CUTTER Business Technology Journal

Management, Innovation, Transformation

Vol. 31, No. 10, 2018

Riding the Next Wave of CLOUD COMPUTING

CUTTER Business Technology Journal

As business models for creating value continue to shift, new business strategies are constantly emerging and digital innovation has become an ongoing imperative. The monthly *Cutter Business Technology Journal* delivers a comprehensive treatment of these strategies to help your organization address and capitalize on the opportunities of this digital age.

Cutter Business Technology Journal is unlike academic journals. Each monthly issue, led by an expert Guest Editor, includes five to seven substantial articles, case studies, research findings, and/or experience-based opinion pieces that provide innovative ideas and solutions to the challenges business technology professionals face right now — and prepares them for those they might face tomorrow. *Cutter Business Technology Journal* doesn't water down or delay its content with lengthy peer reviews. Written by internationally known thought leaders, academics, and practitioners — you can be certain you're getting the uncensored perspectives of global experts.

You'll benefit from strategic insight on how the latest movements in digital innovation and transformation, IoT, big data analytics and cloud, to name a few, are changing the business landscape for both new and established organizations and how cutting-edge approaches in technology leadership, enterprise agility, software engineering, and business architecture can help your organization optimize its performance and transition to these new business models.

As a subscriber, you'll also receive the *Cutter Business Technology Advisor* — a weekly bulletin featuring industry updates delivered straight to your inbox. Armed with expert insight, data, and advice, you'll be able to leverage the latest business management thinking to achieve your organization's goals.

No other journal brings together so many thought leaders or lets them speak so bluntly — bringing you frank, honest accounts of what works, what doesn't, and why. Subscribers have even referred to *Cutter Business Technology Journal* as a consultancy in print and likened each month's issue to the impassioned discussions they participate in at the end of a day at a conference!

Get the best in thought leadership and keep pace with the technologies and business models that will give you a competitive edge — subscribe to *Cutter Business Technology Journal* today!

Founding Editor: Ed Yourdon
Publisher: Karen Fine Coburn
Group Publisher: Christine Generali
Managing Editor: Cindy Swain
Copy Editors: Jennifer Flaxman, Tara Meads
Production Editor: Linda Dias
Client Services: service@cutter.com

***Cutter Business Technology Journal®** is published 12 times a year by Cutter Information LLC, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA (Tel: +1 781 648 8700; Fax: +1 781 648 8707; Email: cbtjeditorial@cutter.com; Website: www.cutter.com; Twitter: @cuttertweets; Facebook: Cutter Consortium). ISSN: 2475-3718 (print); 2475-3742 (online).*

Subscription rates are US $485 a year in North America, US $585 elsewhere, payable to Cutter Information LLC. Reprints, bulk purchases, past issues, and multiple subscription and site license rates are available on request.

NOT FOR DISTRIBUTION
For authorized use, contact Cutter Consortium +1 781 648 8700 or service@cutter.com.

☐ Start my print subscription to *Cutter Business Technology Journal* ($485/year; US $585 outside North America).

Name _____ Title _____

Company Address _____

City _____ State/Province _____ ZIP/Postal Code _____

Email (Be sure to include for weekly *Cutter Business Technology Advisor*) _____

Fax to +1 781 648 8707, call +1 781 648 8700, or send email to service@cutter.com.
Mail to Cutter Consortium, 37 Broadway, Suite 1, Arlington, MA 02474-5552, USA.

Request Online License Subscription Rates

For subscription rates for online licenses, email or call: **sales@cutter.com** or **+1 781 648 8700.**

CUTTER CONSORTIUM
Access to the Experts

Opening Statement

by the *Cutter Business Technology Journal* Team

Cloud computing is no longer the next *big* thing. It is *the* thing. Over the past decade, cloud computing has become the go-to solution for data storage and processing, even as data volume and complexity have grown by leaps and bounds. Cloud's accessibility, low entry costs, and security features have enabled organizations of all sizes to become more agile and decrease time to market. Ultimately, the cloud provides organizations the means to stay competitive as they make significant strides toward their digital transformation goals.

Plus, the demand for and adoption of cloud computing services have increased at such a staggering rate that IT decision makers not opting for the cloud will soon need to justify — if they haven't had to already — why they would *not* use the cloud as their first preference for data storage and processing needs.

As the title of this issue, "Riding the Next Wave of Cloud Computing," implies, the cloud computing market has already achieved some major milestones since its inception. Yet it still remains one of the hottest markets in terms of growth, spending, and revenue generation. Indeed, it is still making its way up a steep growth curve, and as more features and functionality are added that allow businesses to embrace new business models and enable innovation, the growth shows no signs of slowing.

Public cloud allows organizations to tap into a virtually unlimited pool of data storage and processing power, remotely accessed and managed over the Internet, at a lower total cost of ownership than building and operating IT infrastructures themselves. However, the evolution of cloud technologies has not stopped at public clouds; community, private, multicloud, and hybrid cloud solutions are also viable means of storing and processing data, depending on an organization's needs. A hybrid approach to cloud computing, for instance, which uses both public and private cloud resources, allows organizations to search for innovative business models and maximize their operational efficiencies without exceeding their IT budget.

However, as with any business model and process transformation, addressing the organizational culture and change issues to cloud computing are crucial to its success. This issue of *Cutter Business Technology Journal* offers astute evaluations from seven expert authors who outline various opportunities afforded by cloud computing solutions, myriad challenges that we face, and several strategies that will enable us to move forward successfully in our transformation journeys.

The evolution of cloud technologies has not stopped at public clouds; community, private, multicloud, and hybrid cloud solutions are also viable means of storing and processing data, depending on an organization's needs.

In This Issue

Many consider cloud computing's foundation layer — infrastructure as a service (IaaS) — as a utility service, akin to electricity, showing physical similarities in the delivery of the two services. In our first article, authors Priya Sinha, Cutter Consortium Senior Consultant James Mitchell, Jonathan Smith, and David Wallom tell us that "understanding how power markets are structured gives valuable insight into future market development for cloud computing." They leverage the analogy of the electricity market to explain how such features can be replicated for a global IaaS cloud computing market. The authors explore the benefits of the wholesale trading of contracts in delivering cloud services, how we can avoid the mistakes experienced by the electricity markets, and how we can leverage certain features of the electricity market that have led

to increased competition, price transparency, and greater resilience to external shocks. Sinha et al. also examine the current state of the cloud computing market and explore how existing cloud market participants have tried to evolve the market.

> *A hybrid cloud strategy not only prepares an organization for the future but also protects its investment today.*

Next, Prerna Lal discusses how the challenges of moving to a public cloud platform — performance concerns, control, compliance, and security threats — make the case for adopting a hybrid cloud strategy. This strategy entails moving some IT capabilities to the cloud, while maintaining core elements in-house, hosted on-premises. According to Lal, a hybrid model enables organizations "to optimally allocate their resources while keeping their current IT infrastructure operating at low risk." She goes on to say that "a hybrid cloud strategy not only prepares an organization for the future but also protects its investment today." Lal then describes the hybrid cloud market, its benefits, and a four-step process for designing a hybrid cloud strategy.

In our third piece, Łukasz Paciorkowski examines the "product as a service" transformation. He believes that "once products become cloud-connected…, organizations … can maintain a continuous dialogue with customers, provide additional services, and move from being 'just a provider' to an essential business partner." Paciorkowski discusses the impact of product as a service on the business model; the implications for the product lifecycle, supporting processes, and organizational structure; and the inevitable changes to company culture and identity. Although he believes that "a cloud-delivered offering allows for rapid expansion to other markets and unprecedented growth of the end-customer base," Paciorkowski also emphasizes the need to adjust your company culture to the cloud era to avoid the risk of failure.

In our final article, we hear from Cutter Consortium Senior Consultant Claude Baudoin, who takes a very useful and pragmatic view, with a look at lessons learned from the cloud journey to date. Those lessons include the shifting of costs from CAPEX to OPEX, the agility therein gained, the danger of shadow IT, improved rather than weakened security, the value of metrics, and the importance of keeping up with the evolution of cloud technologies. In concluding his lessons learned, Baudoin asserts, "It is now obvious that people who initially doubted the durability of the cloud phenomenon were wrong: the cloud has impacted business and IT much more completely than even the optimists thought possible."

We hope the insight provided in this issue gives you an enlightened perspective on the current and future cloud computing market and the guidance required to make well-informed decisions on the strategies and technologies that will provide your organization a competitive edge.

Upcoming Topics

Fintech: Emerging Trends, Future Directions
Steve Andriole

Business Technology Trends and Predictions 2019
Cutter Consortium

GENERATING "ELECTRICITY"

Journey to a Wholesale Cloud Computing Market

by Priya Sinha, James Mitchell, Jonathan Smith, and David Wallom

Many consider cloud computing's foundation layer — infrastructure as a service (IaaS) — as a utility service, similar to electricity. Various industry books and media sources have drawn a compelling analogy by highlighting the physical similarities in the delivery of the two services. If this analogy holds true, then we can easily apply the approaches to market structure and regulatory oversight that have worked well in electricity markets to the global cloud market.

In this article, we draw attention to the benefits of including a mechanism for the wholesale trading of contracts in delivering cloud services to mirror that of delivering electricity. Plus, we argue that it is possible for the cloud market to avoid the plethora of mistakes experienced in the many deregulated wholesale electricity markets globally. We also detail features from certain electricity markets that have led to attractive market attributes, such as increased competition, price transparency, and increased resilience to external shocks. Moreover, we leverage the analogy to explain how we can replicate such features for a global IaaS cloud computing market.

The Cloud Market

With organizations all over the world rapidly adopting cloud technology, cloud market revenue is growing at an astounding rate. Currently estimated at US $260 billion, the entire cloud market is expected to reach $411 billion by 2020. The vast majority of the public cloud market is dominated by a few hyperscale cloud service providers: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). These hyperscale providers offer their own core cloud infrastructure services, with a myriad of platform as a service (PaaS) and software as a service (SaaS) offerings built on top, in what is primarily a vertically integrated direct-to-consumer sales approach.

