Cutter IT Journal

The Journal of Information Technology Management

Vol. 25, No. 8
August 2012

"Cloud computing represents a game-changing paradigm shift in the industry, with consumers (and the enterprises they serve) crying out for a broadened set of standards in areas where standards have never existed before."

— Mitchell Ummel, Guest Editor

Cloud Computing Standards

CUTTER CONSORTIUM

Opening Statement

by Mitchell Ummel, Guest Editor

The dialogue surrounding standards, standards-setting processes, and the overall benefit of standards is hardly a new one — especially in the field of information technology. Few will argue with that claim — ours is indeed an industry built on standards. In fact, we find the myriad foundational technologies upon which cloud computing services are typically built (including, but not limited to, TCP/IP, HTML, and XML) to be already decades mature in terms of standardization and industry adoption.

Why then the increasing interest in cloud computing standards in 2012? The answer is arguably quite simple: cloud computing represents a game-changing paradigm shift in the industry, with consumers (and the enterprises they serve) crying out for a broadened set of standards in areas where standards have never existed before.

Much is at stake as the cloud computing standards "battlefield" emerges.

BATTLEFIELD PLAYERS

Large, established vendors seek to carve out new (or protect existing) market share as they recast and rebrand their existing computing products and services to be marketecturally[1] "cloud friendly." Small, new vendors seek to enter the marketplace, compete on a level playing field, and find opportunities to either innovate or position themselves as market followers — promoting their wares as "cheaper, faster, better" (pick one or more) than the established competition's.

New markets develop with increasing velocity, as industries of "cloud brokers" or "cloud VARs" arise. We find thousands of startups are able to spin up tens of thousands of offerings on a daily basis with zero up-front investment ("Look, Ma, no VC required!") and ramp up capacity on demand if the offering is the one in 10,000 that goes viral — for a day, a year, or (as the next Google or Facebook) for decades.

We, the consumers of cloud computing services, seek only the promise of value for our enterprise computing dollar, based upon a strong foundation of trust, assurance, and security. Figure 1 illustrates how consumer-driven standards may influence and broaden markets, increase competition and innovation, and lead to products and services with greater consumer value.

TALK ABOUT INTEROPERABILITY …

The analogy between the evolution of the electric energy industry and cloud computing is oftentimes used, and for good reason. It's likely the most applicable predictor of where this industry is heading over the next 10-20 years.

Although slight regional variances exist, it's generally the case that I, as a consumer of electric power, can plug in my appliance anywhere in the world and expect it to work efficiently, safely, and reliably. Standards for voltage regulation, plug/outlet design, and circuit protection are mature and widely embraced, and the electric appliance industry can compete, and innovate, on a level playing field for the benefit of consumers worldwide. The clock radio in my office is one of thousands of such devices in the market today. It is a Japanese brand, with technologies patented in the US and subcomponents manufactured in South Korea. This amazing device, purchased for US $12, is incredibly reliable and runs on "pennies" of electrical power per week.

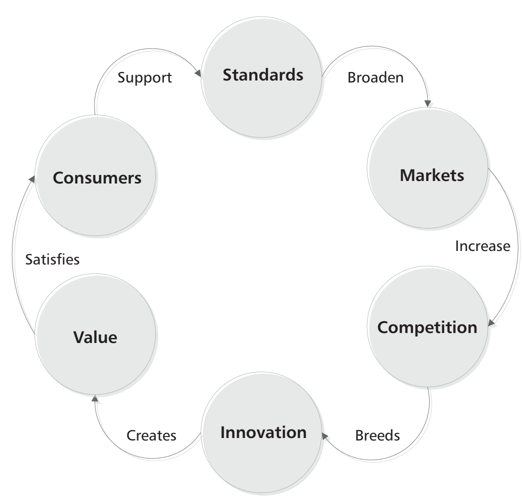

Figure 1 — Standards value chain.

The emerging global cloud computing market space, built on a fabric of underlying standards promoting interoperability, security, assurance, and trust, also promises to amaze consumers in the coming decade. Current computing standards aside, the paradigm shift to the cloud compels us to address new areas of standardization where none have existed before. These include the following:

- Industry-wide agreement and widespread adoption of standard terminology and measures for basic elemental units such as "compute instance," "compute-hour," "MB-month," and so on
- Standard language and provisions guiding the overall structure and framework for a new generation of consumer-friendly cloud computing service-level agreements (SLAs)
- Foundational standards for portability and interoperability of compute instances (virtual machines) across and among participating public cloud computing service providers, as well as private cloud architectures serving an enterprise

CLOUD COMPUTING 2020

One might imagine a not-too-distant future where consumers of standards-based cloud computing services are able to automatically and rapidly (i.e., in near real time) provision and re-provision their enterprise computing load across multiple competing providers based on a number of attributes, including instantaneous pricing, geographical region, temporal (time) shifting, service-level requirements, and so on. Standards will allow consumers to seamlessly move their computing load from their private cloud (enterprise data center) to public cloud providers and back again, as demand or failover requirements dictate.

