Digital/McKinsey: Insights

Navigating risk in the digital world

September/October 2017

The articles in *Digital McKinsey Insights* are written by consultants in the Digital McKinsey practice together with colleagues from 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 technology 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: digitalmckinsey@mckinsey.com.

Editor: Roberta Fusaro

Managing Editors:
Michael T. Borruso, Venetia Simcock

Art Direction and Design:
Todd Buxton, Nicole Esquerre

Editorial Production:
Elizabeth Brown, Heather Byer, Roger Draper, Katie Gilgour, Heather Hanselman, Gwyn Herbein, Katya Petriwsky, John C. Sanchez, Dana Sand, Sneha Vats, Belinda Yu

Cover photo:
© Donald Iain Smith/Getty Images

McKinsey Practice Publications

Editor in Chief: Lucia Rahilly

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

Features

Navigating risk in the digital world

Cybersecurity and risk management have become huge priorities for businesses and governments, as practically all of life goes digital. In this issue of *Digital McKinsey Insights*,[1] we consider what that means for leaders in both the private and public sector as they reinvent their organizations. This collection explores how leaders can mitigate risk and ensure productive and secure interactions with employees, customers, and suppliers.

It starts with dispelling the common myths executives tend to hold about their data-protection programs. As our partners in cybersecurity explain, it's a fallacy that all assets in the organization must be secured in the same way. A tiered approach, giving the most valuable assets the highest levels of protection, has proved to be much more effective at keeping critical business information safe while holding down cybersecurity costs. What's more, this tiered approach can also foster digital resilience in businesses.

Our experts discuss the need for more collaboration and stronger internal training and for protocols when it comes to cybersecurity: "You would get rid of half of your problems as an enterprise if you just train your folks and put controls in place," former IBM CEO Sam Palmisano notes in an interview.

Articles in this issue also explore the emerging rules and technologies that are changing the way companies and government agencies manage risk. What do businesses need to know about forthcoming General Data Protection Regulations and how best to comply with them? How can risk analytics help executives ferret out potential problems before they become full-blown crises?

Read on to find out more.

[1] Editor's note: To longtime readers of *McKinsey on Business Technology*, welcome to our revamped design and direction for the publication. The name, look, and feel are a bit different, but the content remains focused on addressing executives' biggest questions relating to the use of technology. We hope you enjoy the issue.

fotomay/Getty Images

Hit or myth? Understanding the true costs and impact of cybersecurity programs

Jason Choi, James Kaplan, Chandru Krishnamurthy, and Harrison Lung

Cybersecurity is a critical but often misunderstood aspect of companies' technology infrastructures. Here's how business and technology leaders can ensure that assets remain safe.

Companies are using all kinds of sophisticated technologies and techniques to protect critical business assets. But the most important factor in any cybersecurity program is trust. It undergirds all the decisions executives make about tools, talent, and processes. Based on our observations, however, trust is generally lacking in many organizations' cybersecurity initiatives—in part, because of competing agendas. Senior business leaders and the board may see cybersecurity as a priority only when an intrusion occurs, for instance, while the chief security officer and his team view security as an everyday priority, as even the most routine website transactions present potential holes to be exploited.

This lack of trust gives rise to common myths about cybersecurity—for example, about the

types of threats that are most relevant, the amount of spending required to protect critical data, and even about which data sets are most at risk. Perceptions become facts, trust erodes further, and cybersecurity programs end up being less successful than they could be. If incidence of breaches has been light, for instance, business leaders may tighten the reins on the cybersecurity budget until the CIO or other cybersecurity leaders prove the need for further investment in controls—perhaps opening themselves up to attack. Conversely, if threats have been documented frequently, business leaders may reflexively decide to overspend on new technologies without understanding that there are other, nontechnical remedies to keep data and other corporate assets safe.

In our experience, when there is greater transparency about companies' cybersecurity programs and trust among the various stakeholders, companies reap significant benefits. Businesses can make better decisions about their security priorities and response plans, as well as the training and investments required to hold attackers at bay. In this article, we explore four common myths executives tend to believe about cybersecurity, and we suggest joint actions business and IT executives can take to create more transparency and understanding company-wide about the technologies and processes that are most effective for protecting critical business information.

Separating myths from facts

Based on our work with companies across industries and geographies, we've observed that business and cybersecurity leaders fall under the sway of four core myths when discussing or developing protection programs for corporate assets.

Myth 1: All assets in the organization must be protected the same way

Not all data are created with equal value. The customer data associated with a bank's credit-card program or a retailer's loyalty-card program are of greater value than the generic invoice numbers and policy documents that companies generate in-house. Companies don't have endless resources to protect all data at any cost, and yet most deploy one-size-fits-all cybersecurity strategies. When faced with a request from the IT organization for more funding for cybersecurity, C-suite leaders tend to approve it reflexively (particularly in the wake of a recent security breach) without a more detailed discussion of trade-offs—for instance, how much is too much to spend on protecting one set of critical data versus another? Or if the company protects all external-facing systems, what kind of opportunities is it missing by not bringing suppliers into the fold (using appropriate policies and governance approaches)? Indeed, most business executives we've spoken with acknowledge a blind spot when it comes to understanding the return they are getting on their security investments and associated trade-offs.

In our experience, a strong cybersecurity strategy provides differentiated protection of the company's most important assets, utilizing a tiered collection of security measures. Business and cybersecurity leaders must work together to identify and protect the "crown jewels"—those corporate assets that generate the most value for a company. They can inventory and prioritize assets and then determine the strength of cybersecurity protection required at each level. By introducing more transparency into the process, the business value at risk and potential trade-offs to be made on cost would

then be more obvious to all parties. A global mining company, for example, realized it was focusing a lot of resources on protecting production and exploration data, but it had failed to separate proprietary information from that which could be reconstructed from public sources. After recognizing the flaw, the company reallocated its resources accordingly.

Myth 2: The more we spend, the more secure we will be

According to our research, there is no direct correlation between spending on cybersecurity (as a proportion of total IT spending) and the success of a company's cybersecurity program. Some companies that spend quite a bit on cybersecurity are actually underperforming the rest of the market with respect to developing digital resilience[1] (Exhibit 1). In part, this is because those companies were not necessarily protecting the right assets. As we mentioned earlier, companies often default to a blanket approach (protecting all assets rather than the crown jewels). Throwing money at the problem may seem like a good idea in the short term—particularly when an intrusion occurs—but an ad hoc approach to funding likely will not be effective in the long term. Business and cybersecurity leaders instead must come to a shared understanding of costs and impact and develop a clear strategy for funding cybersecurity programs. The business and cybersecurity teams at a healthcare provider, for example, might agree that protecting patient data is the first priority but that confidential financial data must also be secured so as not to compromise partner relationships and service negotiations. They could allocate resources accordingly. Without this shared understanding, business leaders may balk when a data breach occurs after they've funded significant changes in the security infrastructure. The lack of transparency and trust between the C-suite and the IT organization will only get worse.

Myth 3: External hackers are the only threat to corporate assets

It is true that threats from outside the company are a huge concern for cybersecurity teams, but there are significant threats inside corporate walls as well. The very people who are closest to the data or other corporate assets can often be a weak link in a company's cybersecurity program—particularly when they share passwords or files over unprotected networks, click on malicious hyperlinks sent from unknown email addresses, or otherwise act in ways that open up corporate networks to attack. Indeed, threats from inside the company account for about 43 percent of data breaches.[2]

Business and cybersecurity leaders must therefore collaborate on ways to improve internal risk culture. They must educate employees at all levels about the realities of cyberattacks and best practices for fending them off—for instance, holding town meetings, mounting phishing campaigns, or staging war-game presentations to familiarize employees with potential threats and raise awareness. Many of these activities will need to be led by the CIO, the chief security officer, or other technology professionals charged with managing cybersecurity programs. But none will be fruitful if the company's business leaders are not fully engaged in a dialogue with the cybersecurity function and if companies don't build explicit mechanisms for ensuring that the dialogue continues over the long term.

[1] Unless otherwise indicated, statistics relating to the composition and effectiveness of companies' cybersecurity programs are from the 2015 McKinsey Cyber Risk Maturity Survey.

[2] *Grand theft data*, Intel Security, 2015, mcafee.com.

EXHIBIT 1

Companies' spending on cybersecurity does not necessarily correlate with level of protection.

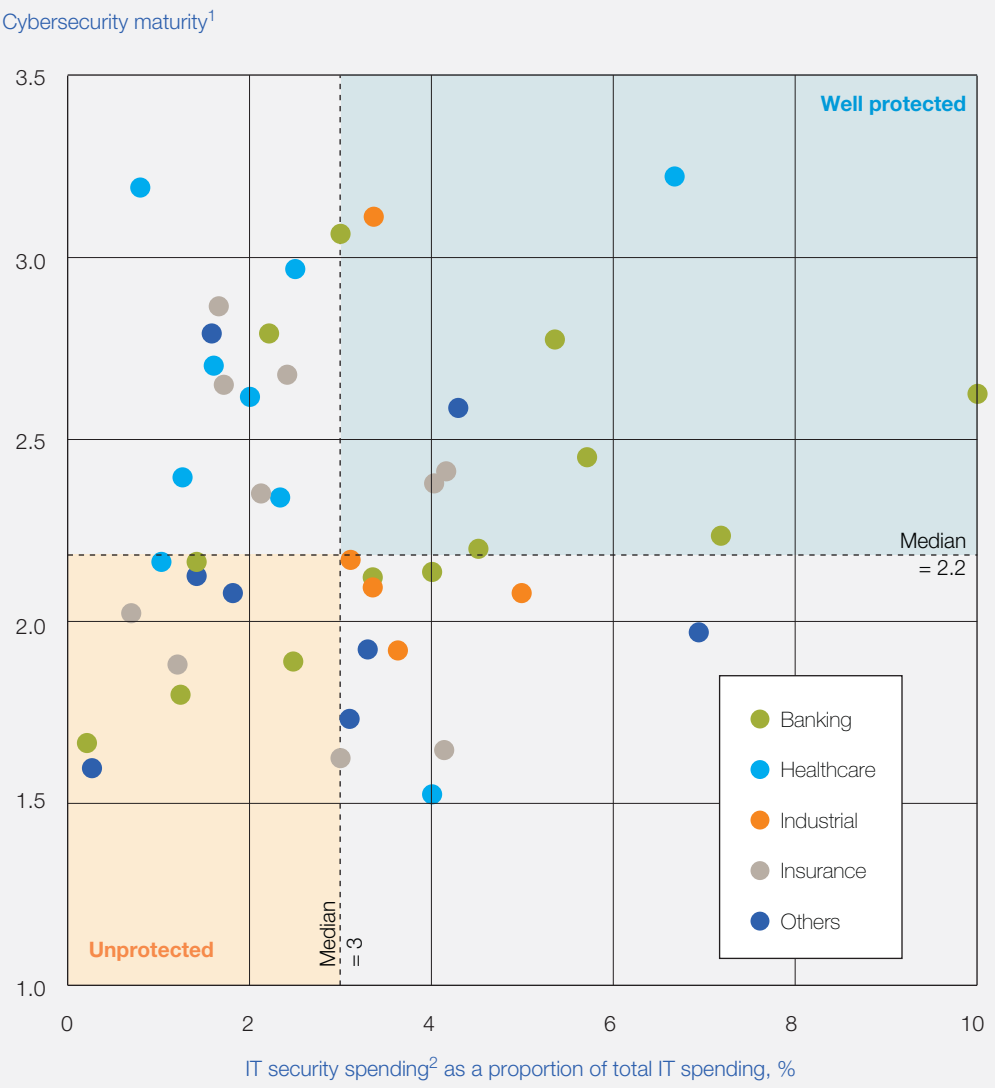

Note: Reflects responses from 45 companies in the Global 500 about their cybersecurity spending and capabilities.

[1]Companies' cybersecurity maturity is rated on a scale of 1 to 4, with 4 being the most mature (highest-level talent and capabilities).
[2]Spending is rated on a scale of 1 to 10; no companies allocated more than 10% of their budget to security.

Source: 2015 McKinsey Cyber Risk Maturity Survey

Business leaders at all levels must realize that they are the first line of defense against cyberthreats, and cybersecurity is never the sole responsibility of the IT department.

Myth 4: The more advanced our technology, the more secure we are

It is true that cybersecurity teams often use powerful, cutting-edge technologies to protect data and other corporate assets. But it is also true that many threats can be mitigated using less advanced methods. After all, most companies are not dealing with military-grade hackers. According to research, more than 70 percent of global cyberattacks come from financially motivated criminals who are using technically simple tactics, such as phishing emails.[3]

When companies invest in advanced technologies but do not understand how best to use them or cannot find properly skilled administrators to manage them, they end up creating significant inefficiencies within the cybersecurity team, thereby compromising the cybersecurity program overall.

Companies must, of course, explore the latest and greatest technologies, but it is also critical that companies establish and maintain good security protocols and practices to supplement emerging technologies—for instance, developing a robust patch-management program[4] and phasing out software for which vendors no longer provide security updates. This sort of foundation can help companies mitigate many of the biggest threats they may face. Consider the following example: a patch covering the vulnerabilities that could be exploited by the WannaCry cryptoworm was released March 14, 2017—some two months before the ransomware worked its way into more than 230,000 computers across more than 150 companies.

Building a culture of resilience

Rather than perpetuate myths, business and cybersecurity leaders should focus on bridging the trust gaps that exist between them. We believe most companies can do that when technology and business leaders jointly train their attention on two main issues of control: how to manage trade-offs associated with cybersecurity, and how to discuss cybersecurity issues and protocols more effectively.

How do we manage trade-offs?

Technology professionals have a role to play in reeducating the C-suite about best practices in cybersecurity spending—specifically, illustrating for them why a tiered approach to cybersecurity may be more effective than blanket coverage for all. The budget cannot grow and shrink depending on whether the company recently suffered a system intrusion. Cybersecurity must be considered a permanent capital expenditure, and allocations should be prioritized based on a review of the entire portfolio of initiatives under way. Business and technology professionals must work together to manage the trade-offs associated with cybersecurity.

When discussing which initiatives to invest in and which to discontinue, business and cybersecurity professionals can use a risk-categorization model with four threat levels denoted, from minor to severe. The cybersecurity team can then engage the

[3] *2017 Data breach investigations report*, Verizon, 2017, verizonenterprise.com.

