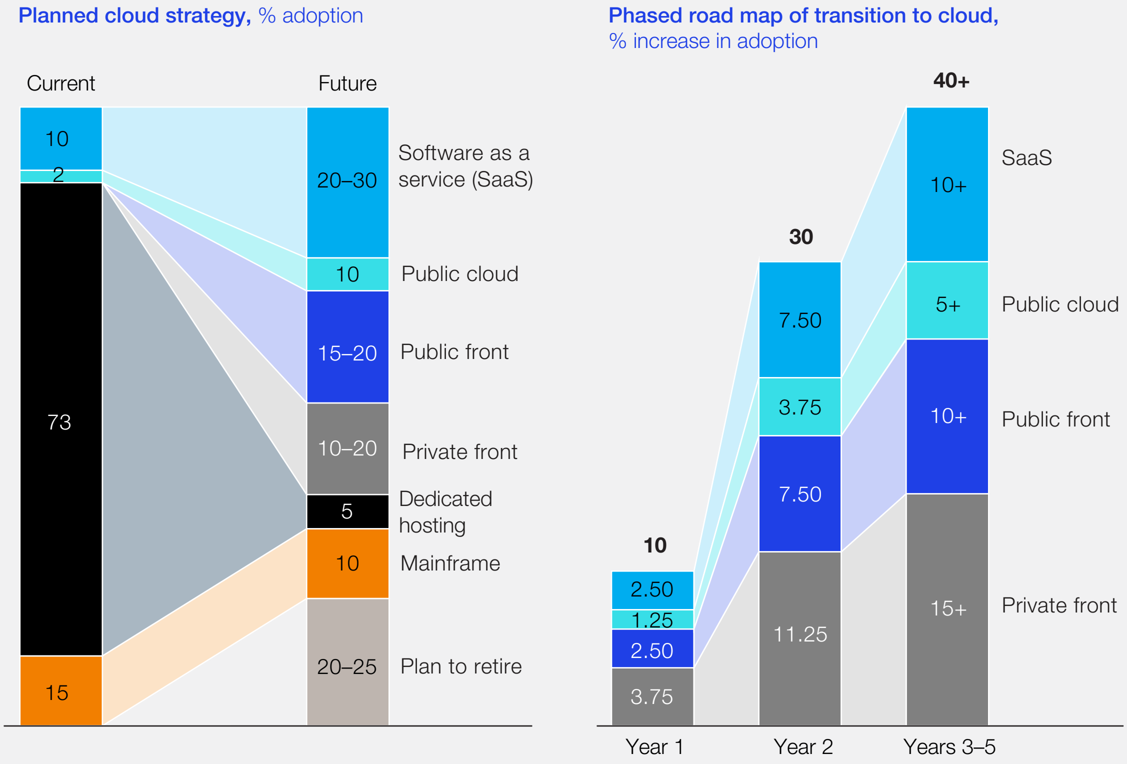

Planned cloud strategy, % adoption

Phased road map of transition to cloud, % increase in adoption

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.

Copyright © 2018 McKinsey & Company. All rights reserved.