

Inevitable architecture

Complexity gives way to simplicity and flexibility

ORGANISATIONS ARE OVERHAULING THEIR IT LANDSCAPES BY COMBINING OPEN source, open standards, virtualisation, and containerisation. Moreover, they are leveraging automation aggressively, taking steps to couple existing and new platforms more loosely, and often embracing a “cloud first” mind-set. These steps, taken individually or as part of larger transformation initiatives, are part of an emerging trend that some see as inevitable: the standardisation of a flexible architecture model that drives efficiency, reduces hardware and labor costs, and foundationally supports speed, flexibility, and rapid outcomes.

IN some companies, systems architecture is older than the freshman tech talent maintaining it. Sure, this legacy IT footprint may seem stable on a day-to-day basis. But in an era of rapid-fire innovation with cloud, mobile, analytics, and other forces implemented on the edge of the business fueling disruption and new opportunities, architectural maturity is becoming a persistent challenge—one directly linked to business problems.¹ Heavy customisation, complexity, security vulnerabilities, inadequate scalability, and technical debt throughout the IT environment have, directly or indirectly, begun to impact the bottom line. In Deloitte’s *2016 Global CIO Survey*, 46 percent of 1,200 IT executives surveyed identified “simplifying IT infrastructure” as a top business priority. Likewise, almost a quarter deemed the performance, reliability, and functionality of their legacy systems “insufficient.”²

Previous editions of *Tech Trends* have examined strategies that CIOs are deploying to modernise and revitalise their legacy core systems, not only to extract more value from them but to make the entire IT footprint more agile, intuitive, and responsive.


Likewise, *Tech Trends* has tracked IT’s increasingly warm embrace of automation and the autonomic platforms that seamlessly move workloads among traditional on-premises stacks, private cloud platforms, and public cloud services.

While these and other modernisation strategies can be necessary and beneficial, they represent sprints in a longer modernisation marathon. Increasingly, CIOs are thinking bigger-picture about the technology stacks that drive revenue and enable strategy. As they assess the capacity of current systems to meet future needs, many are likely asking, “If I could start over with a new IT architecture, what would I do differently to lower costs and increase capabilities?”

In the next 18 to 24 months, CIOs and their partners in the C-suite may find an answer to this question in a flexible architecture model whose demonstrated efficiency and effectiveness in start-up IT environments suggest that its broader adoption in the marketplace may be *inevitable*. In this cloud-first model—and in the leading practices emerging around it—platforms are virtualised, containerised, and treated like malleable, reusable resources. Systems

Figure 1. Infrastructure management tools: Maturity and developer and operations impact

The infrastructure management tools landscape is heavily segmented and complex. While some capabilities such as cloud management and autonomies are at the origin of disruption for IT, technology leaders should also address adjacent capabilities to enable seamless IT.



Deloitte University Press | dupress.deloitte.com

are loosely coupled and, increasingly, automated to be self-learning and self-healing. Likewise, on-premises, private cloud, or public cloud capabilities can be deployed dynamically to deliver each workload at an optimum price and performance point. Taken together, these elements can make it possible to move broadly from managing instances to managing outcomes.

It's not difficult to recognise a causal relationship between architectural agility and any number of potential strategic and operational benefits. For example, inevitable architecture—which also embraces open-source/open-stack approaches—provides the foundation needed to support rapid development and deployment of flexible solutions that, in turn, enable innovation and growth. In a competitive landscape being redrawn continuously by technology disruption, time-to-market can be a competitive differentiator.

Moreover, architectural modernisation is about more than lowering capital expenses by transitioning to the cloud and eliminating on-premises servers. Indeed, its most promising opportunity for cost savings often lies in reducing the number of people required to maintain systems. This does not mean simply decreasing headcount, though in some cases that may become necessary. Rather, by automating mundane “care and feeding” tasks, companies can free critical IT talent to focus on higher-level activities such as developing new products and services, working hand-in-hand with the business to support emerging needs, and deploying a revitalised technology stack that enables rapid growth in an innovation environment.

