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 last 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 fully exploit 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.

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.

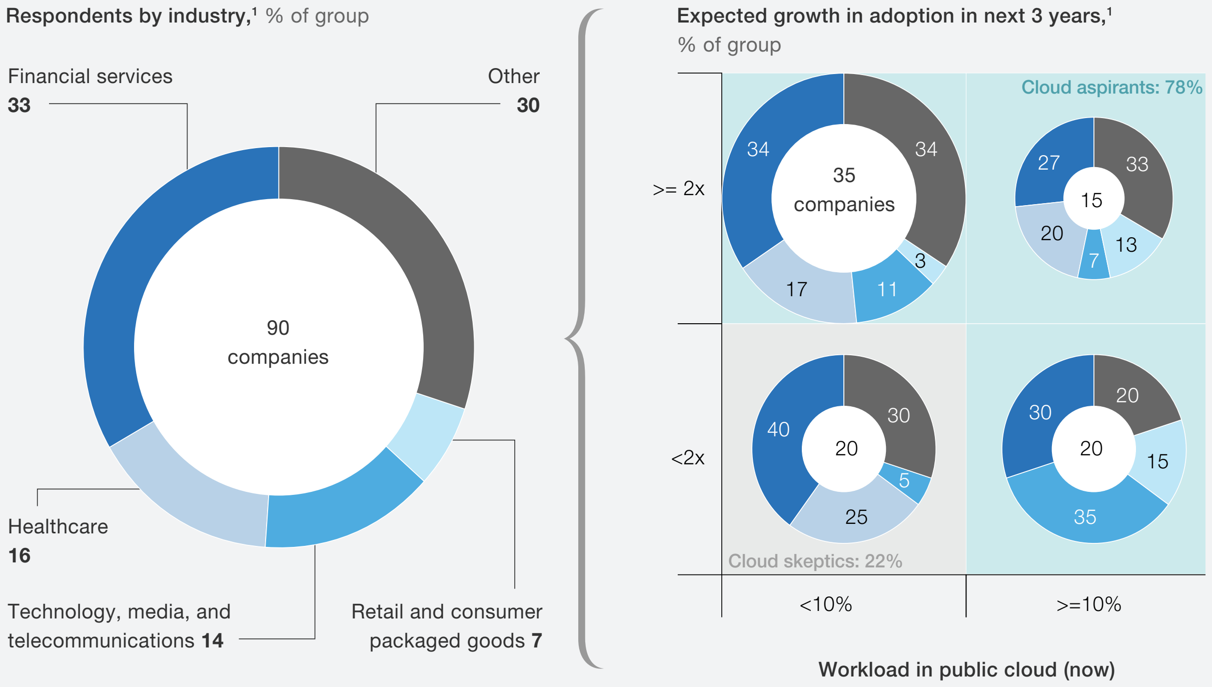

[1]Percentages may not sum to 100% due to rounding.

McKinsey&Company | Source: McKinsey global cloud cybersecurity research, 2017

[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 30 percent are in other industries.

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 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 cybersecurity, compared to what providers will do.** 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.

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

- **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 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.

The most effective approach is to reassess the company's cybersecurity model in terms of 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 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).

Exhibit 2

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

Provider of perimeter-security control

■ Enterprise ■ Cloud-service provider (CSP) ■ Third party

Backhauling: All public-cloud access is through private infrastructure with external gateway.

Adopting CSP controls by default: CSP controls for public cloud only. Separate private security controls.

Cleansheeting: Best-of-breed security controls for public cloud and private cloud.