[4] Patch management is the structured process of acquiring, testing, and installing code changes to an administered computer system.

C-suite in discussions about the most important data assets associated with each part of the business value chain, the systems they reside in, the controls being applied, and the trade-offs associated with protecting higher-priority assets versus lower-priority ones.

At a broader level, technology professionals can help the C-suite create benchmarks for cross-company and multiyear expenditures on cybersecurity initiatives that can be reviewed regularly—for instance, cybersecurity spending as a percentage of overall IT expenditures. The CIO and his or her team could create a capital-expenditure index for security investments to help the C-suite justify cost per risk-adjusted losses or cost per percentage of infrastructure protected. Or, technology and business professionals could jointly develop a formula for quantifying the upside of making improvements to the cybersecurity program. In this way, they can make clear decisions about which tools to buy and add to the existing cybersecurity architecture, which systems to upgrade, and which to retire.

Regardless of the metrics used, it is important to have a comprehensive, formal approval process for planning and reviewing capital expenditures associated with cybersecurity. Priorities must be set from a

EXHIBIT 2

Cybersecurity teams' access to the C-suite is limited.

How many direct reports away is the senior-most cybersecurity executive from the CEO?, % of survey respondents

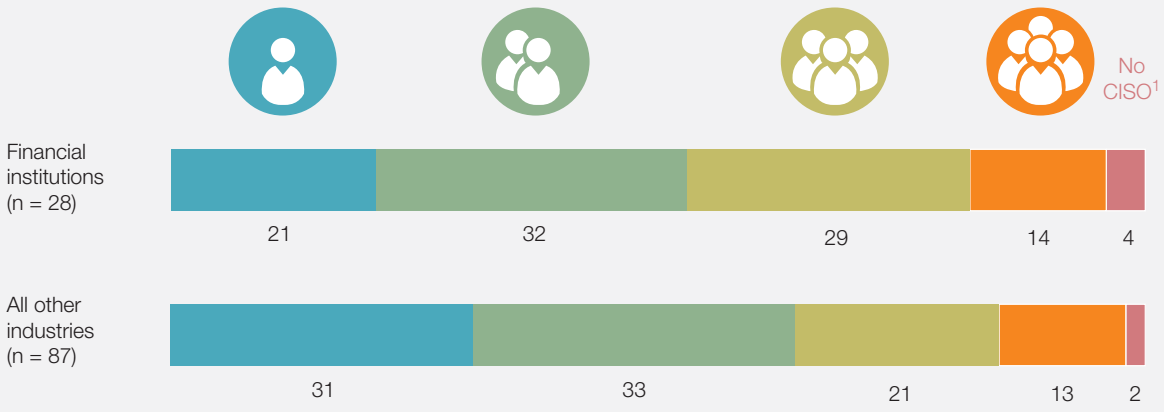

Note: Executives polled included chief information-security officers and other C-suite executives charged with making decisions about cybersecurity investments.

[1]Chief information-security officer.

Source: 2015 McKinsey Cyber Risk Maturity Survey

business perspective rather than a systems perspective. CIOs and chief security officers must collaborate with the business to identify those assets with the potential to generate the greatest amount of value for the business and develop a cybersecurity road map accordingly. The road map would illustrate the distribution of crown jewels across the organization and the greatest surface areas of exposure. It would outline current controls and the sequence for launching new security initiatives, looking two to three years out. Of course, business and cybersecurity executives would need to revisit these plans quarterly or annually to ensure that they are still relevant given changes to the environment. The road map would also define roles and responsibilities, as well as mechanisms by which the C-suite and the leaders in the cybersecurity function could monitor progress made against the plan and revise it accordingly.

How do we talk about cybersecurity?

Weak communication accounts for much of the lack of trust between business leaders and members of the cybersecurity function. Our research indicates that in most companies, cybersecurity professionals are at least two layers from the CEO in the corporate hierarchy, with few opportunities for direct discussion about protection issues and priorities (Exhibit 2). What's more, in about half of the companies we studied, there was little to no formal documentation shared by the cyber function with the C-suite about the status of their defense systems; many companies relied instead on occasional emails, memos, and notes (Exhibit 3).

Furthermore, when business and technology professionals do get in a room together, cybersecurity is usually discussed using highly technical language—for instance, "We already

EXHIBIT 3

Many cybersecurity teams use informal means to communicate with business leaders.

How do you summarize the status of defense systems to the chief information-security officer and business-level executives?, %

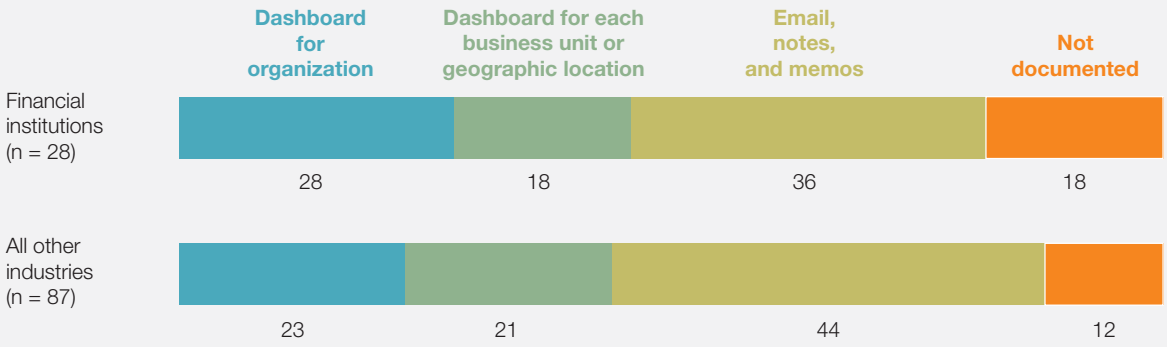

Note: Executives polled included chief information-security officers and other C-suite executives charged with making decisions about cybersecurity investments.

Source: 2015 McKinsey Cyber Risk Maturity Survey

have measures to cover all CVE, however APT is something we need to watch out for. With our current SVM and SIEM infrastructure, there is no way we can defend these advanced attacks."[5] Jargon notwithstanding, the technology and business professionals in the room all understand how critical it is to build a robust cybersecurity program given the potential effects on the bottom line if corporate assets are compromised. But each side is typically only getting half the story.

Instead of reporting that "ten vulnerabilities were remediated," for example, technology professionals can use visual aids and outcomes-oriented language to help business leaders understand potential security threats and ways to address them. A status update might be better phrased in the following manner: "Our cybersecurity team has patched a security hole in our customer-relationship-management system that could have given hackers access to millions of packets of our retail customers' data, creating $100 million in financial damage." Cybersecurity professionals could also clearly delineate and communicate levels of systems access for intended and unintended users—a database administrator would have greater privileges than frontline employees, for instance.

Finding a common vocabulary is important not just for ensuring clear communication between the C-suite and the cybersecurity function but also for raising awareness about potential cyberthreats and risks among employees throughout the company. Members of the cybersecurity function should schedule frequent, regular check-ins with staff at all levels to educate them about relevant cybersecurity topics—how to recognize a phishing email, for example—and to showcase the company's security capabilities. The cybersecurity team at one technology firm conducts "road shows" to demonstrate which systems are being scanned and how they are being monitored. One online retailer, meanwhile, includes details about its cybersecurity efforts in existing financial reports—for instance, reporting on its development of an antimalware scanner to protect the integrity of its recommendation engine, which helps drive advertising. It does this to illustrate that cybersecurity is part of the business process and can help drive revenue.

These discussions should take place regardless of whether the company is facing an imminent threat or not. The cybersecurity team at one company we observed shared with top leadership a simple breakdown of a typical security-event drill (Exhibit 4). The team wanted to give members of the board and the C-suite a step-by-step overview of what would happen in a typical attack—not just to prove the effectiveness of the company's security capabilities but also to familiarize individuals with potential threats so they might recognize them when they encounter deviations from the norm.

◆◆◆

As we mentioned earlier, technology leaders may have to lead the charge in forging direct communications, creating cost transparency, and identifying business priorities. But the tasks suggested will require experience in C-suite-level communication, budgeting, and strategy planning—some of which may be beyond the core skill set of those on the

[5] CVE stands for common vulnerabilities and exposures, APT stands for advanced persistent threat, SVM stands for security and vulnerability management, and SIEM stands for security information and event management.

EXHIBIT 4

Cybersecurity data theft has a pattern of event and response.

1

Insider takes sensitive data via flash drive

A disgruntled employee installs indexing malware in corporate systems and transfers files from servers to USB drive.

Visible hints

- Inquiry is made to senior executives about temp file being created and deleted.
- Slow laptops are reported to IT department and chief information officer.
- Help-desk ticket is sent to IT security lead.

Typical response

- Initially, the IT-security team does not realize that data are being threatened.
- Once the data are breached, the security team tries to determine best way to inform senior executives; the process is ad hoc, because protocols are not clear.

2

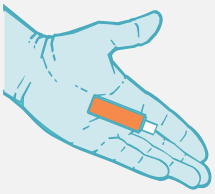

Insider gives or sells employee data to a cybercriminal

Cybercriminal uses old but valid credentials to access company servers and download employee records containing personally identifiable information (PII).

Visible hints

- Data-loss alerts are sent to the security lead in the IT organization.

Typical response

- Team focuses on the forensics of the alert but is not able to connect it to previous notifications.

3

Cybercriminal sells PII data to identity thieves on the black market

Identity thieves buy and use the employee data for fraudulent transactions.

Visible hints

- Based on individuals' and organization's complaints, the FBI detects the data breach and files a report with government affairs.

Typical response

- IT security reactively investigates employee data leak, trying to determine the scope of the breach.
- Team escalates event to privacy team.

4

Sensitive data are published on social media

Online bloggers publish video with references to the sensitive data stolen.

Visible hints

- An online video, found by employees, is sent to the head of communications.

Typical response

- The security team engages the communications group.

Source: 2015 McKinsey Cyber Risk Maturity Survey

cybersecurity team. To come up to speed more quickly, cyber leaders may want to reach out to others with relevant expertise—for example, vendors and partners who can share best practices. In the spirit of agile development, cybersecurity teams may also want to take on these activities in "launch, review, adjust" mode. They could update threat and risk profiles in one- to six-month sprints, thereby ensuring they are responsive to the latest trends and technologies.

Make no mistake, the time to foster greater transparency about cybersecurity is now. The board must have trust in the C-suite and its ability to handle security breaches without dramatically affecting the company's value and brand. The C-suite needs to trust the chief information-security officer's claims that every penny spent on improving the security of IT infrastructure is worth it. The company needs to trust that vendors can properly protect shared data or ensure service stability if breaches occur. And, of course, customers need to trust that their personal data are being carefully safeguarded behind corporate walls.

The C-suite and the cybersecurity function can no longer talk past one another; security must be a shared responsibility across the business units. It must be embedded in various business processes, with the overarching goal of building a culture of resilience. The companies that take steps now to build greater trust between the business and the IT organization will find it easier to foster a resilient environment and withstand cyberthreats over the long term. ◆

Jason Choi is a consultant in McKinsey's Hong Kong office, where **Harrison Lung** is an associate partner; **James Kaplan** is a partner in the New York office, and **Chandru Krishnamurthy** is a senior partner in the Atlanta office.

The authors wish to thank Suneet Pahwa and Chris Rezek for their contributions to this article.

Digital risk: Transforming risk management for the 2020s

Saptarshi Ganguly, Holger Harreis, Ben Margolis, and Kayvaun Rowshankish

Significant improvements in risk management can be gained quickly through selective digitization—but capabilities must be test hardened before release.

Digitization has become deeply embedded in banking strategy, as nearly all businesses and activities have been slated for digital transformations. The significant advantages of digitization, with respect to customer experience, revenue, and cost, have become increasingly compelling. The momentum to adopt the new technologies and operating models needed to capture these benefits continues to build. The risk function, which has seen significant growth in costs over the past decade, should be no exception. Indeed, we are starting to see digital transformations in risk create real business value by improving efficiency and the quality of risk decisions. A digitized risk function also provides better monitoring and control and more effective regulatory compliance.

Experience shows that the structural changes needed to bring costs down and improve effectiveness in risk can be accomplished

much like digital transformations in other parts of the bank. The distinguishing context of the risk environment, however, has important implications. First, risk practitioners in most regulatory jurisdictions have been under extreme pressure to meet evolving regulatory requirements and have had little time for much else. Second, chief risk officers have been wary of the test-and-learn approaches characteristic of digital transformation, as the cost of errors in the risk environment can be unacceptably high. As a result, progress in digitizing risk processes has been particularly slow.

This status quo may be about to change, however, as global banking leaders begin to recognize how substantial value can be unlocked with a targeted digital agenda for risk featuring fit-for-purpose modular approaches. In addition to the objective of capturing value, this agenda incorporates risk-specific goals. These include ensuring the ongoing effectiveness of the control environment and helping the risk function apply technology to better address regulatory expectations in key areas—like risk measurement, aggregation, and reporting.

What is digital risk?

Digital risk is a term encompassing all digital enablements that improve risk effectiveness and efficiency—especially process automation, decision automation, and digitized monitoring and early warning. The approach uses work-flow automation, optical-character recognition, advanced analytics (including machine learning and artificial intelligence), and new data sources, as well as the application of robotics to processes and interfaces. Essentially, digital risk implies a concerted adjustment of processes, data, analytics, and IT, and the overall organizational setup, including talent and culture.

Three dimensions of change: Processes, data, organization

To realize the full benefits of process and decision automation, banks need to ensure that systems, processes, and behaviors are appropriately fitted for their intended purpose. In the risk environment, prioritized use cases are isolated in such areas as credit underwriting, stress testing, operational risk, compliance, and control. In most banks, current processes have developed organically, without a clearly designed end state, so process flows are not always rational and efficient. Operational structures will need to be redesigned before automation and decision support can be accordingly enabled.

Data, analytics, and IT architecture are the key enablers for digital risk management. Highly fragmented IT and data architectures cannot provide an efficient or effective framework for digital risk. A clear institutional commitment is thus required to define a data vision, upgrade risk data, establish robust data governance, enhance data quality and metadata, and build the right data architecture. Fortunately, processes and analytics techniques can now support these goals with modern technology in several key areas, including big data platforms, the cloud, machine learning, artificial intelligence, and natural-language processing.