Beyond start-ups

To date, much of the “tire-kicking” and experimentation done on the inevitable-architecture model has taken place in the IT ecosystems of start-ups and other small companies that are unburdened by legacy systems. However, recently some larger organisations have begun experimenting with one or more inevitable-architecture components. For example, Marriott Hotels, which runs VMware virtual machines on an IBM cloud, recently deployed a hybrid cloud co-developed by the two vendors.

This time-saving solution makes it possible to move virtual workloads created and stored on-premises to the cloud quickly, without having to rewrite or repackage them. Marriott reports that with the new solution, a process that could take up to a day can now be completed in an hour.³

And retail giant Walmart has embraced open-source with a new DevOps platform called OneOps that is designed to help developers write and launch apps more quickly. Walmart has made code for the OneOps project available on GitHub.⁴

In the coming months, expect to see other organisations follow suit by experimenting with one or more of the following inevitable-architecture elements:

Loosely coupled systems: Large enterprises often deploy thousands of applications to support and enable business. In established companies, many of the core applications are large, monolithic systems featuring layers of dependencies, workarounds, and rigidly defined batch interfaces. In such IT environments, quarterly or semiannual release schedules can create 24-month backlogs of updates and fixes. Moreover, making even minor changes to these core systems may require months of regression testing, adding to the backlog.⁵

Over the last several years, the emergence of microservices and the success of core revitalisation initiatives have given rise to a fundamentally different approach to infrastructure design in which core components are no longer interdependent and monolithic but, rather, “loosely coupled.” In this structure, system components make little if any use of the definitions of other components. Though they may interface with other components via APIs or other means, their developmental and operational independence makes it possible to make upgrades and changes to them without months of regression testing and coordination. They can also be more easily “composed” than large monolithic systems.

Containers: Similar to loosely coupled systems, containers break down monolithic interdependent architectures into manageable, largely independent pieces. Simply put, a container is an entire runtime environment—an application plus its dependencies,

libraries and other binaries, and configuration files—virtualised and bundled into one package. By containerising the application platform and its dependencies, differences in OS distributions and underlying infrastructure are abstracted away.⁶ Containers are fully portable and can live on-premises or in the cloud, depending on which is more cost-effective.

Container management software vendors such as Docker, CoreOS, BlueData, and others are currently enjoying popularity in the software world, due in large part to their ability to help developers break services and applications into small, manageable pieces. (Current industry leader Docker reports that its software has been downloaded more than 1.8 billion times and has helped build more than 280,000 applications.)⁷

Though containers may be popular among developers and with tech companies such as Google, Amazon, and Facebook, to date corporate CIOs seem to be taking a wait-and-see approach.⁸ That said, we are beginning to see a few companies outside the tech sector utilise containers as part of larger architecture revitalisation efforts. For example, financial services firm Goldman Sachs has launched a yearlong project that will shift about 90 percent of the company’s computing to containers. This includes all of the firm’s applications—nearly 5,000 in total—that run on an internal cloud, as well as much of its software infrastructure.⁹

Speed first, then efficiency: As business becomes increasingly reliant on technology for emerging products and services, it becomes more important than ever for IT to shrink time-to-market schedules. As such, inevitable architecture emphasises speed over efficiency—an approach that upends decades of processes and controls that, while well intentioned, often slow progress and discourage experimentation. Though it may be several years before many CIOs wholly (or even partially) embrace the “speed over efficiency” ethos, there are signs of momentum on this front. Acknowledging the gravitational pull that efficiency can have on development costs and velocity, streaming service Netflix has created an IT footprint that doesn’t sacrifice speed for efficiency.¹⁰

Open source: Sun Microsystems co-founder Bill Joy once quipped, “No matter who you are, most of the smartest people work for somebody else. . . . It’s better to create an ecology that gets most of the world’s smartest people toiling in your garden for your goals. If you rely solely on you own employees, you’ll never solve all your customers’ needs.” This idea, now referred to as Joy’s Law, informs open-source software (OSS) and open-stack approaches to development, both core attributes of inevitable architecture.¹¹