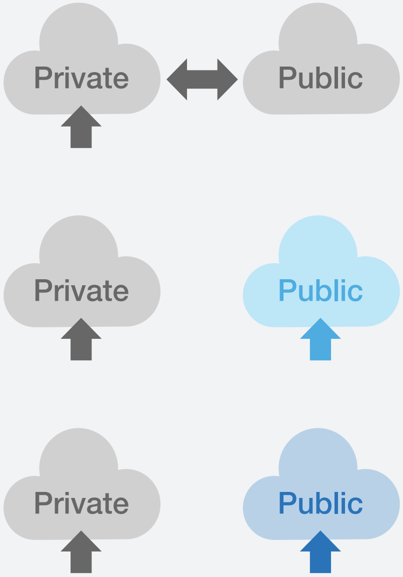

- **Cleansheeting.** Cleansheeting involves designing a “virtual perimeter” and developing cloud-specific controls from solutions offered by various external providers. Used by around 15 percent of cloud-aspirant companies, this approach enables companies to apply the best perimeter-security solutions 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.

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 CSP 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.”

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 encrypting 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. 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, companies need not apply the same archetype to their entire public-cloud profile. It’s possible, even advantageous, to use different archetypes for applications

Exhibit 3

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

Performance of archetype against evaluation criteria

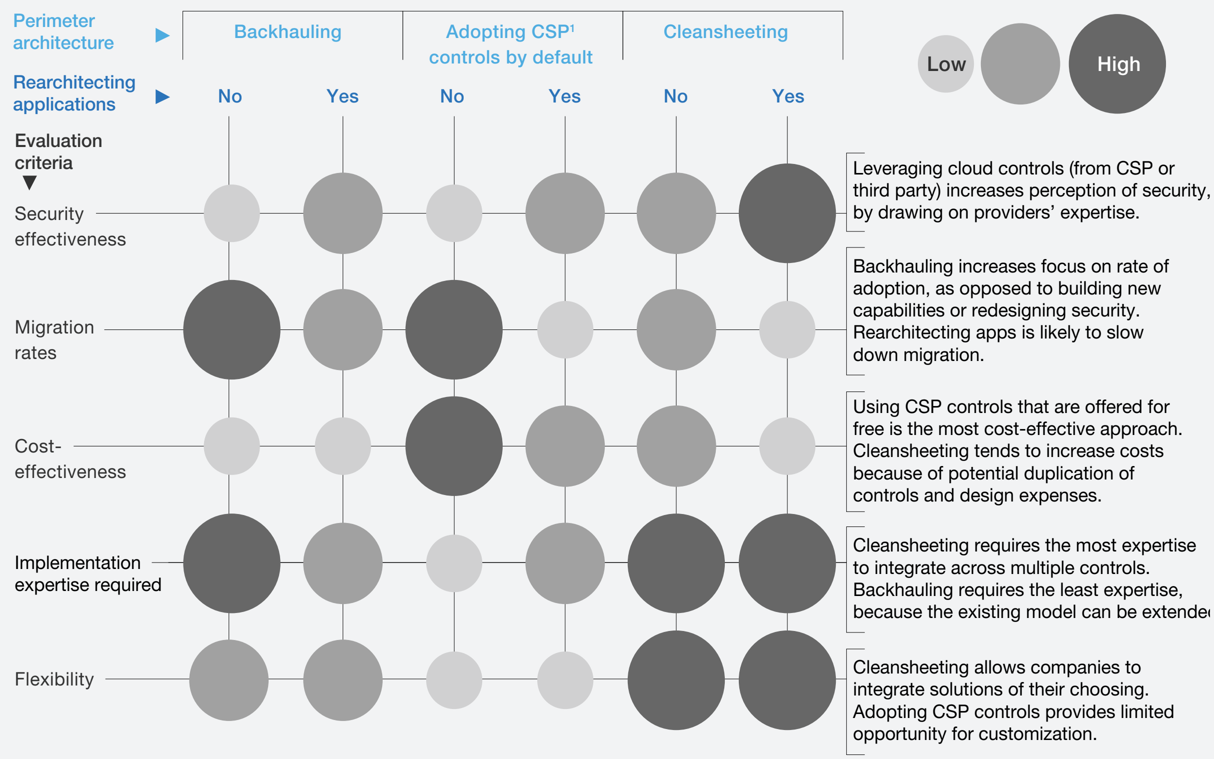

Leveraging cloud controls (from CSP or third party) increases perception of security, by drawing on providers' expertise.