The organization and operating model will require new capabilities to drive rapid digitization. Although risk innovation takes place in a very specific, highly sensitive area, risk practitioners still need to create a robust culture of innovation. This means

putting in place the right talent and nurturing an innovative "test and learn" mind-set. Governance processes must enable nimble responses to a fast-moving technological and regulatory environment. Managing this culture of innovation in a way that is appropriate for risk constitutes a key challenge for the digitized risk function.

Adapting digital change to the risk context

Most institutions are digitizing their risk functions at a relatively slow pace, taking modular approaches to targeted areas. A few have undertaken large-scale transformation, achieving significant and sustainable advances in both efficiency and effectiveness. Either way, in the risk context, care must be taken when adapting test-and-learn pilots commonly used in digital transformations in other parts of the bank. Robust controls must be applied to such pilots, as the tolerance for bugs and errors in risk is necessarily very low. When digitizing processes relating to comprehensive capital analysis and review (CCAR), for example, solutions cannot be introduced into production before thorough testing has convinced designers and practitioners of their complete reliability and effectiveness. In certain other risk areas—such as monitoring and early-warning systems in commercial credit risk—banks can use test-and-learn approaches effectively.

Sizing the opportunity

Our experience suggests that by improving the efficiency and effectiveness of current risk-management approaches, digital risk initiatives can reduce operating costs for risk activities by 20 to 30 percent. Risk management at most global, multiregional, and regional banks is abundant with opportunity. Current processes are resource intensive and insufficiently effective, as indicated by average annual fines above $400 million for compliance risk activities alone (Exhibit 1).

The potential benefits of digital risk initiatives include efficiency and productivity gains, enhanced risk effectiveness, and revenue gains. The benefits of greater efficiency and productivity include possible cost reductions of 25 percent or more in end-to-end credit processes and operational risk, through deeper automation and analytics. Risk effectiveness can be strengthened with superior transparency, gained through better management and regulatory reporting and the greater accuracy of model outputs due to better data. Revenue lift can be achieved through better pricing or an enhanced customer and frontline experience—for example, by reducing the know-your-customer (KYC) cycle time from one week to under one day, or the mortgage-application process to under 30 minutes, from 10 to 12 days. Improved employee satisfaction can also be achieved through focusing talent on high-value activities.

Target risk processes: Credit risk, stress testing, and operational risk and compliance

The possible action areas for digital risk are extensive, but in our view three specific areas are optimal for near-term efforts: credit risk, stress testing, and operational risk and compliance. Alhough no one bank has fully digitized all three of these areas, we are seeing leading banks prioritize digital initiatives to realize discrete parts of the total savings available. The following discussion is based on actual digital risk initiatives across risk types and processes.

EXHIBIT 1

Digital risk management can significantly reduce losses and fines in core risk areas.

Note: Credit risk losses are gross charge-offs; operational and compliance risk losses do not include opportunity costs (such as unearned revenue due to operational risk events); the average total yearly fines are given for banks fined at least once in the period 2009–15.

Source: Bank holding company Y9C reporting forms; *Financial Times*' bank-fines data; McKinsey analysis

Credit risk

Credit delivery is hampered by manual processes for data collection, underwriting, and documentation, as well as data issues affecting risk performance and slow cycle times affecting the customer experience. Digital credit risk management uses automation, connectivity, and digital delivery and decision making to alleviate these pain points. Value is created in three ways: by protecting revenue, improving risk assessments, and reducing operational costs.

To protect revenue in consumer credit, digital risk strengthens customer retention. It improves the customer experience with real-time decisions, self-service credit applications, and instant credit approvals. The improvements are enabled through integration with third parties for credit adjudication and the use of dynamic risk-adjusted pricing and limit setting. One European bank is exploring the potential for digital risk to expand revenue in consumer credit within the same risk appetite. Digitized credit processes will permit faster decision making than the competition while the bank maintains its superior risk assessment.

Value is also created by improving risk assessment. Advanced analytics and machine-learning tools can increase the accuracy of credit risk models used for credit approvals, portfolio monitoring, and workouts. It can also reduce the frequency of judgment-based errors. The integration of new data sources enables better insights for credit decisions, while real-time data processing, reporting, and monitoring further improve overall risk-management capabilities. Operational costs are also reduced as credit processes are digitized. A greater share of time and resources can be dedicated to value-added activities, as inputs and outputs become standardized and paperless.

In addition to improving default predictions, we have seen credit risk improvements in these areas creating a revenue lift of 5 to 10 percent and lowering costs by 15 to 20 percent (Exhibit 2).

Stress testing, including CCAR

Banks find that significant value can be captured through a targeted digitization effort for stress testing, including CCAR. The current approach is highly manual, fragmented, and sequential, presenting challenges with data quality, aggregation, and reporting time frames and capacity. The processes are prime candidates for digital automation and work-flow tools.

The underlying stress-testing process is the starting point. The improvement program will aim at optimizing resources. Dedication of resources will be prioritized based on materiality of risk. Institutions can achieve additional efficiency through parallel processing, centralization, and cross-training of staff, as well as better calendaring. Templates and outputs are standardized, and “golden” sources for data are designated. The resulting process becomes increasingly transparent and effective. Process optimization is supported by digital-automation initiatives for data loading, overlays, Y14A reports, and the end-to-end review-and-challenge process. Real-time visualization and sensitivity analyses are digitally enabled as part of the transformation. In addition to optimizing stress testing directly, banks are also looking for opportunities to harmonize the data, processes, and decision-making models with business planning.

We have seen digitization in CCAR and stress testing bring significant cost improvements and—even more important—free up capacity so that experts can apply more insight and improve the quality and use of outputs (Exhibit 3).

Operational risk and compliance

At many global banks, manual processes and fragmented systems have proliferated across operational risk and compliance controls and activities. In anti–money laundering (AML), for

example, processes and data have become unwieldy, costs have skyrocketed, and efforts have become ineffective. Significant opportunities to increase the effectiveness and efficiency of AML operations lie in thorough end-to-end streamlining of the alert-generation and case-investigation processes.

In alert generation, digital risk improvements ensure that reference data available for use in the analytic engine is of high quality. Advanced-analytics tools such as machine learning are used to test and refine the case-segmentation variables and support "auto-adjudication" where possible. In addition,

EXHIBIT 2

An integrated digital risk program for consumer credit can protect revenue, improve risk assessments, and reduce operational costs.

EXHIBIT 3

There are many ways digitization can improve efficiency and effectiveness of comprehensive capital analysis and review (CCAR) and stress testing.

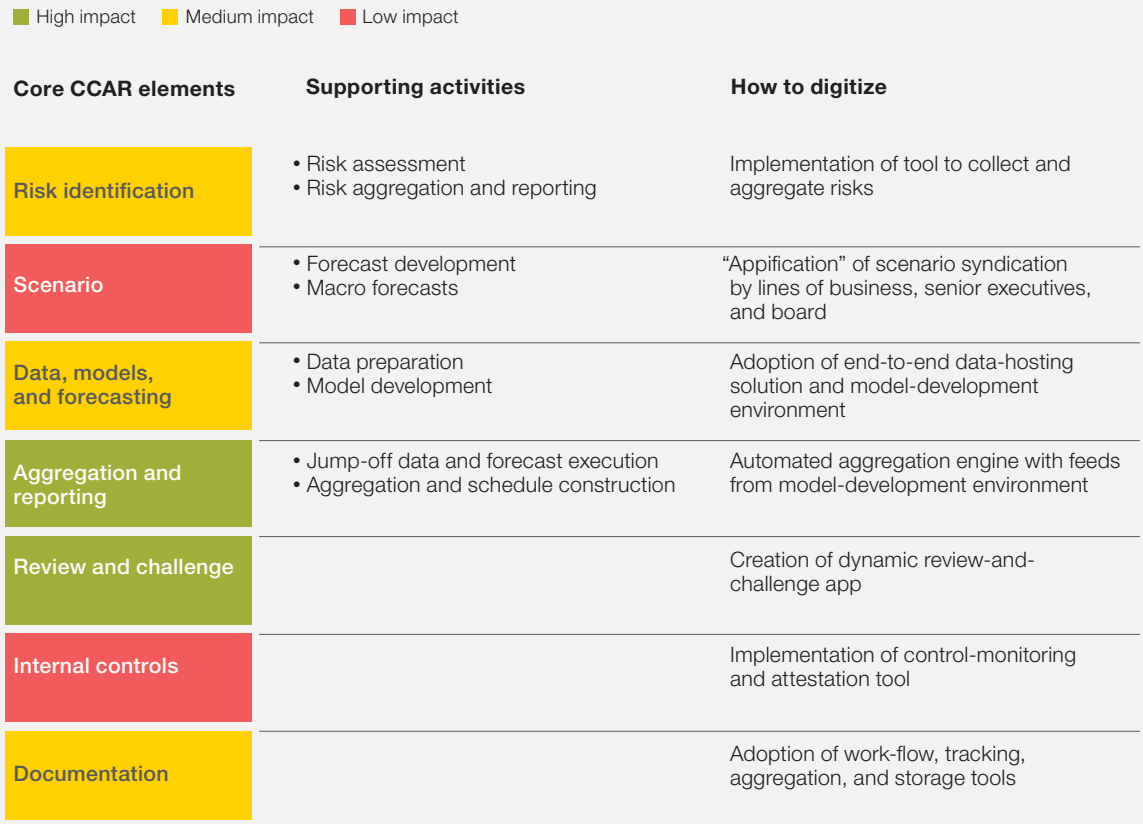

digitization and work-flow tools can support smart investigations and automated filing of suspicious-activity reports, an improvement that enhances the productivity of the investigation units.

Our experience of digital risk initiatives in AML is that they invariably improve effectiveness and efficiency, typically in the range of 20 to 25 percent. The overall impact of such improvement is even greater, however, given the large cost base of this function across institutions and the risk of not identifying bad actors.

Digital risk is different

A digital risk program must be designed in recognition of those aspects of the

Digital risk is never a self-contained effort—it will depend on data from all businesses and functions.... Innovative approaches such as agile and digital labs provide effective options to implement solutions incrementally.

risk function that distinguish it from other functions, such as frontline digital sales. For risk, regulators will not accept the characteristic approaches of traditional digital transformations. A live launch of "minimum viable products" to be tested and refined in production is not an appropriate path for most risk activities. Most approaches to digitization focus on improving the customer experience. Digital risk will involve some actual external customers, such as in credit delivery, but in most areas the focus will be on internal customers, stakeholders, and regulators. Moreover, digital risk is never a self-contained effort—it will depend on data from all businesses and functions. Development thus proceeds at a pace limited by the careful management of these interdependencies. Innovative approaches such as agile and digital labs provide effective options to implement solutions incrementally.

Direct impact will be felt in cost and risk reduction

While digital risk offers clear opportunities for significant cost reduction, the impact on revenue is less obvious but implicitly understood by leaders. Frontline digital transformations are often aimed at direct revenue improvement; proof of this impact from digital risk programs is more elusive, since risk is an enabling function. Faster turnaround times for loan applications is a typical digital risk improvement. This will likely drive higher lending volumes and, consequently, increased revenue—even if the correlation cannot be precisely determined. Given the indirect impact on revenue, digital risk programs should focus primarily on reducing risk and cost. The exception is digital credit, where the case for revenue lift will be clearer.

Designing a program

An effective digital risk program begins with chief risk officers asking the right questions—those that point the institution toward specific initiatives for digital innovation. "Can we reduce the time needed for structured credit approvals to a few minutes?" "How can we increase straight-through processing rates?" "How can we improve the efficiency and streamlining of KYC activities to reduce pain points in the account-opening process?" "How can we make CCAR less sequential and resource intensive?" "How can we improve the timeliness of reporting to meet regulatory objectives?" "What value

can we extract from better use of internal data?" "What is the incremental benefit of including new data sources?" The answers will help shape initiatives, which will be prioritized according to current resource-allocation levels, losses and regulatory fines, and implementation considerations, such as investment and time.

Digital risk programs can incorporate the familiar design features of digital transformations, such as zero-based process and interface redesign and an agile framework. The testing and refinement, however, takes place entirely within a controlled environment. The design approach, which can be modular, must also be comprehensive, based on a thorough review of risk activities, appetite, and policies.

The designs cannot be migrated into production until they have been thoroughly tested and syndicated, often with regulatory bodies. Because of its highly sensitive environment, risk is digitized end to end over a longer timeline than is seen in customer-service areas. Specific capabilities are developed to completion and released discretely, so that risk management across the enterprise is built incrementally, with short-term benefits.

The anatomy of a transformation

A digital risk program can get a running start by capturing high-value opportunities first. The anatomy of the transformation will resemble that of other digital transformations, with the usual three stages: 1) priority initiatives are identified according to the value at stake and the feasibility for near-term implementation, 2) digital solutions are designed to capture that value and tested and revised according to stakeholder input, and 3) the improvement is introduced into production, with continued capability building to embed the design, engineering, and change management into the operating model and invest in the right capabilities and mind-sets.

The opportunities identified in stage one are matched in stage two with digital and other solutions that will reduce waste and optimize resources while improving standardization and quality. These solutions will involve work-flow automation, digital interfaces, and the use of advanced analytics and machine learning. The technology design may use a two-speed architecture to support fast innovation in IT while allowing the main IT infrastructure to operate normally. New functionality is rigorously tested prior to migration into production, to ensure a smooth, error-free transition for critical risk functions. Iterative test-and-learn processes take place within environments featuring higher control standards than typical elsewhere. Stakeholder feedback and often regulator syndication are obtained prior to production release.

In the third stage, where the innovation is introduced into production, the organization focuses on change management. In itself, this is no different from typical digitization programs in other business areas. The focus is on embedding the design into the operating model and continuing to invest in digital capabilities to build momentum for further launches. Having the right talent in place, whether drawn from internal or external sources, is the key to a successful transition to digital risk.