Cloud computing is a generic term given to consumer-focused third-party-operated IT services (i.e., users self-serve the resources they require). The widely adopted National Institute for Standards and Technology (NIST) model of cloud computing describes cloud computing in terms of service types, deployment models, and essential characteristics.[1]

Service types range from tightly constrained, highly specific, but easy-to-use SaaS offerings, through less constrained software or service development environments delivered as PaaS offerings, down to the highly generic, heavily customizable, somewhat harder-to-use IaaS offerings. In this article, we focus on IaaS, as it is commonly the underlying foundational layer for both PaaS and SaaS cloud services.

The deployment model is a significant cost driver, as the attractive economics of cloud are a direct result of the extent of sharing capacity among different users. Where the users are all from the same organization, the cloud service supporting those users is a "private cloud." Where those users come from different organizations, but they all belong to a well-defined community (e.g., academia, public sector, or financial institutions), the cloud service is termed a "community cloud." Where a cloud service accepts any user, then it is a "public cloud." Finally, the term "hybrid cloud" is used as well, typically referring to the use of both public and private clouds for the same application. Our focus here is on public cloud IaaS, including the component used in conjunction with a private cloud IaaS in hybrid cloud use cases.

All the essential characteristics defined by the NIST model should apply to a service for it to be considered a true "cloud." Of these, self-serve usage, scaling up and down elastically, and the option to pay under a pay-as-you-go or on-demand price plan are the attributes that make the analogy with electricity so strong. Understanding how power markets are structured gives valuable insight into future market development for cloud computing.

The Electric Utility

In many of the different electricity markets around the world, electricity is traded as if it were a commodity. Technology and market structure have developed since

the introduction of the electric utility in Godalming, England, in 1881, to enable the impression that electricity is a fully "fungible" (i.e., interchangeable) commodity. However, electricity delivered to the grid by different types of power stations has different characteristics to the electricity consumed by each class of end users, and hence clearly is not actually fungible. The differences include voltage, the current and its frequency, and other factors such as the reactive power component.

A key constraint of an electricity grid is that electricity supply and demand must never be unbalanced for more than a few seconds, at all points on the grid. This means that metering of output from each power generator — and metering of consumption by each end user — is paramount to be able to map the flow of currents in real time and hence forecast locations and times when the network will be under stress. This balancing challenge is managed with a combination of centralized control of some power stations, voluntary control (interruptible contracts) for others, and involuntary control (underfrequency protection relays or other involuntary curtailment measures) on end users.

In a deregulated market, a somewhat abstracted wholesale market layer is superimposed over these activities to minimize the use of them, thereby minimizing disrupted supply. The wholesale market derives both forward and spot prices that motivate actions similar to, but typically more economic and effective than, what a central operator would otherwise have achieved through mandated actions.

Most deregulated markets start off with the creation of three roles: a power-generating company (a "provider"), a power-supplying company (a "supplier"), and a network operator that owns and manages the electricity cable network that carries power from generators to end users. In many cases, network ownership is also separated from the system operation role.

The *provider* is responsible for owning and operating power stations and selling all the power the stations generate to the grid, as measured by their generation meter for each metered period. The *supplier* is responsible for buying all the power taken from the grid by a group of consumers, as measured by their consumption meters. In some power markets, there is a further separation of roles, with a *specialist meter operator* validating and guaranteeing the accuracy of the meter readings, and a *data collector* and a *data aggregator* for collecting and aggregating the metered production and usage, though these are not necessary for the configuration of the market itself. This metered data is held centrally, often by a regulated entity, and is accessible by some market participants. As metering technology evolves, and smart meters are installed in more jurisdictions, so too are the roles evolving. Consequently, there are important questions being raised and addressed relating to data privacy around detailed consumer electricity usage.

The delivery of electrical power is via a shared grid infrastructure, owned by one or more *network operators*, that is generally split into a state-level "transmission system" that transmits the power at high voltage over long distances with minimal losses per kilometer, and a more local "distribution grid" operating at a much lower voltage but with higher electrical losses, to cover the last few miles to each individual power consumer. A key facet of this arrangement is that there is no direct link between the power delivered to the grid and the power extracted from the grid. That is, a consumer cannot, in any physical way, choose a flow of electrons from a specific generator, though some suppliers do sell to customers intimating as such through green tariffs (where contracts rather than a flow of power provide an audit trail to evidence the original source of generation).

Cloud Computing as a Utility

The cloud computing market is evolving in a similar way to the electric utility model. In the same way that a power consumer can outsource the owning and operation of power plants, a cloud consumer can outsource the owning and operation of data centers and the IT services hosted within. This means that, like electricity, the cloud market also includes the distinct roles of provider and consumer.

As in the early days of the electric utility, all major cloud providers started off with a "vertically integrated" direct sales model (i.e., they also played the role of cloud supplier). However, several major cloud providers have started supporting sales through an indirect sales channel, leveraging specialist partners to reach out to the many different industries that could benefit from cloud services. They have allowed those channel partners to step into the billing chain, thus creating a more unique cloud supplier role, much like those electricity suppliers that stepped into the billing chain between electricity providers and electricity consumers.

Cloud metering, or measurement of services consumed, is far more complicated for cloud services than for electricity because there is no common unit of usage (e.g., the kilowatt-hour [kWh]) that can be used for all the myriad services that can be consumed under a cloud model. This has led to each cloud provider retaining the roles of meter operator and, generally, data collector. There are examples of independent third-party data aggregators; however, they rely upon the accuracy of metering performed by the cloud providers, which is a naturally conflicted position.

In contrast to grid-delivered electricity where there is no clearly identifiable flow of electrons from a particular generating station to a particular user, in the cloud it is absolutely possible to selectively utilize one cloud provider rather than another. This is important, as in the cloud market there is the concept of customer stickiness to a particular provider. One driver of this stickiness is that the provider holds large quantities of the consumer's data, which can be expensive to move between providers. This issue is known as "data gravity," and is not present in the electricity markets.

In contrast to grid-delivered electricity where there is no clearly identifiable flow of electrons from a particular generating station to a particular user, in the cloud it is absolutely possible to selectively utilize one cloud provider rather than another.

It is possible for cloud consumers to deliberately recreate for the cloud market the abstraction afforded to the electricity consumer due to the nature of a power grid. A cloud consumer can, in principle, choose to control and consume cloud services from multiple providers via a set of APIs that provide an abstraction layer, making the cloud providers' services fungible, at least in theory, and in practice for certain types of application. This area of development is evolving rapidly and is strongly supported by the European Commission, among others.

The two primary benefits of this abstraction are price competition and improved service quality due to reduced impact from service outages, as when one provider has an outage, it is possible for the cloud consumer to failover to an alternate provider. The resulting increased competition has the potential to drive innovation and disruption in a self-reinforcing cycle in well-designed markets.

In both electricity and cloud, providers invest billions of dollars in building infrastructure with the capacity to provide services. If vertically integrated, these providers sell services directly to end users who consume variable but small amounts at a time on short-term contracts. This leads to providers having a large exposure to the future prices of their services, together with exposure to the future volume of services they can sell based on their ability to attract users to consume and pay those prices for services in the future. The electricity market manages this price and volume risk by "hedging" with a range of potential risk mitigation instruments. These include highly customized, long-term "power purchase agreements" that contractually agree to usually fixed prices for the electricity generated by a specific generation station (including, in some cases, the purchaser holding the physical "outage risk").

There are less tailored methods of hedging available, too, including selling fungible "forward contracts" into the wholesale market for the electricity the generation station forecasts to produce in the future. Such approaches can be more flexible for the electricity provider, but rely on the existence of a liquid-traded wholesale market for electricity, since if the forecast proves to be inaccurate, the electricity provider must buy contracts to match what it has sold up front.

The cloud market offers solutions for the equivalent risk management challenges using what the major cloud service providers have variously termed "reserved instances," "reserved virtual machine instances," and "committed use discounts," which are different forms of forward agreements to buy agreed volumes of cloud services over specified future delivery terms with a pre-agreed pricing methodology.

Local Electricity vs. Global Cloud

There are many differently structured electricity markets around the world; generally, one in each country, with multiple markets in large, federated countries such as the US. Otherwise distinct electricity markets are connected in many cases by physical interconnection in the form of high-voltage alternating current (AC) power lines on land, and for markets separated by modest gaps, such as the English Channel or the Bass Strait between the Australian mainland and Tasmania, subsea direct current (DC) power lines. These now-interconnected markets were developed

separately and built specifically to serve the geographic area in question, based upon the physical characteristics of the generators, networks, and consumers in that area.

Over time, neighboring areas with different market designs have found opportunities to build some interconnection, and a range of operating models to flow electricity between markets has developed. Often, these models are also supported by the trading of electricity wholesale market products. In the absence of government interventions in the form of policies or subsidies that drive systematic electricity wholesale price differences between adjacent countries, individual power markets export power only to nearby markets because if power generation is not close to demand, then resistance within the power cables causes significant losses over larger distances.

This is not an issue for the provision of cloud services. The closest parallel for cloud services, which is not an issue for electricity, is latency (i.e., the time it takes for a user action to trigger an effect "in the cloud"). Even latency is not an issue for many cloud computing use cases, however, and this has enabled the hyperscale cloud providers to grow so quickly. Cloud providers can offer certain cloud services to latency-insensitive end users scattered all over the world, initially from a single data center that could be located anywhere on the planet.

In contrast to electric utilities that have remained predominantly national or regional (mainly due to physical losses and country-specific market design and retail rules), the cloud utility has rapidly become a single global market to an extent that is almost unrivaled for other perishable, capacity-based services. For each new incremental service, ready access to fungible customers spread across the globe has led to a small number of hyperscale cloud providers dominating the market. These have begun with IaaS offerings such as compute, storage, and networking as a service, and increasingly include platform services that form a layer on top of infrastructure, ranging from machine learning to video transcoding and everything in between.

