

Modernizing IT for digital reinvention

Digital McKinsey: Insights
July 2018



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, Belinda Yu

Cover photo:
© Mehau kulyk/science photo library/Getty Images

McKinsey Practice Publications

Editor in Chief: Lucia Rahilly

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

Copyright © 2018 McKinsey & Company. All rights reserved.

This publication is not intended to be used as the basis for trading in the shares of any company or for undertaking any other complex or significant financial transaction without consulting appropriate professional advisers.

No part of this publication may be copied or redistributed in any form without the prior written consent of McKinsey & Company.

Table of contents

Introduction



Modernizing IT for digital reinvention

Features




IT's future value proposition




Toward an integrated technology operating model



Perpetual evolution: The management approach required for digital transformation



Ten trends redefining enterprise IT infrastructure



Learning from leaders in cloud-infrastructure adoption



Making a secure transition to the public cloud



Five ways to unlock win-win value from IT-services sourcing relationships



Why you need a digital data architecture to build a sustainable digital business

In brief



Executives on IT modernization



Modernizing IT for digital reinvention

IT's future is as a source of innovation, not just operations and enablement. That point came through loud and clear in a recent McKinsey Global Survey of executives. At present, though, few respondents think their IT functions make meaningful contributions in areas that promote strong business performance. To get there, IT needs to modernize.


Part of the problem is the excessive complexity of incumbents' technology organizations and processes. The bulky IT architectures that enterprises have assembled over decades can complicate efforts to install advanced technologies. Waterfall-style IT processes prevent established companies from updating vital capabilities as quickly as customers might like. And when incumbents attempt to modernize IT systems, they sometimes pursue changes haphazardly, rather than prioritizing activities to yield maximum benefits.

Overcoming these impediments and habits is seldom easy, but some established companies are showing the way. At a technical level, businesses can start by developing a comprehensive view of their digital capabilities and linking them with an enterprise

architecture that is built for continual updates, or what we call "perpetual evolution." In many cases, they can also benefit from tapping into cloud services and setting up flexible, multilayered data architectures.

Companies should also pay attention to the human aspects of IT modernization. Boosting the technology function's effectiveness involves synchronizing the operations of new digital and legacy IT teams so that they collaborate effectively on complex, high-value technology projects. Businesses can also reframe their relationships with IT-services providers to make full use of their know-how and innovative capacity. That helps IT managers capitalize on important trends and guard against emerging risks.

Success in the age of digital disruption will belong to companies that digitize their core businesses, launch new business models, and apply state-of-the-art technologies—activities that are possible only with sophisticated IT architectures and well-coordinated technology teams. In this edition of *Digital McKinsey Insights*, we plot the moves that IT organizations can take to reach the leading edge.



Jean-François Martin

IT's future value proposition

Naufal Khan, Jason Reynolds, and Christoph Schrey

Many executives expect IT will play a growing role in driving business results, according to a new survey. For that to happen, CIOs must broaden their profiles and prove IT's effectiveness in areas such as digital and innovation.

IT is poised to play a new, more strategic role in companies, one that moves beyond support to create business value through technology-based business innovation and digital initiatives. But according to a McKinsey Global Survey on business technology,¹ IT organizations continue to struggle with performance issues, both in conventional IT and in areas that are critical for the future. As a result, technology leaders aren't often the clear owners of technology-related activities

and capabilities, and many respondents—especially on the business side—see their IT organizations as replaceable by third-party providers. For IT and its leaders to become business partners, the results suggest that CIOs must raise their skills and influence within the organization, leverage technology to move the business's innovation agenda forward, and address the strategic, operating-model, and talent problems that underlie IT's ineffectiveness.

¹ The online survey was in the field from October 11 to October 21, 2016, and garnered responses from 709 participants. Of these, 395 have a technology focus, and the remaining 314 are C-level executives representing other functions. The participants represent the full range of regions, industries, company sizes, and tenures. To adjust for differences in response rates, the data are weighted by the contribution of each respondent's nation to global GDP.

A shifting value proposition for IT

According to the results, many executives—both in and out of IT—expect IT’s value proposition to change meaningfully in the coming years (Exhibit 1). Currently, the largest shares of respondents say IT creates the most value through more traditional business enablement and operational support. But they predict that in the next few years, technology will drive business results. Respondents are most likely to expect that IT will contribute most through innovation and through integrating technology solutions in support of business results—a dramatic shift from IT’s current role.

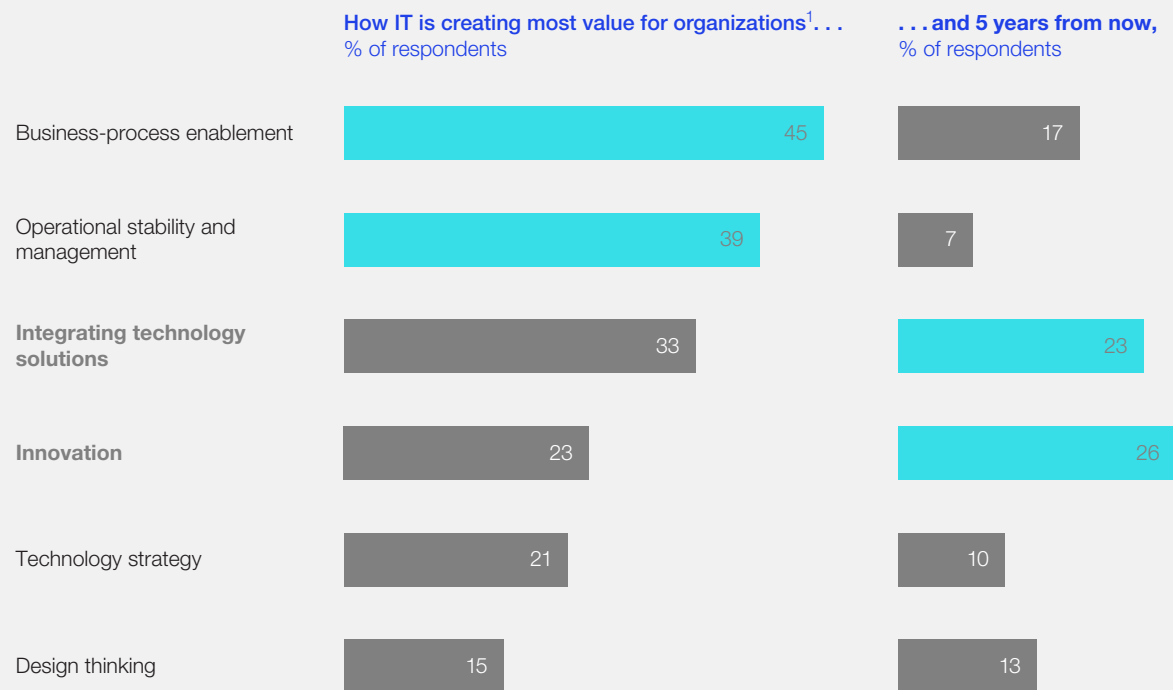
As part of that, respondents also believe that IT should be playing an important role in shaping strategy around digitization. Roughly 80 percent say that business and technology should collaborate on digital strategy, compared with only 55 percent who say they do so now.

Other responses reinforce the merit of IT’s contributions, both current and potential. When technology leaders are involved in shaping business strategy, IT’s ability to create value is greater. As we’ve seen in previous surveys,² respondents report greater IT effectiveness

² Pedja Arandjelovic, Libby Bulin, and Naufal Kahn, “Partnering to shape the future—IT’s new imperative,” May 2016, McKinsey.com; “Why CIOs should be business-strategy partners,” February 2015, McKinsey.com.

EXHIBIT 1

Executives expect that IT’s value proposition will shift dramatically, away from enablement and operations and toward integration and innovation.



¹When asked about the ways that IT creates the most value currently, respondents could select up to 2 options; when asked about the value they expect IT will create 5 years from now, they could select only 1 option. Respondents who answered “don’t know” are not shown; n = 709.

when their CIOs are very involved in overall business strategy. With engaged CIOs, digital initiatives do better as well: 43 percent of respondents with very involved CIOs report significant business impact from their digital initiatives, compared with 23 percent of all others who say the same.

IT struggles to perform, and its future is uncertain

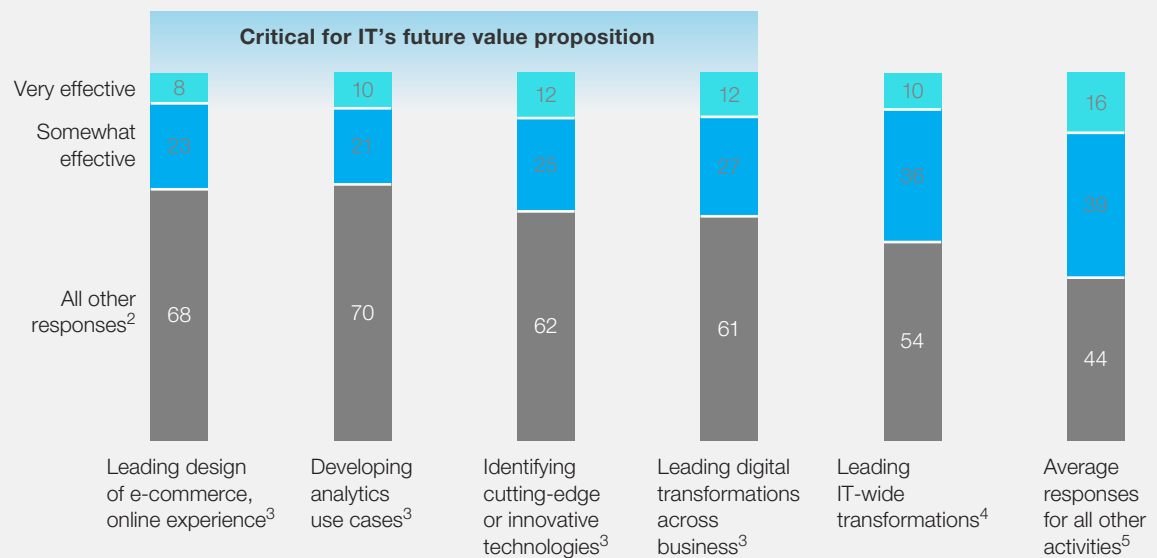
Despite this opportunity for IT, this year's results continue a long-standing pattern of performance concerns in the IT function, even among IT respondents themselves. What's worse, perceptions are especially negative in

the areas that are most critical to IT's future value proposition (Exhibit 2). Just 12 percent of all respondents say their IT organizations are very effective at leading digital transformations across their business, and only 8 percent say IT is very effective at the design of e-commerce and online experience. When organizations have undergone major IT transformations (the modernization of infrastructure, for example), few business leaders have even noticed. Fifty-one percent of IT respondents report having undergone major transformations in the past two years, while just 36 percent of their business peers say the same.

EXHIBIT 2

Few respondents characterize IT's performance as very effective, especially in areas that are critical to IT's future value proposition.

IT organizations' effectiveness at functional activities, % of respondents¹



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

² Includes respondents who said "neutral," "somewhat ineffective," "very ineffective," "don't know," or "not applicable; our IT organization isn't involved."

³ n = 709.

⁴ n = 387.

⁵ Includes 6 additional activities, out of 11 the survey asked about. For these 6, the shares saying "very effective" range from 17% to 34%. Three of the activities ("leading transformations within IT organization," "managing enterprise-wide data architecture/infrastructure," and "managing IT's performance") were asked only of IT respondents.

For some of these same capabilities that will be critical to IT's future, respondents also report a lack of clear ownership. Executives are more likely to say there's no clear owner for activities such as e-commerce design and technical delivery than they are to say these capabilities are the CIO's responsibility. With regard to who should be leading these activities in order for the organization to use technology most effectively, neither the CIO nor the CTO is cited by a majority of respondents.

Moreover, many executives can imagine replacing IT with external vendors or service providers. About one-third of all respondents, including 43 percent of business leaders, describe IT as significantly or fully replaceable by vendors and third-party providers (Exhibit 3). Since the previous survey, the gap between business and IT executives who say technology is substitutable has also grown considerably.

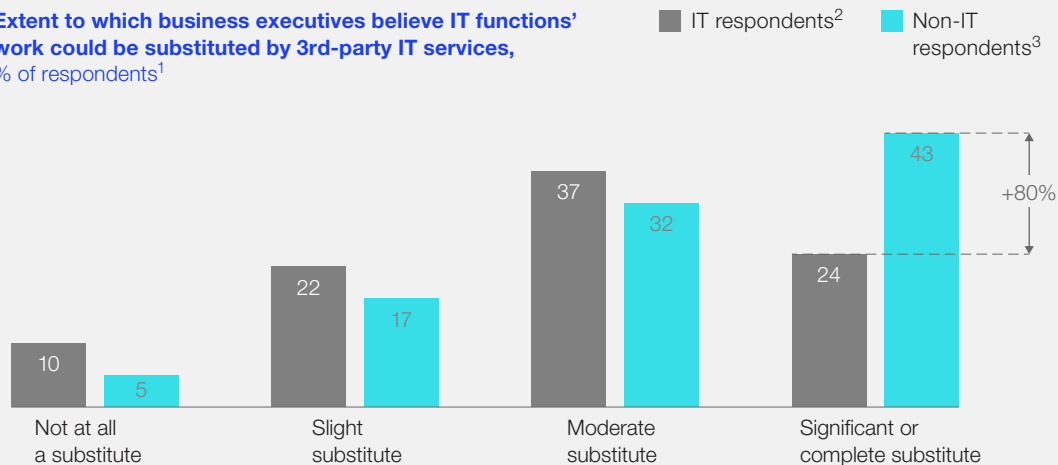
And while nine in ten respondents agree that, over the next five years, their central IT organizations will undergo some fundamental changes (for example, a change of 30 percent or more to overall budget or resources), the jury is still out on what, exactly, those differences will be. Respondents are nearly as likely to expect that the IT function's responsibilities will increase as technology becomes more central to the overall business as they are to predict that the business side will execute most of the work that IT does now.

An imperative to improve

For CIOs and technology leaders to strengthen IT's value proposition and relevance in the digital era, they must make meaningful contributions to growth and innovation. The results suggest that a greater leadership role for the CIO and improved alignment and ways of working are critical to this success.

EXHIBIT 3 More than four in ten business executives believe IT is significantly or fully replaceable by third-party services.

Extent to which business executives believe IT functions' work could be substituted by 3rd-party IT services, % of respondents¹



¹ Respondents outside of IT were asked to share their own perspective on IT's potential to be replaced by 3rd-party work. Respondents in IT were asked to answer on behalf of business executives at their organizations, and how they believe their business counterparts perceive IT's potential to be replaced. Those who answered "don't know" are not shown.

² n = 387.

³ n = 322.