Backhauling increases focus on rate of adoption, as opposed to building new capabilities or redesigning security. Rearchitecting apps is likely to slow down migration.

Using CSP controls that are offered for free is the most cost-effective approach. Cleansheeting tends to increase costs because of potential duplication of controls and design expenses.

Cleansheeting requires the most expertise to integrate across multiple controls. Backhauling requires the least expertise, because the existing model can be extende

Cleansheeting allows companies to integrate solutions of their choosing. Adopting CSP controls provides limited opportunity for customization.

[1] Cloud-service provider

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. The eight control areas are listed below, along with observations from our research.

- **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.

Moving into the next generation of IAM

A Fortune 500 healthcare company we spoke with has redesigned its 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 CSP-provided identity services, 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."

- **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]

- **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 web gateway, web application firewall, and network monitoring from different third parties that support multiple clouds.

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

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 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."

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

- **User end points.** Moving workloads onto 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 apply an appropriate security model archetype (such as backhauling or cleansheeting) to determine the best security controls and their scope. 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 shortlist and prioritize controls that can be standardized and automated, and implement them in agile iterations.

Clarifying internal responsibilities for cybersecurity, compared to 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 them 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 and 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 their logs readily available in a format 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 they 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 receives 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 methods requires companies to foster a culture 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 also helps 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.

Exhibit 4

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

Cloud-deployment process with secure DevOps

Architecture & design
Enhancements:
- Developers with architecture-security expertise design more secure architectures from project inception
- Architectures are approved for implementation faster, without the need for security team's oversight

Implementation
Enhancements:
- Developers with secure coding expertise introduce fewer vulnerabilities
- Modular security components "snap in," without separate design and implementation
- Milestones achieved faster, without the need for security team's oversight

Security challenge eliminated: No need for design, implementation, and code reviews to be performed by developers with specialized security knowledge

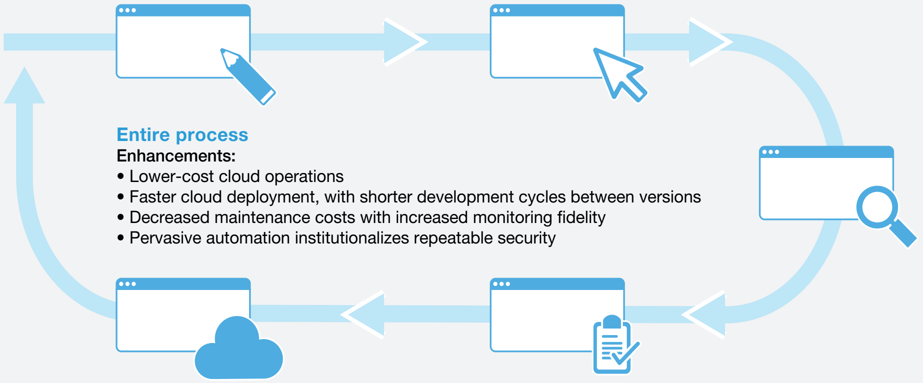

Entire process
Enhancements:
- Lower-cost cloud operations
- Faster cloud deployment, with shorter development cycles between versions
- Decreased maintenance costs with increased monitoring fidelity
- Pervasive automation institutionalizes repeatable security

Code review
Enhancements:
- Secure code scanners conduct automated code reviews for common vulnerabilities
- Developers with secure coding expertise locate and eliminate vulnerabilities before they can be accepted into code base

Deployment
Enhancements:
- APIs for cloud-environment creation include functions to specify secure configuration
- Configurations are done securely by default, with strong encryption and authentication pre-selected

Testing
Enhancement:
- Security test cases are created and automated by the team's own developers, without the need for outside assistance from the security team

Security challenge eliminated: No need for separate testing, because cloud environments are configured to security standards by default and instrumented before deployment into products

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 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 process. 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 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 to integrate them with other existing solutions.** This requires the company to gain a full understanding of CSP's 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 to 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 onto 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 into 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 specialist 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 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.