Many large companies still have old-world open-source policies featuring bureaucratic approval processes that ultimately defeat the entire purpose of OSS.¹² Yet, to quote reluctant Nobel laureate Bob Dylan, the times they are a-changin’. One recent study of open-source adoption estimates that 78 percent of companies are now using OSS. Even the generally risk-averse federal government is getting in on the act, having recently outlined a policy to support access to custom software code developed by or for the federal government.¹³

Fault expecting: For decades, developers focused on design patterns that make systems fault-tolerant. Inevitable architecture takes this tactic to the next level by becoming “fault expecting.” Famously illustrated in Netflix’s adventures using Chaos Monkey, fault expecting is like a fire drill involving real fire. It deliberately injects failure into a system component so that developers can understand 1) how the system will react, 2) how to repair the problem, and 3) how to make the system more resilient in the future. As CIOs work to make systems more component-driven and less monolithic, fault expecting will likely become one of the more beneficial inevitable-architecture attributes.

Autonomic platforms: Inevitable architecture demands automation throughout the IT lifecycle, including automated testing, building, deployment, and operation of applications as well as large-scale autonomic platforms that are largely self-monitoring, self-learning, and self-healing. As discussed in *Tech Trends 2016*, autonomic platforms build

upon and bring together two important trends in IT: software-defined everything’s climb up the tech stack, and the overhaul of IT operating and delivery models under the DevOps movement. With more and more of IT becoming expressible as code—from underlying infrastructure to IT department tasks—organisations now have a chance to apply new architecture patterns and disciplines. In doing so, they can remove dependencies between business outcomes and underlying solutions, while redeploying IT talent from rote low-value work to higher-order capabilities.

Remembrance of mainframes past

For many established companies, the inevitable-architecture trend represents the next phase in a familiar modernisation journey that began when they moved from mainframe to client-server and, then, a decade later, shifted to web-centric architecture. These transitions were not easy, often taking months or even years to complete. But in the end, they were worth the effort. Likewise, successfully deploying and leveraging inevitable architecture’s elements will surely present challenges as the journey unfolds. Adoption will take time and will likely vary in scope according to an organisation’s existing capabilities and the incremental benefits the adopting organisation achieves as it progresses.

Once again, the outcome will likely justify the effort. It is already becoming clear that taken together, inevitable architecture’s elements represent a sea change across the IT enablement layer—a modern architecture up, down, and across the stack that can deliver immediate efficiency gains while simultaneously laying a foundation for agility, innovation, and growth. Its potential value is clear and attributable, as is the competitive threat from start-ups and nimble competitors that have already built solutions and operations with these concepts.



Future stack: Louisiana builds an IT architecture for tomorrow

With some legacy IT systems entering their third decade of service, in 2014 the state of Louisiana launched the Enterprise Architecture Project, an ambitious, multifaceted effort to break down operational silos, increase reusability of systems and services, and most important, create the foundation for the state’s next-generation architecture. The end state was clear: Accelerate the delivery of flexible, secure building blocks and the capabilities to speed their use to produce value for Louisiana citizens. Changes to how IT procures, assembles, and delivers capabilities would be not only needed but inevitable.

“We deliver end-to-end IT services for 16 executive branch agencies,” says Matthew Vince, the Office of Technology Services’ chief design officer and director of project and portfolio management. “With a small IT staff and that much responsibility, we didn’t see any way to take on a project like this without jumping into inevitable architecture.”

Given the age and complexity of the state’s legacy architecture—and that Louisiana’s government takes a paced approach in adopting new technologies—it was important to define a clear strategy for achieving short- and longer-term goals. The team

focused agencies on architecture principles and the “so what” of the inevitable modern capabilities it would supply: enabling components and reuse over specifying “microservices,” how automation could help agencies over specifying a particular IT product, and the need for workload migration and hybrid cloud enablement over specifying a particular form of virtualisation.