◆◆◆

The path to digital risk will be a multiyear journey, but financial institutions can begin to capture significant value within a few months, launching tailored initiatives for high-value targets. As the risk function becomes progressively digitized, it will be able to achieve higher levels of efficiency, effectiveness, and accuracy. In the future, risk management will be a lean and agile discipline, relieving cost pressures, improving regulatory compliance, and contributing to the bank's ability to meet escalating competitive challenges. The first steps toward that future can be made today. ◆

Saptarshi Ganguly is a partner in McKinsey's Boston office, **Holger Harreis** is a partner in the Düsseldorf office, and **Ben Margolis** is an associate partner in the New York office, where **Kayvaun Rowshankish** is a partner.

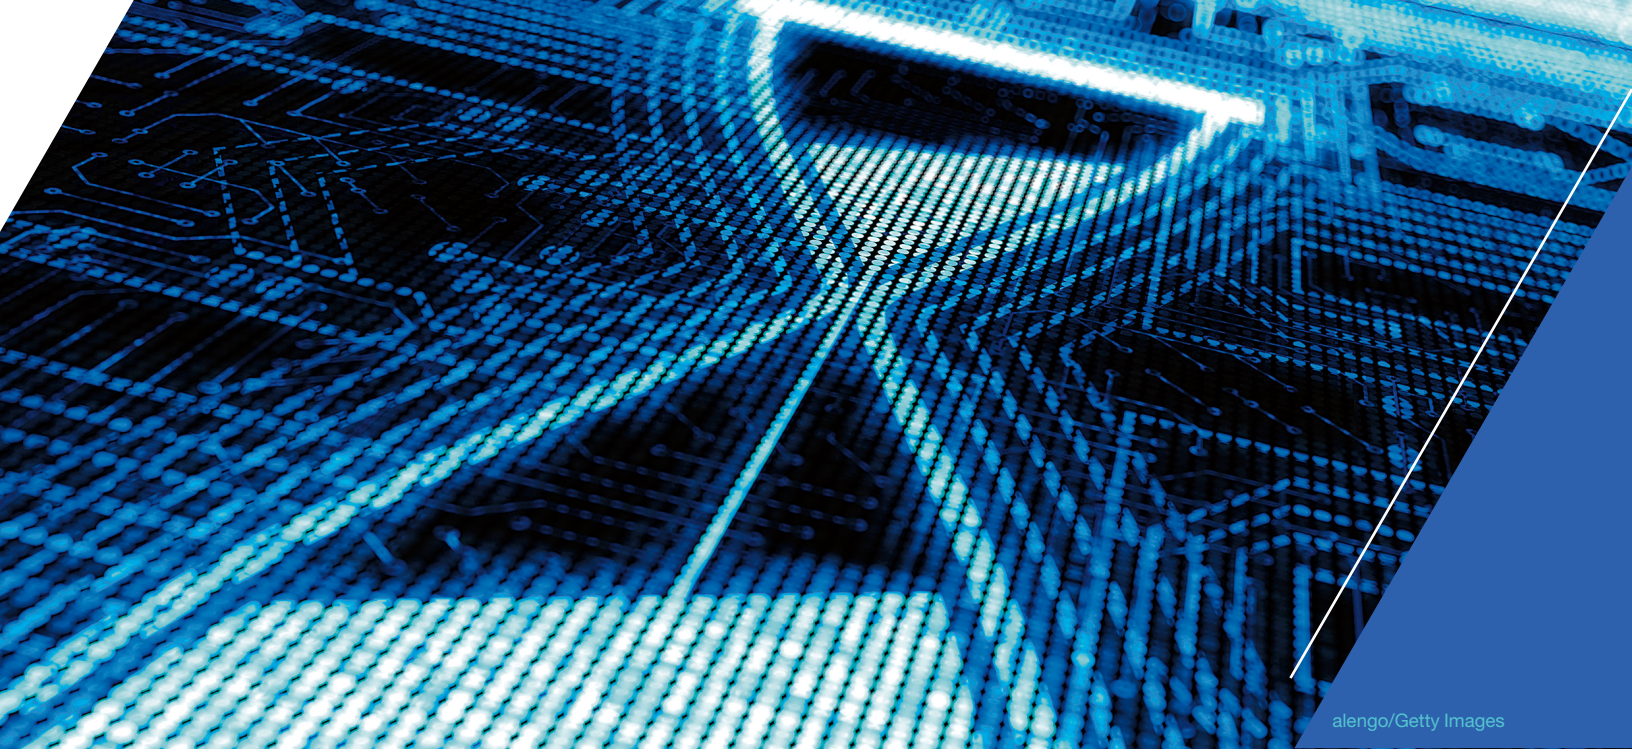

Tackling GDPR compliance before time runs out

Daniel Mikkelsen, Kayvaun Rowshankish, Henning Soller, and Kalin Stamenov

Data protection has always been important. Now it's becoming urgent. Here's a primer on how companies can adapt to the new rules.

Europe is on the brink of a sea change in its data-protection laws. In fact, when the General Data Protection Regulation (GDPR) takes effect on May 25, 2018, the effects will reverberate far beyond the continent itself. The GDPR goes further than harmonizing national data-protection laws across the European Union and simplifying compliance; it also expands the reach of EU data-protection regulation and introduces important new requirements. It seeks to ensure that personal data are protected against misuse and theft and to give individuals in the European Union control over how data relating to them are being used. Any entity that is established in the European Union or that processes the personal data of individuals in the European Union in order to offer them goods or services or to monitor their behavior—whether as customers, employees, or business partners—will be affected. Any failure to comply with the regulation could incur severe reputational

damage as well as financial penalties of up to 4 percent of annual worldwide revenues (see sidebar, "The GDPR: Key facts," for a synopsis of the new rules).

After an initial wait-and-see approach, many companies in Europe and beyond—including those in Asia, the Middle East, and the United States—are starting to set up sizable compliance programs. Yet our recent surveys of major companies revealed that a third of the executives in the sample felt their organizations still had a long way to go on the road to compliance.[1] As the GDPR is based on principles rather than rules, the onus is on individual companies to determine implementation in their particular context (exhibit). This process is fraught

[1] We surveyed 19 executives at McKinsey's European General Data Protection Regulation (GDPR) Roundtable in February 2017; most were chief information-security officers. In May 2017, we conducted an informal online poll of eight US executives who were leading GDPR efforts.

EXHIBIT

The General Data Protection Regulation sets out guiding principles for data protection.

Principle	Explanation
Lawfulness	Data should be processed only when there is a **lawful basis** for such processing (eg, consent, contract, legal obligation)
Fairness	The organization processing the data should provide data subjects with **sufficient information** about the processing and the **means to exercise their rights**
Transparency	The **information provided** to data subjects should be in a **concise and easy-to-understand format** (eg, the purpose of consent should not be buried in a lengthy document of terms and conditions)
Purpose limitation	Personal data may be collected only for a **specific, explicit, and legitimate purpose** and should not be further processed
Data minimization	The processing of personal data should be **adequate, relevant, and limited to what is necessary** in relation to the purposes for which those data are used
Accuracy	Data should be **accurate and kept up to date**
Storage limitation	**Data should not be held any longer than necessary** in a format that **permits personal identification**
Security	Data should be processed in a manner that **ensures security and protection against unlawful processing, accidental loss, damage, and destruction**
Accountability	The **data controller** is responsible for **demonstrating compliance**

Source: Regulation (EU) 2016/679 of the Council of the European Union, European Commission, and European Parliament

The GDPR: Key facts

The scope of the General Data Protection Regulation (GDPR) is broad, covering any information that can be linked to an identifiable individual (such as search-engine entries, employee authentication, payment transactions, closed-circuit-television footage, and visitor logs) in any format (structured or unstructured) and in any medium (online, offline, or backup storage). The regulation introduces stringent consent requirements, data-subject rights, and obligations on organizations that gather, control, and process data. Its core requirements cover the following:

Documentation. Organizations should maintain a record of data-processing activities and be ready to present it to the regulator at any time.

Legal basis. All data processing should have a legal basis, such as the consent of the data subject or the need to fulfill a contract or legitimate business purpose.

Rights of data subjects. Organizations should implement rights such as the right to be forgotten (or, more accurately, to data erasure), the right to data portability, the right to object, the right to revoke consent, and the right to restrict processing.

Security. Organizations should protect data through means such as encryption or "pseudonymization" and have effective operational procedures and policies for handling them safely.

Third-party management. Vendors and suppliers, including outsourcing partners, should be required to protect personal data and should be monitored to ensure that they do so.

Privacy by design. Any organization planning a new technology, product, or service should consider data-protection requirements from the beginning of the development process.

Breach notification. Data breaches likely to result in high risk to individuals' rights and freedoms should be reported to the authorities within 72 hours, and subsequently to the data subjects as well in certain cases.

The new regulation will be enforced via national supervisory authorities within the European Union that are granted wide-ranging enforcement powers and sanctions, such as the power to ban data processing. The fines for failure to comply will be high, as much as 4 percent of annual worldwide revenues. The GDPR also allows individuals to seek civil actions (including class-action lawsuits) against organizations that violate their data-protection rights.

with uncertainty, and many companies are struggling to understand how they can best interpret, measure, and monitor compliance. Below we examine some of the main stumbling blocks and identify the steps that successful companies are taking to overcome them.

Why businesses are struggling with GDPR compliance

From our survey and conversations with executives, we have identified a number of ways that compliance efforts are falling short:

- **Underestimating the scope of the regulation.** Some of the executives who responded to our survey were not fully aware of the breadth of the GDPR, regarding it as merely an enhancement to existing regulations. Conversely, others felt that complying with the new provisions—especially the business and IT implementation of data-subject rights—would be onerous for their organization, and they were doubtful they would reach full compliance by May 2018. Indeed, only one of the 19 participants in our European survey believed his/her company would fully comply by the deadline.

- **Uncertainty about how to interpret the requirements.** The GDPR sets out a number of principles that organizations should observe in processing personal data, but most companies have yet to decide how to put these principles into practice. For instance, under the principle of lawfulness, any organization processing personal data must have either the consent of the individuals concerned or some other lawful basis for that processing. Although the GDPR provides guidance on what might constitute a lawful basis—such as to carry out a contract, to comply with a legal obligation, or to serve the legitimate interest of the data controller or a third party—that guidance leaves a great deal of room for interpretation. In practice, we see organizations taking very different views on issues such as the extent to which new consents are required from customers. In all these matters, companies will need to consult with lawyers. And lawfulness is not the only principle in the GDPR where there is uncertainty over interpretation. Take the accuracy principle, for example: it requires organizations to keep personal data up to date and take every reasonable step to rectify inaccuracies, but it is left to the organizations themselves to decide what steps they consider reasonable.

- **Slowness in identifying the additional security measures needed.** As the GDPR uses similar language to the current directive, many organizations are relying on their existing security measures, including protocols for particular customer segments, for compliance. However, as they build their records of processing activities, they will need to ensure that these measures are proportionate to the risks pertaining to different types of personal data. This calls for a structured approach to defining data risk and the measures necessary for mitigation—"pseudonymization," anonymization, encryption, deletion, and so on.

- **A struggle to build and maintain a comprehensive inventory of all their personal data–processing activities.** To satisfy this requirement, most of the banks we spoke with are relying initially on

manual methods, typically using an internal survey to identify relevant data-processing activities within their organization. Such an approach may suffice for creating the inventory in the first place, but it is unlikely to be adequate to the task of keeping the inventory current and readily available to the regulator on demand. Sustainable processes and tools for maintaining detailed records have proved elusive so far for many organizations.

- **Lack of capabilities to fulfill their obligations.** Many companies are struggling to identify and develop the capabilities they will need to execute data subjects' rights in a timely manner. Consider, for example, the right to data portability. If a wealth-management firm receives a request from a customer to hand over all of her personal data to a different institution, what capabilities will it need to compile these data and transmit them to the new wealth manager? Under the GDPR, the data covered by the portability requirement are not confined to the personal data an individual provides and the transactions they perform, but includes observed data, such as search history, location, and so on. Building IT capabilities to fulfill these requirements may require banks to consolidate data from disparate systems, create new authentication methods, and introduce external APIs.

Steps to a successful GDPR effort

Drawing on our industry observations and regulatory experience, we have identified a number of actions that contribute to a successful GDPR effort and can help overcome some of the difficulties outlined above. Our advice is to check whether your institution is already taking these steps, and, if not, act now while there is still time.

- **Assign ownership of the program to a cross-functional task force.** A typical GDPR program does not have a natural owner in the organization; the challenge of ensuring compliance requires an approach that cuts across functions and businesses. All of the teams involved—legal, compliance, the business, IT, risk, and others—must commit to and share responsibility for a road map for change. Senior leadership approval and buy-in is vital so that the program is securely anchored in a company's overall strategy.

- **Define the scope of your GDPR program and use a business lens to determine what should be ready for May 2018.** Most of the companies we surveyed believed they would not be fully compliant by the implementation date, so it is important to identify which aspects of the regulation and which data assets are critical to compliance and make them a priority. This means not only understanding legal requirements but also defining what risks the business is willing to accept, and what value it seeks to extract from the program.

- **Develop an in-house interpretation of GDPR requirements** that identifies the big strategic questions they pose and seeks to address them early on. The approach should reflect the most likely scenario, take the industry view into account, and neither downplay nor exaggerate the impact of the regulation. Adopting a black-or-white legalistic approach may not be helpful, so it will

Companies are using GDPR-inspired reforms as an opportunity to build greater flexibility into their data platforms so that they not only comply with the new provisions but also respond more readily to future regulatory changes.

be important to stay close to peers as well as regulators and see what practical steps they are taking to comply. As your program progresses, take regular pulse checks to keep it on track. Given the heavy IT requirements, your program validation should be performed well before the second quarter of 2018 to allow time for course correction, if needed.

- **Assess your GDPR readiness to uncover any gaps** and plan measures to fill them, whether that involves modifying marketing processes to secure customer consent, developing new in-house data-protection measures, or carrying out vendor evaluations. Bear in mind that adopting manual solutions to satisfy requirements such as ensuring data portability can lead to high ongoing running costs. Building an automated solution at the outset—such as APIs for data transfer—could simplify compliance and reduce costs in the long run if you believe there will be sufficient demand (for instance, for data portability) to justify the investment involved.