Regulation of Oligopoly Electricity Markets

There are electricity markets around the world that are at least as dominated by a few providers as the global cloud computing IaaS markets. In France, for example, prior to deregulation in 2015, the incumbent controlled 85% of the market.[2] Each dominant electricity provider is subject to heavy government regulatory oversight, primarily because electricity is a service consumed by domestic end users, who are heavily protected by competition law. This has put a lot of pressure on the government to get regulation right, with mixed success.

In contrast, cloud IaaS is mainly sold as a business-to-business (B2B) service (i.e., it is almost exclusively consumed by businesses, including SaaS vendors). As such, cloud IaaS has not yet received the direct attention of the competition authorities and national regulators, which tend to focus on any examples of anti-competitive behavior in the business-to-consumer (B2C) SaaS layer. This may be about to change, however. The Financial Stability Board, the global organization tasked with identifying systemic risks to the world's economies, has highlighted to its members that there is a growing concentration risk as financial institutions all select vendors from the same small group of leading IT service providers.[3]

The Current Cloud Market

In this section, we consider the shape and distribution of the current market for IaaS cloud computing and explore how existing cloud market participants have themselves tried to evolve the market.

Market Domination by Hyperscale Providers

In the current global market for IaaS cloud computing, AWS has the largest market share, followed by Microsoft Azure, Alibaba Cloud, and GCP. The market share is not straightforward to quantify, as the major cloud providers each defines its addressable markets differently and ever more broadly as time goes by, as they build ever more targeted services on top of their core infrastructure services, competing in more markets, thereby diluting their market share even as they further dominate their core market. Revenue is really the only way to measure market share for cloud, as there is no universally accepted unit of measure for cloud (i.e., nothing akin to the kWh for electricity). Instead, a bewildering number of different cloud resources are metered, and mostly charged for, by the cloud providers themselves.

A typical rule of thumb used by competition authorities to measure market concentration is the Herfindahl-Hirschman Index (HHI).[4] The Competition and Markets Authority (CMA) in the UK, for example, regards markets with HHI < 1,000 as "unconcentrated," HHI

1,000-2,000 as "concentrated," and HHI > 2,000 as "highly concentrated."[5] The US Department of Justice (DOJ) has higher thresholds; namely, HHI < 1,500 as "unconcentrated," HHI 1,500-2,500 as "moderately concentrated," and HHI > 2,500 as "highly concentrated."[6] By way of comparison, the UK's energy regulator, Ofgem, calculated HHI for the UK power market to be 1,599 in June 2017 (i.e., it is "concentrated" or "moderately concentrated," but not "highly concentrated").[7]

We have calculated the HHI for the global IaaS public cloud services market as between 2,050 and 2,100, depending on how many market participants the category "others" is broken into. In other words, the global IaaS public services cloud market should be deemed by competition authorities to be "highly concentrated" by the CMA's definition and "moderately concentrated" by the DOJ guidelines. It is highly unlikely, however, that every local market is shared in the same way; thus, it can reasonably be assumed that in AWS's largest markets (such as the US, where Alibaba Cloud is likely to have a lower market share than its enormous Chinese market share), the localized HHI value could well be much higher. Unfortunately, there is no reliable data publicly reported that allows us to confirm this.

The DOJ observes that "mergers resulting in moderately concentrated markets that involve an increase in HHI of more than 100 points potentially raise significant competitive concerns and warrant scrutiny."[8] By our calculations, the acquisition by AWS of a competitor with as low a market share as 1.15% would increase HHI by 100 points and hence trigger this recommendation of scrutiny.

Market Evolution

It is vitally important to point out, at this juncture, that Amazon has not built this dominant market share by acquisition but by organic growth. It was an early pioneer for the public cloud industry, enjoying incredible growth while delivering ever-lower prices for customers. It is much harder to justify a forced breakup using anti-trust legislation of a dominant provider that has achieved that position through organic growth.

In stark contrast to AWS's staggering organic growth is the acquisitive way in which IBM, which used to hold a dominant position in the IT industry, has tried to compete. It replaced its own SmartCloud offering with the acquisition of Softlayer, and most recently has announced its $34 billion acquisition of Red Hat.[9]

Amazon, Microsoft, and, most recently, Google, have voluntarily embraced the notion of selling cloud indirectly to customers through a resale channel. Cloud resellers are functionally equivalent to suppliers in the electricity industry (i.e., they step into the billing chain such that they pay one bill but are paid by many individual customers). Interestingly, this was exactly the solution that AT&T requested, and which was granted when US District Judge Harold Greene ordered the 1984 breakup of AT&T's US telecom monopoly.[10] This at least provides long-term competition at the point of price setting to end-user customers, even if the underlying cloud provider is the same company. We can consider this as equivalent to having a dominant electricity generation technology being sold by the manufacturer of that technology to various competing power generation companies. That is, the introduction of resale should be seen as extremely positive by competition watchdogs, and the introduction of any restraints on that resale should be monitored closely. It is in fact a very positive signal that the leading public cloud providers are embracing channel sales voluntarily.

Despite the implementation of resale, society is still left with a black swan risk that there may be fundamental flaws in any underlying technology used by these hyperscale providers that may be common among them. However, the technology companies have, thus far, done a good job of managing this kind of risk and dealing with any issues that have resulted. There are also related risks, such as the heavily publicized design fault in Intel chips[11] that has affected most cloud service providers in a way that greater public cloud market fragmentation would not have mitigated. It has been notable that the public cloud service providers acted to remedy the situation from the Intel chip design flaw far more quickly than would most enterprises operating their own private clouds.

While the major cloud service providers do have formal cloud reseller programs in place, a properly sophisticated resale market is not yet fully implemented, at least not in the manner that one would recognize as analogous to that found in sophisticated electricity markets. The cloud market can therefore be reasonably assumed to be operating below optimal efficiency, both for cloud providers and for cloud buyers. However, we remain optimistic that the cloud market is evolving

in a direction that is reducing black swan risks, not exacerbating them.

Acknowledgments

The work of Priya Sinha and David Wallom has been supported through Innovate UK Knowledge Transfer Partnerships.

James Mitchell would like to thank John Woodley for his years of mentorship in the areas of market development and price competition.

Endnotes

[1]Mell, Peter, and Timothy Grance. "The NIST Definition of Cloud Computing." Special Publication 800-145, National Institute of Standards and Technology (NIST), September 2011 (https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-145.pdf).

[2]Bayer, Edith. "Report on the French Power System." Agora Energiewende, October 2015 (https://www.agora-energiewende.de/fileadmin2/Projekte/2014/CP-Frankreich/CP_France_1015_update_web.pdf).

[3]"Financial Stability Implications from FinTech." Financial Stability Board, 27 June 2017 (http://www.fsb.org/wp-content/uploads/R270617.pdf).

[4]"Herfindahl-Hirschman Index." US Department of Justice (DOJ), 31 July 2018 (https://www.justice.gov/atr/herfindahl-hirschman-index).

[5]"Merger Assessment Guidelines: CC2/OFT1254." Competition Commission (CC2) and Office of Fair Trading (OFT), September 2010 (https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/284449/OFT1254.pdf).

[6]Brandenburger, Rachel, and Joseph Matelis. "The 2010 US Horizontal Merger Guidelines: A Historical and International Perspective." *Antitrust*, Vol. 25, No. 3, Summer 2011 (https://www.justice.gov/atr/2010-us-horizontal-merger-guidelines-historical-and-international-perspective).

[7]"Electricity Supply Market Shares by Company: Domestic (GB)." Ofgem, October 2018 (https://www.ofgem.gov.uk/data-portal/electricity-supply-market-shares-company-domestic-gb).

[8]Brandenburger and Matelis (see 6).

[9]"IBM to Acquire Red Hat, Completely Changing the Cloud Landscape and Becoming World's #1 Hybrid Cloud Provider." IBM News Room, 28 October 2018 (https://newsroom.ibm.com/2018-10-28-IBM-To-Acquire-Red-Hat-Completely-Changing-The-Cloud-Landscape-And-Becoming-Worlds-1-Hybrid-Cloud-Provider).

[10]Pinheiro, John. "AT&T Divestiture & the Telecommunications Market." *Berkeley Technology Law Journal*, Vol. 2, No. 2, September 1987 (https://scholarship.law.berkeley.edu/cgi/viewcontent.cgi?article=1064&context=btlj).

[11]Williams, Chris. "Kernel-Memory-Leaking Intel Processor Design Flaw Forces Linux, Windows Redesign." *The Register*, 2 January 2018 (https://www.theregister.co.uk/2018/01/02/intel_cpu_design_flaw).

Priya Sinha is a Data Science Research Associate–Cloud Ecosystem at Oxford e-Research Centre, University of Oxford, UK, working with Prof. David Wallom and Dr. Neil Caithness. From 2016-2018, she was a KTP Fellow at the University of Oxford. Dr. Sinha is an accomplished Indian classical dancer (Bharatnatyam) and a recipient of the Centre for Cultural Resources and Training scholarship from the Ministry of Culture, India. She earned a PhD in experimental condensed matter physics from the University of Leeds, UK, with a Marie Curie Fellowship; a dual master's degree in nanoscience and nanotechnology from the University of Joseph Fourier, France, with a French government Scholarship (Égide-CFIU), and the University of Delhi, India; and a bachelor of science degree with honors in physics from Sri Venkateswara College, University of Delhi, India. She can be reached at sinha.priyasmita@gmail.com.