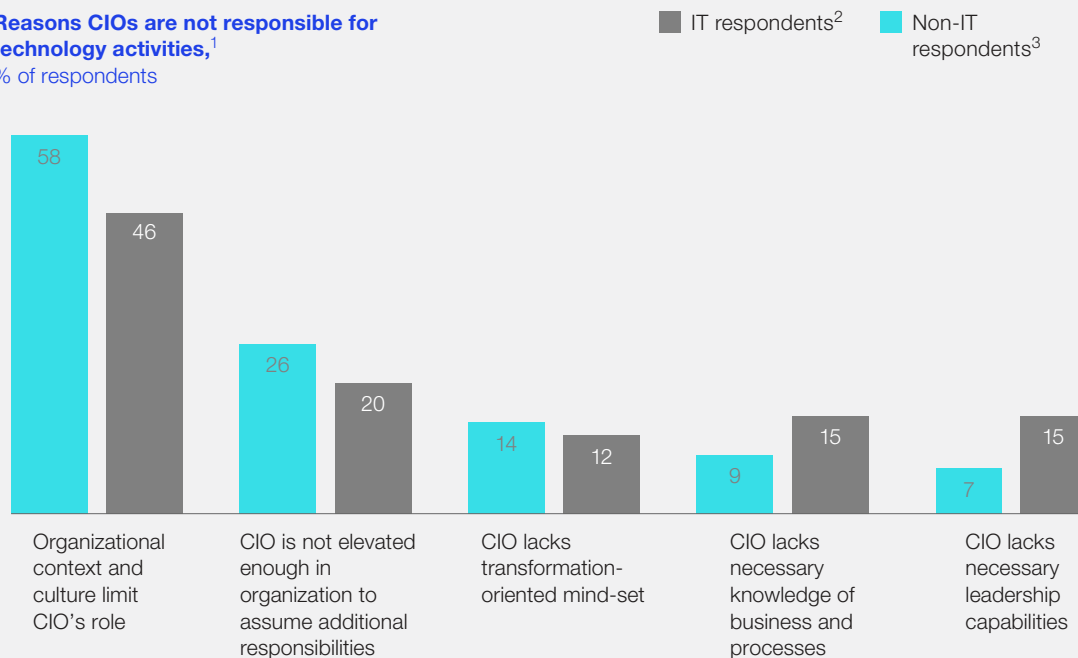
First, CIOs need to establish themselves as genuine business leaders and partners. At organizations where CIOs don't have responsibility for key technology activities, IT and non-IT respondents alike tend to say it's because the organizational context and culture limit the CIO's role (Exhibit 4). They are much less likely to say the main reason is a lack of leadership skills or limited knowledge of business processes, though business respondents identify these factors almost twice as often as their IT peers do. What's more, in several surveys now, we've confirmed the importance of CIO leadership. The more involved a CIO is in shaping overall business

strategy, the better the IT function performs, both overall and on digital strategy (Exhibit 5). The latest results show that elevating the CIO's role, both structurally and culturally, could be the key to achieving this. When the CIO reports directly to the CEO, rather than the CFO or other senior roles, respondents are 2.5 times likelier than others to say their CIOs are very involved in company strategy (Exhibit 6).

Second, the root causes of IT's ineffectiveness must be addressed. According to IT respondents, the most significant problems are a lack of clear priorities for the IT function, weakness in IT's operating model, and talent

EXHIBIT 4 Business and IT respondents agree on what holds back their CIOs: organizational culture and a limited role.

Reasons CIOs are not responsible for technology activities,¹
% of respondents



¹This question was asked only of respondents who cited a role other than the CIO as owner of at least 1 of the following technology activities: design of e-commerce and online experience, technical delivery of e-commerce and online experience, developing analytics use cases for insight generation, identifying cutting-edge innovative technologies for the business, running organization's online e-commerce business, and digital marketing. Respondents were asked to select up to 2 reasons, and respondents who answered "other" or "don't know" are not shown.


²n = 375.

³n = 291.

EXHIBIT 5

More CIO involvement in the business correlates with an IT function that's more effective, both overall and on digital strategy.

Likelihood that IT organization is very effective,¹
% of respondents



Likelihood that business impact of digital initiatives is significant,²
% of respondents



¹Respondents were asked to rate IT's effectiveness in 11 different areas, and this analysis includes the average rating across these areas.

²This question was asked only of respondents who said their organizations have pursued a digitization transformation (ie, a large-scale change effort that is more comprehensive than a short-term improvement program) in the past 2 years; n = 360.

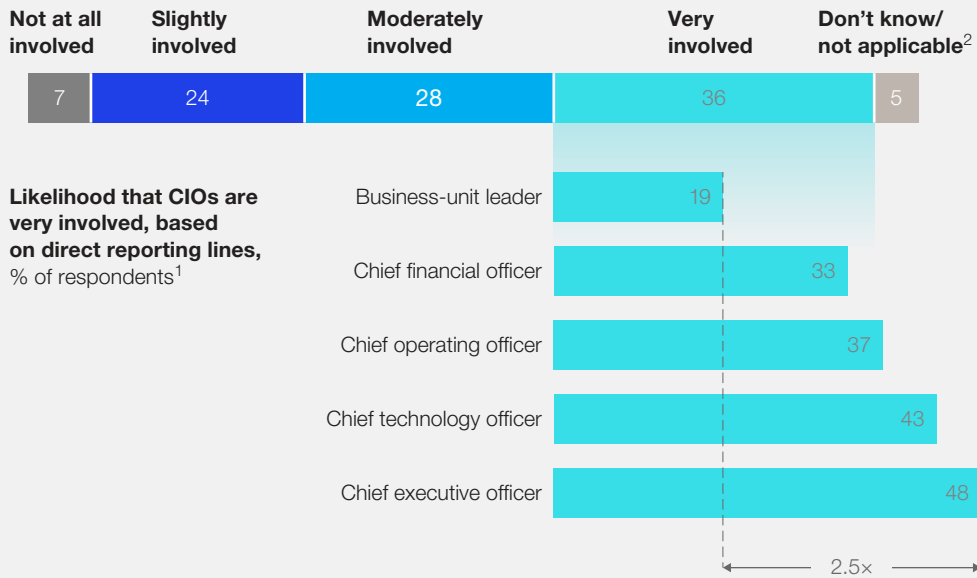
³"Not involved" includes respondents who say their CIOs are not at all involved, somewhat involved, and moderately involved in business strategy at their organizations.

⁴"Involved" includes respondents who say their CIOs are very involved in business strategy.

EXHIBIT 6

Of the CIOs who are involved in business strategy, nearly half report directly to their CEOs.

CIO involvement in shaping organization's business strategy and agenda,
% of respondents¹



¹n = 709.

²Figures may not sum to 100%, because of rounding.

issues. In fact, talent has actually grown as a root cause; respondents were twice as likely to cite talent issues as they were in 2015. With the operating model, the key challenge for CIOs to solve is inefficient governance and work-intake processes. Just after that, IT respondents most often cite weak alignment between business and IT, unclear roles and responsibilities (both of which affect the clarity of IT's priorities), and lack of a hybrid digital-IT operating model (which is needed for digital initiatives and solutions to work). By addressing these ongoing issues, IT leaders and their organizations will resolve three-quarters of the main reasons why IT isn't performing effectively.

Looking ahead

In response to the challenges that the survey results revealed, here are three steps that can help CIOs and IT organizations strengthen their value proposition and contributions to the broader business. These steps are mutually reinforcing, so taking all three together will increase the success and impact of each.

- **CIOs must rewrite their job descriptions.** Despite performance concerns and an uncertain future for IT, CIOs will need to increase expectations for themselves and the IT function. They must also work hard to elevate their role within the organization, developing both their leadership and business muscles while building a more direct reporting line to the CEO. To do so, they will need to write a more ambitious job description that reflects their organizations' broader aspirations for growth and innovation. This could mean taking on newer responsibilities around customer engagement, such as omnichannel design, design and oversight of analytics, and the centralization and automation of core business functions. CIOs will also need to focus on developing

both the functional skills (such as digitization and delivery) and the leadership skills necessary to gain credibility as a true business partner, and they must ensure that the IT organizations they lead are meeting—or even surpassing—expectations.

- **Address nagging causes of IT ineffectiveness.** The results point to three critical areas of IT ineffectiveness—a lack of priorities, operating-model weaknesses, and issues related to talent—on which organizations must make quick progress. The first requires a frank discussion with business leaders to close the gap between perceived and actual priorities. Agreeing on priorities will help IT play a clear, focused role in the organization, ensure visibility and appreciation for the technology-related transformations IT is leading, and let IT leaders shift their time and resources to the areas the business values most, such as innovation and integration. The second—strengthening IT's operating model—has been a top-two cause of poor performance for two years in a row and is especially crucial for organizations pursuing digital transformations. These organizations will need to move to a more unified and flexible operating model to support large-scale digital efforts that will inevitably span disparate technologies (legacy and next-generation) and delivery practices (agile and traditional methodologies). Finally, the search for top IT talent must include new approaches to workforce planning, attraction, evaluation, and development, as well as the culture of the IT organization.
- **Integrate technology across the enterprise.** Another opportunity for CIOs is the role of integrator. Respondents report a wide variety of technology-leadership roles at their organizations, and that technology

is touching upon the work of many business functions. CIOs, then, are in a unique position to observe these activities at their organizations and serve as a central architect to help manage the technology-enabled innovations and capabilities. To do so, they will need to strengthen their own transformation muscles by freeing up change-minded technology leaders from

their day-to-day activities and building transformation-leadership capabilities within their teams. They will also need to connect more closely with committed business partners who understand the long-term journey of transformation via technology and are willing to help navigate the organization through potential disruptions. ♦

The contributors to the development and analysis of this survey include **Naufal Khan, Jason Reynolds,** and **Christoph Schrey**, a senior partner, partner, and associate partner, respectively, in McKinsey's Chicago office.

Copyright © 2018 McKinsey & Company. All rights reserved.



JamesBrey/Getty Images

Toward an integrated technology operating model

Naufal Khan, Gautam Lunawat, and Amit Rahul

Companies may be able to get digital transformations off the ground by separating digital from conventional IT, but that approach is not sustainable. Here's a better way.

Technology organizations are now expected to play a central role in helping companies capitalize on new digital capabilities—connectivity, advanced analytics, and automation, for instance. These capabilities can help them build deeper relationships with customers, launch new business models, make processes more efficient, and make better decisions.

To a greater degree than before, technology groups must focus on integrating these new

digital tools and approaches with existing legacy systems and methodologies—a task that isn't always as straightforward as it sounds. Companies have introduced costly, complicated initiatives designed to deploy digital tools and approaches organization-wide, only to see such programs fall short of their potential or stall completely. The evidence? Rich data sets are accessible only to a few groups of privileged users. Innovative processes used in one business unit are never shared

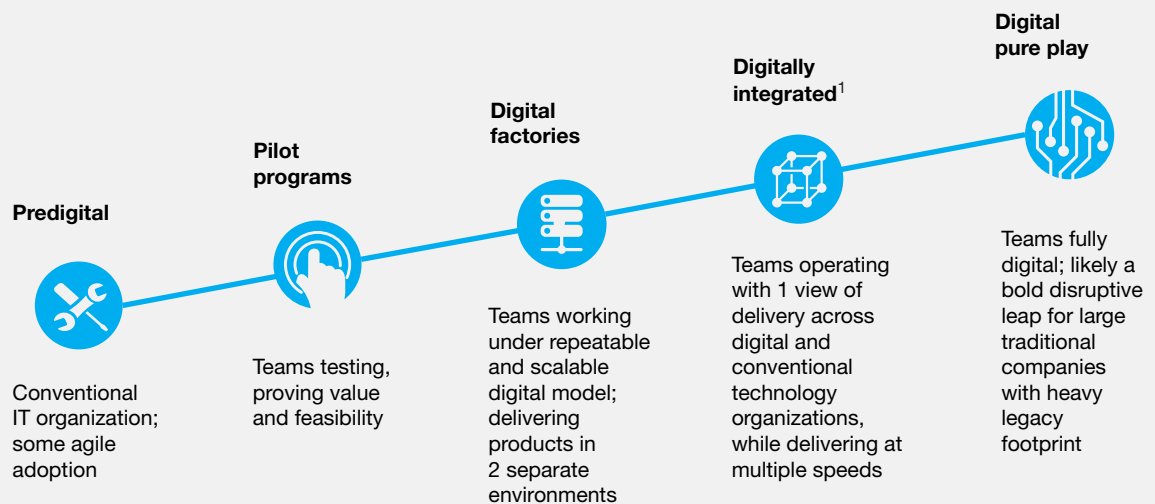
across the company, and the impact of digitization remains small and isolated.

A critical factor in these shortfalls is the lack of a common operating model for digital and IT teams. The digital factory model that most businesses tend to use to launch their digital programs can undeniably speed up a company’s pace of innovation in the short term. Skunk Works digital teams working outside the purview of a conventional IT organization can quickly tackle pilot projects that they can then turn into innovative products or customer experiences. For their part, most senior business leaders often decide to stay the course with this approach, with separate digital and IT units adhering to different operating and service-delivery models. They recognize that

a shift to the “pure play” model of digitization pioneered by the likes of Amazon and other Internet companies might be overly ambitious or disruptive in the near to midterm (Exhibit 1).

In our experience, however, at least 60 percent of the highest-value technology projects companies pursue require collaboration and delivery from multiple technology groups across both digital and IT teams. The lack of a common operating model can thwart such cooperation. What’s more, fragmented technology stacks can put pressure on overall system stability, scalability, and resilience. The physical split between digital and IT groups can create confusion among business stakeholders about which team is handling which tasks. Even within technology groups

EXHIBIT 1 Companies’ digital programs tend to follow a common evolutionary path.



¹Integrated operating model.

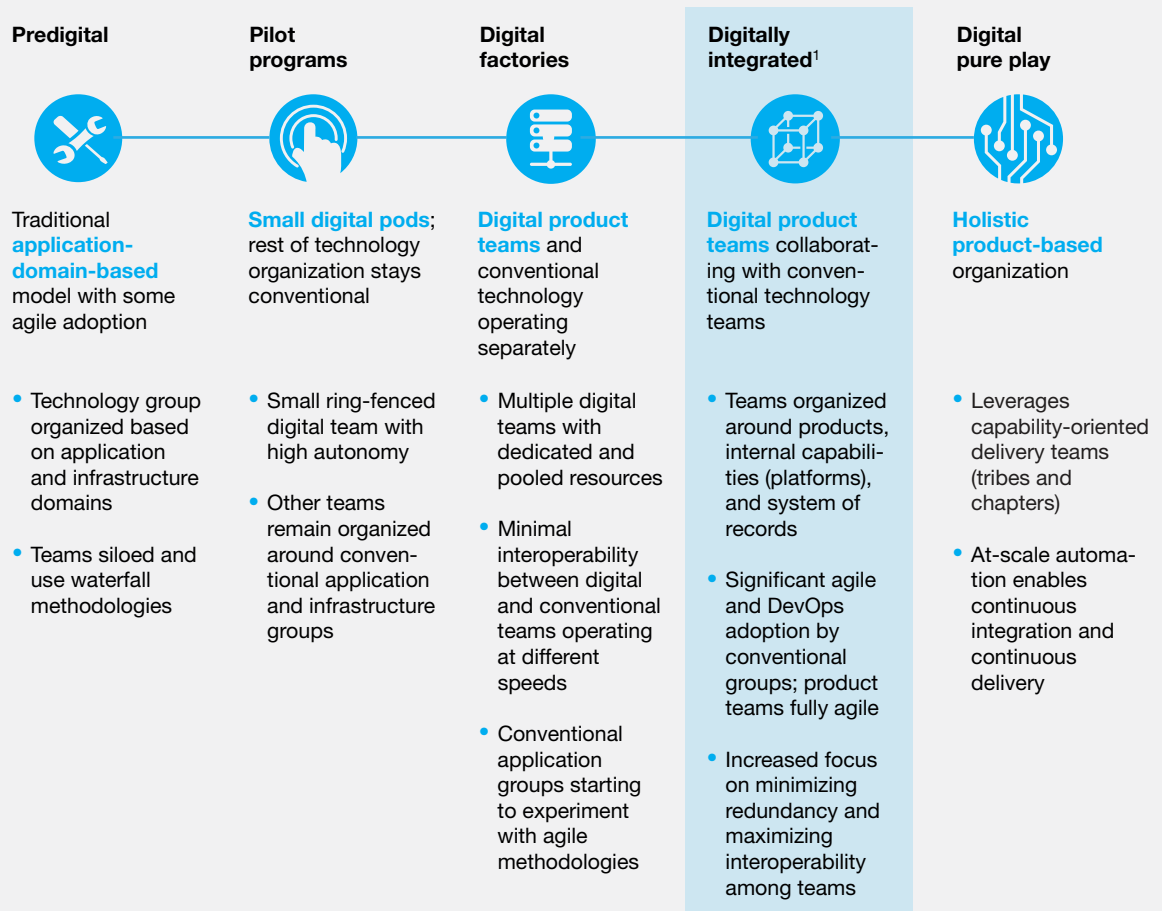
Source: McKinsey analysis

themselves, the culture can become bifurcated as employees identify with either old or new ways of working.