- **Begin building a "golden record" of every personal data–processing activity** in the organization to ensure compliance and traceability. This goes beyond documenting the system inventory and involves maintaining a full record of where all personal data come from, what is done with them, what the lawful grounds for processing are, and whom the data are shared with. Map business or functional activities that use personal data and get the owners of these activities to complete a detailed questionnaire about the data processing involved. In parallel, work with vendors and internal IT experts to build tools and processes to maintain the inventory in steady state. This can be done as part of your software-development life cycle and data-protection impact assessments. Some companies adopt special data tools to discover personal-data assets and provide compliance reporting, but these tools have yet to be proven at scale in the marketplace.

- **Define your organizational setup for data protection.** Designating a data-protection officer (DPO) is not enough.

Companies also need to weigh the pros and cons of different organizational set-ups in order to arrive at a reporting structure that enables the DPO to operate independently; to interact effectively with the chief information-security officer, chief privacy officer, and heads of legal, compliance, and risk; and to report to the highest level of management. Having decided on the new structure, companies then need to determine the resources required to support it and fulfill their data-protection responsibilities more broadly.

- **Define the uncertainties in interpreting the GDPR requirements, and identify unacceptable risks to your business and IT.** Many aspects of GDPR will be gradually resolved through industry practices and codes of conduct, regulatory guidance, or the court system. Interpretations of what is legally acceptable may also change over time. Frequent interactions with legal and business partners on compliance, legal issues, cybersecurity, application development, third-party vendor management, operations, marketing, and so on will help companies build a shared understanding of what they need to do. Beyond pure compliance, IT and the business should work together to define where the program should go the extra mile to minimize reputational risk, maintain customer trust, and avoid last-minute IT scrambles. This may involve implementing more stringent consent requirements, prominently announcing opt-out possibilities, implementing tougher-than-necessary security measures, and setting a high bar for sending personal data to third parties.

- **Consider strategic value.** Half the chief information-security officers in our sample regarded GDPR primarily as a hindrance to their business. Undoubtedly, the regulation will impose a burden on organizations, and with a matter of months to go before implementation, companies are racing to limit any negative impact it may have. However, what many leaders miss are the benefits that can be realized through a GDPR program. A well-conceived program can help an organization to build customer trust, improve customer relationships, establish better data controls, and improve internal data handling and availability. One company is taking advantage of its GDPR program to reengineer its master data-management platform so that all parts of the organization have a complete picture of all personal data on any given customer. Other companies are using GDPR-inspired reforms as an opportunity to build greater flexibility into their data platforms so that they not only comply with the new provisions but also respond more readily to future regulatory changes. Seen in this light, a GDPR program can be an opportunity to embark on a wider data transformation that will benefit the whole business.

◆◆◆

The steps above will help any institution get on the right track to meet next year's implementation date. GDPR should not be taken lightly. Organizations that fail to comply could face high fines, civil actions, and reputational damage, while those that use their GDPR program to spur a broader data transformation may be able to capture additional business flexibility and value. These are compelling reasons to treat the

new regulation as a high priority for the whole organization, not just the risk, legal, and compliance functions. And with the implementation date imminent, companies need to act fast. ◆

Daniel Mikkelsen is a senior partner in McKinsey's London office. **Kayvaun Rowshankish** is a partner in the New York office, where **Kalin Stamenov** is a consultant. **Henning Soller** is an associate partner in the Frankfurt office.

The authors wish to thank Malin Strandell-Jansson for her contributions to this article

Hoxton/Tom Merton/Getty Images

Protecting your critical digital assets: Not all systems are created equal

Piotr Kaminski, Chris Rezek, Wolf Richter, and Marc Sorel

Top management must lead an enterprise-wide effort to find and protect critically important data, software, and systems as part of an integrated strategy to achieve digital resilience.

The idea that some assets are extraordinary—of critical importance to a company—must be at the heart of an effective strategy to protect against cyberthreats. Because in an increasingly digitized world, protecting everything equally is not an option. The digital business model is, however, entirely dependent on trust. If the customer interface is not secure, the risk can become existential. Systems breaches great and small have more than doubled in the past five years, and the attacks have grown in sophistication and complexity. Most large enterprises now recognize the severity of the issue but still treat it as a technical and control problem—even while acknowledging that their defenses will not likely keep pace with future attacks. These defenses, furthermore, are often designed to protect

the perimeter of business operations and are applied disjointedly across different parts of the organization.

Our research and experience suggest that the next wave of innovation—customer applications, business processes, technology structures, and cybersecurity defenses—must be based on a business and technical approach that prioritizes the protection of critical information assets. We call the approach "digital resilience," a cross-functional strategy that identifies and assesses all vulnerabilities, defines goals on an enterprise-wide basis, and works out how best to deliver them. A primary dimension of digital resilience is the identification and protection of the organization's digital crown jewels—the data, systems, and software applications that are essential to operations.

Burgeoning vulnerabilities, finite resources, fragmented priorities

In determining the priority assets to protect, organizations will confront external and internal challenges. Businesses, IT groups, and risk functions often have conflicting agendas and unclear working relationships. As a result, many organizations attempt to apply the same cyber-risk controls everywhere and equally, often wasting time and money but in some places not spending enough. Others apply sectional protections that leave some vital information assets vulnerable while focusing too closely on less critical ones. Cybersecurity budgets, meanwhile, compete for limited funds with technology investments intended to make the organization more competitive. The new tech investments, furthermore, can bring additional vulnerabilities.

The work to prioritize assets and risks, evaluate controls, and develop remediation plans can be a tedious, labor-intensive affair. Specialists must review thousands of risks and controls, and then make ratings based on individual judgment. Some organizations mistakenly approach this work as a compliance exercise rather than a crucial business process. Without prioritization, however, the organization will struggle to deploy resources effectively to reduce information-security risk. Dangers, meanwhile, will mount, and boards of directors will be unable to evaluate the security of the enterprise or whether the additional investment is paying off.

All data and systems are not created equal

In any given enterprise, some of the data, systems, and applications are more critical than others. Some are more exposed to risk, and some are more likely to be targeted. Critical assets and sensitivity levels also vary widely across sectors. For hospital systems, for example, the most sensitive asset is typically patient information; other data, such as how the emergency room is functioning, may even be publicly available. Risks to priority data include breach, theft, and even ransom—recall that a Los Angeles hospital paid a $17,000 Bitcoin ransom to a hacker that had seized control of its systems. An aerospace-systems manufacturer, on the other hand, needs to protect intellectual property first and foremost, from systems designs to process methodologies. A financial-services company requires few controls for its marketing materials but is vulnerable to fraudulent transactions; its M&A database, furthermore, will need the best protection money can buy. Attackers can be

individuals or organizations, such as criminal syndicates or governments with significant resources at their command. The attacks can be simple or sophisticated, the objectives varying from immediate financial reward to competitive or even geopolitical advantage.

Cybersecurity spending: When more is less

In the face of such diverse threats, companies often decide to spend more on cybersecurity, but they are not sure how they should go about it.

- A global financial-services company left cybersecurity investments mainly to the discretion of the chief information-security officer (CISO), within certain budget constraints. The security team was isolated from business leaders, and resulting controls were not focused on the information that the business felt was most important to protect.

- A healthcare provider made patient data its only priority. Other areas were neglected, such as confidential financial data relevant to big-dollar negotiations and protections against other risks such as alterations to internal data.

- A global mining concern focused on protecting its production and exploration data but failed to separate proprietary information from information that could be reconstructed from public sources. Thus, broadly available information was being protected using resources that could have been shifted to high-value data like internal communications on business negotiations.

These examples illustrate the need for a unified, enterprise-wide approach to cyber risk, involving the business and the risk, IT and cybersecurity groups. The leaders of these groups must begin to work together, identifying and protecting the organization's critical digital assets as a priority. The process of addressing cyber risk will also have to become technologically enabled, through the implementation of work-flow-management systems. Cybersecurity investment must be a key part of the business budget cycle and investment decisions must be more evidence-based and sensitive to changes.

The business-back, enterprise-wide approach

The key is to start with the business problem, which requires a consideration of the whole enterprise, and then to prioritize critical risks. This work should be conducted by a company-wide team composed of key people from the business, including those in product development and in the cybersecurity, IT, and risk functions. The team's main tasks are to determine which information assets are priorities for protection, how likely it is that they will be attacked, and how to protect them. To function, the team must successfully engage the leaders of several domains. They need to work together to discover what is most important—no mean challenge in itself. The best way to get started is to found the team on the agreement that cyber risks will be determined and prioritized on an enterprise-wide "business back" basis. In other words, the team will first of all serve the enterprise. Critical risks, including the impact of various threats and the likelihood of occurrence, will be evaluated according to the dangers they pose to the business as a whole.

Guiding principles

The following principles can help keep companies on track as they take the unified approach to prioritizing digital assets and risk:

- **Start with the business and its value chain.** The effort should be grounded in a view of the business and its value chain. The CISO's team, particularly when it is part of the IT organization, tends to begin with a list of applications, systems, and databases, and then develop a view of risks. There are two major flaws to this approach. First, it often misses key risks because these can emerge as systems work in combination. Second, the context is too technical to engage the business in decision making on changes and investments. By beginning with the business, the team encourages stakeholder engagement naturally, increasing the likelihood that systemic exposures will be identified.

- **The CISO must actively lead.** In addition to being a facilitator for the business's point of view, the CISO should bring his or her own view of the company's most important assets and risks. By actively engaging the business leaders and other stakeholders as full thought partners, the CISO will help establish the important relationships for fully informed decision making on investments and resource allocation. The role of the CISO may thus change dramatically, and the role description and skill profile should be adjusted accordingly.

- **Focus on how an information asset can be compromised.** If an information asset is exposed by a system being breached, the vulnerability of this system should be considered, even if the system's primary purpose does not relate to this information asset.

- **Focus on prioritization, not perfect quantification.** The team needs only enough information to make decisions on priority assets. It does not need highly precise risk quantifications—these would be difficult to produce and would not make a difference in deciding among investment options.

- **Go deeper where needed.** The same level of analysis is not needed to quantify all risks. Only for particularly high-impact or complex risks should the team invest in deeper analyses. It should then decide on and acquire the information needed to make more informed investment decisions.

- **Take the attacker's view.** Risk reviews and vulnerability analyses must not focus solely on the value of the information to the company and the ascertainable gaps in its defenses. The profiles of potential attackers are also important: Who wants the organization's information? What skills do they possess? Thinking about likely attackers can help identify new gaps and direct investment to protect the information that is most valuable to the most capable foes.

A flexible, systematic process with a designed platform

The object of the enterprise-wide approach is to identify and remediate gaps in existing control and security systems affecting critical assets. The solution, in our experience, will be an end-to-end process, likely requiring multiple development iterations, including a detailed account of hundreds of assets. A work-flow system and asset database would be an ideal tool for supporting this complex process, allowing focus on prioritizing risks. A flexible, scalable, and secure online application can be easy to use while managing all the inventory and mapping data, the

rigorous risk and control evaluations, sector-specific methodologies, and rationales for each risk level. The platform can also support detailed data to be used when needed as the team undertakes analysis of the priority assets and gaps and makes the recommendations that will shape remediation initiatives.

In developing this approach for clients, McKinsey experts defined the following five key steps:

1. **Identify and map digital assets,** including data, systems, and applications, across the business value chain. This can be accelerated by applying a generalized-sector value chain and a common taxonomy for information assets and then customizing these to the organization.

2. **Assess risks for each asset,** using surveys and executive workshops. By basing this analysis on the business importance of the asset, the organization will have identified its crown jewels.

3. **Identify potential attackers, the availability of assets, and current controls and security measures** protecting the systems through which access can be gained to the assets, using surveys and workshops similar to those in step two.

4. **Locate where security is weakest** around crown-jewel assets and identify the controls that should be in place to protect them, by comparing the results of these assessments using dashboards.

5. **Create a set of initiatives to address the high-priority risks and control gaps.** Implementation will involve a multiyear plan, including timelines for follow-up reviews. Once the initial assessment is complete, this plan becomes a living document, regularly refreshed to reflect new data, systems, applications, risks, and mapping, as well as progress to remediate known vulnerabilities (see sidebar, "An institution's progress").

The process promotes cyber-risk transparency, answering key stakeholder questions: What are our inherent information risks? Where is our organization vulnerable? How big (and where) is the residual exposure? What remediation actions should we prioritize? How do we know if what we did is working? Information-risk trade-offs can be defined based on a perspective on value at risk across the company. This helps the C-suite and board discuss information-security risk with regard to enterprise value, providing transparency on what risks they are willing to accept and why.

Results inform budget and investment decisions, helping to satisfy both regulatory and shareholder expectations. With investments targeted to best protect the most sensitive digital assets, costs are held down as the digital resilience of the organization is elevated. To build digital resilience into their operations, furthermore, the process can help organizations create periodic assessments to highlight trends and new gaps. Risk managers can then develop new initiatives prioritized according to the enterprise's global needs.

◆◆◆

Organizations in sectors with higher digital maturity will benefit the most from this approach, including financial services, manufacturing, and healthcare. They face the tough task of fully protecting their most

An institution's progress

One financial institution that used our approach was able to identify and remediate gaps in the control and security systems affecting its critical assets. The change program began with a risk assessment that had highlighted several issues. Business and IT priorities on cybersecurity spending were found to be somewhat out of alignment, while communication on risks and risk appetite between risk management and businesses was less than optimal. The lack of agreement among stakeholder groups consequently stalled progress on a mitigation plan for cyber risk.

In response, the company established a unified group, which developed a work plan to protect critical data. The team inventoried all systems and applications in all business units, validating the results with key stakeholders to ensure completeness. It then identified critical data and performed a risk assessment with input from the stakeholders. The team was now able to identify the critical information assets based on potential risk impact. The level of control in each system was also evaluated, as the team mapped information assets to the systems and applications where they reside and isolated gaps between current and needed controls.

The critical data assets requiring additional protection were identified globally and by business unit. The systems and applications holding critical data that needed remediation could then be addressed. The team developed a series of detailed scenarios to reveal system vulnerabilities and help stakeholders understand what could happen in a breach. A comprehensive set of prioritized initiatives and a multiyear implementation plan was then created. The data resulting from this process are continually updated and provide regular guidance on budgeting decisions and board reviews.

important assets while not stifling business innovation. To achieve this balance, the business, IT, risk, and other functions will have to work together toward the same, enterprise-wide end—to secure the crown jewels so that the senior leaders can confidently focus on innovation and growth. ◆

Piotr Kaminski is a senior partner in McKinsey's New York office, **Chris Rezek** is a senior expert in the Boston office, **Wolf Richter** is a partner in the Berlin office, and **Marc Sorel** is a consultant in the Washington, DC, office.