James Mitchell is a Senior Consultant with Cutter Consortium's Business Technology & Digital Transformation Strategies practice; CEO of Strategic Blue, a financial cloud brokerage firm offering for cloud computing, services more commonly seen in the commoditized energy markets; and a Cutter Consortium Summit *keynoter. For the past five years, Dr. Mitchell has been pioneering various cloud brokerage concepts that are now gaining widespread acceptance and adoption based on his background as an originator of structured physical commodity transactions at Morgan Stanley. Featured in books and reports as the leading authority on how the IT market will change because of the move to a utility (or "cloud") delivery model, he is a frequent speaker at events and often successfully predicts how the cloud market will develop over the next five years. Dr. Mitchell has a doctorate in DNA nanostructures from the University of Oxford, UK, and a first-class master's degree in experimental and theoretical physics from Cambridge University, UK. He can be reached at jmitchell@cutter.com.*

Jonathan Smith has more than 11 years' experience within the energy wholesale and retail markets. During his career, he worked as an originator at Morgan Stanley in the UK power and gas wholesale markets, before heading up an independent energy retailers' group-wide electricity and gas forecasting, trading, and pricing activities. There, Dr. Smith designed hedging strategies to manage key risks, with trade execution on spot and forward markets in the UK and Germany. He oversaw a more than tenfold growth of the hedging book to several 100 million pounds of annually delivered commodity. Dr. Smith earned a master's degree in physics and a PhD in semi-conductor physics, both from the University of Oxford, UK. He can be reached at jonathansmith1@gmail.com.

David Wallom is Associate Professor/Associate Director of Innovation, Oxford e-Research Centre, University of Oxford, UK, where he leads three research activities: energy and ICT, advanced e-infrastructure and cloud computing, and volunteer computing. He has led over 30 research projects in areas such as cloud utilization, smart energy grids, research data management, green IT, ICT security, and institutional repositories. Prof. Wallom is a member of the Centre of Excellence in Cyber Security (recognized by the Government Communications Headquarters) and the UK Space Agency Ground segment Advisory Group. He can be reached at david.wallom@oerc.ox.ac.uk.

LOOKING TO THE FUTURE

Designing a Hybrid Cloud Strategy for Competitive Advantage

by Prerna Lal

Cloud computing, one of the most disruptive trends in the IT market over the past decade, has fueled digital disruption and has enabled businesses to pursue digital transformation. Today's competitive business environment needs constant connectivity — the lifeblood of digital transformation — between customers, devices, and organizations. Consequently, organizations of all sizes are exploring the myriad ways that disruptive technologies, particularly cloud, can enable them to stay connected with customers through any device, irrespective of location and time. The results of this exploration are providing endless opportunities for organizations to reinvent themselves by creating new markets and new business models. But to keep pace with ever-growing competition, organizations need to deploy an IT infrastructure that is agile and can rapidly align with fast-changing business needs. Cloud computing is that infrastructure. For several years now, we have witnessed the great migration of IT from traditional on-premises locations to the cloud platform, with organizations leveraging cloud benefits in terms of degree of scalability, pay-per-use flexibility, and agility.

To keep pace with ever-growing competition, organizations need to deploy an IT infrastructure that is agile and can rapidly align with fast-changing business needs.

There has been a rapid increase in the number of organizations adopting cloud computing. Moreover, organizational spend on cloud-specific solutions is expected to grow at more than six times the rate of general IT spending through 2020.[1] As cloud computing became ever-more prevalent, private and hybrid cloud models received less consideration than public clouds, largely due to the popularity of public cloud IT services being offered by industry leaders (i.e., Amazon, Microsoft, and Google), which users can easily and quickly access. Initially, these IT services were mainly related to test and development environments in need of fast deployment and with less criticality than production workloads. But over the years, organizations have realized that concerns over performance, control, compliance, and security threats outshine the benefits of moving to the public cloud.

In view of these challenges, many organizations are now focusing on a hybrid cloud strategy: moving part of their IT capabilities to the cloud, while maintaining core elements in-house, hosted on-premises. The hybrid model is becoming immensely customary among organizations, as it enables them to optimally allocate their resources while keeping their current IT infrastructure operating at low risk. A hybrid cloud strategy not only prepares an organization for the future but also protects its investment today.

Defining Hybrid Cloud

In simple terms, a hybrid cloud model integrates an organization's private cloud infrastructure with one or more public clouds. Different organizations, as well as cloud service providers, have explained the term "hybrid cloud" in several ways; Table 1 summarizes a few definitions. The definition from the National Institute of Standards and Technology (NIST) is the most accepted: a hybrid cloud is "a composition of two or more distinct cloud infrastructures … that remain unique entities but are bound together by standardized or proprietary technology that enables data or application portability."[2]

Hybrid cloud is sometimes confused with *multicloud*, but there is a difference in each of these strategies. The multicloud strategy involves using a mix of cloud services from different cloud vendors for specific nonconnected workloads, while the hybrid cloud strategy creates a mix of on-premises private cloud and third-party public cloud services, which may be from different vendors, with orchestration between the two.

	Hybrid Cloud Definition	Source
National Institute of Standards and Technology	"A composition of two or more distinct cloud infrastructures … that remain unique entities, but are bound together by standardized or proprietary technology that enables data or application portability."	Mell, Peter, and Timothy Grance. "The NIST Definition of Cloud Computing." Special Publication 800-145, National Institute of Standards and Technology (NIST), September 2011 (https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-145.pdf).
Amazon Web Services	"In cloud computing, hybrid cloud refers to the use of both on-premises resources in addition to public cloud resources."	"Hybrid Cloud Architectures with AWS." Amazon Web Services (AWS) (https://aws.amazon.com/enterprise/hybrid/).
Microsoft Azure	"A hybrid cloud is a computing environment [that] combines a public cloud and a private cloud by allowing data and applications to be shared between them."	"What Is a Hybrid Cloud?" Microsoft Azure (https://azure.microsoft.com/en-in/overview/what-is-hybrid-cloud-computing/).
IBM	"The hybrid model is primarily a private cloud that allows an organization to tap into a public cloud when and where it makes sense."	"Going Hybrid: Best of Both Worlds in Cloud Computing." *Forward View* (https://www.ibm.com/midmarket/us/en/forwardview/article_Cloud1_1204.html).

Table 1 — Defining hybrid cloud.

One of the most attractive benefits of the multicloud strategy is the risk avoidance of vendor lock-in, as switching to a different cloud vendor for a specific workload will not impact other cloud-based applications. Data sovereignty may be another reason to select a multicloud strategy, since with certain countries, laws and regulations require organizations to store data in the country in which they operate.

Hybrid Cloud Market

The size of the global hybrid cloud market, which was valued at US $36 billion in 2017, is expected to reach $171 billion by 2025, growing at a compound annual growth rate of 21.7% from 2018 to 2025.[3] The global hybrid cloud market can be segmented by component (solution or service), service model (infrastructure as a service, platform as a service, or software as a service [SaaS]), organization size (small, medium, or large), industry vertical (healthcare, BFSI [banking, financial services, and insurance], retail, transportation and logistics, IT, and telecom), or geographic region. Microsoft, Hewlett Packard Enterprise (HPE), VMware, Amazon Web Services (AWS), Rackspace, Dell Technologies, Google, and Alibaba Cloud are among the key players in the global hybrid cloud market.

Hybrid Cloud Benefits

In the following sections, we look at the benefits of hybrid cloud, including security and compliance, scalability, and cost-effectiveness.

Security and Compliance

Cloud service providers claim to ensure the security of data and applications for all cloud models, but recent data breaches of public clouds have raised concerns.[4] Such breaches have caused organizations to revisit their cloud model choice as they seek to ensure better security. Indeed, security is the key reason for organizations migrating their workloads and applications from public cloud to either on-premises or private cloud.[5] Hybrid clouds not only provide a higher level of security for sensitive data but also ensure implementation of compliance requirements for a specific industry. A majority of the organizations adopting hybrid cloud belong to banking and financial services organizations, a sector in which data security and compliance is of paramount importance. Organizations implementing hybrid cloud continue to enjoy the various benefits of public cloud–based solutions for better customer engagement while ensuring compliance with regulations and control of customer data, which is kept

securely inside the bank's or financial services organization's firewall using on-premises or private cloud.

Scalability

Scalability refers to a system's ability to expand output on demand when resources are added.[6] Hybrid clouds allow organizations to leverage their legacy IT investments and to seamlessly scale them up to the cloud based on the fluctuating demands of the business, all without compromising data security.

Cost-Effectiveness

Scalability and cost-effectiveness go hand in hand with hybrid clouds. Hybrid cloud implementation saves organizations from investing a huge amount on creating IT resources to meet fluctuating demand. Hybrid cloud implementation results in using cost-effective, on-premises IT resources along with premium-priced, flexible, public cloud resources to handle short-term spikes in demand. Thus, organizations have the flexibility of paying for the use of public cloud–based IT resources only for a short period of time, as needed, instead of making a huge capital investment in expanding their on-premises or private cloud, which could remain idle for long periods of time.

Designing a Hybrid Cloud Strategy

Successful implementation of a hybrid cloud strategy requires alignment with the business strategy. Hybrid cloud strategy is not limited to the incorporation and balance of *new* versus *legacy* technology. Rather, it should be designed to help organizations leverage emerging technologies for rapid business transformation in today's fast-paced digital world, ensuring operational efficiency and keeping costs under control. Let's take a look at a four-step hybrid cloud strategy.

Step 1: Align Hybrid Cloud Goals to Business Goals and Objectives

Designing a hybrid cloud strategy is not only about understanding technical issues but also requires a thorough understanding of the business perspective of the organization. Understanding the business perspective helps identify how to use a hybrid cloud model to optimize IT asset utilization and to respond rapidly to changing business needs while maintaining necessary controls.

The very first step is to identify business needs as well as pain points. Business needs and pain points may be related to any of the functional areas, be it marketing, sales, HR, finance, or manufacturing. The next step is to identify whether cloud-based IT solutions can help solve business problems or reduce the impact of business pain points. Assuming the answer is yes, the final step is to select the hybrid cloud deployment model (i.e., whether to go for a traditional data center and/or private cloud and a public cloud or choose a managed private cloud and a public cloud). This critical decision depends on various factors, such as the size and maturity level of the organization, the size and capabilities of the in-house IT team, the organization's financial strategy with respect to IT investments, customer location, access to the applications, compliance and regulations related to data management, flexibility of the service-level agreement, and the organization's preference around the control of data and intellectual property.