Companies should instead consider shifting to an integrated digital IT operating model in which there is one operating model and

one view of how technology capabilities are delivered by both digital and conventional IT groups (Exhibit 2). Under this model, teams organize around technology capabilities rather than specific technology assets and functions, and they often use agile methodologies to speed up the provision of IT services.

EXHIBIT 2 Companies should pursue an integrated digital operating model for their technology groups.



¹Integrated operating model.

Source: McKinsey analysis

According to our research, companies that pursue an integrated IT operating model can realize greater process efficiencies, often through the elimination of redundant roles and initiatives, and they can deliver products and services to customers more quickly.

A plan for integration

The journey toward an integrated model is neither easy nor quick. It can take years to complete depending on a company's starting point and digital aspirations. It therefore requires a commitment from the business and technology groups (both digital and conventional IT teams) to reconsider existing ways of working and collaborate on devising a new path. Business leaders must show a willingness to "test and learn," and technology leaders will need to become active thought partners to the business units.

Organizations will, of course, need to address issues relating to core technology. For instance, they will need to design flexible, perpetually evolving enterprise architectures, with lightweight connections, that can support the development and deployment of new business capabilities. They will also need to develop agile data-management practices—that is, centralizing the collection and storage of data and allowing employees across the company to access critical business information from multiple systems.

Perhaps the most critical changes associated with making a successful shift to an integrated digital IT operating model, however, are those relating to processes and people—that is, rethinking the composition of the technology organization, the methods for providing IT services, and the management of technology talent. Let's take a closer look at these three factors.

Rethinking the technology organization

To successfully pursue an integrated digital IT operating model, companies should reconsider how digital and conventional technology groups are organized and governed: What processes does each group currently follow, and how could those processes be standardized to ease collaboration? What governance structures do they use, and what modifications could be made to improve decision making? Under an integrated model, the digital and conventional teams would jointly pursue the company's digital agenda and may work under a single overall technology leader—likely from the technology group—to ensure accountability at the top. They would also need to take the following steps:

- **Redefine critical roles in technology leadership.** As part of the integrated organization design, companies will need to redefine leadership roles associated with the construction of products—for instance, product managers and designers, engineers, data managers, and IT architects. New roles may be required. Those in existing roles may need to develop new skill sets and areas of expertise. The nature and extent of those redefinitions will depend on a number of factors, including the company's digital goals, its corporate culture, and its existing technology capabilities. Many leadership roles will likely need to become "hybrids"—incorporating both digital and conventional IT perspectives. A large B2C company undergoing integration of its digital and IT organizations created a role under the CIO called head of consumer technology. This individual is responsible for the development of all digital and conventional customer-facing applications regardless of the channel (online, mobile, and stores).

- **Centralize IT-architecture and IT-infrastructure teams.** In an integrated organization, common resources for digital and IT teams, such as technology architecture and infrastructure, will need to be centralized. By combining the teams managing these resources, companies can eliminate redundant tasks, facilitate standardization of processes, and deliver benefits more broadly to the business units. For instance, one manufacturer is convening an end-to-end technology IT-architecture function that would be responsible for making critical decisions relating to both digital and conventional IT assets. Senior leaders believe this new structure will help prevent system proliferation, a perennial issue for the company, and that it will ensure that new technology capabilities are acquired or built based on company-wide needs, rather than according to business-unit or functional needs.
- **Deploy agile, user-centric product-development teams.** Technology staffers should be encouraged to move in and out of cross-functional product or project teams. These self-organizing teams would come together to offer specific customer and end-user experiences or capabilities and then disband when objectives have been met. The leaders of these teams would work directly with business stakeholders to jointly define priorities and identify areas where technology could significantly enhance business processes. The technology team at an online retailer came up with an idea for enhancing payment processes, and it collaborated with the business team to find funding for the project and to design and build the prototype software that would support the

process change. Pilot tests were mounted quickly, with frequent input from the business, and the full process change was implemented within six months.

- **Revisit funding and portfolio-management processes.** IT organizations' funding and portfolio-management processes would also need significant changes under an integrated model. Staged venture-capital-style funding could be applied to projects that involve both digital and conventional IT team members. Funding decisions for those projects could be contingent upon the integrated teams successfully meeting certain milestones during the development cycle. They could also be tied to business outcomes. Meanwhile, business and technology leaders should jointly review all technology initiatives under way—meeting quarterly or biannually—to ensure balanced investments in initiatives that are critical for supporting day-to-day operations as well as those needed to fuel business innovation and growth. In this way, foundational technology investments, such as the modernization of aging IT platforms, which are nonetheless relevant for supporting end-to-end digital capabilities, wouldn't get lost in all the conversation about cutting-edge technology pilots and experiments.

Rethinking technology provision

IT organizations typically manage three major archetypes of work: purely digital projects (creating a mobile application interface, for instance), purely conventional projects (making enhancements to a mainframe application, for example), or hybrid projects that affect both digital and conventional assets (developing a self-checkout application for in-store customers,

for instance). When digital and conventional IT teams' systems and mechanisms for providing technology support remain separate, hybrid projects may be particularly compromised. Such initiatives can be delayed and deadlines missed when conventional IT teams do not anticipate the number and frequency of changes made by digital IT groups, which are typically operating under the test-and-learn principles of agile development.

An integrated delivery model would ensure joint planning on such projects—involving both digital and conventional IT teams at the very start of the life cycle of a project—which would help reduce delays and create more transparency. Companies could take the following steps to help digital and IT groups find common ground and deliver products and services more efficiently. Some of these actions may seem obvious, but it is surprising how many companies take them sporadically, or not at all.

- **Conduct regular planning sessions** to ensure that digital and conventional IT groups are aware of their commitments to project objectives and deadlines and that all potential risks have been evaluated early on. The IT infrastructure team within a conventional IT group, for instance, could agree to allocate some capacity each quarter to address just-in-time requirements from digital teams (working them in between maintenance tasks).
- **Designate a decision-making body** to help remove bottlenecks for hybrid projects. This is not unlike the job done by a traditional project-management office, which imposes standards and processes to ensure that projects stay on track. Indeed, some companies may choose to rely on

their existing project-management offices to meet this need. But others may install a steering committee of stakeholders from the business units and from digital and conventional IT groups to meet and decide periodically on primary issues and risks associated with hybrid projects.

- **Encourage partnerships among IT-support teams** to address the business units' requests more dynamically. In both conventional and digital IT groups, there are teams whose sole purpose is to support development efforts—focusing on quality assurance, infrastructure management, and production efficiencies, for instance. When these groups adopt an agile mind-set—collaborating early in development phases, for instance, and sharing feedback on product and process iterations—they can reduce the turnaround time expected of them in hybrid projects. One company's digital IT group welcomed representatives from the conventional IT group—members of the infrastructure team—in daily meetings associated with the development of a new web feature. Normally, the digital team would have relied on a ticketing system to communicate with the infrastructure team and set work-flow priorities. Instead, it was able to prioritize and convey its requests directly in the meetings. In doing so, the digital team was able to launch the feature quickly, and service completion time dropped 30 percent.
- **Adopt DevOps capabilities** to reduce digital teams' wait time on components from conventional teams. DevOps is a phrase from the world of enterprise software development used to describe the agile relationship between a

company's software-development and IT-operations teams. The methodology advocates for better and more frequent communication and collaboration between these two groups. Under an integrated operating model, the conventional IT team could use DevOps capabilities to gain easy access to the critical assets needed to automate processes for building, testing, and deploying new products and services. The conventional IT team could make its software code available to the digital team quickly and frequently to match its release cycles, thus increasing the speed of development for hybrid projects.

- **Use microservices** to increase the technology organization's ability to provide cross-unit and cross-application functions. Microservices refers to the development of software applications as a package of independent components, each of which can be deployed on its own or in tandem with others, and each of which runs a unique computing process. Through the use of microservices, conventional and IT groups could take advantage of applications and assets previously available to only one group or the other, and could improve their collaborations on hybrid projects that involve both groups' assets.

Revitalizing your talent strategy

The increasing rate of digitization in companies means nearly every business today must make a radical shift in its talent-management strategies. Companies will need to adapt their cultures in ways that will appeal to both next-generation digital workers, who can bring fresh perspectives and innovation to companies, and conventional IT workers, who often carry with them years of valuable institutional knowledge. Specifically, business and IT leaders should focus on making changes in the following areas:

- **Attracting talent.** Companies will need to evaluate their pools of digital and conventional talent and identify any skill gaps that could hinder the pursuit of their digitization goals. As they begin reaching out to possible job candidates, hiring managers will need to work with recruiters to create tailored roles and customized candidate-vetting experiences. Some companies have established standard hiring archetypes (based on the type of talent being targeted) and then crafted ideal requirements and development journeys for people who fit each persona. Thus, the recruiting and onboarding experience for a developer who is fresh out of college, for instance, would be structured differently from that of an IT architect with more than ten years of experience. Companies may also need to make certain cultural changes to attract a millennial cohort that seems to perform best in less bureaucratic, more innovative environments.
- **Retaining talent.** Companies need to ensure that they have the right elements in place to motivate and retain members of the integrated technology organization. The majority of the technology workforce may perceive digital work to be more desirable, making it difficult to keep conventional IT teams motivated. To keep both sides engaged, businesses may want to establish incentives that reward employees based on the scope of their influence within the technology organization, the impact they are having on business outcomes, and their ability to collaborate across teams. In this way, both digital and conventional IT staffers will be motivated to do their best to ensure high-quality customer experiences and successful business outcomes. At

one company, for instance, digital and conventional IT teams jointly created a real-time analytics product that helped to streamline the customer-purchasing experience. Members of both teams were rewarded equally for their success with this hybrid project.

- **Building capabilities.** One of the core benefits of establishing an integrated technology organization is that employees of all stripes, working side by side under one operating model, will gain a greater appreciation of their colleagues' work. They may also find new advantages and opportunities in both digital and conventional areas—thereby expanding the company's talent pool while ensuring the free flow of ideas. Companies can augment this dynamic further by creating skill-development opportunities where expert practitioners can train and coach workers in real-world assignments. Such programs can go a long way toward reducing the cultural friction between the digital and conventional technology groups.



For those incumbents that are trying to catch up to digital-native companies, digital transformation of core products and processes is essential. But the transformation cannot succeed or sustain momentum when the digital technology group is not integrated with the rest of the technology function. The digital factory model will only take companies so far, especially if they aspire to bring all their technology assets to bear in building innovative customer experiences.

Companies must instead pursue an integrated digital IT operating model. Regardless of the rollout plan, the overarching goal should be to minimize the divide between digital and conventional IT groups, thereby assuring business stakeholders that the integrated teams are supporting common strategic objectives and that they are investing in the systems, processes, and talents that can ensure future success. ◆

Naufal Khan is a senior partner in McKinsey's Chicago office, where **Amit Rahul** is a consultant; **Gautam Lunawat** is an associate partner in the Silicon Valley office.

The authors wish to thank Anusha Dhasarathy, Erez Eizenman, and Jason Reynolds for their contributions to this article.

Copyright © 2018 McKinsey & Company. All rights reserved.



kentoh/Getty Images

Perpetual evolution: The management approach required for digital transformation

Oliver Bossert and Jürgen Laartz

Companies that commit to continually updating their enterprise architectures can deliver goods and services as fast as Internet-born competitors do.

Internet retailers can make crucial changes to their e-commerce websites within hours, while it takes brick-and-mortar retailers three months or more to do the same. Cloud-based enterprise-software suppliers can update their products in days or weeks. By contrast, traditional enterprise-software companies need months.

Why can't established companies move as quickly as their Internet-born competitors?

In part, because they are limited by their enterprise architecture, which is the underlying design and management of the technology platforms and capabilities that support a company's business strategies.

The enterprise architecture in traditional companies typically reflects a bygone era, when it was not necessary for companies to shift their business strategies, release new products and services, and incorporate new

business processes at hyperspeed. Consider that until this decade, mobile devices, the Internet of Things, and big data and analytics platforms weren't crucial for competing in the marketplace. Companies did not have an acute need to continually infuse new IT-enabled business capabilities into their operations.

They do now.

To compete against digital-born companies, traditional companies need to adopt a much different approach to designing and managing enterprise architecture—a model we call “perpetual evolution,” because it emphasizes continual changes to and modular design of business capabilities as well as the technologies behind them. This approach encompasses a range of widely known enterprise-architecture frameworks but links them together in a new way. It compels executives to take a comprehensive view of their digital capabilities and technologies but to manage them in a way that mitigates or removes interdependencies and emphasizes speed. Indeed, our work with companies exploring digital transformations suggests that a shift to the perpetual-evolution model can result in faster product-development cycles and greater operational efficiencies—outcomes that are in sync with customers' expectations.

An enterprise architecture built for perpetual evolution differs from a traditional one in six important ways. When considering business processes and activities, IT and business leaders emphasize end-to-end customer journeys rather than discrete product- or service-oriented processes. They use multiple operating models rather than one. When considering the application landscape, IT leaders design and develop applications to be

modular and work independently rather than being tightly coupled with other applications or systems. The enterprise architecture features a central integration platform that boasts lightweight connections rather than a heavyweight bus.¹ The IT organization deploys an application-development model in which developers and IT operations staffers work closely to test and launch new software features quickly (DevOps). And the general view of information and communications technology is as a commodity rather than a strategic factor (Exhibit 1).

In this article, we compare the perpetual-evolution model with existing approaches to designing and managing enterprise architecture, and we explore what's required to shift to this newer approach. The companies that do can unburden themselves of their legacy business processes and mind-sets. They can build the systems and capabilities required to thrive in this era of digitization, enhanced service delivery, and dramatically reduced software-release cycles.