“For example, think about containers,” says Michael Allison, the state’s chief technology officer. “Having the ability to jump into Docker or Kubernetes today is not as important as being able to adopt container orchestration solutions in the future. Our immediate goal, then, is to lay the foundation for achieving longer-term goals.”

Laying that foundation began by aligning IT’s goals and strategies with those of agency leaders and business partners. IT wanted to determine what tools these internal customers needed to be successful and what options were available from their suppliers and partners before planning, prototyping, and providing working solutions to meet those needs.

After careful consideration of business and IT priorities, budgets, and technical requirements, IT teams began developing a next-generation software-defined data centre (SDDC) to serve as the platform of choice for the enterprise architecture project and future line-of-business applications. This SDDC was put into production at the end of

2016 and now supports efforts in several EA project areas, including:

Security and risk management: Replaces network segregation with software-defined data security policies.

Cloud services: Uses both public and enterprise cloud services, including commercial infrastructure-as-a-service offerings.

Consolidation and optimisation: Consolidates and optimises data centre services, service-oriented architecture, and governance to increase efficiencies and lower costs.¹⁴

Progress in these areas is already having a net positive impact on operations, Vince says. “Traditionally we were siloed in service delivery. Suddenly, barriers between teams are disappearing as they begin working through the data centre. This has been one of the greatest benefits: Lines between traditionally segregated skill sets are becoming blurred and everyone is talking the same language.

We are seeing SDDC technology can be used across the enterprise rather than being confined to use by a single agency,” he continues. “A priority in our move to the SDDC is integrating systems, maximising reusability, and—when legislation allows us to—sharing data across agencies. Going forward, everything we do will be to meet the needs of the entire state government, not just a single agency.”¹⁵

The keys to Citi-as-a-service

Responding to a seismic shift in customer expectations that has taken place during the past decade, leading financial services company Citi is transforming its IT architecture into a flexible, services-based model that supports new strategies, products, and business models.

“The sophistication and generational change in customer expectation in our industry as well as the delivery requirements impacts everything we do now, from behind the scenes items such as how we run our data centres to the way we design branch offices, mobile and ATM capabilities, and the financial instruments we offer,” says Motti

Finkelstein, managing director, CTO-Americas and Global Strategy and Planning at Citi. “Due to the dynamic nature of this change, the whole design of IT architecture needs to be more agile. Infrastructure should be flexible—compute, network, storage, security, databases, etc.—with elastic real-time capabilities to meet business demands in a time-efficient manner.”

Citi actually began working toward its ultimate architecture goal more than a decade ago, when the organisation took initial steps to virtualise its data centres. Other virtualisation projects soon followed, eventually leading to an effort to standardise and focus architectural transformation on building a broad everything-as-a-service (XaaS) environment. Finkelstein estimates that ongoing efforts to transform Citi’s architecture have reduced infrastructure costs significantly in the last decade. “We are continuing on a trajectory of not only expediting time to market delivery but becoming more cost-efficient as we move our optimisation efforts up the stack with technologies such as containerisation paired with platform-as-a-service,” he says. Container technologies are rapidly becoming enterprise ready, and may inevitably become the bank’s target architecture. Barriers to adoption are diminishing with industry standards coming into place and the up-front investment required to re-architect applications onto these platforms steadily reducing.

Leaders also found that while these new technical building blocks drove cost and efficiency savings for IT, equally important were the new capabilities they enabled. “While we were investing in virtualisation, automation, and XaaS, we are also building our development capabilities,” Finkelstein says. “Development will become much more efficient and productive as our engineers are designing and implementing XaaS environments, reducing time to market and improving adoption times for innovation in technology.” Citi has engineered a PaaS environment with predefined microservices to facilitate rapid and parallel deployments; this enables the use of clear standards and templates to build and extend new services both internally and externally, which marks a significant change from only a few years ago.