In the future, we might even see a virtual "stock market" develop for brokers or aggregators of cloud computing services — where compute instances, capacity, and options are traded, in real time, much like stocks are traded on a stock exchange. Enterprises that continue to invest in "bare metal" data center capacity *and* subscribe to industry cloud computing standards for interoperability may even have the opportunity to sell excess capacity back to the "grid," in much the same way that a homeowner with a residential wind power generator can sell back excess power to the grid in utility markets that provide "net metering" capabilities.

TRUTH IN CLOUD COMPUTING

Akin to truth in lending requirements for financial institutions offering mortgages, or truth in labeling requirements for the food and drug industry, consumers of cloud computing services may soon demand a truth in computing "label" on services they consume. Such a label would allow for easy comparison of pricing among competing providers, and assurance of elemental service-level capability and interoperability, as a means to promote and enforce a basic level of transparency and accountability among vendors in the marketplace.

IN THIS ISSUE

How does this all play out? Our contributing authors tackle this topic head on in this month's issue of *Cutter IT Journal*. Now let me introduce these authors, each of whom contributes to our understanding of this broad subject area through their insights and unique domain knowledge.

Carlos Viniegra's thoughtful article sets the stage with a foundational discussion of standards setting in general. He writes that "economic history shows us that consumers will take advantage of savings and drive markets toward their optimal economies of scale." Viniegra further explores the motivation for stakeholders involved in standardization and the behavior types that typically exist in the context of "agency theory." He then presents three "information asymmetry" problems that, according to him, will increase the pressure for government regulation in this market.

Next, coauthors Duff Bailey and Jeffrey Wu promote the expansion of beneficial standards in the cloud computing space. While they recognize that government regulation and certification will continue to play a role, they also argue that it will be primarily market forces that drive even the largest of cloud vendors to standardization.

In our third article, David Bernstein presents a core theme that a standard is not a document but a process. Within his article, he reflects upon past standards-setting-as-a-process experiences from the early days of Ethernet leading up to and including WLAN and related wireless security protocols. He further concludes that we are stuck in a cloud computing "early adopter" stage with consumer confidence, in the absence of practical standards, sorely lacking.

Next, Cutter Senior Consultant and cloud computing expert Claude Baudoin promotes a consumer/customer-driven (vs. industry-driven) standards-setting focus. He also provides a rundown of the numerous complementary, and sometimes competing, standards-setting initiatives already underway. Baudoin argues that, due to the rapidly moving cloud adoption cycle, we the users must become active, making our voices heard through initiatives such as the OMG-affiliated Cloud Standards Customer Council (CSCC), an organization that is attempting to prioritize and harmonize cloud standards-setting activities among all standards bodies that are active today.

Our fifth author, Beth Cohen, explores the interrelation between cloud computing standards efforts and the open source movement. Her deep dive into the frameworks and technologies that make up "de facto" cloud standards suggests that the open source community is best positioned to address the standards "gap" that currently exists in cloud computing. Her bold article, entitled "An Open Source Approach to IaaS Cloud Standards," is a technical delight.

And last but not least, Aditya Watal leads us through the cloud computing "stack" from SaaS, through PaaS, to IaaS, and in the process presents a unique vantage point on the viability of standards at each layer. Suggesting that the greatest opportunities for standardization exist at the PaaS layer, Watal profiles the top six PaaS offerings on the market today, assessing them against his proposed Cloud Architecture Reference Model for Applications (CARMA). He then goes on to describe "what an ideal application sitting on such a standards-compliant platform would look like."

CALL FOR ACTION

We at Cutter Consortium are seeing an increasing interest by our readers in learning more about cloud computing standards setting, how to become involved in those efforts, and how best to practically embrace cloud computing within their organizations as part of an enterprise cloud computing strategic roadmap. We are interested in your feedback and success stories in this area and would seek to include your experiences in future research published by Cutter.

ENDNOTE

[1]According to Wikipedia, "Marchitecture (or Marketecture) is a portmanteau of the words 'marketing' and 'architecture.' The term is applied to any form of electronic architecture perceived to have been produced purely for marketing reasons. It may be used by a vendor to place itself in such a way as to promote all [its] strongest abilities whilst simultaneously masking [its] weaknesses."