Comparing old and new management approaches

A good way to understand the evolution of enterprise architecture is to consider how companies have traditionally treated its core elements—business operations, business capabilities, the IT-integration platform, IT-infrastructure services, and the underlying information and communications technologies. How would those elements look different under a perpetual-evolution model?

Business operations

Companies have typically designed their business operations using technologies and methodologies with an eye toward simplifying

¹ A connection layer that contains most of the business logic (or rules of computing).

EXHIBIT 1

Digital transformation requires a different model for managing enterprise architecture.

Traditional model	Elements of enterprise architecture	Perpetual-evolution model
Focus on product- or service-centered processes	Business operations	Focus on customer-centric journeys
Reliance on one operating model	Business capabilities	Use of multiple operating models (working at 2 speeds)
Emphasis on interdependency	Business applications	Emphasis on decoupling applications
Use of heavyweight bus ¹	IT integration platform	Use of lightweight connections
Software development managed centrally	Infrastructure services	Software developers and IT operations jointly build new products and features (DevOps)
Managed as precious asset	Information and communication technology	Managed as commodity
		Technology stack

¹Connection layer that contains most of the business logic (or rules of computing).

Source: McKinsey analysis

internal processes. They may build systems that automate internal transactions such as “order to cash” and “service inquiry to resolution,” for instance, and only update those systems incrementally.

Under a perpetual-evolution model, business operations and digital systems must be designed with an outward-facing view—that is, focused on the customer experience online and offline. Priorities have changed. The customer used to be an element in a product- or company-centered process; now the products and services are an element in the customer journey. To be sure, companies’ inward-focused view isn’t obsolete. Enterprises need to maintain core transactional processes and systems, whether they are accounts payable and receivable, order management, procurement, or something else. And they must also make sure those business processes and technologies remain efficient.

However, businesses’ operations and IT systems must now reflect all phases of and elements within the customer journey—not just the exact moment of purchase. And the experience must be continually updated. Individual companies are becoming part of larger industry ecosystems that are focused on supporting end-to-end customer journeys. In the old world of TV manufacturing, for example, companies designed their business operations and IT systems to follow the product to retailers. Today’s digital TVs have become platforms for manufacturers to provide a range of TV-related services to the home, such as identifying shows consumers might want based on their viewing habits, targeted advertising, and more. As a result, TV manufacturers’ business operations and IT systems must encompass the end user’s TV viewing experience, not just the retailers’ requirements. And because end-user preferences will be ever-changing, business

operations and activities must be adapted on the fly.

Note that B2B companies are not immune to this trend, especially those that embed digital technologies into their products to sell predictive maintenance, performance improvement, and other services—for example, construction equipment, aircraft engines, power turbines, and drilling equipment. Companies' enterprise architecture must be able to support customers for the entire time in which they use products and services, even in real time.

Business capabilities

As we mentioned earlier, until this decade, companies have not had an acute need to continually infuse new IT-enabled business capabilities into their operations—for instance, identifying the product a customer is most likely to buy next. Rather, they introduced these capabilities into their enterprise architectures slowly and periodically. Business applications that support these capabilities, such as enterprise-resource-planning (ERP), product-life-cycle-management (PLM), and customer-relationship-management (CRM) systems, were managed as tightly coupled systems; making changes in one often required making big changes in others.

In today's fast-changing digital world, however, companies must be able to continually improve business capabilities without fear of disrupting entire systems. One way to do so is to group processes and systems into two categories: digital business capabilities that are differentiating for the customer experience, and those that support transactional capabilities. We call this a two-speed architecture, and it is a critical element of the perpetual-evolution

model because it helps companies direct their resources appropriately.²

Consider a retail chain that sells a growing proportion of its products through its website. The company cannot take months to enhance its product-recommendation engine when a digital-born competitor can do that in days or weeks. It must have an architecture that makes business capabilities systems-agnostic. It shouldn't matter, for example, what kind of core systems the retailer has; its new or enhanced product-recommendation approach should be able to be implemented and changed easily. These digital business capabilities become the basis on which to compete in an online world.

IT-integration platform

The first two elements of enterprise architecture we have discussed are focused on front-end operations and activities, whereas the other four involve consideration of the back end of companies' enterprise architectures.

Under a traditional model of enterprise-architecture management, companies' IT-integration platform would typically feature a single heavyweight enterprise service bus. This setup can make it difficult for companies to operate digitally in real time. The number of connections increases exponentially in a digital environment, and when all service calls have to pass through the heavyweight bus, the connection layer can become a bottleneck. So companies can have a hard time, for instance, offering website visitors faster page-loading times. Such delays can represent billions of dollars in lost revenue.

The perpetual-evolution model, by contrast, emphasizes lightweight connections to improve

² Oliver Bossert, Martin Harrysson, and Roger Roberts, "Organizing for digital acceleration: Making a two-speed IT operating model work," October 2015, McKinsey.com.

transmission performance and address the problem of latency—the time it takes for companies to deliver web pages to online customers who demand instant responses at every click. The functional elements of the purchasing experience, such as payment- or promotion-management applications, can be decoupled from one another—although when a change does not affect a single service but the entire platform it can still be managed on a slower development track. In this way, companies can upgrade core applications within CRM, ERP, PLM, and supply-chain-management systems module by module (or service by service) without having to make whole-system replacements. The application-migration process can happen faster, and any risks—of downtime, for instance, or the introduction of system bugs—can be kept to a minimum.

IT-infrastructure services

In most traditional companies, IT-infrastructure services (the hardware, software, and network resources required to support an enterprise IT environment) are centrally managed by an independent team. After application developers and code testers finish their tasks, they turn over their assignments to a production team, whose complex testing and handover processes could delay the delivery of a new system to the market for weeks or months.

Under a perpetual-evolution model, DevOps becomes central to a company's ability to test new digital business capabilities and bring them to market rapidly. The concept of DevOps has firmly taken hold in many companies. It involves bringing together IT developers with IT-operations staffers to codevelop new software products and features. Because both sides have skin in the game—with no organizational siloes or middlemen between them—they can

address problems proactively. Under this approach, companies are seeing increased productivity within their software-development teams, faster release of digital products and services, and improved customer experiences. Our experience suggests, for instance, that companies can reduce the average number of days required to develop code and move it into live production from 89 days to 15 days, a mere 17 percent of the original time.

Information and communications technology

Information and communications technologies (ICT)—the combination of all the company's audiovisual, telephone, and computing networks—have tended to be costly. Companies deployed them carefully as expensive (but necessary) assets. However, advances in connectivity, cloud computing, and other technologies have made it easier for companies to adopt a perpetual-evolution mind-set and model for managing ICT. They can use cloud-technology services, for example, to turn IT into an affordable resource, regardless of company size. Indeed, even start-up companies can get up to speed in their target markets quickly by renting computing power and storage space from cloud vendors. ICT is now a commodity, and prior investments are no longer necessarily a big competitive advantage or barrier to market entry.

Establishing a perpetual-evolution architecture

Managing changes systematically across all elements of the technology stack will enable companies to move to an architecture of perpetual evolution. Most companies, however, still view each as a separate system or capability rather than as critical interconnected components of architecture. We have found five principles to be critical for changing this mind-set:

- Free up development teams from unnecessary dependencies.
- Be consistent; focus on change across all areas of the enterprise architecture.
- Break down silos ...
- ... but maintain a strict separation of the platform team from other teams.
- Recognize that transformation of enterprise architecture must be an ongoing process.

Free up development teams from unnecessary dependencies

Companies must be able to change elements of their digital products and processes quickly, thus keeping up with competitors' ability to generate new and innovative customer experiences on demand. To do that, companies must free up their development teams from unnecessary dependencies (Exhibit 2). They can do this by deploying DevOps models and decoupling applications from larger platforms. Teams would no longer have to wait for sign-offs, handoffs, and preparation of test environments when writing code. Those tasks would be managed within the team, with immediate input from development and operations specialists. Such freedom could help development teams reduce their software-release times from months to hours.

Eliminating dependencies is crucial if companies want to design and sell new digital capabilities to ever-more targeted customer segments, each of which will have different needs. Let's use the example of an auto manufacturer that has embedded digital technologies into its cars that enable customers to make online updates to navigation, infotainment, and other systems.

To ensure perpetual evolution, the automaker needed to design those systems so it could isolate the business capabilities it wants to offer customers—for example, a certain navigation capability or a specific new feature of the infotainment system—and so it could change or update these elements independent from one another.

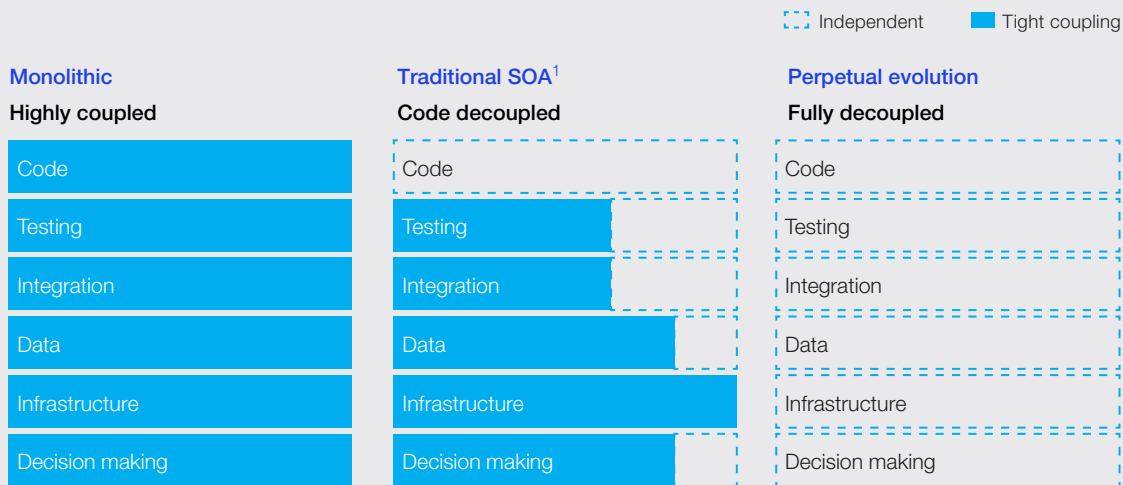
Be consistent; focus on change across all areas of the enterprise architecture

Coding isn't the only place to worry about dependencies. Dependencies also crop up in testing, integration, data, infrastructure, and decision making. By the latter, we mean the individuals who must sign off on the implementation of new business capabilities—is it the team chartered to build and enhance them, or senior management? If, after the capabilities are developed, senior management must approve them before they are put into the marketplace, you can bet it will take those new capabilities a long time to come to market.

Such dependencies are a feature, in effect, of earlier approaches to enterprise architecture. All the elements of the enterprise architecture were tightly coupled. Different modules used the same code base, so a change in one area prompted time-consuming dependency checks to determine how other areas might be affected. The installation of new software depended on the schedules of software testers and resources. Even when developers decoupled software functionality, they often coupled the data, which created dependencies. And when developers intended to decouple the integration layer from applications, teams still too often hardwired business logic into the heavyweight bus, also creating dependencies. When software was ready to move into production, the handover from the development team to the

EXHIBIT 2

The perpetual-evolution model eliminates dependencies among elements of the technology stack.



¹Services-oriented architecture.

Source: McKinsey analysis

infrastructure team often slowed things down. They were now working on the production team’s schedule, competing against a long queue of software releases. Perhaps most important, awaiting senior management’s approval for a new software system or functionality upgrade before it went into production could set things back by weeks.

To be sure, companies’ movement over the past few decades toward services-oriented architecture (SOA) plus a decoupling of code from the other five elements of the enterprise architecture have been major advancements. Companies can now design web services around specific business capabilities. Yet in most companies, the testing, integration, data, infrastructure, and decision-making activities remain tightly coupled. Companies must explore the use of web services so that new software features can be launched independent of any others, and independent of any piece in the IT stack. In fact, their ultimate

goal should be just that, rather than to create a focused service.

Break down silos ...

IT architects have often been stereotyped as “people drawing funny boxes in charts.” For their part, software developers have been viewed as the people who write code for the modules that those “funny boxes” represent. This division of labor has all too often led to both groups operating in their own worlds rather than working closely together. A company that wants to be digitally competitive will need enterprise architects more than ever. However, those architects can no longer maintain an arm’s-length relationship with developers. They must work closely with them to make sure the architectural rules of perpetual evolution—not just the code—are written into software. Architects need to be part of the teams focused on a business capability or group of related capabilities. They will find themselves working alongside product

managers, developers, marketers, testers, production people, legal help, and others.

... but maintain a strict separation of the platform team from other teams

Every company informally manages a part of their IT architecture as a platform, and it organizes other parts according to business capabilities (for example, microservices associated with customer onboarding or marketing campaigns). To shift to a perpetual-evolution architecture, companies must draw explicit boundaries between these two parts of the architecture. Then they must enforce those boundaries through strict oversight and other governance processes.

A company's digital business capabilities enable it to make rapid changes to products and processes, therefore IT professionals must shift their focus along these lines as well. They should define the parts of the IT platform according to the business capabilities they support, rather than as technologies. Defining an IT capability as "service integration" will help the company identify the technologies in the organization with comparable functionality. It will also help the company create more meaningful roles, such as "service-integration architect," rather than "XYZ product architect."

Recognize that transformation of enterprise architecture must be an ongoing process

By drawing clear boundaries between business capabilities and technology platforms, companies will be able to isolate the fast-moving parts of their infrastructure (the business capabilities) from the slower-moving


ones (the platforms). Nonetheless, they cannot ignore the need to continuously improve their platforms. Companies must make sure they can update pieces of their platforms continuously as well. For many on the senior-leadership team, this will require a significant change in mind-set; traditionally they have been focused on requesting and approving "big bang" system changes. IT leaders and enterprise architects will need to educate the C-suite about the benefits of the perpetual-evolution model, which emphasizes continual monitoring and continual renewal, across all elements of the technology stack. They may need to introduce new forms of reporting and communications, for instance, to help business executives understand the need and to keep track of outcomes.



To stay competitive in a world in which providing a great customer experience has become paramount, companies in nearly every industry must continually innovate digital products and services, as well as the business processes that support those products and services. They can gain greater agility if they abandon rigid enterprise-architecture-management practices of the past and adopt a new approach that enables perpetual evolution—changing out elements of enterprise architecture quickly, adding new parts in no time, and incorporating the latest and greatest functionality. This shift in methodology can help traditional companies keep pace with digital-born competitors.

Oliver Bossert is a senior expert in McKinsey's Frankfurt office, and **Jürgen Laartz** is an alumnus of the Berlin office.

Copyright © 2018 McKinsey & Company. All rights reserved.



tcareob72/Getty Images

Ten trends redefining enterprise IT infrastructure

Arul Elumalai, Kara Sprague, Sid Tandon, and Lareina Yee

The IT infrastructure landscape is evolving rapidly. What will it look like in 2020?

When people think of enterprise IT infrastructure, they often imagine racks of hardware locked away in data centers and basements. But it is actually a focal point of disruption and innovation in every area, from servers and storage to networking and software.