Inevitably, speed to market and rapid outcomes are attracting the right kind of attention: “When we started working to create an XaaS environment, we thought it would be based on policy requirements,” Finkelstein says. “Today, everyone—including the business—understands the need for this development process and the benefit of templates standardisation to build new products and services. This model being DevOps-operated needs to have a lot of visibility into the process; we are starting to see different organisations become more involved in building those templates for success. At this point, it becomes more about tweaking a template than about starting from scratch.” (DevOps is the term used for agile collaboration, communications, and automation between the development and operations technology units.)¹⁶

Inventing the inevitable

Google has been passionate about the piece-parts of inevitable architecture—a cloud-based, distributed, elastic, standardised technology stack built for reliability, scalability, and speed—since its earliest days. The company’s head of infrastructure, Urs Hölzle (Google’s eighth badged employee), recognised the inherent need for scale if the company were going to succeed in building a truly global business allowing people to search anything, from everywhere, with content delivered anywhere. With seven of its products each having more than a billion active users, that early vision has not only been the foundation of Google’s growth—it has driven the very products and solutions fueling the company’s expansion.

According to Tariq Shaukat, president of customers at Google Cloud, “We recognised that our technology architecture and operations had to be built to fulfill a different kind of scale and availability. We invented many of the core technologies to support our needs . . . because we had to.” That led to the development, deployment, and open sourcing of many technologies that help fuel the growth of inevitable architecture—including

MapReduce, Kubernetes, TensorFlow, and others. These technologies have become a complement to Google’s commercial product roadmap, including the Google Cloud Platform, across infrastructure, app, data, and other architectural layers.

At the heart of Google’s technology landscape is a systems mind-set: deeply committed to loosely coupled but tightly aligned components that are built expecting to be reused while anticipating change. This requires that distributed teams around the globe embrace common services, platforms, architectural principles, and design patterns. As development teams design and build new products, they begin with an end-state vision of how their products will operate within this architecture and, then, engineer them accordingly. Not only does this approach provide a common set of parameters and tools for development—it also provides a consistent approach to security and reliability, and can prevent architectural silos from springing up around a new application that doesn’t work with all of Google’s global systems. Extensions and additions to the architecture happen frequently but go through rigorous vetting processes. The collaborative nature of capabilities such as design and engineering is coupled with tight governance around non-negotiables such as security, privacy, availability, and performance.

Google’s architectural efforts are delivering tangible benefits to both employees and customers. The company’s forays into machine learning are one example: In a period of 18 to 24 months, it went from an isolated research project to a broad toolkit embedded into almost every product and project. Adds Shaukat, “The reason it was possible? A common architecture and infrastructure that makes it possible to rapidly deploy modules to groups that naturally want to embrace them. It was absolutely a grassroots deployment, with each team figuring out how machine learning might impact their users or customers.” By rooting agility in a systems approach, the company can make an impact on billions of users around the globe.¹⁷

MY TAKE

**JOHN SEELY BROWN,
INDEPENDENT CO-CHAIRMAN,
CENTRE FOR THE EDGE**

DELOITTE LLP

The exponential pace of technological advancement is, if anything, accelerating. Look back at the amazing changes that have taken place during the last 15 years. At the beginning of the naughts, the holy grail of competitive advantage was scalable efficiency—a goal that had remained mostly unchanged since the industrial revolution. Organisations and markets were predictable, hierarchical, and seemingly controllable. The fundamental assumption of stability informed not only corporate strategies and management practices but also the mission and presumed impact of technology. By extension, it also informed IT operating models and infrastructure design.

Then the Big Shift happened. Fueled in part by macro advances in cloud, mobile, and analytics, exponential advances across technology disciplines spawned new competitors, rewired industries, and made obsolete many institutional architectures. Rather than thinking of our business models as a 200,000-ton cargo ship sailing on calm open waters, we suddenly needed to mimic a skilled kayaker navigating white water. Today, we must read contextual currents and disturbances, divining what lies beneath the surface, and use these insights to drive accelerated action.