Step 2: Design Technology Portfolio

Technical aspects that can help with the decision as to whether a workload/application should reside on an on-premises/private cloud or on a public cloud include an evaluation on four parameters: performance, security, integration, and data volume. Workloads with significant performance needs, security requirements, multiple back-end integrations, or large data volume are better on private clouds. On the other hand, workloads with minimal performance integration or storage requirements tend to be better suited for public clouds and possibly SaaS solutions. Other important factors include scalability requirements, customizations, and support needs. Analyzing these factors can help produce a clear understanding of which applications/workloads should move to/remain on an on-premises/private cloud and what can be managed on the public cloud.

Step 3: Plan Migration

Migration to the hybrid cloud should be done in stages. Initial stages may involve migrating noncritical workloads, while later stages may include business-critical workloads. The organization needs to carry out a detailed assessment of the application portfolios

moving to the cloud. The results of this assessment, based on parameters such as security, performance, and so on, determine migration priority.

An organization may select different approaches for migrating different workloads, taking into consideration how to take maximum advantage of capabilities available only on a cloud platform. The most common approaches for cloud migration include rehost, replatform, repurchase, refactor, retire, and retain, commonly known as the "six Rs." The selection of a migration strategy will be on a case-by-case basis. For example, if business pain points are related to management of data servers, then refactoring may be the best option; on the other hand, employee-leave reporting issues may be best solved through rehosting.

It is important to include key employees from multiple departments because migration to a cloud platform may require changes in business processes across the organization. These employees can help the migration team assess the workload importance and identify the specific security and regulatory requirements to which they must adhere. This will help the team not only prioritize the workload for migration but also design a better hybrid cloud strategy.

During migration planning, the organization should also assess the skill sets both of employees who will manage cloud-based applications and of employees who will use them. It is important to ensure that the in-house team that may need to manage the on-premises/private cloud part of the hybrid cloud have the required skills and knowledge. In the event of a skills gap, the organization will need to decide whether to provide training or to hire new resources.

The organization must provide communication regarding the migration to the hybrid cloud to every stakeholder. It is imperative to sensitize every stakeholder to the benefits of the hybrid cloud deployment; benefits that may include an increase in operational efficiency, improved accessibility of applications, reduction in costs, and so on. Communication is critical, as well-informed stakeholders who understand the benefits will appreciate the efforts involved in the migration process, will have less resistance to any changes in business processes, and will accept the hybrid cloud model faster. Finally, testing of the hybrid cloud before going live, including data exchange between the on-premises/private cloud and the public cloud, is of utmost importance to ensure that there is no room for error.

Step 4: Manage Hybrid Cloud

Hybrid cloud management is complex, as both the public and private cloud have their own native APIs and resources for managing storage, networking, provisioning, security, and so forth. Thus, organizations with diverse IT portfolios require a single interface that can combine information from the different clouds to reduce complexity and provide a greater sense of control when managing a hybrid cloud environment.

The process of implementing a hybrid cloud is an evolving one. Organizations may start migration with a small number of workloads or applications and expand the portfolio later after evaluating outcomes. In addition, new advancements in cloud-based technology may persuade organizations to change their hybrid cloud strategy.

Conclusion

Today's wave of digital disruption is unprecedented. Hybrid cloud, in particular, is changing the management and delivery of IT services, while enabling organizations to search for innovative business models and maximize their operational efficiencies without exceeding the IT budget. The pace at which organizations are migrating to a hybrid cloud model is remarkable. However, hybrid cloud implementation poses challenges. Though cloud vendors claim to ensure a smooth migration to a hybrid cloud environment, the integration of an on-premises/private cloud with a public cloud still remains a key concern. Some challenges that organizations face while implementing a hybrid cloud, and which must be carefully considered while designing a hybrid cloud strategy, are management of the complexity of the technology portfolio to be migrated, creation of a governance structure to manage the hybrid cloud, compatibility issues, and, the most critical challenge, ensuring security. Despite these concerns, the future of hybrid cloud has a silver lining, as cloud service vendors are constantly innovating and collaborating in order to provide cloud-based solutions that are safe, scalable, and available globally. It will be interesting to witness how technology revolutionizes the next wave of cloud computing.

Endnotes

[1]Gantz, John F., and Pam Miller. "The Salesforce Economy: Enabling 1.9 Million New Jobs and $389 Billion in New Revenue Over the Next Five Years." IDC, 2016 (https://cdn2.hubspot.net/hubfs/2148167/Nagarro-Sep2016-Theme/PDF_Files/IDC-salesforce-economy-study-2016.pdf?t=1499578990603).

[2]Mell, Peter, and Timothy Grance. "The NIST Definition of Cloud Computing." Special Publication 800-145, National Institute of Standards and Technology (NIST), September 2011 (https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf).

[3]Rake, Rachita. "Hybrid Cloud Market." Allied Market Research, October 2018 (https://www.alliedmarketresearch.com/hybrid-cloud-market).

[4]See, for example: Dawn-Hiscox, Tanwen. "Human Error Causes Australian Broadcasting Corporation Breach." Data Centre Dynamics, 20 November 2017 (https://www.datacenterdynamics.com/news/human-error-causes-australian-broadcasting-corporation-breach); and "Misconfigured Servers Lead to Multiple Major Breaches — Are You at Risk?" Foresite, 10 April 2018 (https://foresite.com/misconfigured-servers-lead-multiple-major-breaches-risk).

[5]Haranas, Mark. "Businesses Moving from Public Cloud Due to Security, Says IDC Survey." *CRN*, The Channel Company, 13 August 2018 (https://www.crn.com/businesses-moving-from-public-cloud-due-to-security-says-idc-survey).

[6]Nielsen, Christian, and Morten Lund. "Building Scalable Business Models." *MIT Sloan Management Review*, Winter 2018, 7 December 2017 (https://sloanreview.mit.edu/article/building-scalable-business-models).

Prerna Lal is Assistant Professor in Information Management at International Management Institute, India. She has more than 14 years' experience in academics and research in the areas of data warehousing/data mining, business analytics, management information systems, software project management, IT service management, cyber law, and cloud computing. Ms. Lal has been published in various journals and publications. She is SAP-certified and has an ITIL Foundation certificate in IT service management. Ms. Lal earned an MBA from the Indian Institute of Technology Roorkee and a PhD in cloud computing from Banasthali University, India. She can be reached at prernalal@yahoo.com.

COMING TO TERMS WITH REALITY

The Product Is Dead. Long Live the Service!

by Łukasz Paciorkowski

We live in times when renting is often better than owning, and services are overtaking products. Consider, for example, how BMW through its DriveNow car-sharing program allows registered customers to pay for driving hours rather than buy a new car; how Rolls-Royce, known for its jet engines, sells the engines "as a service," charging fees on the basis of flight hours; and that Michelin offers tire leasing programs priced per mile. Moreover, Michelin commits to its customers' seeing a reduction in fuel consumption; if that target is unmet, the consumer gets a reduced price. Today's digital technologies, including cloud, Internet of Things (IoT), and others, enable offerings as a service. Importantly, moving from product to service has transformational effects on the organizations executing such change.

For almost 20 years, cloud and related technologies have been transforming entire industries. Interestingly, most of the discussions around transformation tend to focus on the technical aspects, technical solutions, platforms, and various as-a-service offerings. Although the technology perspective is very important, it's not the most critical. Truly transformational projects require a much broader perspective, covering the impact on the business model; product lifecycle, processes, and organizational structure; and company culture and identity. A truly successful digitization project will change a company to its core. Thus, product-to-service transformation is probably the best example of the pervasiveness of digital technologies.

Cloud-Enabled Transformation: Beyond Technology

At my company, we are convinced that every business is a digital business today. As experienced architects, engineers, and digital designers, we are fascinated with technology and the opportunities it brings. But, based on our experience, we also know that digital transformation is mainly about business, people, and culture. Companies that failed during their cloud-enabled digital transformation journey often struggled with the cultural and operating model change. The "soft" part of the transformation affected their performance and overall success more than the technical aspects related to implementation of the transformation.

Examples of failed digital transformations show many factors contributing to poor performance, missed targets, and loss of credibility. GE is one example where digital change did not go exactly as planned.[1] Instead of becoming one of the world's top software companies, GE is in the process of selling its digital child Predix in an attempt to redefine its transformational roadmap.[2] Let's explore some less obvious implications of the transformation toward a cloud-centric digital company.

Companies that failed during their cloud-enabled digital transformation journey often struggled with the cultural and operating model change.

Implications for the Business Model

Moving to cloud-based solutions goes beyond a technical migration of workloads (applications) to a data center managed by Amazon or Google. For your organization to realize the full potential of such a change, you must consider how your business will (or can) change. Will you still sell products? Or maybe it is time to move to services? Will you charge customers full price for your goods, or will you move to a service fee charged periodically? What will your pricing strategy be? How will a new pricing model impact your liquidity?

In early 2000, Michelin made a very bold move going from being "just" a tire producer to a fleet service provider. Becoming a service provider was a huge change to the operating paradigm of the entire company. Initial trials were unsuccessful because the value proposition was not clearly articulated. But Michelin

quickly learned from its mistakes. After investing further in cloud-enabled IoT solutions, the company launched EFFIFUEL, an offering that uses IoT sensors, cloud, big data, and analytics to help customers optimize their transportation fleet costs.[3] So, today, the value proposition is straightforward. The World Economic Forum, in fact, reports that "a reduction in fuel consumption of 2.5 liters per 100 km represents annual savings of €3,200 [about US $3,660] for long-haul transport (at least 2.1% reduction in total cost of ownership and 8 tonnes in CO_2 emissions)."[4]

A cloud-delivered offering allows for rapid expansion to other markets and unprecedented growth of the end-customer base.