What are the trends that are giving rise to such disruption and innovation? And what are the implications for business-technology strategy? Both IT infrastructure providers and customers must answer these questions as they plan their futures. We have identified ten trends that are already having a major impact

on IT infrastructure and will bring even more disruption over the coming years. Here is a look at what is changing and how companies can respond.

Familiar trends at a faster pace and greater scale

The following trends will not be news to anyone, but their recent acceleration and the scale of their impact might come as a surprise.

1. “As-a-service” consumption for everything from software to hardware.

Enterprise buyers increasingly prefer consumption-based pricing models—a

phenomenon that started with software and has now moved into hardware. This shift from capital expenditures to operational expenditures helps reduce risk, frees up capital, and provides increased flexibility. From 2015 through 2016, revenues for infrastructure as a service (IaaS) and platform as a service (PaaS) rose by 53 percent, making them the highest-growth segments in cloud and infrastructure services.¹ Considering that a unit of compute/storage in the cloud can be up to 40 to 50 percent cheaper in total cost of ownership than a unit on premises, the shift to as-a-service models is striking.² In addition to moving from on premise to cloud, IT providers and customers are experimenting with annuity-based payments for traditional hardware.

2. The public cloud goes mainstream.

While companies have been moving their workloads to the public cloud for years, there has recently been a sea change at large enterprises. Capital One, GE, Netflix, Time Inc., and many others have drastically reduced or even eliminated their private data centers, moving their operations to the cloud.³ In fact, cloud providers are expected to account for about 80 percent of shipped server and storage capacity by 2018.⁴

Amazon is the leader in IaaS, with about 40 percent market share.⁵ Microsoft is a clear second, followed by Google and IBM. Together these players account for approximately 65 percent of the IaaS market today.⁶ With the decline of on-premises data centers, they could account for almost half of all IT infrastructure provisioning by 2020. If that is the case, only companies with significant capital-investment capabilities could compete with them. One potential candidate would be Alibaba, which has recently experienced triple-digit year-over-year cloud-related revenue growth, driven largely by cloud adoption in China.⁷

3. Increased use of open-source offerings, up and down the stack.

Approximately 65 percent of companies increased their use of open-source software from 2015 to 2016, according to the 2016 Future of Open Source Survey conducted by Black Duck and North Bridge. Major IT providers now rely on programs such as Apache Spark, Kubernetes, and OpenShift. Moreover, Airbnb, Airbus, eBay, Intel, and Qualcomm are among the many large companies using TensorFlow, Google's open-source library of machine-learning code.⁸ Facebook's Open Compute Project, which aims to make hardware more efficient,

¹"2016 review shows \$148 billion cloud market growing at 25% annually," Synergy Research Group, January 3, 2017, srgresearch.com.

² Nagendra Bommadevara, James Kaplan, and Irina Starikova, "Leaders and laggards in enterprise cloud infrastructure adoption," October 2016, McKinsey.com.

³ Julie Bort, "Netflix, Juniper, and Intuit explain how Amazon is eating the \$3.5 trillion IT industry," *Business Insider*, January 13, 2016, businessinsider.com.

⁴ Arul Elumalai, Irina Starikova, and Sid Tandon, "IT as a service: From build to consume," September 2016, McKinsey.com.

⁵"Microsoft, Google and IBM public cloud surge is at expense of smaller providers," Synergy Research Group, February 2, 2017, srgresearch.com.

⁶ *Ibid.*

⁷ Ron Miller and Jon Russell, "Ambitious Alibaba takes aim at the kings of cloud computing," *TechCrunch*, February 27, 2017, techcrunch.com.

⁸ Scott Carey, "What is TensorFlow? How are businesses using it?," *Computerworlduk*, November 17, 2017, computerworlduk.com; TensorFlow, tensorflow.org.

flexible, and scalable, has helped extend the open-source movement into the data centers of companies that are participating members, such as AT&T, Deutsche Telekom, and Goldman Sachs.⁹

4. Cybersecurity remains a major concern.

Cybersecurity continues to be a top C-suite and board-level priority. Across all industries, attacks are growing in number and complexity, with 80 percent of technology executives reporting that their organizations are struggling to mount a solid defense.¹⁰ Many companies cannot recruit the internal talent needed because there is a shortage of cybersecurity experts, leading them to invest in managed security services. Cloud-based security offerings are also becoming more attractive to companies, with McKinsey estimating that they will comprise 60 percent of security products by 2020, up from 10 percent in 2015.

5. Mainstream comfort with “white box” hardware.

Traditionally, IT infrastructure providers have relied on assembling branded systems for their server, storage, and networking offerings. To do so, they outsourced hardware manufacturing to original-design manufacturers (ODMs). However, this model is becoming obsolete because customers are increasingly unwilling to pay for assembly. Instead, customers go directly to ODMs, using designs for servers obtained from sources such as Facebook’s Open Compute Project to customize their data-center configurations. Open Compute Project member companies that have taken this route include IBM, Fidelity

Investments, and Verizon.¹¹ As discussed later in this article, many of these ODMs are located in Asia, which is driving more hardware business to that region. By 2020, IDC estimates that “self-built” servers will comprise half the hyperscale-server market.

6. Internet of Things business applications are ready for adoption.

McKinsey estimates that business-to-business applications will account for nearly 70 percent of the value that will flow from the Internet of Things (IoT) in the next ten years.¹² According to our 2017 Enterprise IoT Executive Survey, 96 percent of companies expect to increase their IoT spending over the next three years, with some planning to devote as much as a quarter of their IT expenditures to IoT-related capabilities. The most popular use cases for enterprise IoT involve increasing visibility into operations, optimizing operational tasks, or assisting with the development of new business models. The upshift in adoption is even occurring in industries that have traditionally been slow to adopt new technologies, such as oil and gas. The growth of enterprise IoT will vastly increase demand for the compute-and-storage infrastructure, augmenting demand for hyperscale resources and IoT-specific PaaS solutions.

BI Intelligence predicts that more than five billion IoT devices, such as inventory-control and safety-monitoring tools, will require edge solutions by 2020 because they must collect and process data in real time.¹³ Edge solutions allow information processing at the

⁹ Membership directory, Open Compute Project, opencompute.com.

¹⁰ Tucker Bailey, James M. Kaplan, and Chris Rezek, “Repelling the cyberattackers,” *McKinsey Quarterly*, July 2015, McKinsey.com.

¹¹ Membership directory, Open Compute Project, opencompute.com.

¹² Jacques Bughin, Michael Chui, and James Manyika, “An executive’s guide to the Internet of Things,” *McKinsey Quarterly*, August 2015, McKinsey.com.

¹³ Nicholas Shields, “Microsoft brings IoT to the edge,” *Business Insider*, May 12, 2017, businessinsider.com.

device or gateway level, rather than within the cloud or a data center, reducing both latency and connectivity dependencies. Of the \$500 billion in growth expected for IoT through 2020, McKinsey estimates that about 25 percent will be directly related to edge technology. Edge computing will help improve data compression and transfer in the connectivity layer of the technology stack, reducing network bandwidth and making a wider range of IoT applications possible.

New trends to watch

In addition to the acceleration of familiar trends, several new developments are altering the IT infrastructure landscape for both providers and customers. These include the shift to Asia in hardware, the use of DevOps for software and hardware, container-first architectures, and the growth of artificial intelligence and machine-learning-optimized stacks.

7. The shift of the hardware infrastructure market to Asia. Asian original-equipment manufacturers have been making inroads in the IT infrastructure market dominated by US-based providers. Consider two examples in the server market:

- Huawei plans to shore up its position in the server market by spending about \$1 billion of its annual \$9 billion R&D budget on equipment for data centers.¹⁴
- Lenovo acquired IBM's x86 server business in 2014, helping to expand its footprint in large enterprises globally.¹⁵

An equally important shift involves Asian ODMs, which have also increased their share of the hardware market as white-box systems become more popular. Taiwan-based Quanta Computer's cloud-computing revenue from server, storage, switch, and IoT devices has been strong. Several Asian ODMs now provide servers to some of the top global hyperscale cloud providers, including Amazon, Facebook, and Google, all of which are investing heavily in expanding their data-center infrastructure.¹⁶ As noted earlier, initiatives such as Facebook's Open Compute Project are accelerating with this shift, since they allow members to obtain plans and designs for servers, storage, and networking. Some Asian ODMs are also offering off-the-shelf products based on open-source designs. If current trends continue, Asian ODMs may increase their revenue share of the hardware market two- or threefold by 2020.¹⁷

8. DevOps for software and hardware.

IT departments have to deliver new features even faster. Meanwhile, companies now expect greater availability from them—24-hour coverage every day of the week. DevOps can help achieve both goals by fostering a high degree of collaboration along the entire IT value chain.

The new DevOps business model extends beyond application development to encompass application operations and IT infrastructure. Within DevOps, all three groups work as one. Many organizations understand the benefits of this model and are moving in

¹⁴ Barb Darrow, "Chinese Giant Huawei to attack server market," *Fortune*, September 12, 2016, fortune.com.

¹⁵ "The new X factor, Lenovo," September 2014, lenovo.com.

¹⁶ Agam Shah, "Made-in-China servers attracting more buyers," *PC World*, March 9, 2016, pcworld.com.

¹⁷ IDC Worldwide Quarterly Server Tracker; IDC Worldwide Quarterly Disk Storage Tracker.

this direction. In McKinsey's 2017 IT-as-a-Service Survey, 80 percent of respondents stated that they had implemented DevOps practices in some part of their organization. In addition, 53 percent of respondents stated that they would apply these practices across their entire organization by 2020, up from 37 percent today.

In keeping with these trends, demand for DevOps talent will surge over the next few years. Companies may have trouble finding staff to fill all roles, since 40 percent of survey respondents stated that a lack of internal talent and skills was the primary factor preventing DevOps from becoming mainstream.

9. Container-first architectures. No longer confined to niche development environments, containers are on the path to overtake virtual machines and become the primary unit of deployment in the cloud. Atlassian's 2016 report, *Software development trends and benchmarks*, revealed that 34 percent of software professionals have adopted containerization in their development teams.

What is most remarkable about containerization is the speed of its growth. In RightScale's 2016 *State of the cloud* report, only 18 percent of respondents reported deploying containers in production environments. In McKinsey's 2017 survey, by contrast, respondents stated that Docker was their most frequently used DevOps tool. The growth of containerization has been occurring in tandem with the proliferation of microservice architecture—the development of software applications in small, independent units. As developers refine microservices, they

are also addressing many of the challenges that prevented containerization's growth, including inadequate security, problems with management or orchestration, and scalability.

In parallel with these trends, the next logical step in application atomization is emerging. It involves the abstraction of compute resources, in which functions become a unit of deployment, or function as a service. This will eliminate the need to provision infrastructure or manage compute resources for these functions.

10. Artificial intelligence and machine-learning-optimized stacks. After many years of refinements, artificial intelligence (AI) is delivering benefits to companies across industries.¹⁸ Consider, for instance, how AI helps utilities forecast electricity demand, or how it allows automakers to create self-driving cars. Various developments are encouraging this new wave of AI, including increased computation power and the availability of more sophisticated algorithms and models. Perhaps most important, data volume is exploding, with network devices collecting billions of gigabytes every day.

The McKinsey Global Institute (MGI) estimates that the entrepreneurial activity unleashed by AI drew between \$26 billion and \$39 billion in investment in 2016—three times the amount attracted in 2013.¹⁹ Most AI investment comes from large digital natives, such as Amazon, Baidu, and Google, which are exploring innovations in semiconductors, infrastructure software, and systems. Some companies are building new computing paradigms that incorporate tensor processing units from

¹⁸ Dorian Pyle and Cristina San José, "An executive's guide to machine learning," *McKinsey Quarterly*, June 2015, McKinsey.com.

¹⁹ For the full McKinsey Global Institute report, see "How artificial intelligence can deliver real value to companies," June 2017, on McKinsey.com.

Google, graphics processing units from Nvidia, and field-programmable gate arrays from Xilinx. The large hyperscale providers are also offering AI and machine-learning capabilities to enterprises through the cloud.

As enterprises gain increased access to leading-edge AI and machine-learning technologies, automation will increase. According to MGI, about half of all the activities people are paid to do in the world's workforce could be automated, accounting for almost \$15 trillion in wages.²⁰




The scale of disruption in the technology-infrastructure landscape is unprecedented, creating huge opportunities and risks for industry players and their customers. Executives at technology infrastructure companies must drive growth by transforming their portfolios and rethinking their go-to-market strategies. They should also build the fundamental capabilities needed for long-term success, including those related to digitization, analytics, and agile development. All of these ambitious steps will require more capital and capacity, but customers in the new IT infrastructure landscape will reward their efforts. ♦

²⁰ For the full McKinsey Global Institute report, see “Harnessing automation for a future that works,” January 2017, on [McKinsey.com](https://www.mckinsey.com).

Arul Elumalai and **Sid Tandon** are partners in McKinsey's Silicon Valley office; **Kara Sprague** is an alumna of the San Francisco office, where **Lareina Yee** is a senior partner.

Copyright © 2018 McKinsey & Company. All rights reserved.



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. Last, you have much better reliability and security for your applications. Those servers can be either internal for your enterprise—and we call this 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. Second, 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 past 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 that 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 and what are the right standards 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 is clearly a group of companies that are emerging as leaders in cloud adoption, and we are calling them cloud savvy. They have achieved a much higher adoption of the 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 was 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 is 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 of 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 have been wide reports of a number of security breaches both 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 are the available tools 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, this notion 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 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 center

in the private cloud because of legacy, cost, or security reasons. What I also believe is that that 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 the 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 about 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 the operating model is the hardest thing to get done right when migrating to the cloud at scale.

Even companies that anticipated that that 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. ♦

Roberta Fusaro is a senior editor at McKinsey Publishing in the North American Knowledge Center. **Irina Starikova** is a partner in McKinsey's Silicon Valley office.

Copyright © 2018 McKinsey & Company. All rights reserved.



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 cybersecurity² 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

¹ For more, see Nagendra Bommadevara, James Kaplan, and Irina Starikova, “Leaders and laggards in enterprise cloud infrastructure adoption,” October 2016, McKinsey.com; 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.

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.

While adoption of the public cloud has been limited to date, the outlook for the future is markedly different. Just 40 percent of the

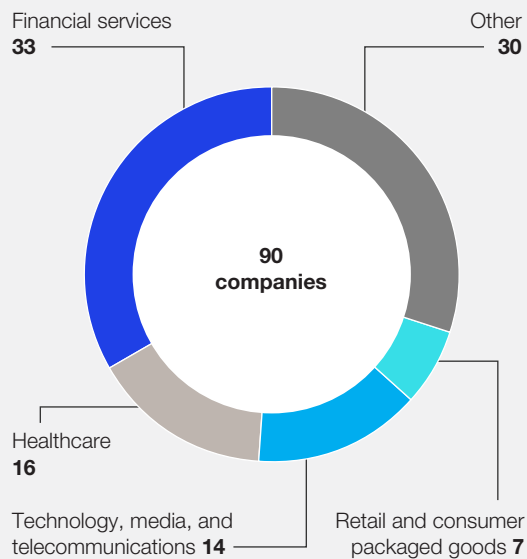
companies we studied have more than 10 percent of their workloads on public-cloud platforms; in contrast, 80 percent plan to have more than 10 percent of their workloads in public-cloud platforms in three years or plan to double their cloud penetration (Exhibit 1).³ We refer to these companies as “cloud aspirants.” They have concluded that

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.