At a lower level, new principles have emerged in the world of computer science. Eventual consistency is now an acceptable alternative to the transactional integrity of two-phase commit. Tightly aligned but loosely coupled architectures are becoming dominant. Mobile-first mind-sets have evolved, giving way to “cloud first”—and in some cases “cloud only”—landscapes. Distributed, unstructured, and democratised data is now the norm. Moreover, deep learning and cognitive technologies are increasingly harnessing data. Open standards and systems are providing free capabilities in messaging, provisioning, monitoring, and other disciplines. The proliferation of sensors

and connected devices has introduced a new breed of architecture that requires complex event correlation and processing.

Collectively, these advances form the underpinnings of inevitable architecture. They are redefining best practices for IT infrastructure and are providing new opportunities to use this infrastructure to empower the edge. This, in turn, makes it possible for companies to rapidly bring complex new ideas and products to market.

“WE MUST DIVINE WHAT LIES
BENEATH THE SURFACE.”

Moreover, inevitable architecture supports experimentation at the pace of overarching technological advancements. Importantly, it also provides opportunities to automate environments and supporting operational workloads that can lead to efficiency gains and concrete cost savings while also enabling a new kind of agility. Companies can use these savings to fund more far-reaching inevitable architecture opportunities going forward.

In describing Amazon’s relentless focus on new offerings and markets, Jeff Bezos once famously said, “Your margin is my opportunity.” The gap between exponential technologies and any organisation’s ability to execute against its potential is the new margin upon which efficiencies can be realised, new products and offerings can be forged, and industries can be reshaped. If every company is a technology company, bold new approaches for managing and reimagining an organisation’s IT footprint become an essential part of any transformational journey. They can also help shrink the gap between technology’s potential and operational reality—especially as the purpose of any given institution evolves beyond delivering transactions at scale for a minimum cost. Inevitable architecture can provide the wiring for scalable learning, the means for spanning ecosystem boundaries, and the fluidity and responsiveness we all now need to pursue tomorrow’s opportunities.

Overhauling IT infrastructure by embracing open source, open standards, virtualisation, and cloud-first strategies means abandoning some long-held systems architecture design principles. Doing so may require rethinking your approach to security. For example, if you have been a network specialist for 30 years, you may consider some kind of perimeter defense or firewall foundational to be a good cyber defense strategy. In a virtualised, cloud-first world, such notions aren't necessarily taken as gospel anymore, as network environments become increasingly porous. Going forward, it will be important to focus on risks arising from new design philosophies. Open standards and cloud-first aren't just about network design—they also inform how you will design a flexible, adaptable, scalable cyber risk paradigm.

Also, as companies begin overhauling their IT infrastructure, they may have opportunities to integrate many of the security, compliance, and resilience capabilities they need into the new platforms being deployed by embedding them in templates, profiles, and supporting services. Seizing this opportunity may go a long way toward mitigating many of the risk and cyber challenges often associated with a standardised, flexible architecture model. Moreover, designing needed cybersecurity and risk management capabilities into new systems and architectures up front can, when done correctly, provide other long-term benefits as well:

Speed: Retrofitting a new product or solution with cyber capabilities may slow down development initiatives and boost costs. By addressing cyber considerations proactively as they begin overhauling IT infrastructure, companies may be able to accelerate the product development lifecycle.

Product design as a discipline: When risk and compliance managers collaborate with designers and engineers from the earliest stages of product development and transformation initiatives, together they can design “complete” products that are usable and effective and maintain security and compliance standards.

Effectiveness: By establishing practices and methodologies that emphasise risk considerations—and by taking steps to confirm that employees understand the company's approach to managing different types of risk—organisations may be able to move more efficiently and effectively. Just as infrastructure transformation initiatives offer designers the opportunity to incorporate risk and cybersecurity into new products and services, they offer the same opportunity to instill greater technology and risk fluency in employees throughout the enterprise.

“GOING FORWARD,
IT WILL BE IMPORTANT TO
FOCUS ON RISKS ARISING
FROM NEW DESIGN
PHILOSOPHIES.”

However, with these benefits come questions and challenges that CIOs may need to address as they and their teams build more flexible and efficient architecture. For example, they may have to create additional processes to accommodate new cyber and risk capabilities that their organisations have not previously deployed.