Companies using cloud-related technologies, like IoT, to enhance their offerings with an "as a service" model are expanding their operations to the full product lifecycle. Before "connected" products, companies had little knowledge about how their products were used, where they were used, and who used them. Once products become cloud-connected, however, organizations like Michelin can maintain a continuous dialogue with customers, provide additional services, and move from being "just a provider" to an essential business partner.

A cloud-delivered offering allows for rapid expansion to other markets and unprecedented growth of the end-customer base. Well-known examples are Apple Store or Google Play. Software delivery has never been easier. Entering new markets in countries around the world is greatly simplified. Once you publish your application to the cloud marketplace, it can quickly reach millions of potential customers. For example, Uber has shown how fast it can roll out operations around the world. Indeed, the cloud eliminated the need for heavy investment in IT infrastructure in each region where it operates. Local operations can be reduced to the bare minimum. Thus, cloud-based solutions enable speed and agility.

Another very important aspect of becoming an Internet-first, "cloudified" company is the customer engagement model. You (and your company) are no longer the only marketing channel for your products and offerings; your customers are also a channel. Even if you do not have a presence on Facebook, LinkedIn, or Twitter, it doesn't mean that your company is not being discussed on such a platform. Being able to respond quickly to changes in customer sentiment is critical to company/product/offering success. When social media exploded with the recent incident of a United Airlines passenger being dragged off one of the airline's planes, the airline, unprepared to address the customer issue in a timely and public manner, quickly faced a PR nightmare.[5] So you need to ask yourself if your company has the relevant capability to handle digital customer relationships; if it is lacking in that area, you need to develop it.

Implications for the Product Lifecycle, Supporting Processes, and Organizational Structure

Selling product as a service requires a continuous maintenance and improvement strategy. Your product lifecycle strategy will have to change drastically. Instead of a one-, two-, or seven-year release cycle, for instance, consider update releases every one to six months. Remember that the leaders in cloud design, like Netflix, are pushing hundreds of changes into production every day! Is it even possible to assure such update frequency for your physical product?

For example, take Tesla. Tesla cars are known to receive major overnight updates. This is possible using cloud infrastructure, connectivity, and physical product configurability managed through software.[6] Artificial intelligence and machine learning allow modern car manufacturers to optimize the operational parameters of your car, leading to improved performance without the need to visit the service garage. The ability to instantly upload and download vast amounts of relevant digitized data from (and to) the vehicles would be very difficult without the supporting cloud infrastructure.

Consider what needs to change in your back-end processes and supporting organization to be able to fix bugs, test updates, and deploy patches to your connected products. Here again, software takes center stage. Delivering changes with speed requires a lot of supporting IT infrastructure, software configuration, and automated procedures implementation. Only then can you shorten the develop-test-deploy process from weeks, or even months, to days and hours. Thankfully, DevOps, SysOps, and continuous delivery are disciplines native to cloud environments. Many products and solutions (e.g., Jenkins, Puppet, ELK, or Consul) are

available to support you in automating the process of updates, upgrades, and bug-fix solutions. Be sure to look at the market leaders — Netflix, Amazon, or Facebook — to fully understand that IT automation is a business differentiator.[7]

As for organizational structure, let's consider again the example of BMW's DriveNow service.[8] The team tasked with the creation of a new car-sharing platform originated from the product-based company — BMW — but its setup was very different from that of its parent organization. A large, multidisciplinary team of designers, architects, engineers, UX specialists, and data analysts was split into small, nimble teams. The operating method of choice was Agile, which helped introduce continuous release cycles with adjustable speed of delivery (called *cadence* in Agile). Some ideation-to-production sprints lasted a few days, others a few months. Eventually the DriveNow service was successfully launched across cities in Europe and is expanding to other areas.

As we can see from this example, turning toward cloud-related technologies most likely will impact the organizational structure in your company. Some IT functions will be centralized; others can be completely outsourced and provided from any place in the world. Business and operations will require access to cloud-based apps, and, therefore, the company network will need to extend way beyond typical intranet boundaries. If we add IoT to cloud-based apps, your network might even extend to your customers' premises. Furthermore, if most of your IT infrastructure is in the cloud, consider a working-from-anywhere approach versus colocation. Limiting office space is one way to lower operational costs. But are you ready to allow your employees to work remotely? Here's where culture comes into play.

Implications for Company Culture and Identity

Digital transformation just about always triggers a discussion around the company's culture and identity. Many product-based companies came into existence due to the founders' love of mechanical engineering. Such companies thrive on production excellence, unparalleled quality, and a culture of engineers working with their hands. Mechanical and electrical engineers are these companies' differentiating asset and constitute a major part of their workforce. The new breed of connected products, which heavily leverage software and a cloud infrastructure, require the same (or an even greater) number of programmers and IT architects. An engineering company is suddenly becoming a software company, and physical products are giving way to services. But what is such a company's core business? Should it outsource or insource software development or partner with another organization? What is its market differentiator — the physical product, a software platform, or a service around the offering?

Changes in a company's business model, operations, and workforce demographics necessarily mean that the internal culture of the company will also change. If you neglect the cultural aspect in your digital transformation journey, there is a high risk that you will fail. Corporate culture can be changed, but it is a long and difficult process. Many people will not get it or not agree with it, so prepare for increased staff turnover. Nevertheless, adjusting your company culture to the cloud era is a necessary step. Explore examples and lessons learned from others. Stories of failed digital transformation programs from GE, Proctor & Gamble, LEGO, and others provide useful tips on how to influence a shift in a company's culture.[9]

Changes in a company's business model, operations, and workforce demographics necessarily mean that the internal culture of the company will also change.

Conclusion

Cloud-enabled transformations should never be discussed purely in a technology context. Digital transformation starts in the minds and hearts of people. If you do not address the concerns and cultural challenges, even the best cloud-based solution will not help you excel and thrive in today's digital economy. Especially in the context of product-to-service transformation, increased attention to the soft areas of digital change is critical. If your colleagues and employees do not understand the new, service-based digital reality, how can your customers buy into it? It is worth remembering that working with digital means working mainly with people.

Endnotes

[1]Colvin, Geoff. "What the Hell Happened at GE?" *Fortune*, 24 May 2018 (http://fortune.com/longform/ge-decline-what-the-hell-happened).

[2]Tozer, Alice. "GE Digital Could Be the Latest Cast-Off from General Electric's Ever-Shrinking Portfolio." *Fortune*, 31 July 2018 (http://fortune.com/2018/07/31/general-electric-ge-digital-on-sale/).

[3]Nadim, Emad. "Michelin: Tires-as-a-Service." HBS Digital Initiative, Harvard Business School, 17 November 2016 (https://rctom.hbs.org/submission/michelin-tires-as-a-service).

[4]"Michelin Solutions." World Economic Forum (http://reports.weforum.org/digital-transformation/michelin-solutions/).

[5]Cooper, Gael Fashingbauer. "United's PR Disaster Lands in Twitter's Unfriendly Skies." CNET, 10 April 2017 (https://www.cnet.com/news/united-airlines-dragged-passenger-off-twitter-memes-unfriendly-skies).

[6]O'Kane, Sean. "Tesla Can Change So Much with Over-the-Air Updates That It's Messing with Some Owners' Heads." The Verge, 2 June 2018 (https://www.theverge.com/2018/6/2/17413732/tesla-over-the-air-software-updates-brake).

[7]Null, Christopher. "10 Companies Killing It at DevOps." TechBeacon (https://techbeacon.com/10-companies-killing-it-devops).

[8]Gottschkes, Sebastian. "Agile by Example: DriveNow." DZone, 2 August 2016 (https://dzone.com/articles/agile-by-example-drivenow).

[9]Colon, Tony. "The First Step to Shift to a Digital Culture." *Forbes*, 21 March 2018 (https://www.forbes.com/sites/forbestechcouncil/2018/03/21/the-first-step-to-shift-to-a-digital-culture/#384d25a11c7).

Łukasz Paciorkowski is Principal Architect and cofounder of Architects for Business, an independent group of technical professionals on a journey to digitize the world. He focuses on translating business visions into technical solutions, particularly in the areas of IoT, cloud, mixed reality, and digitization. Previously, Mr. Paciorkowski worked at Accenture as Senior Manager and IoT Domain Lead, where he combined strategy, business, and technology views to create end-to-end engagement plans for global corporations from multiple industries; and at IBM as Executive Architect and Cloud Advisor, focusing on digital transformation enabled by the cloud. He is also a former member of the IBM Academy of Technology Leadership Team, where he was coleader of the Organic Innovation program. Mr. Paciorkowski is coauthor of multiple publications, has one patent pending, and is a frequent conference speaker and lecturer. In a visual, unconventional, and sometimes provocative way, he encourages listeners to think outside the box and look at technology topics from an unusual perspective. Mr. Paciorkowski holds a master's degree in artificial intelligence from the WIT Academy, Poland. He can be reached at lukasz.paciorkowski@a4bee.com.

FACING THE NEW NORMAL

Cloud Lessons Learned

by Claude Baudoin

Since the birth of the phrase "cloud computing" about 12 years ago, we have seen many papers and articles covering it with similar subtitles or taglines: "Cloud Challenges and Opportunities," "Cloud Computing: Myths and Realities," or something along those lines. This author is admittedly guilty as charged. Responding to another call for papers on the benefits, barriers, risks, and so forth, of the cloud is therefore a challenge: what can be useful to the reader that has not been said already?

This article tries to take a very pragmatic viewpoint: what are the things we have learned during this journey? What do most reasonable analysts and users now agree on, as opposed to questions on which the jury is still out? What should you spend time worrying about, and what should you consider settled, for good or for bad? Finally, with various lessons learned, what should you educate your managers or clients about, so they don't waste their time or yours?

My answers to those questions have been influenced by several things:

- A historical perspective: utility computing is as old as mainframe timesharing — and was predicted by John McCarthy as early as 1961.
- Cloud services are a very broad domain; we use many of them as individuals without much hesitation.
- What I have heard from my consulting clients.
- My work since 2011 within the Cloud Standards Customer Council (CSCC), recently rebranded as the Cloud Working Group of the Object Management Group (OMG), which I cochair.