The authors wish to thank Oliver Bevan and Rich Cracknell for their contributions to this article.

luismmolina/Getty Images

Finding a strategic cybersecurity model

Protecting critical and sensitive information is of paramount importance in business and government, but plans must be in place to handle inevitable breaches too.

Cybersecurity has become one of the biggest priorities for businesses and governments, as practically all of life migrates its way to data centers and the cloud. In this transcript of an episode of the McKinsey Podcast, recorded at the Yale Cyber Leadership Forum in March, Sam Palmisano, chairman of the Center for Global Enterprise and the retired chairman and CEO of IBM, and Nathaniel Gleicher, head of cybersecurity strategy at data-and-cloud-security company Illumio, speak with McKinsey about how governments and companies can vastly improve their cyberprotections.

Roberta Fusaro: First up from the forum is Sam Palmisano, who, in this wide-ranging conversation with McKinsey's Marc Sorel, makes the case that strong cybersecurity programs are critical for improved innovation and economic growth.

Sam, thank you for joining us today. I want to talk a little bit about your work on the Commission on Enhancing National Cybersecurity. What was the original mandate? What was the process by which you came up with your findings? And what were some of the most surprising results?

Sam Palmisano: Thank you, Roberta. The thing was that President Obama had reached the conclusion that the digital economy or the Internet is so fundamental now to economic growth and society that something needed to be done to make some recommendations to enhance it or strategically position it for the future. A great example is the Internet of Things, because it's no longer just phones and desktop computers. It's everything in life. It's self-driving cars, it's thermostats, it's music players, it's cameras.

Now, you take this infrastructure and you're making billions of things that are computers, which are smart devices. But that's what they are, they're chips with software with all the vulnerabilities, unless you design for security from the beginning. And you've taken this problem, and you've put it on steroids.

The complexity there is one of getting consensus to go fast and address the issues prior to billions of things being out there that aren't secure, which is the path we're headed down.

Marc Sorel: How do you think about what the private sector, and to some extent the social sector, need to do now to be part of that?

Sam Palmisano: We need to form a private–public collaboration. The reason for it: the government doesn't have the skills to do this itself. We spent nine months crawling through their statements of skill. They can argue all they want. They don't [have the skills]. That doesn't mean that elements of government don't have some skill. To take the intelligence agencies out of this discussion and get to that commercial side, they don't have the capability. They need the capability, so they had to form a partnership. The skills exist in the academic community and in the research universities and in the technology community.

Marc Sorel: Did you all, as a commission, see a model in the market today for what that collaboration could look like?

Sam Palmisano: There are established entities within government that are a combinationof academic, private sector, government, and technical. A lot of the technical communities come together.

General Keith Alexander ran the Cyber Command Center. There were probably 20 of us that met once a quarter for five, six years. The same guys that were running IBM, Google, Dell, Microsoft, HP, and Verizon, plus all the government-appropriate people, would meet quarterly. The technical people would meet even more often to tackle some of these issues, and it was self-funding. We solved problems just by pitching in because it was in the best interest of everyone to solve some of these issues, and in the best interest of the industry, because you wanted to expand and grow.

To really do this, though, this was going to require funding. To solve the problem we're talking about, it's going to require some amount of money and research, like a DARPA[1] or related fund, pick something like that as the funding source that government can coordinate, and then convene this body. Then do the work as we suggest. Now, the work is going to get complicated. Because there's two pieces to it. One is, let's say, for example, to come up with a standard for the Internet of Things that you would put in this device, this object. Then within that object, you'd have

[1] Defense Advanced Research Projects Agency.

this standard. Then you'd also have a nutrition label on the standard. We called it the Cyber Star. It's like the health seal that says, "OK, if you're the manufacturer and you've complied with these standards, you get the star." You get the Cyber Star.

There were also guys that recommended a thing called secure, they call it clean pipes. With clean pipes, there are a lot of policy implications, a lot of criminal-justice-systems implications. But technically, you could create a clean path and you could have a secure path, and you could argue for certain areas where life is threatened.

In the autonomous vehicles or drones or things where people could actually be seriously injured or die, you'd want a secure, clean path. You don't want this on the open Internet.

Marc Sorel: So you're talking about creating a separate secure environment for these privileged parts of the ecosystem.

Sam Palmisano: Right. Think of it as a commercial virtual private network but beyond that. Put that on steroids from an encryption and security perspective. For all these Internet of Things devices. Health, heart monitor, things you're putting in your body. Pacemakers, et cetera. Defibrillators. Those kinds of things. Not Fitbits that you wear on your wrist, but serious things that could do serious harm like stop your heart. You want to have that information flowing in a secure way. In an encrypted, secure way. That doesn't mean everything should be that. If you're sharing your photos with friends, I don't think you need that level or cost associated with those kinds of technologies.

Marc Sorel: You're basically saying at some level, there should be a tiering of Internets to acknowledge the degree of security required for different pieces of the ecosystem to communicate.

Sam Palmisano: That is a solution to the problem. Now you have to make it commercially viable, which gets you into things like net neutrality. But if you were to technically solve the problem, you would begin to architect portions of the Internet. You can't go re-create the past. It's just too old, it's too cobbled together. Let that be what it is.

But anything that's life threatening or takes down the infrastructure or the world economy. Let's just start there. The premise or the assumption is that you can't solve this in the Internet as it exists today. It just was too complicated. It's too convoluted. It's too open by design. That's why it was so successful, because it was an open architecture. We had all these debates, all of the technical guys, and said, "Look. We used to do this 40 years ago." ATMs never got hacked. Money didn't start spitting out on the curb and stuff because it was a secure connection. It was a proprietary network. We know how to do it technically.

But there are people that did these things for years. We've moved onto an open innovative system, which is terrific because it drives innovation at a much more rapid pace. It also gives people more economic opportunity to participate. That's a big plus. But in certain areas where you're dealing with, let's say, major societal issues, we ought to go back to some of the classical approaches to how you design the systems.

Roberta Fusaro: Most people today would say, "If I had to place a bet on who's going to gain ground on whom and put space between themselves, it's the attackers that are going to continue to distance themselves in terms of capability from the defenders in terms of their capabilities." Do you agree with that?

Sam Palmisano: Eighty percent of the cybersecurity issues that have occurred in the commercial world are internal process and people. It's not just the disgruntled employees who got fired and therefore they gave somebody their access codes. It's also people who didn't protect their access codes or they tape it to their computer. Or they leave it in the top drawer of their desk, and the cleaning people can go get the stuff. You would get rid of half of your problems as an enterprise if you just train your folks and put controls in place.

It's a combination of monitoring, process training, audit people. Did you follow the process? So there's an accountability in the system. That'll clean up a lot of the stuff in the commercial world. Password authentication and end points. If the civilian side of government, .gov, did those things, they would clean up probably 95 percent of their problems and save a ton of money, too.

We also talked about this idea, which never got traction in the commission report, but we thought it was a good idea, where you basically would create a national ID like a credit bureau. You could create this national ID foundry where you get your birth certificate. You also get your digital identity at birth, and that digital identity is secure and protected. Now, you can modify for simple things—sharing your photos on the Internet—or you can modify it for very sophisticated things like financial transactions, your health information.

Marc Sorel: Why didn't it catch on?

Sam Palmisano: In the commission itself?

Marc Sorel: Yeah.

Sam Palmisano: What we did was say, further studies should take place, and we recommended that Treasury would look at, further look at creating this kind of an entity. We also looked at commercial insurance as well, and the purpose of commercial insurance.

The purpose of commercial insurance was that if you agreed on the standards, and therefore you complied with those standards, you should be able to get higher liability coverage at a lower rate than somebody who didn't.

Our view was that would drive up the adoption rate because people are going to want to find an insurance policy for cyber. That's going to happen. How do you get these companies to make the investments to move up the risk-protection curve? Well, you make it to their advantage by having insurance that says, "We could audit those standards. And if you've complied with those standards, like burglar-alarm systems or fire alarms in your home, you're going to get higher liability coverage at a lower rate." That's to make it an economic-based system versus a government-mandated system. The commission was very biased toward private-sector solutions versus government-mandated solutions. You need a private sector or an economically driven set of motivations to solve the problem.

Roberta Fusaro: This has been a fascinating conversation. Thank you, Sam, for taking the time to be with us today.

Sam Palmisano: Oh, thank you. It was great being with you.

Next up from the forum is Nathaniel Gleicher, who describes how businesses can learn a lot from the model of protection used by the US Secret Service.

Roberta Fusaro: Welcome, Nathaniel. Thank you for joining us today for the McKinsey Podcast.

Nathaniel Gleicher: No problem. Glad that I could join.

Roberta Fusaro: Your company has been providing cyber options for four or five years now, and I'm wondering how you've seen the market change over that time in terms of what customers are looking for or technologies that have emerged.

Nathaniel Gleicher: There used to be a perception that cybersecurity was black magic, particularly outside of the technical community, and that outside of that community, people would sort of say, "I don't understand this. Just make it work." As long as you don't hear anything, no news is good news. The increasing scope and scale of breaches and the degree to which organizations are moving into these exposed environments has changed that. If you look at business leaders, I think they are focused on how do you quantify the risks that you face, and how do you measure the benefit that you're getting from the solutions you invest in? It's a much more quantification-driven industry than it used to be. I don't know that we're very good at quantification yet. But the desire to quantify is an important change.

Roberta Fusaro: Apart from quantification, are there other hot topics in cyber that you're seeing or managing right now?

Nathaniel Gleicher: Sometimes I think we do cybersecurity like fourth graders play soccer. Chase the ball across the field, the whole group runs. There are always hot topics. What's interesting to me is that we've known for a while there are a few steps that if you took them, environments would be much more secure.

If you think about encrypting data, using strong passwords, whitelisting your applications, segmenting your environment, patching your vulnerabilities, and people generally haven't done that because it's been hard to figure out how to do that at scale across these large organizations.

One of the biggest challenges that we face in cybersecurity today is that we don't really have a single, coherent strategic model to describe how to protect an environment. There are a lot of tactical models, so if you look at the SANS[2] top 20, if you look at NIST,[3] if you look at some of these other frameworks, they will tell you, you should be investing in encryption. You should be investing in segmentation. You should be investing in certain kinds of detection. They'll tell you all

[2] SysAdmin, Audit, Network, and Security.

[3] National Institute of Standards and Technology.

the tools you should use, and you can think about how to line them up, but it's very tactical. It's hard to find a model that lets you pull back and think about the threat as a whole.

I'm starting to see groups of companies trying to solve that problem, trying to think, how do you do these steps that don't seem all that sexy but that actually drive to security.

Roberta Fusaro: What are some of the potential remedies?

Nathaniel Gleicher: If you look at security disciplines through the ages, whether it's law enforcement, executive protection, physical security for locations, military security, any of these sort of well-built disciplines, the foundation of every security discipline is understanding the environment you're protecting and exerting control over that environment.

In cybersecurity, we are not good at understanding the environment we're defending. Most organizations don't understand the network. They don't understand what's connected and what's communicating with what. Because of that, they have relatively few options to control that environment. I mentioned before a few simple things people could do to strengthen their environment. Those are all about control, and what I mean by control, people often think there's prevention, keeping the bad guys out, and then there's detection and response, catching them once they get in.

Those are both important components. In general today, people would tell you, you can't invest all in one or the other, that prevention by itself isn't enough. People are going to get in. What people miss in that debate is the reason detection and response works is because you understand your environment, and you control it.

If you don't know where your high-value assets are, and if you don't know what connects to them, how someone would

> "The president [of the United States] is a lot like a high-value asset in a data center, in that he's very valuable, very targeted, and also very exposed.... The job [of the Secret Service] is really about managing risk, which is similar to the way we're protecting assets in the data center."

access them, it's incredibly hard to know what you need to protect. If you don't have the resources to control that, you're defending an open field. So you have hundreds and hundreds of paths you need to defend, potential connections you need to worry about, and the attacker gets to move first. On the flip side, if you invest to understand your environment first, and control your environment first, it actually makes detection and response better.

Roberta Fusaro: What are some ways to identify the crown jewels, the things that really do matter? I can imagine that could be an incredibly difficult task, given all the assets that companies manage.

Nathaniel Gleicher: It's different for every organization, to some degree, but it's about understanding business risk. The question is, what are the assets that I defend, or that my business relies on, such that if they were exposed or compromised, it would fundamentally harm the way I do business?

Whether that's healthcare data about your customers, or customer information, whether that's the systems on which your business runs, whether that's the exchanges across which you connect, every business has a different set of factors they need to judge. But often, if you think in terms of business risk, we're pretty good at figuring that out, because businesses have been measuring and concerned about risk for quite some time. It's just a question of translating that and understanding the technical implications.

A model that I like to use when I think about this is the way the Secret Service protects the president. The president is a lot like a high-value asset in a data center, in that he's very valuable, very targeted, and also very exposed. The Secret Service doesn't get to take the president, put him in a box somewhere, and have him not talk to anyone. He's constantly talking to people, so the job is really about managing risk, which is similar to the way we're protecting assets in the data center.

When the Secret Service is protecting the president, if you imagine the president speaking in an auditorium, the Secret Service shows up months before the president is going to be there. The first thing they do is they map the auditorium to understand that if the president's going to be here, speaking on this stage, here are all the attack vectors. Here are all the ways someone could reach the president. An auditorium is built for openness, so there are going to be a lot. The Secret Service tries to control that environment, to shrink the number of attack vectors. The reason they do this is, as we said before, if you have to watch 100 attack vectors, it's really expensive, and you're really spread out thin. If you have to watch 20, you're in much better shape as a defender. So you can say we don't leave this doorway open, and no one's going to sit in this portion of the auditorium. You can close things down to simplify your environment. That's important for a lot of reasons, but the biggest reason is it makes detection much easier.

If there's a section of the auditorium where no one is supposed to sit, that doesn't necessarily mean no one will show up there. People always do strange things. But if someone does, you know they've broken a policy. It's not a false positive. There's no risk of confusion. You can simply react, and it lets the Secret Service act much more quickly

because rather than basing their actions on uncertain analysis, they're basing it, they create firm boundaries. When someone breaks a boundary, they know what to do. If the Secret Service wanted to, they have a lot of resources, they could put a metal detector at every seat in the auditorium.