Likewise, new platforms often come with risk and security capabilities that can be deployed, customised, and integrated with existing solutions and business process. While companies should take advantage of these capabilities, their existence does not mean you are home free in terms of risk and cybersecurity considerations. Risk profiles vary, sometimes dramatically, by organisation and industry. Factory-embedded security and risk capabilities may not address the full spectrum of risks that your company faces.

Embracing open source and open standards also raises several risk considerations. Some may question the integrity and security of code to which anyone can contribute, and one challenge that organisations using open source code may face is mitigating whatever risk does come from using code written by anonymous strangers. CIOs should consider performing broad risk and capability assessments as part of an architecture transformation initiative. Likewise, they should evaluate what IT and legal support should be put in place to address potential issues.

Where do you start?

In organising systems and processes to support an inevitable architecture (IA) model, start-ups often have the advantage of being able to begin with a blank canvas. Larger companies with legacy architectures, on the other hand, will have to approach this effort as a transformation journey that begins with careful planning and preparatory work. As you begin exploring IA transformation possibilities, consider the following initial steps:

- **Establish your own principles:** Approaches to inevitable architecture will differ by company according to business and technology priorities, available resources, and current capabilities. Some will emphasise cloud-first while others focus on virtualisation and containers. To establish the principles and priorities that will guide your IA efforts, start by defining end-goal capabilities. Then, begin developing target architectures to turn this vision into reality. Remember that success requires more than putting IA components in place—business leaders and their teams must be able to use them to execute strategy and drive value.
- **Take stock of who is doing what—and why:** In many companies, various groups are already leveraging some IA components at project-specific, micro levels. As you craft your IA strategy, it will be important to know what technologies are being used, by whom, and to achieve what specific goals. For example, some teams may already be developing use cases in which containers figure prominently. Others may be independently sourcing cloud-based services or open standards to short-circuit delivery. With a baseline inventory of capabilities and activities in place, you can begin building a more accurate and detailed strategy for industrialising priority IA components, and for consolidating individual use case efforts into larger experiments. This, in turn, can help you develop a further-reaching roadmap and business case.
- **Move from silos to business enablement:** In the same way that CIOs and business teams are working together to erase boundaries between their two groups, development and infrastructure teams should break free of traditional skills silos and reorient as multidisciplinary, project-focused teams to help deliver on inevitable architecture's promise of business enablement. The beauty? Adopting virtualisation and autonomics across the stack can set up such a transition nicely, as lower-level specialisation and underlying tools and tasks are abstracted, elevated, or eliminated.
- **CIO, heal thyself:** Business leaders often view CIOs as change agents who are integral to business success. It's now time for CIOs to become change agents for their own departments. Step one? Shift your focus from short-term needs of individual business teams to medium or longer-term needs shared by multiple groups. For example, if as part of your IA strategy, seven of 10 internal business customers will transition to using containers, you will likely need to develop the means for moving code between these containers. You may also need to bring additional IT talent with server or container experience on staff. Unfortunately, some CIOs are so accustomed to reacting immediately to business requests that they may find it challenging to proactively anticipate and plan for business needs that appear slightly smaller on the horizon. The time to begin thinking longer-term is now.

Bottom line

In a business climate where time-to-market is becoming a competitive differentiator, large companies can no longer afford to ignore the impact that technical debt, heavy customisation, complexity, and inadequate scalability have on their bottom lines. By transforming their legacy architecture into one that emphasises cloud-first, containers, virtualised platforms, and reusability, they may be able to move broadly from managing instances to managing outcomes.