With this background in mind, I will attempt to provide guidance on the following topics:

- The benefits of cloud adoption — in particular, is cost reduction a real argument and what are the other justifications?
- What is the impact of the cloud on how IT relates to the rest of the business?
- How do you address security in the cloud?
- How do you measure success?
- How do you select a cloud strategy given the increasing variety of deployment options?

Cost Shifting, Not Cost Reduction

One of the original arguments made to justify "going to the cloud" was that it would significantly reduce IT costs. We had heard this before; it was one of the major arguments for outsourcing in the 1990s. Cloud vendors emphasized that argument because it was an obvious way to sell their offering. Some customers really believed it, while others may have doubted it, but cynically thought that "it may or may not be true, but that's the only argument I can make to sell this to my boss or to the board."

We have learned that the impact of the cloud on IT costs is more complex than a simple reduction. Over the total lifetime of a system or application, renting the service in the cloud could very well cost more than it would to operate it on premises. However:

- **The nature of the cost is different.** Instead of a CAPEX that will be depreciated over three or five years, it is a "pay per use" monthly cost. Thus, there is a risk that you will keep paying a monthly fee beyond the point where the two curves crossed, but frankly, how often have we used equipment or software beyond its depreciation lifetime?
- **Depreciation is a nice accounting trick, but you still need to pay the price up front.** For companies with cash flow limitations, paying by the month is a lot easier to manage.

- **The cloud reduces or eliminates the risk of ending up with costly unused (or underutilized) systems.** If you want or need to get off a particular cloud service, you can stop paying, as long as you have ensured that the contract had reasonable termination clauses.

In fact, the cloud has completely upset the traditional total cost of ownership calculation because the initial purchase or licensing costs are low or zero (you still have to count the labor costs incurred to adopt the solution, including such things as data migration and user training), while the recurring costs are higher than the traditional 15% or 18% annual support costs incurred for on-premises systems.

Lesson Learned

You should explain to decision makers that the key effects of cloud adoption are to shift costs from CAPEX to OPEX and to reduce the risk of write-offs of unused assets — which is likely to be worthwhile even if the total lifetime cost of a solution is not necessarily lower.

> *The cloud allows you to try something, fail, and move on to something else with relatively less severe consequences.*

It's About Agility

So if cost reduction is not the key motivation or benefit of going to the cloud, what other key factor(s) should genuinely influence that decision?

We have learned that a big part of the answer is agility. Call it flexibility or scalability if you want, but "agility" is a bit broader and evokes the principles of the Agile movement. The cloud allows you to try something, fail, and move on to something else with relatively less severe consequences. The cloud does not require you to calculate with high accuracy the amount of resources you will need, or the length of time you will need them. You should be able to scale up and down, to add and remove services, to get the latest updates, and so on — without having to do the work yourself.

Interestingly, while we can be reluctant to use agility or flexibility as a primary motivating factor in the enterprise, it is exactly what we consider on a personal basis when we buy a new smartphone and a communications plan for it. Think about this: I would never have dreamed of paying $750 for a phone — yet I will fairly happily pay $25 per month for three years instead. During that time, I know that I can trade the phone in for a new model, add or remove services, and so forth.

I am not arguing that you should lightly choose the first solution you hear without studying your requirements. After all, switching from one CRM solution in the cloud (to take an example) to another one is costly in terms of configuration, data migration, and learning curve. But choosing the wrong system, or sizing it inaccurately, is no longer a five-year sentence if it is cloud-based.

Lesson Learned

A primary business justification for "going to the cloud" should be the agility it gives the enterprise to adapt to change through faster IT sourcing and the ability to scale up or down, or to change solutions, whenever necessary.

The Cloud and "Shadow IT"

The IT department used to own the keys to the computing resources, literally and figuratively. The data center was the crown jewel of IT. To most users, it was impressive and mysterious, and a symbol of the relationship between IT and its users.

Because users in a line of business (LOB) had to go through the corporate IT department to obtain any IT resources, they had to follow the process defined by IT. This made the CIO or IT manager both powerful and resented. Few IT people understood the users' plight of preparing justifications, listening to technical jargon they did not understand, and sitting in countless meetings until finally (i.e., after their project had been delayed by lack of the required capability) receiving a deliverable that may or may not have been what they needed. And then they got the bill.

What we have learned is that the cloud has enabled a *shadow IT* to emerge. That sounds scary (mostly to the IT people), right? But we have also learned that shadow IT is not totally a bad thing, as long as there is communication, coordination, architecture, and governance. Let's discuss this more specifically.

Shadow IT means that each LOB, or each function (e.g., HR, marketing, sales, field service) is making its own decisions and paying its own costs for certain IT capabilities that it deems necessary, without going through the centrally controlled process on which we just heaped sarcasm. As a result, the agility benefit of cloud solutions can now directly and immediately benefit end users. Since many of these departments do not have the skills needed to perform a good study of their requirements and carefully select a solution, the decision process has shortcomings and can border on the arbitrary. Here are three typical ways in which a shadow IT solution is selected:

1. A department employee who seems to be pretty good with technology (often a junior one) is told to research and propose a solution.

2. A manager who has used a certain solution at home — for his homeowners' association, say, or heard about it from a buddy at the gym — decides that it must be good enough for his business need (that's how so many confidential documents end up being stored in Google Docs).

3. A consultant is hired to do the study. A very cheap consultant is chosen because the cost of this study needs to fly below the financial controls radar. In terms of the suitability of the recommendation, you get what you pay for.

What is typically *not* done is to go to the CIO and say, "Look, we need a solution for X, and we need it quickly. We're not willing to go through a protracted process, and we heard that you guys typically take way too long. But we do want to make sure that what we choose will not create a mess; we want to make certain we can exchange data with other systems — even though we may not yet know what those are; we realize that we may need some form of support at some point; and we want to remain reasonably good friends. Will you help us achieve those goals? Call it a "proof of concept" if that sounds better. And guess what, you're even going to learn something useful in the process, which could be applied elsewhere in the organization."

Seen this way (i.e., very optimistically), shadow IT projects don't quite deserve their name; they are no longer happening in complete darkness. They try to achieve a balance: work fast without having to jump through bureaucratic hoops, but not risk rejection by the host organism.

Some progressive CIOs understand the need to monitor, facilitate, and even embrace these processes rather than ignore or fight them. One way to do this (usually in a large enterprise because of the resource allocation that it requires) is to create a small "rapid reaction team" of IT specialists who work in quasi-startup mode. Their job is to quickly build prototypes for the internal clients, trying free or cheap solutions in the cloud, allowing decisions to be made in the course of days or weeks, not months, while keeping the IT organization informed and ensuring that the selected solutions can be integrated and supported. A key challenge is to protect this team from being "recaptured" inside mainstream projects when a resource crunch or an emergency happens.

Lesson Learned

If IT does not embrace agility in delivering solutions to the business, the business will go ahead and select solutions in the cloud without involving IT. But if the business procures such solutions on its own, integration will become very difficult.

Shadow IT projects don't quite deserve their name; they are no longer happening in complete darkness.

Follow the Data

One of the key reasons organizations invoke for staying away from public clouds is security, and specifically the risk that unauthorized parties will access data.

From a technology perspective, placing data in a public cloud may actually *improve* security rather than weaken it. There are four simple reasons for this:

1. A cloud service provider (CSP) is likely to have many more security specialists than any one of its clients and can therefore do a better job at protecting its infrastructure and reacting to an attack.

2. If a CSP suffers a damaging attack, the business consequences can be devastating. Therefore, a CSP is even more highly motivated to protect its infrastructure than its clients are to protect theirs.

3. Multi-tenancy, which is a usual property of a public cloud, makes it harder to find a particular client's data. If I stored my data on premises, attackers could find my IP addresses through the domain name system and would know exactly where to attack. If my documents are in a public cloud instead, attackers must find a needle in a haystack.[1]

4. Most security breaches are committed by insiders. Within a company, it is relatively easy to copy data without leaving fingerprints. Retrieving data from a public CSP, however, is likely to leave a trace.

> *As if the issue of personal data protection was not bad enough, we need to superimpose on it the problem of data residency.*

Statistics about security breaches suffered by public CSPs are misleading. Companies prefer to keep quiet about internal security incidents, which means that claims of higher security for on-premises infrastructure are suspect. (The same arguments, by the way, can be made to debunk the claims that data is less available when placed in the cloud.)

However, we need to balance the technology arguments with the management perspective. Indeed, a combination of lack of education about the above considerations and issues of risk management, regulatory compliance, and legal exposure may result in staying away from the cloud. Personal survival may also come into play. If there is a successful attack against a bank and credit card numbers are stolen, the CEO can hope to escape personal blame (often by firing the CIO immediately after the initial mop-up effort). But if the CEO approved a proposal to move to the cloud, and a similar attack against the CSP results in a loss of confidential information, then the CEO's head is also likely to roll.

The punitive measures included in the General Data Protection Regulation (GDPR) of the European Union (EU) have heightened, for managers of large companies, fear about losing control of data. If data contains personal information about citizens of the EU, a company may be fined as much as 4% of its annual revenue (or €20 million, whichever is *higher*) if this data is leaked "intentionally or negligently." It is only if a data leak could not reasonably be expected, by virtue of having implemented strict security measures such as those detailed in Article 32,[2] that the organization can avoid severe penalization.

Other countries are now considering the adoption of measures similar to those of the GDPR, which is often considered "model legislation."[3] This is the case in Canada with the Personal Information Protection and Electronic Documents Act (PIPEDA).[4] And while it is unlikely to become law in its exact current form, a bill was recently introduced in the US Congress that would impose not only GDPR-like penalties but also jail time on executives of companies that fail to protect consumers' personal data.