They could put one at every single seat. They could get the best metal detector in the world. The problem is, they would never do that. They would get thousands and thousands of alerts and lots of them would be because someone had a particularly heavy watch on, or had change in their pocket. Whatever it might be. To test those alerts, they would have to send Secret Service agents out into the auditorium to check each one. And Secret Service agents are really expensive, and they're rare. It takes a long time to train them. They're hard to find. What you really want to do is take your precious resource, your Secret Service agents, and you want to direct them at the hardest, smallest slice of the problem.

So take that and apply it to the data center. If you are detecting everything everywhere, and you don't have control over the environment, you're going to get a lot of alerts. The statistics we see right now back that up. Organizations get 500, 1,000 critical alerts a day, which is a huge number of alerts that supposedly you have to deal with.

On average, organizations say they have the capacity to investigate something like 1 percent of them. So you're investigating 1 percent of all these critical alerts. Quickly, you start to turn things off because that data is dirty. If you're following the model, you would do the same thing the Secret Service does. You don't put a metal detector everywhere. What you do is you control the environment. You limit the places people can be, the paths they can take, so you know where to watch. So you know if this is my high-value asset in my data center, then if anything strange happens there, obviously it should be my highest priority. If anything strange happens in something connected to it that might be a secondary priority. You can start to prioritize these alerts and focus on the problems that matter more.

Roberta Fusaro: What are some of the policies or regulations that are emerging that business executives need to concern themselves with?

Nathaniel Gleicher: In a lot of ways, 2017 will be a year of regulation in cybersecurity. Not exactly the regulation people think about. I don't know that it'll come from DC. SWIFT,[4] the financial-transactions organization, recently put out controls that all of its members need to comply with to segment and protect their SWIFT application.

This is in response to all the criminal activity targeting SWIFT applications. That's one. The New York State Department of Financial Services, the financial regulator, put out controls around cybersecurity quite recently. The European Union recently put out a new general data-protection regulation, which has a whole range of controls built into it, but there are specific pieces around where is data stored, and how is it stored, which raise serious concerns for companies.

There are a lot of pieces coming out from different places that, depending on what industry you sit on, you need to watch. The pattern that I'm seeing, though, is

[4] Society for Worldwide Interbank Financial Telecommunication.

each of these has components that require organizations to do a better job exerting control over the data in their possession.

Organizations have said, "My data just pools in all these places. I don't even know where it is. It moves through these systems too fast for me to follow." It has been acceptable for companies not to know answers to these technical questions. You're seeing these regulations start to come out that push back on that. There's this increasing requirement on organizations to understand what's happening in those systems, and where that data's going.

Roberta Fusaro: How might this increased oversight affect companies' ability to innovate? So many new business models are data and analytics driven.

Nathaniel Gleicher: There's this old apocryphal joke that if we built cars like we built computers, cars would go 500 miles an hour, get 500 miles a gallon, and blow up once a week. We've made this choice, historically, around computer and Internet innovation that the consequences of unreliability aren't all that high.

We'd rather have rapid innovation, but what's happening now is more and more, you see the technical world, the Internet world, colliding or reconnecting with the physical world, whether it's autonomous cars, whether it's health innovation like you're seeing, whether it's integrating smart solutions into the home, whether it's integrating smart solutions into our transportation framework.

There are more and more opportunities integrating technology and smart solutions into the financial systems that our society runs on. There are more and more opportunities for surprisingly small bugs to cause very big chain effects in the physical world. The push and pull that you're seeing is how do you maintain the pace of innovation that has been so valuable, and such an engine of economic growth, an engine of competitive edge for us, while still mitigating the risks of all of these autonomous systems, and more and more sophisticated systems that are impacting the physical world.

Roberta Fusaro: What are the opportunities for VCs and start-ups in this changing environment?

Nathaniel Gleicher: There are huge opportunities in pointing artificial-intelligence solutions and orchestration solutions at problems that are incredibly hard to do at scale for large organizations. We tend to think of cybersecurity as a technology solution because that's convenient.

The truth is, it's really an organizational solution. If you only have one computer, obviously, anyone can make a computer secure by turning it off. But if you have one computer, if you have one system, a sophisticated defender is going to be much better able to protect that than if you have a thousand systems and hundreds of employees, or 10,000 systems, and hundreds or thousands of employees.

The challenge is getting large organizations to operate in a coherent fashion, when large organizations are made up of people, and we aren't always good at operating in a coherent fashion. What organizations really need, and where there's real potential, is how do you

make it so those things we talked about at the beginning—encryption, strong passwords, segmentation, whitelisting applications, patching vulnerabilities—can be done reliably, consistently, and at scale, because if we can do that, we would solve a large chunk of our security problem.

Roberta Fusaro: Nathaniel, thank you so much for joining us today.

Nathaniel Gleicher: Thank you for having me. ◆

Nathaniel Gleicher is the head of cybersecurity strategy at Illumio, and **Sam Palmisano** is chairman of the Center for Global Enterprise and retired chairman and CEO of IBM. **Roberta Fusaro** is a senior editor at McKinsey Publishing, and **Marc Sorel** is a consultant in McKinsey's Washington, DC, office.

Laszlo Podor/Getty Images

Risk analytics enters its prime

Rajdeep Dash, Andreas Kremer, Luis Nario, and Derek Waldron

All the ingredients are in place for unprecedented advances in risk analytics. Now it's up to banks to capture the opportunities.

With the rise of computing power and new analytical techniques, banks can now extract deeper and more valuable insights from their ever-growing mountains of data. And they can do it quickly, as many key processes are now automated (and many more soon will be). For risk departments, which have been using data analytics for decades, these trends present unique opportunities to better identify, measure, and mitigate risk. Critically, they can leverage their vast expertise in data and analytics to help leaders shape the strategic agenda of the bank.

Banks that are leading the analytical charge are exploiting both internal and external data. Within their walls, these banks are integrating more of their data, such as transactional and behavioral data from multiple sources, recognizing their high value. They are also looking externally, where they routinely go beyond conventional structured information, such as credit-bureau reports and market information, to evaluate risks. They query unconventional sources of data (such as government statistics, customer data from utilities and supermarket loyalty cards, and geospatial data) and even new unstructured sources (such as chat and voice transcripts, customer rating websites, and social media). Furthermore, they are getting strong results by combining internal and external data sets in unique ways, such as by overlaying externally sourced map data on the bank's transaction

information to create a map of product usage by geography. Perhaps surprisingly, some banks in emerging markets are pioneering this work. This is possible because these banks are often building their risk database from scratch and sometimes have more regulatory latitude.

The recent dramatic increases in computing power have allowed banks to deploy advanced analytical techniques at an industrial scale. Machine-learning techniques, such as deep learning, random forest, and XGBoost, are now common at top risk-analytics departments. The new tools radically improve banks' decision models. And techniques such as natural-language processing and geospatial analysis expand the database from which banks can derive insights.

These advances have allowed banks to automate more steps within currently manual processes—such as data capture and cleaning. With automation, straight-through processing of most transactions becomes possible, as well as the creation of reports in near real time. This means that risk teams can increasingly measure and mitigate risk more accurately and faster.

The benefits—and challenges—of risk analytics

Banks that are fully exploiting these shifts are experiencing a "golden age" of risk analytics, capturing benefits in the accuracy and reach of their credit risk models and in entirely new business models. They are seeing radical improvement in their credit risk models, resulting in higher profitability. For example, Gini coefficients of 0.75 or more in default prediction models are now possible.[1]

Exhibit 1 lays out the value that analytics can bring to these models.

Some banks are expanding their risk models to new realms. A few have been able to automate the lending process end to end for their retail and small-and-medium-size-enterprise segments. These banks have added new analytical tools to credit processes, including calculators for affordability or preapproval limits. With this kind of straight-through processing, banks can approve up to 90 percent of consumer loans in seconds, generating efficiencies of 50 percent and revenue increases of 5 to 10 percent. Recognizing the value in fast and accurate decisions, some banks are experimenting with using risk models in other areas as well. For example, one European bank overlaid its risk models on its marketing models to obtain a risk-profitability view of each customer. The bank thereby improved the return on prospecting for new revenue sources (and on current customers, too).

A few financial institutions at the leading edge are using risk analytics to fundamentally rethink their business model, expanding their portfolio and creating new ways of serving their customers. Santander UK and Scotiabank have each teamed up with Kabbage, which, using its own partnership with Celtic Bank, has enabled these banks to provide automated underwriting of small-business loans in Canada, Mexico, and the United Kingdom, using cleaner and broader data sets. Another leading bank has used its mortgage-risk model to provide a platform for real-estate agents and others providing home-buying services.

[1] Gini coefficients measure variation or randomness in a set of values, where 0 is completely random and 1 is perfectly ordered. In a model that predicts default, a Gini coefficient of 0 would indicate that the model is no better than a coin toss, and 1 would indicate that the model's output perfectly predicted the eventual defaults.

EXHIBIT 1 **Analytically enhanced credit models can improve banks' returns in four ways.**

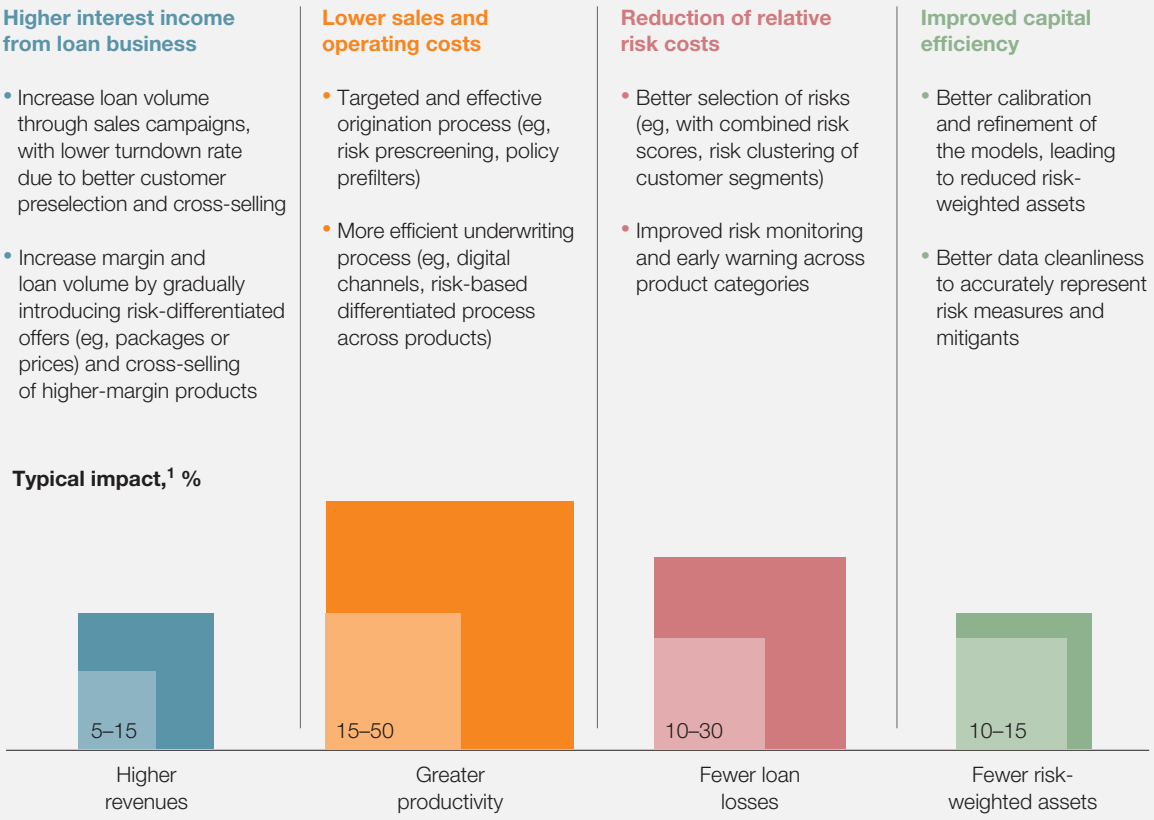

[1]Impact not additive and depends on the bank's portfolio.

Source: McKinsey analysis

Realizing the potential

For many banks, the advantages of risk analytics remain but a promise. They see out-of-date technology, data that are difficult to clean, skill gaps, organizational problems, and unrelenting regulatory demands. The barriers seem insurmountable. Yet banks can get things moving with some deliberate actions (Exhibit 2).

Perhaps the most salient issue is that risk analytics is not yet on the strategic agenda. Bank leaders often don't understand what is really at stake with risk analytics—at times because the analytics managers present highly complex solutions with a business case attached as an afterthought. Lagging banks miss out on the benefits, obviously, and also put other programs and activities at risk. Initiatives to grow revenue and optimize pricing can founder if imprecise risk assessment of customer segments leads to poor choices. In lending, when risk models

EXHIBIT 2

Several factors keeping banks from realizing the potential promise of risk analytics should be reexamined.

	Perceived barrier	A better way to think about it
Strategic agenda 	• Risk analytics is disconnected from business strategy and often seen as only a technology or regulatory-compliance initiative	• Risk analytics is at the heart of many strategic topics (eg, digital, capital productivity, loan-book health, market entry)
Data and technology 	• Unclean, unmatched data means waiting for that never-ending, "nearly complete" data transformation	• The data available can generate high value, often in combination with external data
	• Technological landscape is so complex that a simplification and upgrade is required before doing anything	• The "art of the possible" can produce high-value projects
Skills and organization 	• IT group doesn't have the authority to enforce data-management policies	• The business can take responsibility for data quality, integrity, and access, supported by a strong IT organization
	• Building analytics means hiring scarce, expensive data engineers and scientists	• Banks can move quickly through inorganic growth and partnerships
Regulations 	• Regulatory burden does not allow banks to focus on anything else, including analytics	• Analytics business cases can tease out surprising synergies between regulatory needs and business aspirations
	• Regulators would not agree with use of advanced models and more advanced data	• Sophisticated, value-generating models can be built even within constraints established by the Basel Committee and the European Union
Change 	• Building a model is relatively easy and can be done any time	• Digital economy has "winner takes all" economics; first movers have a huge advantage

Source: McKinsey analysis

underperform, banks often add business rules and policies as well as other manual interventions. But that inevitably degrades the customer experience, and it creates an opening for fintechs to capture market share through a better experience and more precise targeting. Taken to its logical conclusion, it is conceivable that banks might be relegated to "dumb pipes" that provide only financing.