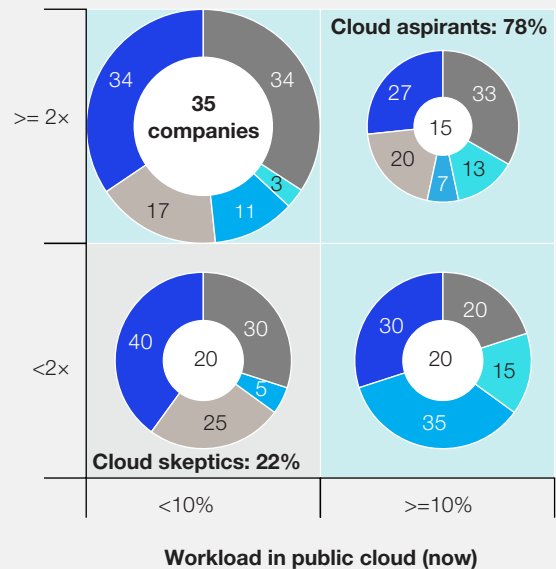
EXHIBIT 1

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.

Respondents by industry,¹ % of group



Expected growth in adoption in next 3 years,¹ % of group



¹Figures may not sum to 100%, because of rounding.
Source: McKinsey global cloud-cybersecurity research, 2017

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 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 the security resources of cloud-service providers (CSPs) 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 with 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 according 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

EXHIBIT 2 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-cloud security controls.



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



the applications. These two key choices also inform each other. 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).
- **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

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 (CSP) controls where available and third-party controls for areas that CSPs do not support. The executive said, "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."

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 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 (Exhibit 3). 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 in implementing new security programs, and the flexibility they desire from their security architectures.

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 is 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.

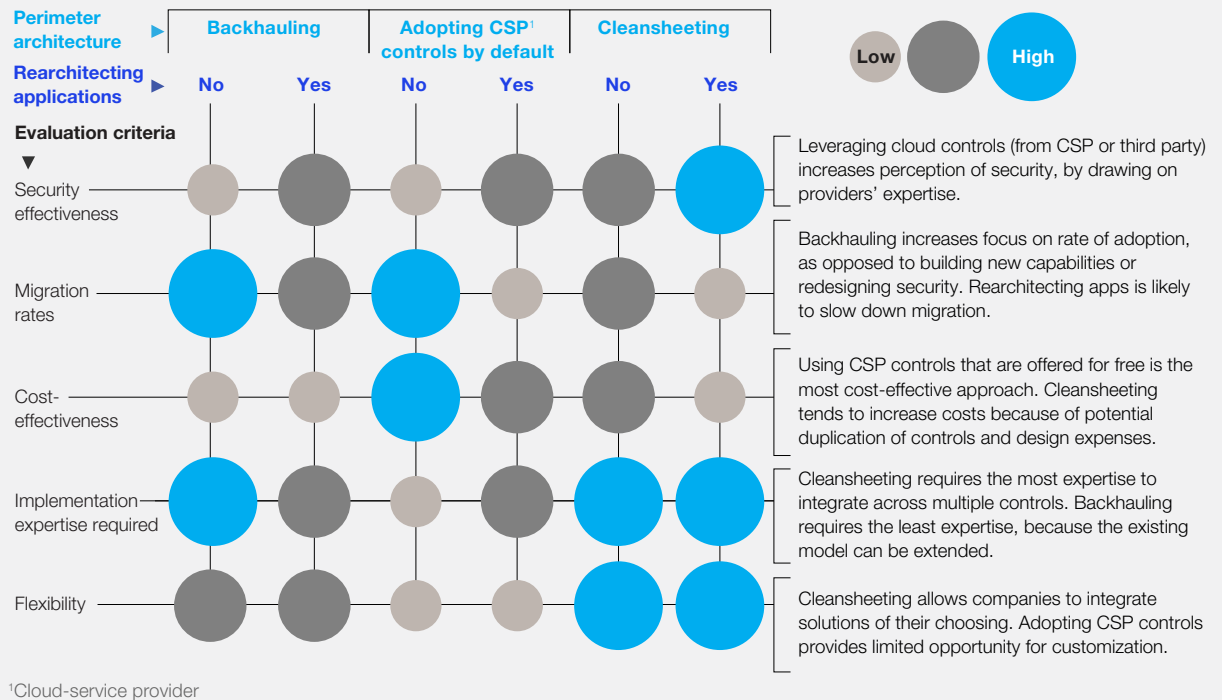
Redesigning a full set of cybersecurity controls for the public cloud

Once enterprises have decided on a security

EXHIBIT 3

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

Performance of archetype against evaluation criteria



Moving into the next generation of identity and access management

A Fortune 500 healthcare company we spoke with has redesigned its identity-and-access-management 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 cloud-service providers (CSPs), 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."

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. Organizations need to think about the following eight areas of cybersecurity controls, which we list along with observations from our research, in combination:

- **Identity and access management (IAM).** IAM solutions for cloud-based applications and data are gradually shifting into the cloud (see sidebar “Moving into the next generation of identity and access management”). 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

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 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.”

- **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.
 - **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 more 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 with what providers will do

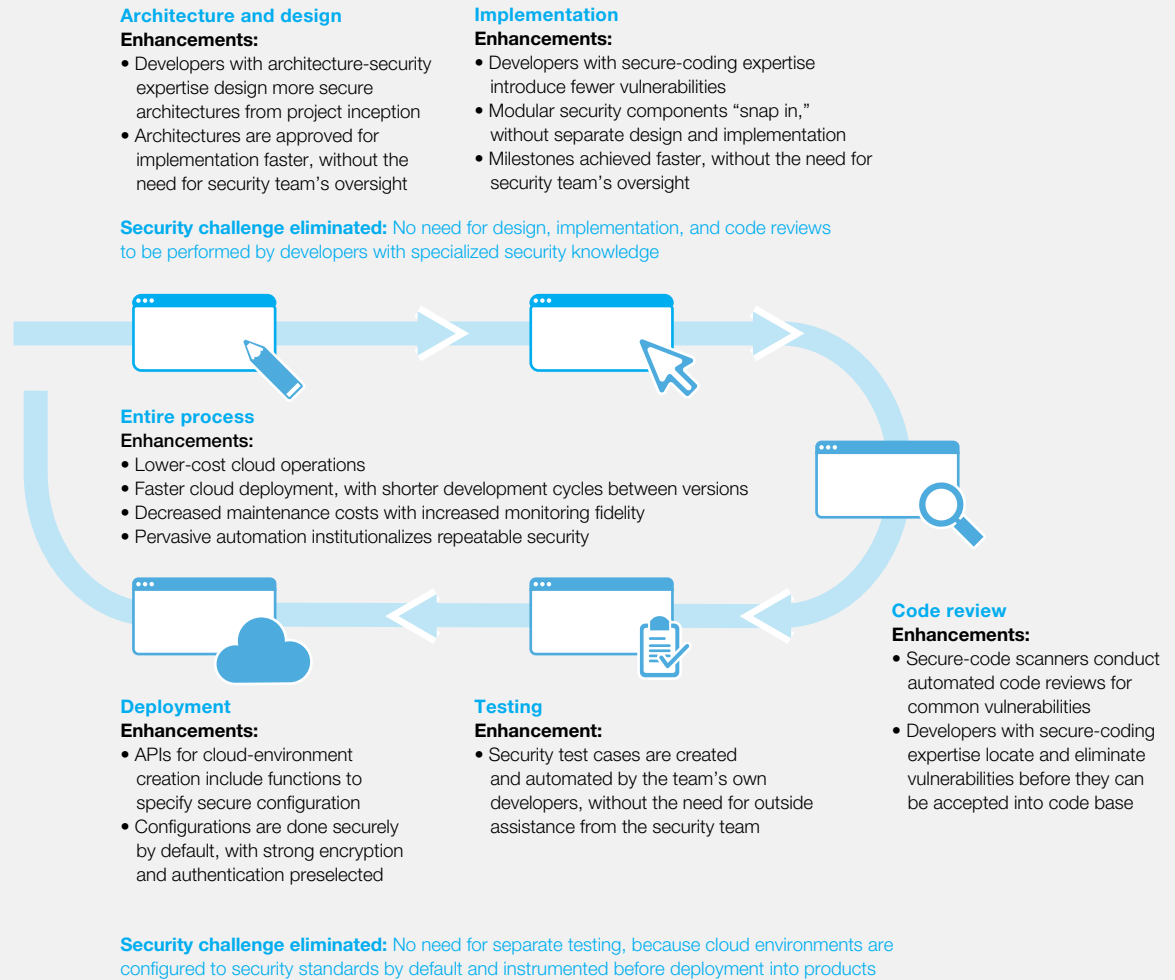
When enterprises migrate applications and data to the public cloud, they must depend on CSPs and third-party providers for some security controls—but they should not depend on them to provide all of the necessary controls. Unless companies and CSPs clearly divide all the responsibilities for cybersecurity 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, the following 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 as well as 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.

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

Cloud-deployment process with secure DevOps



- Multicloud IAM capabilities.**
 Companies should insist that CSPs provide native multifactor authentication (MFA). Those that use identity-as-a-service 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.

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. The following 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 as well as 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 MFA, 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 the 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.** 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 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.

Copyright © 2018 McKinsey & Company. All rights reserved.



Sashkinw/Getty Images

Five ways to unlock win–win value from IT-services sourcing relationships

Rahil Jogani, Aditya Pande, and Vikrant Shirdade

IT purchasers and providers can achieve win–win outcomes by altering their sourcing routines. Here’s how they can be more strategic about the principles and practices they follow.

Global companies in all industries typically acquire a significant portion of their IT services from external providers. Annual global spending on external IT services is about \$900 billion, with companies procuring IT consulting, systems integration, application development and maintenance, and IT-infrastructure-management services, among others.¹

We can expect spending on IT services to increase as companies explore digital products and business models, and require more and different types of technology support. Indeed, digital transformations typically require higher-order IT capabilities and involve strategic partnerships with providers that supply critical knowledge and

¹ Forecast alert: IT spending, worldwide, 4Q16 update, Gartner, January 2017, gartner.com.

expertise along with new technologies and support services.

For a number of reasons, however, IT-sourcing relationships have been difficult to get right. Given the speed at which new technologies and software-development approaches emerge, IT purchasers often struggle to understand how to set realistic objectives and incentives; how to balance multiple priorities relating to cost, efficiency, quality, and innovation; and how to structure governance arrangements to benefit both sides. For their part, IT providers wrestle with similar issues: how best to meet a range of customers' expectations, how to prioritize objectives and resources to help customers meet their individual needs, and how to create next-generation improvements and innovations for customers rather than just carrying out immediate tasks.

To help executives understand how to answer these questions, we conducted reviews of hundreds of contracts over the past three years.² The contracts covered IT-sourcing relationships across multiple industries and regions. We analyzed the contracts along three main dimensions: general terms and conditions, commercial terms and conditions, and governance structure (Exhibit 1).

In many of the contracts we reviewed, the sourcing relationship was not meeting its full potential (Exhibit 2). For instance, greater innovation was a desired goal on both sides but often was lacking, according to the executives with whom we spoke. Such performance gaps exist because of shortfalls on both sides of the partnership.

Our research revealed five obvious but often overlooked changes IT purchasers and providers can make to their sourcing routines that could bridge these gaps and create win-win outcomes. Specifically, they must develop a shared understanding of business outcomes, emphasize the long term, actively collaborate on critical IT architecture decisions, pursue transformation with clear planning and relentless "grit," and devise win-win contract mechanisms (Exhibit 3). Businesses and IT providers should address all five of these areas if they want to achieve a full spectrum of benefits from IT contracts, beyond just cost. Based on our experience, the value gained by both sides could be between two and four times that of pursuing traditional contracting approaches (Exhibit 4). We believe the best way to break from status quo practices and relationships is to fully recognize the dynamics at play and devise clear plans to alter them.

Develop a shared understanding of business outcomes

For many IT purchasers and providers, it can be hard to achieve a shared understanding of business objectives. In more than 60 percent of the contracts reviewed, teams had not followed a thorough process for internally discussing desired business outcomes.

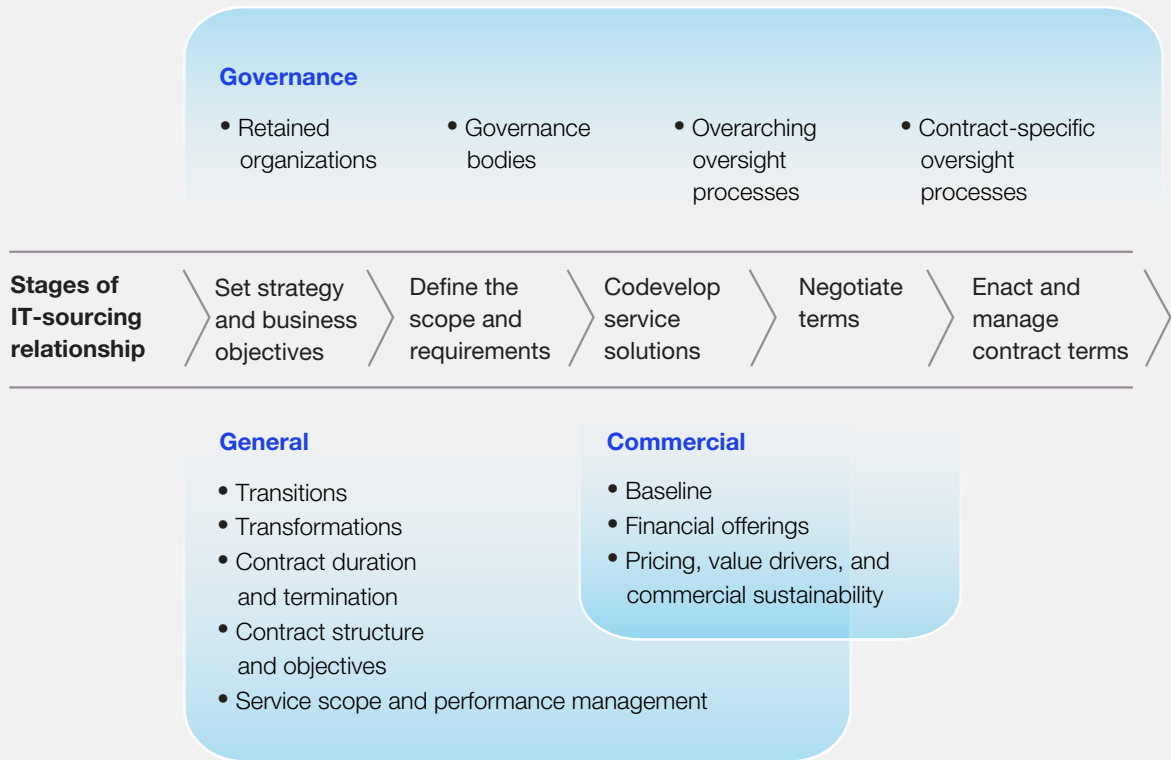
IT purchasers faced hard internal deadlines for developing and finalizing contracts with providers. They felt they did not have enough time to engage all relevant business stakeholders in defining the full potential value to be gained from investment in external IT services—whether it be cost savings, increased productivity, or more agility and innovation. As a result, they were often unsure

² Our observations are drawn from a sample of more than 50 companies and about 200 live contracts. We examined an average of 35 to 40 data points per contract. In performing this research, we employed the strict confidentiality procedures that govern our work with purchaser-provider relationships.