As if the issue of personal data protection was not difficult enough, we need to superimpose on it the problem of *data residency*. This term refers to the set of issues and challenges posed by the location of data — especially if it ends up located in (or even just passing through) a jurisdiction that exposes it to greater risks, including legal demands for access, prohibition of strong encryption, or state-sanctioned spying.[5]

Rational decisions about using or not using a public cloud are made more difficult by the fact that some IT managers will exaggerate the risks in order to preserve their fiefdom. Claiming that a public cloud is inherently less secure is a way to preserve the data center and staff from which they derive their status.

Lesson Learned

Much of what is feared about security in the cloud is a myth. Whether to place data outside of the enterprise or not is ultimately a cross-functional risk management decision, involving business management (including the legal department), IT, and the CSP itself, which must specify in the cloud service agreement the security measures it has deployed and its commitment to prompt reporting and remediation of attacks.

Measuring the Impact

It should be clear by now that I do not propose to measure the success of a cloud transformation strategy through cost savings, since the nature of the costs is changing, and the ultimate gain or loss can only be assessed over the entire lifetime of an application or system.

We have learned that security and availability metrics (e.g., number of incidents, time to repair) can be improved rather than reduced by a move to the cloud, but that the issue is *clouded* in myth. Therefore, it seems essential to firmly base this discussion on facts.

Baseline metrics about the "as is" environment are often missing. Before embarking on a cloud transformation, an organization should have data about its current on-premises performance levels, especially in terms of security and reliability. Only then will it be able to determine later whether the cloud provides better results — allowing fact-based decisions to change policy or to change providers.

Since we highlighted, early on, the importance of agility, it must be reflected in metrics (and a baseline measurement), such as the following:

- Time taken (including management and administrative tasks) to upgrade IT resources (e.g., to add storage space)
- Number of incidents related to lack of scalability (e.g., degraded operations due to a saturated resource)
- Time required to fulfill an internal customer request, from the study of the internal customer's requirement until its satisfaction
- User satisfaction with the IT sourcing process derived from surveys

A well-managed cloud adoption program should lead to improvements in all these metrics — unless IT was doing a fabulous job on all counts, which, frankly, is very rare.

Establishing goals for these metrics, before a cloud migration, should be possible in many cases — and a CSP can help. It should have statistics from its other customers, and its service-level agreement should specify such things as how long it takes to bring a new server or a new disk online, or how many hours of downtime are needed to upgrade an application.

Lesson Learned

Measure what matters to the business, not only to IT. And do not compare the performance of a CSP to ideal on-premises statistics of 100% reliability and 0% security incidents, because that's not the truth.

The Increasingly Complex Cloud Deployment Options

At the beginning of the cloud, things were relatively simple. Thanks to the National Institute for Standards and Technology (NIST) Cloud Computing Reference Architecture from 2011,[6] we could divide the cloud world into three service models (software as a service [SaaS], platform as a service [PaaS], and infrastructure as a service [IaaS]) and four deployment models (public, private, hybrid, and community). Most initial adoptions focused on public clouds, and either on SaaS or IaaS — none of which required a great deal of technical complexity to adopt.

In 2018, things have become much more complex. Here are some things we have learned since NIST's original work:

- Hybrid clouds and multiclouds have become much more frequent than we first expected, and they require significant tooling in order to manage the complex environment they create.
- To avoid being locked into a relationship with a single provider, ways to package the payload (especially in the case of applications; data storage is easier) have been developed to achieve portability across providers.

We now have a proliferation of deployment technologies: containers, container as a service (CaaS), cloud-native applications, microservices, function as a service (FaaS), bare metal servers, virtual machines (VMs), nested VMs … the list seems to increase each quarter.[7]

This evolution means that if an organization is migrating to the cloud some of its internal applications (or legacy commercial ones that are not delivered in the cloud by the vendor), it needs people with new technical knowledge to either package or refactor these applications. While retraining existing IT personnel may be possible, the learning curve can be steep and not everyone can climb it. In the short term, an IT organization probably needs to hire specialized consultants or contractors, which does not come cheap.

Lesson Learned

Do not underestimate the pace of evolution of cloud technologies. If no one in the organization watches (and understands) the new cloud delivery technologies, you

may choose the wrong one, or you may be misled by a vendor that will fail to mention a better option because the vendor does not offer it.

Conclusion: Cloud Is the New Normal

The "accelerating pace of IT" has become a trite phrase. We should be quite used to these ever-shortening cycles of transformation, so it should not be a surprise that in the course of a decade the cloud has changed the landscape of IT so quickly for so many organizations.

It is now obvious that people who initially doubted the durability of the cloud phenomenon were wrong: the cloud has impacted business and IT much more completely than even the optimists thought possible. McCarthy's prediction of an age of "utility computing" came true half a century after it was made. Even with some real risks and persistent misunderstandings (what I have described as myths), the cloud is clearly the "new normal" of IT management.

If your business is not to supply computing resources, then owning such resources is not important — it is even a distraction. Most companies don't operate their own power generator or water treatment plant. Except in some very specific cases (e.g., real-time constraints, massive data access rates), why should they operate their own servers?

What is important, instead, is managing this environment and the resulting relationships. The IT organization becomes the broker in charge of supplying appropriate and scalable resources in a timely manner to its internal customers. Thus, in conclusion, we have learned that the key capabilities of IT in the cloud era will not be software development or system administration, but the ability to:

- Elicit and understand user requirements
- Keep track of technology and market trends
- Implement a solid (but rapid) sourcing process
- Manage relationships with providers
- Measure performance and service levels and react to changes

Endnotes

[1]This argument was made as far back as July 2010 by my late Cutter Consortium colleague Mitch Ummel; see: "Cloud Computing: Separating the Hype from the Reality." Webinar, Cutter Consortium, 7 July 2010 (https://www.cutter.com/event/cloud-computing-separating-hype-reality-webinar-391736).

[2]"General Data Protection Regulation (GDPR): Chapter 4, Article 32 — Security of Processing." Intersoft Consulting (https://gdpr-info.eu/art-32-gdpr).

[3]Remarks by Satya Nadella, Chairman of Microsoft, during an interview at the Viva Technology conference in Paris, 24 May 2018. From the author's notes.

[4]"The Personal Information Protection and Electronic Documents Act (PIPEDA)." Office of the Privacy Commissioner of Canada (https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda).

[5]The nature of the risks, particularly their potential negative impact on the cloud industry, are documented in a paper by the Cloud Standards Customer Council (CSCC); see: "Data Residency Challenges: A Joint Paper with the Object Management Group (OMG)" (https://www.omg.org/cloud/deliverables/data-residency-challenges.htm).

[6]Liu, Fang, et al. "NIST Cloud Computing Reference Architecture." National Institute of Standards and Technology (NIST), 8 September 2011 (https://www.nist.gov/publications/nist-cloud-computing-reference-architecture).

[7]To bring some clarity to this complexity, the Cloud Working Group of OMG is starting a new guide on this topic, due out in the first quarter of 2019.

Claude Baudoin is a Senior Consultant with Cutter Consortium's Business & Enterprise Architecture and Data Analytics & Digital Technologies *practices. He has been owner and Principal Consultant of cébé IT and Knowledge Management, a boutique consulting practice, since 2009. Prior to that, Mr. Baudoin spent 35 years in IT and software management roles in industry, including 26 years at Schlumberger, the global oilfield services company, in France and the US. Since 2015, he has served as an energy domain advisor to the Object Management Group and the Industrial Internet Consortium and has sat on the Steering Committee of the Cloud Standards Customer Council. Mr. Baudoin has published two books on software engineering,* Méthodes de Programmation *and* Realizing the Object-Oriented Lifecycle, *and numerous papers and conference presentations, including a number of Cutter Consortium reports, advisors, and* Cutter Business Technology Journal *articles. He holds two patents related to IT infrastructure and security. Mr. Baudoin's undergraduate engineering degree is from École Polytechnique in Paris, and he holds a master of science degree in computer science from Stanford University. He can be reached at cbaudoin@cutter.com.*

Cutter Consortium Membership
The Ultimate Access to the Experts

Cutter Consortium Membership opens up multiple avenues to interact with Cutter's experts to brainstorm and gain guidance to transform your organization and boost success.

Like everything business technology, one size does not fit all.

That's why we encourage you to choose the Membership that's right for your organization. Whether you choose Digital Transformation & Innovation, Enterprise-wide, Practice-specific, or CIO Membership you'll see a strong return.

Contact us at +1 781 648 8700 or sales@cutter.com to arrange a sample inquiry call with a Cutter expert and see for yourself how quickly your return on Membership can be realized.

What Do You Get from Cutter Membership?

- Get guidance in leveraging new strategies, emerging technologies, and business management practices to enable digital transformation and boost competitive advantage
- Learn how to mine data to create new products and services and improve customer experience
- Get input on how to reduce expenses through more cost-effective strategies
- Gain insights and get ideas on achieving sustainable innovation, successful change management, and prudent risk management
- Get coaching and insights on leadership and team-building practices that boost productivity
- Discover vendor-agnostic advice to ensure unbiased purchasing decisions

CUTTER Business Technology Journal

About Cutter Consortium

Cutter Consortium is a unique, global business technology advisory firm dedicated to helping organizations leverage emerging technologies and the latest business management thinking to achieve competitive advantage and mission success. Through its research, training, executive education, and consulting, Cutter Consortium enables digital transformation.

Cutter Consortium helps clients address the spectrum of challenges technology change brings — from disruption of business models and the sustainable innovation, change management, and leadership a new order demands, to the creation, implementation, and optimization of software and systems that power newly holistic enterprise and business unit strategies.

Cutter Consortium pushes the thinking in the field by fostering debate and collaboration among its global community of thought leaders. Coupled with its famously objective "no ties to vendors" policy, Cutter Consortium's *Access to the Experts* approach delivers cutting-edge, objective information and innovative solutions to its clients worldwide.

For more information, visit www.cutter.com or call us at +1 781 648 8700.