CONTACTS



JOHN STARLING
Partner, Technology Consulting
Deloitte MCS Limited
020 7303 7225
jstarling@deloitte.co.uk



STEPHEN MERCER
Head of Technology
Consulting
Deloitte MCS Limited
0161 455 6836
stmercer@deloitte.co.uk



RICHARD FORREST
Director, Technology Consulting
Deloitte MCS Limited
020 7303 8547
rforrest@deloitte.co.uk

AUTHORS

RANJIT BAWA

US Cloud and Infrastructure leader, Deloitte Consulting LLP

SCOTT BUCHHOLZ

Managing Director, Technology practice, Deloitte Consulting LLP

KEN CORLESS

Principal, Technology Strategy and Architecture practice, Deloitte Consulting LLP

JACQUES DE VILLIERS

Cloud Services director, Deloitte Consulting LLP

EVAN KALINER

Solutions Network practice leader, Deloitte Consulting LLP

ENDNOTES

1. Scott Buchholz, Ben Jones, and Pavel Krumkachev, *Reimagining core systems: Modernising the heart of the business*, Deloitte University Press, February 24, 2016, <https://dupress.deloitte.com/dup-us-en/focus/tech-trends/2016/reimagining-core-systems-strategy.html>.
2. Khalid Kark, Mark White, Bill Briggs, and Anjali Shaikh, *Navigating legacy: Charting the course to business value*, Deloitte University Press, November 10, 2016, <https://dupress.deloitte.com/dup-us-en/topics/leadership/global-cio-survey.html>.
3. Alex Konrad, "IBM and VMware deepen hybrid cloud partnership for customers like Marriott," *Forbes*, August 29, 2016, www.forbes.com/sites/alexkonrad/2016/08/29/ibm-vmware-deepen-hybrid-cloud-partnership-for-customers-like-marriott/.
4. Frederic Lardinois, "Walmart launches OneOps, an open-source cloud and application lifecycle management platform," *TechCrunch*, January 26, 2016, <https://techcrunch.com/2016/01/26/walmart-launches-oneops-an-open-source-cloud-and-application-lifecycle-management-platform/>.
5. Buchholz et al., *Reimagining core systems*.
6. Paul Rubens, "What are containers and why do you need them?" *CIO Magazine*, May 20, 2015, www.cio.com/article/2924995/enterprise-software/what-are-containers-and-why-do-you-need-them.html.
7. "Ben Golub on what digital containers hold for the future," *Wall Street Journal*, February 10, 2016, www.wsj.com/articles/ben-golub-on-what-digital-containers-hold-for-the-future-1455103832.
8. Clint Boulton, "CIOs aren't ready for Docker and container technology," *CIO Magazine*, February 15, 2016, www.cio.com/article/3033167/virtualisation/cios-aren-t-ready-for-docker-and-container-technology.html.
9. Steven Rosenbush and Steven Norton, "Big changes in Goldman's software emerge from small containers," *Wall Street Journal*, February 24, 2016, <http://blogs.wsj.com/cio/2016/02/24/big-changes-in-goldmans-software-emerge-from-small-containers/>.
10. Patrick Nommensen, "Adopting microservices at Netflix: Lessons for team and process design," *Dzone*, April 13, 2015, <https://dzone.com/articles/adopting-microservices-netflix-0>.
11. Greg Satell, "Why open beats closed," *Forbes*, September 26, 2016, www.forbes.com/sites/gregsatell/2016/09/26/why-open-beats-closed/.
12. Black Duck Software, *The Tenth Annual Future of Open Source Survey*, 2016, www.blackducksoftware.com/2016-future-of-open-source, accessed November 20, 2016.
13. Tony Scott, "The people's code," White House Blog, August 8, 2016, www.whitehouse.gov/blog/2016/08/08/peoples-code.
14. State IT Recognition Awards, National Association of State Information Officers, Louisiana Department of Administration, Office of Technology Services, Enterprise Architecture and Software Defined Datacentre, 2016, www.nascio.org/portals/0/awards/nominations2016/2016/2016LA5-20160518%20LA%20EA%20NASCIO.pdf, accessed December 18, 2016.

15. Interview with Michael Allison, state chief technology officer, and Matthew Vince, chief design officer and director, project and portfolio management, Office of Technology Services for the state of Louisiana, December 5, 2016.
16. Interview with Motti Finkelstein, chief technology officer, Americas and Global Strategy Planning, Citi, December 7, 2016.
17. Interview with Tariq Shaukat, president of customers at Google Cloud, January 9, 2017.