EXHIBIT 1

Our research examined general, commercial, and governance provisions in 200 live IT contracts.

Terms and conditions reviewed



Source: McKinsey analysis

of priorities when setting new contracts or resetting existing ones. In 100 percent of the contracts we reviewed where business outcomes were not clearly defined, key indicators of performance were not exhaustive; they tended to be focused on cost. Hence, they were inadequate for measuring desired business outcomes.






Because they had incomplete information about purchasers' business priorities, providers were unable to determine how best to allocate talent and resources. And

because providers were required to follow purchasers' standard contracting structures and processes, they were less likely to bring new ideas to the table.

To ensure a shared understanding of objectives, purchasers and providers will need to actively break from time and process constraints. Purchasers should involve end users and business-unit leaders in contract discussions with providers early in the process. Desired outcomes should be captured in a minimum viable contract, with

EXHIBIT 2

Our research on contracting practices identified missed opportunities for both IT purchasers and providers.

<p>60%</p> 	<p>of contracts lacked well-defined, shared business objectives</p>	<p>Purchasers and providers faced</p> <ul style="list-style-type: none"> • Unclear definitions of quality of service • Limited tracking and control on business and financial targets
<p>90%</p> 	<p>of contracts had commercial terms and conditions that were not designed to achieve mutually advantageous benefits</p>	<p>Purchasers and providers demonstrated</p> <ul style="list-style-type: none"> • A heavy focus on near-term benefits, often at the cost of long-term value • Few incentives for joint innovation; therefore, limited evidence of these initiatives
<p>90%</p> 	<p>of contracts had detailed governance plans in place, but providers were still mostly kept at arm's length</p>	<p>Purchasers and providers showed</p> <ul style="list-style-type: none"> • Limited evidence of collaboration, especially on IT-architecture-related topics • An inability to surface innovative ideas, because of providers' limited involvement in internal forums
<p>75%</p> 	<p>of contracts listed cost reduction as a primary transformation objective</p>	<p>Purchasers and providers had not</p> <ul style="list-style-type: none"> • Specified transformation plans or measured outcomes beyond cost • Defined day-to-day key performance indicators
<p>67%</p> 	<p>of contracts did not allow for transparency on pricing structure for both parties</p>	<p>Purchasers and providers lacked</p> <ul style="list-style-type: none"> • Collaborative, comprehensive, value-based negotiations on price • A total-cost-of-ownership approach to pricing that incorporates all possible costs and consumption patterns in different scenarios • Mutual incentives and gain-sharing mechanisms

Note: We developed our Contract Maturity Score (CMS) by comparing top- and bottom-decile contracts. The CMS is the weighted average score of all data points examined for a contract on a scale of 1–10. Our observations are drawn from a sample of more than 50 companies and 200 live contracts. We examined between 35 and 40 data points per contract.

the understanding that there will be further collaboration and refinement on terms over time. At one asset-management firm, IT leaders insisted that the CEO, other C-suite executives, and members of the board be involved in the provider-selection process.

Senior business executives were asked to attend supplier visits, and the board received periodic updates on the selection process. When properly informed in this way, business executives and other stakeholders can help contract teams set meaningful short-term and

IT purchasers and providers need to reframe the sourcing relationship.



Develop a shared understanding of business outcomes

- Actively break from time and process constraints **to get clear on business outcomes** expected by critical stakeholders in the business, the IT organization, and among providers
- Reimagine the request-for-proposal process as a **request for solutions**—an opportunity to jointly identify critical problems and define possible fixes
- Capture expected business outcomes in a **minimum viable contract** and continuously collaborate on and refine specific conditions



Emphasize the long term

- Record and share the long-term vision on both sides; for instance, purchasers can involve providers in regular **hackathons** to uncover value opportunities, and they can attend providers' **customer-council** meetings
- Outline an annual **zero-based reset** process to update contracts with estimated volumes and changes to operating or commercial models
- Maintain a continuously updated **backlog of requirements** to assess ongoing changes in mutual expectations and to redefine terms



Actively collaborate on critical IT-architecture decisions

- Create an empowered **provider-success team** to ensure that goals are met
- Bring subject-matter experts from both purchaser and provider to newly created forums—for instance, an **architecture-enablement board**, where IT-architecture and platform innovation are discussed
- Follow Toyota's **genchi genbutsu**¹ approach: coordinate **go-and-see** sessions for purchasers and providers



Pursue transformation with clear planning and relentless 'grit'

- Jointly design a transformation road map and include cost as well as **nonfinancial** goals
- Jointly support the road map with detailed planning to get to an **operating model**, not just a target cost level



Devise win-win contract mechanisms

- Pursue a balanced set of economic incentives; be transparent about **underlying sources of cost** and set mutual targets
- Jointly adopt a **total-cost-of-ownership** approach to pricing rather than emphasizing unit prices
- Actively monitor the adherence to agreed-on outcomes in a **balanced-value scorecard**

¹*Genchi genbutsu* means "actual place, actual thing," and it is a key principle of the Toyota Production System. It suggests that in order to truly understand a situation, one needs to go to the "real place" where work is done.

EXHIBIT 4

Companies can reap significant benefits from win-win contract-management practices.



1.5x–2.5x

Improvement in service quality and customer-satisfaction scores



30%–50%

Reduction in time-to-market outcomes on projects, due to agile delivery



2x–3x

Better than existing baseline on mutually agreed upon road map for transformation targets



2x–4x

Increase in contract-maturity score, leading to improved governance and relationships

long-term objectives and ensure that enforcing mechanisms are embedded in contracts.

IT purchasers and providers will also need to reimagine the request-for-proposal (RFP) process as a “request for solution” process—an opportunity to jointly identify critical problems and define potential solutions. We observed one provider challenge traditional

boundaries: in response to an RFP, the provider went beyond a simple outline of services it could provide; it suggested ways that the purchasing company could form a joint venture with the provider and go to market with a joint service offering. This proposal reframed service provision as an opportunity to generate revenue—an approach that turned heads and established the provider’s credentials with

the purchaser. The joint-venture idea was not approved, but the provider eventually was rewarded with the services deal.

Emphasize the long term

Almost 80 percent of the contracts we examined included service requirements that were narrow in scope. They were valid for a particular point in time but left little room to incorporate future needs. Contract teams had estimated the volume of work flow required but had built in few or no options to reset contract parameters midstream based on actual usage. Additionally, it was often the case that purchasers had not considered future IT needs in any depth. About the same percentage of contracts we reviewed (80 to 90 percent) did not include adequate mechanisms for encouraging long-term innovation.

The conventional thinking among purchasers is that contracts implicitly compel providers to innovate and adopt a long-term focus. The reality is that the success of the sourcing process will be measured according to savings in the first year rather than value created beyond that. This was evident in our reviews. Ninety percent of the contracts included commercial terms and conditions that were not mutually advantageous, implying that discussions around these terms and conditions had been focused on the here and now rather than on how to activate longer-term innovation and value.

To unlock long-term vision, commitment, and innovation, purchasers should delineate a range of service requirements and expected volumes at the outset, based on real-world business objectives, but then refine and codevelop the service requirements with providers through regular “hackathons”—gatherings aimed at surfacing new ideas and determining the resources required to act on them.

IT purchasers also need to give providers more visibility into the business and its big-picture goals: What capabilities and technologies could providers bring to the table to help the business achieve its objectives? The agreed-on contract could include a continuously updated backlog of requirements and options for both sides to shift to alternative work-flow requirements and management approaches when delivery models, the volume of work, or the scope of work changes. Both sides might also agree to a series of checkpoints at which they could reassess and redefine terms. Purchasers could also join providers’ customer-council meetings to understand and influence their long-term direction.

A consumer-electronics company negotiated for such flexibility from core IT providers as it determined the best ways to shift to a cloud-based infrastructure. The company did not want to upend existing operations all at once, so it developed a three-year transition plan with providers. The providers would continue to deliver traditional infrastructure services in the first year while gradually shifting the company’s work flow to a cloud platform in the second and third years. Frequent reviews were built into the process, and both sides agreed that pricing mechanisms would be changed accordingly. The transformation is still under way, but as a result of this arrangement, the purchaser has made significant progress in migrating applications to the cloud.

Actively collaborate on critical IT architecture decisions

Most of the contracts we reviewed contained comprehensive governance plans—detailed descriptions of multilevel forums convened by purchasers, designed to gather input from providers or to review and refine day-to-day processes and tasks. But in practice, providers were mostly kept at arm’s length,

with no direct input into important technology and innovation forums—for instance, the architecture-review board.

In our observation, this dynamic occurred, in part, because contract-governance mechanisms were generally designed and agreed upon by teams that were not accountable for day-to-day execution and outcomes—for instance, procurement and legal. The teams’ primary focus understandably tended to be on anticipating potential catastrophes and retaining control rather than on creating long-term value from the sourcing relationship.

Win–win relationships cannot exist when IT purchasers do not treat IT-service providers as strategic partners. To facilitate this partnership, companies can activate a provider-success team that includes sourcing experts, technicians, and business leaders from both the purchaser and the provider. This team should create a seat at the table for IT providers in forums such as an architecture-enablement board, where the purchaser discusses ideas on IT architecture and underlying platform innovation jointly with providers.

In our observation, transformations are less likely to succeed—and changes are less likely to be implemented on the ground—when providers do not have input into purchasers’ decisions about the underlying architecture. In about 10 percent of the contracts reviewed, we saw evidence that IT purchasers were beginning to design contracts to actively involve providers in these types of forums and explicitly asking providers to be consulted or informed in architecture-review boards. Companies could improve collaboration by

following Toyota’s *genchi genbutsu* approach—that is, scheduling go-and-see sessions for purchasers and providers.³

A global telecommunications company works with multiple IT providers that support the company’s infrastructure management and delivery of software applications. The telco’s contracts with these providers explicitly define partnership-based governance principles. There is a built-in expectation that subject-matter experts from the IT providers will actively participate in conversations about architecture and innovation. In turn, IT providers have gained greater visibility and have shown greater commitment to ensuring that the IT purchaser meets program-level measures of success—for instance, improvement in service levels or direct impact on fees. This approach has enabled the overall environment to function smoothly, with an emphasis on collaboration and transparency.

Pursue transformation with clear planning and ‘grit’

About 75 percent of long-term IT-sourcing contracts reviewed had language indicating that purchasers expected the sourcing relationship to result in a significantly altered IT landscape. But in most cases, there was limited evidence of a transformation plan or a strategy to get the required investments for changing the underlying operating model. For example, 75 percent of the contracts focused on cost control as a significant transformation objective but had less focus on how to measure progress against their transformation objectives. Any dialogue about transformation plans was typically led and owned by IT representatives rather than business stakeholders. The latter were informed as needed. For their part, providers often were not privy to plans

³ *Genchi genbutsu* means “actual place, actual thing” and it is a key principle of the Toyota Production System. It suggests that in order to truly understand a situation one needs to go to the “real place” where work is done.

regarding the architecture road map, so they felt they could have little influence on outcomes from transformation efforts. In most cases, neither side was completely clear about the starting point for change, the target state, and the exact path to get there.

IT transformations cannot succeed without two things: a shared commitment to creating change, including the underlying IT-architecture choices, and something we call transformation “grit”—or rigorous and relentless attention paid to planning and execution. Purchasers and providers must jointly design a transformation road map; it is critical that it be codeveloped and co-owned by the business-unit leaders. Purchasers and providers must support this road map with detailed planning to get to a new target operating model. In this way, they can build a baseline against which to measure outcomes from any sort of IT transformation. Contract teams can monitor costs, but they can also track nonfinancial performance-based metrics (for both business and IT activities) such as asset-refresh rates and service-delivery times. Under this approach, IT purchasers can manage change in the IT landscape more effectively (with input from business stakeholders), and IT providers can be assured that their share of gains will be based on a solid business case rather than an ambiguous definition of success.

The contract team at one investment bank set aggressive transformation goals at the outset of its conversations with sourcing partners. The contract team sought a 40 percent reduction in costs in the first two years and an additional 40 percent reduction over the following six years through the use of the provider’s services. These numbers galvanized the bank and the IT provider. They jointly considered unconventional ideas for transforming the bank’s operations—for instance, moving

to cloud-based services and retiring some applications—and implemented each according to a detailed plan they had drawn up early in the process. The bank, its stakeholders, and the IT provider managed to reset existing relationships and worked together to push for significant technology improvements over the long term.

Devise win-win contract mechanisms

The reality of IT-services sourcing is that in most contracting relationships, negotiations are treated like miniature battles, typically with each side focused on achieving the best price rather than mutual value. Indeed, in most of the contracts we reviewed, commercial mechanisms were not designed to be mutually advantageous, particularly during large contract resets. Certainly, cost reductions are an important objective for IT purchasers, but those reductions cannot come solely at the expense of providers’ margins if sourcing relationships are to flourish.

Both sides should instead pursue a balanced set of economic incentives. They could ensure sustainable economics by being transparent about underlying sources of cost and by setting mutual targets. Instead of focusing solely on unit prices, they could adopt a total-cost-of-ownership approach to pricing that, at the outset, incorporates all possible costs, consumption patterns, and other factors, given different scenarios. There could be an open discussion about mutual incentives in a scorecard that tracks implementation of gain-sharing mechanisms in the contract. Both sides should also consider the value of the contract over its duration rather than just at the beginning. Some contracts are structured so that one side can meet its business and financial objectives in the first year while the other party benefits in the ensuing years. We found balanced incentives in only about 10 percent of the contracts reviewed.

An automaker struck a deal with the provider of IT-infrastructure services under which it agreed to a certain minimum in annual spending. It also agreed to provide reliable forecasts of service needs for a small group of business activities that were stable and thus easier to predict. In return, the provider offered the automaker commercial discounts and agreed to a gain-sharing arrangement in which both sides would benefit from aggressive productivity improvements.

Similarly, a European telecommunications company struck a win-win deal with the provider of its infrastructure-management and application-development services. Under the terms of the deal, the purchaser would benefit from an aggressive and predefined schedule of productivity improvements that would push unit prices down, thereby counterweighting the effects of inflation. Instead of offering outright volume discounts, the provider changed unit rates every year based on actual consumption by the telco compared with the previous year's numbers. This helped sustain the volume of work flowing to the provider and allowed it to correctly gauge and meet the required levels of service each year. As a result of these win-win economics, the provider was always willing to

address any pain points for the telco during the contract term. The purchaser received reliable service. In turn, the telco's IT users gave the provider high customer-satisfaction scores.



IT sourcing is not going away. Companies in all industries lack the bandwidth necessary to maintain all their IT capabilities in-house. They must rely, to one degree or another, on external service providers to get even the most basic tasks done. And as more companies attempt to digitize their products and operations, IT sourcing becomes even more critical.