Some nimble risk groups are finding ways through these problems, however. Our

analysis suggests these teams have six common behaviors:

- **Take it from the top,** lifting risk analytics to the strategic agenda. For example, four out of ten strategic actions that HSBC Bank laid out in 2015 rely heavily on risk analytics.

- **Think big and apply analytics** to every material decision. Capital One is well known for applying analytics to every decision that it makes, even when hiring data scientists.

- **Go with what you have.** If data are messy or incomplete, don't wait for a better version or for a "data-lake nirvana." Use the data you have, or find a way to complement them. When Banco Bilbao Vizcaya Argentaria (BBVA) wanted to lend to some clients but lacked information, it partnered with Destacame, a utility-data start-up, to provide data sufficient to support a way to underwrite the customers.

- **Accumulate skills quickly,** through either rapid hiring or acquisitions and partnerships. Then retain your talent by motivating people with financial and nonfinancial incentives, such as compelling projects. Banks such as BBVA, HSBC, Santander, and Sberbank have launched funds of $100 million and more to acquire and partner with fintechs to add their market share, sophisticated technologies, and people.

- **To succeed, be willing to fail and iterate quickly** through a series of "minimum viable products" (MVPs) while also breaking down traditional organizational silos. One bank building a fully digital lending product went through six MVPs in just 16 weeks to get to a product it could roll out more broadly.

- **Use model validation to drive relentless improvement.** Validation teams can be the source of many improvements to risk models, while preserving their independence. The key is for teams to style themselves as the guardian of model performance, rather than the traditional activity of merely examining models.

If banks can master these elements, significant impact awaits. Risk analytics is not the entire answer. But as leading banks are discovering, it is worthwhile in itself, and it is also at the heart of many successful transformations, such as digital risk and the digitization of key processes such as credit underwriting.

Risk-analytics leaders are creating analytic algorithms to support rapid and more accurate decision making to power risk transformations throughout the bank. The results have been impressive. An improvement in the Gini coefficient of one percentage point in a default-prediction model can save a typical bank $10 million annually for every $1 billion in underwritten loans.[2] Accurate data capture and well-calibrated models have helped a global bank reduce risk-weighted assets by about $100 billion, leading to the release of billions in capital reserves that could be redeployed in the bank's growth businesses.

[2] Assuming a base Gini coefficient of 0.50 and an observed default rate of 5 percent.

Leveraging the six successful behaviors

Nothing succeeds like success. The behaviors we have observed in successful risk-analytics groups provide the guidance.

Take it from the top

Stress testing and regulatory oversight following the 2008 financial crisis have vaulted risk management to the top of the management agenda. Nine years later, and after significant investment, most big banks have regained a handle on their risks and control of their regulatory relations. However, leading banks, recognizing the value from risk analytics, are keeping these programs at the top of their strategic plans, and top leaders are taking responsibility.

Top-management attention ensures commitment of sufficient resources and removal of any roadblocks—especially organizational silos, and the disconnected data sets that accompany these divides. Leaders can also keep teams focused on the value of high-priority use cases and encourage the use of cross-functional expertise and cross-pollination of advanced analytical techniques. Good ideas for applications arise at the front line, as people recognize changing customer needs and patterns, so banks must also build and maintain lines of communication.

Think big and apply analytics

For some time, analytics has played an important role in many parts of the bank, including risk, where a host of models—such as the PD, LGD, and EAD[3] models used in the internal ratings–based approach to credit risk—are in constant use. What's new is that the range of useful algorithms has greatly expanded, opening up dozens of new applications in the bank. Many small improvements to material decisions can really add up. An obvious example is algorithmic trading, which has transformed several businesses. Already by 2009, for example, it accounted for 73 percent of traded volume in cash equities. An expansion of automated credit decisions and monitoring has allowed banks to radically improve customer experience in residential mortgages and other areas. Banks in North and South America are using advanced-analytics models to predict the behavior of past-due borrowers and pair them with the most productive collections intervention.

These and other important examples are shown in Exhibit 3. What's important is that leading banks are putting analytics to work at every step of these and many other processes. Any time a decision needs to be made, these banks call on risk analytics to provide better answers. Even as they expand the applications of risk analytics, however, leading banks also recognize that they need to strengthen their model risk management to deal with inherent uncertainties within risk-analytics models, as these make up the largest share of risk-related decisions within banks.

Go with what you have

Messy, repetitive, and incomplete databases are a reality—but need not be an excuse. Rather than waiting for improvements in the quality, availability, and consistency of the bank's systems and the data they produce, leading risk-analytics teams ask what can be done now. This might involve using readily available data in the bank to immediately build a core analytic module, onto which new modules are integrated as new data sources

[3] Probability of default, loss given default, and exposure at default.

EXHIBIT 3

Rapid innovation in eight use cases is powered by advanced analytics.

	Credit risk	Description	Use cases
1		Underwriting	**Make better underwriting decisions** by using deep-learning algorithms to process vast amounts of data and more accurately quantify the risk of default
2		Credit-line management	**Reduce charge-off losses** by offering an optimal line to each client that is determined by machine-learning algorithms using the latest information about the client (eg, credit score) and local market (eg, home values)
3		Collections	**Increase recoveries** by making the right offer, at the right time, and through the right channel, with a recommendation engine and decision flow powered by 4 machine-learning algorithms
	Operations risk		
4		Payment fraud detection	**Identify and review high-risk payments before they are executed** by using input from fraud investigators to tune powerful machine-learning algorithms that pinpoint the highest-risk transactions
5		Anti–money laundering	**Quickly suspend money-laundering** operations using a longitudinal view of payment pathways to identify the patterns most indicative of money laundering, and accelerate reviews with powerful investigative tools
	Trading risk		
6		Contract compliance	Automate the **extraction and storage of data** from millions of trading contracts for regulatory compliance using leading-edge image-recognition and machine-learning algorithms
7		Trade surveillance	**Identify high-risk traders** by monitoring their behavior with sophisticated natural-language-processing algorithms that recognize themes in trader communications that are markers of conduct risk
	Model risk		
8		Model validation	Apply **rigorous and efficient model-validation processes** for traditional and advanced models that meet regulatory expectations and adhere to industry bench-marks for model risk management

Source: McKinsey analysis

become available. Alternatively, integrating two or more of the data sets on hand can generate significant value. These approaches hasten new analytical models to market, while at the same time helping the bank gather information as it forms a credit relationship with customers.

Furthermore, leading banks supplement their resources with external data—once they have established that this offers clear additional value. Some US fintechs, for example, obtain customer permission to comb financial data and create a sanitized database that banks can use to make accurate risk decisions based on

cash-flow patterns. A bank in Central America built a credit-approval system for unbanked customers based on data collected from supermarket loyalty cards. The bank used data such as frequency of shopping and the amount that customers typically spent per visit to estimate customers' ability to repay debt. Even better for banks, many external data are free. In some markets, micromarket information such as house prices by postal code or employment by district is available, and can be mined for insights into the creditworthiness of customers, especially small businesses. Conducting geospatial analytics on this information can also provide valuable insights (for example, proximity to a coffee-chain outlet would reveal foot traffic for a retail shop). Banks have also started analyzing unstructured data sets, such as news articles, feedback sites, and even social-network data.

Leading banks apply two tests before acquiring external data: Will it add value, typically through combination with other data sets? And does it conform with the bank's regulatory and risk policies? Consumer-protection regulations restrict the type of data that banks can use for risk-analytics applications, such as lending and product design.

While the practices outlined here will yield fast impact from messy, repetitive, and incomplete databases, most banks would still benefit from establishing sound data governance in parallel (and sometimes are required to do so under data regulations such as BCBS 239).

Accumulate skills quickly

Strong risk-analytics teams use several roles to develop solutions and integrate them into business processes (Exhibit 4).

Recognizing that they might not have the time to build the whole arsenal of skills, leading banks have acquired companies, outsourced some analytical work, invested in fintechs, and entered into formal partnerships with analytic houses. JPMorgan Chase has partnered with OnDeck to lend to small businesses; Bank of America has committed $3 billion annually to fintech investment and joint innovation. Other leading banks have entered into partnerships with digital innovators to better understand customer behavior and risk profiles. Even when leading banks have acquired talent at scale in these ways, they still work to define roles and build skills in the risk-analytics team.

To succeed, be willing to fail and iterate quickly

Speed is as important as completeness in realizing value from risk analytics. A winner-takes-all dynamic is emerging in the race to better serve customers. Banks, fintechs, and platform companies are getting better at locking in customers quickly with highly personalized and desirable offerings. The offerings are dependent on customer data, which get richer and deeper with every new development of risk-analytics capabilities. To reach and exceed the speed at which this race is moving, leading banks rely on quick, narrowly defined experiments designed to reveal the value (or the futility) of a particular hypothesis. When they succeed, they constitute a minimum viable product—something good enough to take to market, with the expectation that it will be soon improved. These experiments take weeks to conduct, rather than the more traditional months-long efforts commonly seen in risk-analytics functions (and that's not even considering the validation process). One form such experiments have taken are

EXHIBIT 4

Strong risk-analytics teams are using these roles to develop solutions and integrate them into business processes.

Structuring a strong risk-analytics team

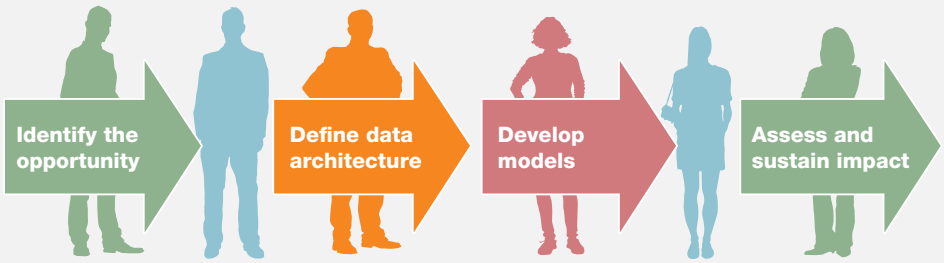

Data engineers and **data scientists**
These roles are already common. What is new is that they encompass new techniques beyond traditional statistics and econometrics. Analytics teams now use such methods as graph theory to analyze supply-chain risk or machine learning to develop highly sensitive early warning systems.

Translators
This new role requires a keen business sense and an understanding of the rationale behind the models. It also requires an entrepreneurial spirit to promote risk analytics throughout the bank.

Business leaders and experts
Business leaders and experts are also involved in developing solutions, taking responsibility for embedding the risk model in current practices.

Source: McKinsey analysis

"hackathons"—coding sessions with analysts and others that have produced promising applications in compressed time frames.

Use model validation to drive relentless improvement

The banks that are developing a competitive edge through analytics constantly improve their current models, even as they build new ones. They make full use of their independent model-validation framework, moving beyond providing regulatory and statistical feedback on risk models every year to a more insightful and business-linked feedback loop. Validation departments can achieve this without losing their independence by changing from a mind-set of "examiners of models" to "guardians of model performance."

To introduce a degree of experimentation into model validation, leading banks incorporate business and model expertise into bursts of rapid development and testing and accept that not all results will be as expected. In this way, the model benefits from a continual 360-degree review, rather than being buried in the risk-modeling team and understood only by the model owner. To be sure, as they do this work, banks must also respect regulatory constraints and explain to supervisors how they are utilizing advanced techniques. But leading institutions do not use regulatory oversight as an excuse not to move forward in an agile fashion. As shown by the multiple examples in this article, even large banks can make significant changes to improve outcomes and customer experience.

Getting started

We have outlined the reasons leading banks see considerable near-term promise in improved risk analytics, and the behaviors and principles that are distinguishing more successful players from the rest. This raises a logical question about what comes next: How can banks develop and execute a long-term, bankwide risk-analytics strategy? While a full discussion is beyond the scope of this article, we see five immediate actions for the chief risk officer (CRO) to maximize the value of existing investments and prioritize new ones. These actions are all consistent with the six successful behaviors discussed above, but distilled into immediate high-payoff steps.

- Assess the current portfolio of risk-analytics projects, assets, and investments and take a hard look at any that cannot answer the following questions satisfactorily:
  - Is the initiative business driven? Does it address one of the biggest business opportunities and define an analytics use case to deliver it? Or is the initiative a hammer looking for a nail?
  - Does the initiative have a clear plan for adoption and value capture? Or is it only a "model building" project?
  - Is the initiative structured to generate quick improvements as well as longer-term impact?
- Make an inventory of your talent, teams, and operating model for each initiative. Success requires multidisciplinary co-located teams of data engineers, data scientists, translators, and business experts. Prioritize actions to find the talent you need, rather than stretching the talent you have to the point of ineffectiveness.
- List your data and technology choke points—the weakest links in the system. Then determine the work-arounds you can develop to get high-priority initiatives moving (such as using external or alternative internal data or vendor solutions). Where no work-around is possible, ensure that precious resources do not lay idle waiting for resolution.
- Explain what you are doing to senior leaders, including business heads, the chief operating officer, and the chief investment officer. Work with them as needed to adjust priorities and redirect the program, but then proceed full steam ahead.

In our experience, risk leaders can take these steps quickly, given the right level of determination and focus. CROs should not hesitate to pull critical people into the exercise for a couple of weeks—it's typically a worthwhile investment that pays off by redirecting a much larger body of work toward maximum impact. ◆

Rajdeep Dash is senior expert in McKinsey's London office, **Andreas Kremer** is a partner in the Berlin office, and **Luis Nario** and **Derek Waldron** are partners in the New York office.

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 digital.McKinsey.com.

This McKinsey Practice Publication meets the Forest Stewardship Council® (FSC®) chain-of-custody standards. The paper used in this publication is certified as being produced in an environmentally responsible, socially beneficial, and economically viable way.

Printed in the United States of America.

FSC Logo here
63% BLACK

Digital/McKinsey

September/October 2017
Designed by Global Editorial Services

mckinsey.com
@digitalmckinsey
facebook.com/DigitalMcKinsey