Companies should continually assess the efficacy of their strategic partnerships—they must evaluate not just technologies provided or service-level agreements forged but also the expertise obtained and innovations achieved. Our research points to five ways to strengthen IT-sourcing dynamics. There are likely other areas for improvement as well. What's clear is that both sides will need to view their interactions differently—as true win-win partnerships, where more value from IT is created together than apart. ♦

Rahil Jogani is an associate partner in McKinsey's Chicago office, **Aditya Pande** is a partner in the Silicon Valley office, and **Vikrant Shirdade** is a specialist in the Gurgaon Knowledge Center.

The authors wish to thank Gaurav Mittal for his contributions to this article.

Copyright © 2018 McKinsey & Company. All rights reserved.



Chris Clor/Getty Images

Why you need a digital data architecture to build a sustainable digital business

Sven Blumberg, Oliver Bossert, Hagen Grabenhorst, and Henning Soller

Companies that succeed at meeting their analytics objectives let business goals drive the technology. Here's how they structure a data architecture that works.

Data architecture has been consistently identified by C-suite executives as a top challenge when preparing for digitizing business. Leveraging our experience across industries, we have consistently found that the difference between companies that use data effectively and those that do not—that is, between leaders and laggards—translates to

a one-percentage-point margin improvement for leaders. In the apparel sector, for instance, data-driven companies have doubled their earnings-before-interest-and-taxes margin as compared with their more traditional peers.

Using data effectively requires the right data architecture, built on a foundation of business

requirements. However, most companies take a technology-first approach, building major platforms while focusing too little on killer use cases.

Many businesses, seeing digital opportunities (and digital competition) in their sectors, rush to invest without a considered, holistic data strategy. They either focus on the technologies alone or address immediate, distinct use cases without considering the mid- to long-term creation of sustainable capabilities. This goes some way toward explaining why a 2017 McKinsey Global Survey found that only half of responding executives report even moderate effectiveness at meeting their analytics objectives. The survey found the second-largest challenge companies face (after constructing a strategy to pursue data and analytics) is designing data architecture and technology infrastructure that effectively support data-and-analytics activities at scale. We found that eight out of ten companies embark on digital data enablement by making their IT departments responsible for the data transformation—with very grand implementation programs—and a small set of business use cases.

This strategy is quite different from that employed by next-generation digital leaders, who typically embark on transformation from a business perspective and implement supporting technologies as needed. Doing the technology first produces more problems than successes:

- **Redundant and inconsistent data storage.** Only two in ten banks we've looked at have established a common enterprise data warehouse, which is

essential for creating a single source of truth for financial and customer data.

- **Overlapping functionality.** Every bank we've analyzed has at least one business function supported by three different technological systems.
- **Lack of sustainability.** The solutions at which financial institutions typically arrive are often quick fixes that ignore the enterprises' larger aspirations for datafication. For example, one insurance company extracted and replicated data from its warehouse each time it was needed rather than building data architecture that would allow it to store each customer element only once, thereby reducing costs and eliminating inefficiencies.

These problems have real business consequences. Meeting leading-edge business requirements, such as real-time customer and decision support, and large-scale analytics requires the integration of traditional data warehousing with new technologies.

The two-speed data-architecture imperative

Today, enterprises must cope with increasingly large and complex data volumes (worldwide, data storage doubles every two years) coming from diverse sources in a wide variety of formats that traditional data infrastructures struggle, and most often fail, to operationalize. Developing new business capabilities—such as individual pricing for customers based on real-time profitability, as some insurance companies have done, automating credit decisions that lead to improved outcomes for banks and greater customer satisfaction,

The traditional data warehouse, through which the organization gains stability and financial transparency, must be scaled down and integrated with the high-speed transactional architecture that supports new products and services, as well as real-time reporting.

or running automated, more cost-effective strategic marketing campaigns as we've seen in the chemicals sector—demands new ways of managing data.

This does not mean, however, that legacy data and IT infrastructures must be trashed, or that new capabilities need to be bolted on. It does mean that the traditional data warehouse, through which the organization gains stability and financial transparency, must be scaled down and integrated with the high-speed transactional architecture that gives the organization the capability to support new products and services (as well as real-time reporting). This is the two-speed principle.

This new, complex technical environment requires companies to closely examine business use cases before making costly technology decisions, such as needlessly ripping out and replacing legacy architectures. Instead, it is preferable to use a capability-oriented reconceptualization of data management as an enabler of digital applications and processes.

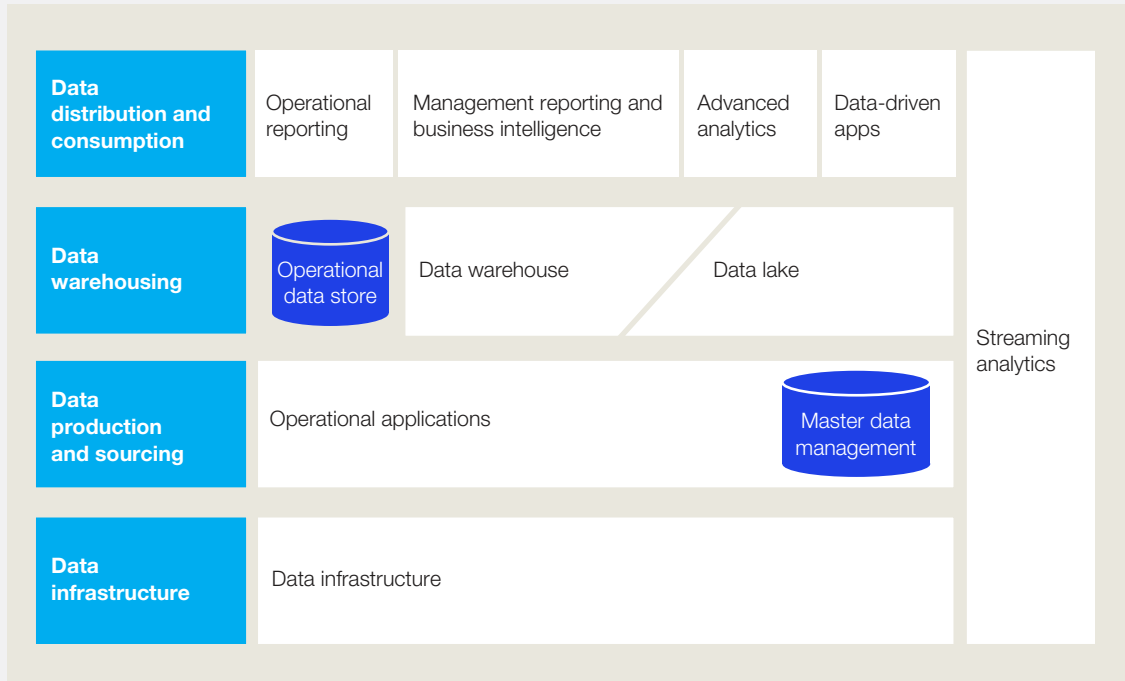
To implement an end-to-end digital data architecture, an enterprise needs first to develop a point of view on its current and, if possible, future business requirements, sketch its desired, flexible data-management architecture, and create a road map for implementation. To begin, one must identify the key business use cases.

To do this, we recommend a thorough review of best-practice use cases across industries that address common value drivers (financial transparency, customer satisfaction, rapid product development, real-time operational reporting, and so on). Then, the company should compare those use cases with its market position and strategic direction, prioritizing those that best reflect the company's situation and aspirations. Once those reference use cases are identified, the company can begin to define target data-architecture capabilities. In this process, the business leads and technology follows.

A layered data architecture has been applied successfully by many organizations

EXHIBIT

In a best-practice use case, a layered data architecture accommodates new digital capabilities.



across various industries, especially finance (exhibit). It can extend to accommodate new digital capabilities, such as the collection of unstructured data, real-time data processing, and streaming analytics.

Such an architecture combines both the traditional requirements of financial transparency via a data warehouse and the capability to support advanced analytics and big data. In a phrase, it's a two-speed approach.

The two-speed architecture adheres to three core principles:

1. a limited number of components with a clear demarcation of capabilities to manage complexity while providing the required functionalities, such as advanced analytics and operational reporting
2. layers that enable the transparent management of data flows and provide a single source of truth to protect against siloes and data inconsistencies (through the data warehouse, which models, integrates, and consolidates data from various sources)

3. integration of state-of-the-art solutions with traditional components, such as the data warehouse, to satisfy such new requirements as real-time processing, and an operational data store based on new database technologies

Companies have used this model in the following ways:

- to help think through and evaluate their options on an architectural level before discussing concrete technical solutions
- to map technology components against capabilities to manage and avoid redundancies while identifying gaps
- to create plans for stepwise transformations driven by business value while limiting business disruption

Getting physical with digital

For example, one of the largest banks in Scandinavia, understanding the business potential of advanced analytics, big data, and better data management to improve fraud detection and prevention, ATM location, and other initiatives, was eager to begin its digital data journey. It was facing intense competition and was considering making a massive, multimillion-dollar investment in its IT and data architecture.

A lot was riding on what the bank decided to invest in, where it decided to invest it, and how. It began by identifying key use cases that reflected the organization's most compelling strategic requirements: improved fraud detection, optimized location and allocation of branches, and more granular customer segmentation.

Based on this determination, the bank outlined a target architecture, founded on the best-practice reference model, that would enable the capabilities the bank desired and assess available solutions. Instead of ripping out its entire IT infrastructure, the bank decided to add a single Hadoop solution that allowed for storage and distributed processing of the bank's extremely large and frequently unstructured data sets across thousands of individual machines. This was especially useful in scaling the bank's high-frequency requirements for its online fraud-detection processes.

For branch location, allocation, and optimization, a Hadoop data lake (a management platform that processes flat, nonrelational data) used the bank's geospatial and population-growth data to determine where best to locate new branches and ATM machines. To improve its customer segmentation, the bank tested a new customer algorithm on the Hadoop database before rolling it out on its legacy data warehouse. This eliminated the typically costly and time-consuming back-and-forth process of develop, pilot, assess, validate, tweak, and pilot again that characterizes traditional data developments.

In this way, the bank achieved its primary business goals. It added new, differentiating capabilities, such as real-time analytics, and created real enterprise value with a relatively small technology investment, not the massive one originally contemplated. This was achieved by deciding what to invest in, where to invest it, and how—before buying systems and software that might not have served it nearly as well. Crucially, instead of first buying the technology, the bank built an in-house analytics team,

skimming off the cream of the local talent in the process.

Today, the bank is considered the leader in financial analytics in its market and sells analytics services to other financial institutions.

The bank knew that the time was ripe to get serious about digital transformation, made it a priority, and in doing so achieved what may well be an enduring competitive advantage, all without disrupting its business with a big-bang

technological transformation. It started with a clear view of its business goals, kept them front and center, and created a two-speed data architecture that worked.

The lesson here is that for many companies, it is both doable and cost effective to add analytics capabilities to an existing IT environment. But that requires a sound data architecture and a well-grounded approach to data management. ♦

Sven Blumberg is a partner in McKinsey's Düsseldorf office, and **Oliver Bossert** is a senior expert in the Frankfurt office, where **Hagen Grabenhorst** is a consultant and **Henning Soller** is an associate partner.

Copyright © 2018 McKinsey & Company. All rights reserved.



IN BRIEF

Executives on IT modernization

In these excerpts from McKinsey interviews, business and technology leaders explain how their companies approached the task of modernizing IT.

Stuart McGuigan, CIO, Johnson & Johnson

I proposed a leapfrog strategy to our leadership team. Instead of repairing the technology we had on the floor—incrementally upgrading servers and storage—we would bypass all that and move our work flow to a hybrid cloud environment. Of course, in 2012, that was much more of a bold statement than it is now. The company needed to understand that moving to the cloud didn't mean losing control and just giving people credit cards and allowing them to buy capacity with cloud-service providers. It meant rethinking our overall computing model. It meant taking advantage of cloud technology and agile development to shift from long product-planning cycles and a capital-intensive IT infrastructure to a highly variable infrastructure and cost structure.

From "Healthcare giant shares prescription for digital reinvention," April 2017

Ted Colbert, CIO, Boeing

Analytics will take billions off the bottom line if you figure out how people across the entire organization can grasp the opportunity—and how to democratize the capability. That can be tricky, because what you don't want is people trying to go create their own data platforms all over the place. It's that fragmentation that went wrong in the IT world 20 years ago and that makes it so hard today to get at data. So you need to keep working on projects that prove the power of data analytics and at the same time, in the background, plan the foundational architecture and work toward a common platform.

From "Data as jet fuel: An interview with Boeing's CIO," January 2018



© Photographer is my life/Getty Images

Leon Bedaux, head of digital IT, KPN

In the past, we would build IT looking out five to ten years, according to a very structured plan. Now, we build IT for only 18 months to two years out. The technology is always getting faster, and automation simpler. We need to keep up with that and the business use cases and requirements that may emerge. In this case, we replaced every platform and phased out traditional software—although in the beginning we kept the legacy technology running until we had fully replaced the required functionality.

From “KPN dials up a new digital strategy,” March 2017

Piyush Gupta, CEO, DBS

When we first started out along this road, we compared ourselves with emerging fintechs and the start-up world and concluded that we really had to digitize completely, not just by putting on digital “lipstick.” We made killing paper a big mantra in the organization, for instance, and were determined to go beyond just tacking on a bunch of digital apps at the front end—that’s the easy bit. We wanted to go all the way through to middleware and the back end.

A company like Uber has reimaged its processes and digitized everything from end to end, and that’s what we have done. This has required rethinking our technology architecture—hard for banks or any company sitting on legacy applications that are 30, 40, or 50 years old—so as to make it API-based and integratable with other applications, maybe open source.

From “The digital reinvention of an Asian bank,” March 2017

Ahmad Azhar Yahya, chief digital officer, Telekom Malaysia Berhad

Going digital is not about one big idea—it’s about solving 1,000 small problems together as one synchronized company. Initially, our main consideration in the digital transformation was moving from a traditional IT platform to a two-speed architecture where one part is customer facing, fast, and flexible and the other is a stable back end for transactions and business support. The two are connected via software that acts as a bridge. This bifurcation enables the development team to preserve core systems while making frequent changes to the front end based on customer feedback.

From “How a large established company built a digital culture,” September 2017

Eric Musser, managing director, robotics and workforce intelligence, Pega

Artificial intelligence, robotics, and automation are important, but do not forget the overall transformation of your underlying systems. Many carriers are still sitting on green-screen and client-server applications, leaving you to interact with multiple third-party browsers and different generations of Java-based applications. Just adding robots on top of this legacy environment is not the entire picture.

From “Automation at scale is driving transformative change across insurance,” June 2017 ♦

The complete interviews are available on [McKinsey.com](https://www.mckinsey.com).

Copyright © 2018 McKinsey & Company. All rights reserved.

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


July 2018

Designed by Global Editorial Services

Copyright © McKinsey & Company

mckinsey.com

 [@digitalmckinsey](https://twitter.com/digitalmckinsey)

 [linkedin.com/showcase/digital-mckinsey/](https://www.linkedin.com/showcase/digital-mckinsey/)