Mitchell Ummel is a Senior Consultant with Cutter's Business & Enterprise Architecture *practice. He is President of UmmelGroup International, Inc., a US-based business and technology management consulting firm. Mr. Ummel is a visionary who is most well known for championing practical application of innovative but lightweight IT process improvement methodologies into today's culture within large organizations. In recent years, he has enjoyed advising state governments and private sector enterprises in their planning and architecture for large, multimillion-dollar business and technology transformations.*

Mr. Ummel's IT experience spans 25 years and includes service in a variety of CIO, executive management, training, coaching, mentoring, and consulting roles for government, law enforcement and justice departments, and telecommunications, electric/gas utility, monitored home security, health insurance, and Internet-based product or service companies. Earlier in his career, he served as a state of Kansas agency CIO, where he contributed to the inception of the Kansas statewide technical architecture, helped define next-generation project and portfolio governance processes, and led technology initiatives enabling electronic health surveillance records management, vital records security/privacy, national bioterrorism response, GIS integration, and public health and environmental informatics among local, state, and federal stakeholders.

Mr. Ummel holds advanced degrees in mathematics and computer science from Fort Hays State University. He is a graduate of, and contributor to, Kauffman Foundation's Entrepreneurial Leadership and New Tech Venture Development Programs, is a frequent speaker at national business and technology conferences, and is a member of the Project Management Institute. He can be reached at mummel@cutter.com.

Keep It Simple! Framing Cloud Computing with Agency Theory

by Carlos Viniegra

A problem clearly stated is a problem half solved.

— Dorothea Brande

Viewed from the perspective of economics, the cloud computing phenomenon can be understood as the change in scale of an industry that, given new technical developments, can move to larger optimal economies of scale. The current debate over the need to establish cloud computing frameworks either by government regulation or industry self-regulation hasn't acknowledged that most of the problems arising from this industry can be framed and understood with tools out of economics and political science.

In this article, I propose the use of agency theory as a means of classifying the main cloud computing issues as information asymmetry problems. These problems, known as moral hazard, conflict of interest, and adverse selection, have been used for decades to describe and solve governance issues for large-scale utility industries. As the utilities examples show, when participants of an industry become a vital component of the economy or too big to fail, they also come to be a source of systemic risk. This becomes evident when crises arise, mainly because information asymmetry problems are left unchecked. When we analyze the projected size and current practices of the cloud computing industry, significant questions arise as to whether self-regulation mechanisms will suffice to manage the systemic risk this industry will add to the economy. This, in turn, shows that there are two probable development pathways for the utility computing business: one in which poor governance progresses to bubble-and-bust cycles, and another in which lessons learned in other utility markets allow for the early adoption of sound governance practices that permit the stable growth of this important market.

GRAPPLING WITH GOVERNANCE

At international cloud computing conferences organized by industry, governments, and international organizations, the discussion almost invariably revolves around the governance structure this developing industry should have. On one side of the discussion, industry movers and shakers make the case for self-regulation as a means of sustaining innovation, while government officials and NGOs point toward the need for government intervention due to the industry's perceived lawlessness and the risk to the general public. The interesting part of these exchanges is that, on both sides of the equation, participants have tried to frame the conversation within the boundaries of engineering and business-specific language. Although we live in a time when value production often depends on holistic approaches from different fields of knowledge, we still have opportunity gaps at the crossroads of engineering and economics. This means that many economists feel too baffled by the fast pace of technical advances to engage the digital revolution, while many engineers may feel themselves fully capable of building workable solutions for the problems that arise in society around the digital environment they are creating.

Regardless of the tacit armistice between engineers and economists, though, the time has come to strip the engineering jargon out of the digital revolution and analyze new developments through the lens of economics. This will enable us to simplify the discussion and open roads toward new solutions.

ECONOMIES OF SCALE, SYSTEMIC RISK, AND UTILITIES

Around the world there is a constant ratio regarding the size and occurrence of small and large companies. Small and medium enterprises (SMEs) represent more than 90% of all companies by number while large enterprises (LEs) constitute less than 10%. Today there is an international debate regarding the size and power of LEs, yet these enterprises have a key role to fulfill in any economy. While SMEs create the lion's share of employment — breaking the size of batches supplied by LEs and personalizing delivery of goods and services — LEs are needed to achieve economies of scale and "commoditization" of goods and services such as energy, transport, finance, and raw materials. Also there's the special case of LEs that make large-scale investments in innovative and cutting-edge technologies such as computer chip and new drug development. Without the efficiencies made possible by LEs, large portions of humanity wouldn't have access to basic goods and services, and many SME business models wouldn't even be possible.

Conventional electricity generation offers an example of optimal scale when the ideal economy of scale for a plant (the size where production cost is minimal) nears one gigawatt. This means that if regulations artificially limited or expanded the size of all plants to suboptimal scale, consumers would have to pay for the inefficiency on their electric bills. Another example is the Concorde airplane, which was retired in 2003. By overreaching the optimal economy of scale that lay *below* the speed of sound, the market failed, spelling the demise of supersonic commercial flight. Thus, economic history shows us that consumers will take advantage of savings and drive markets toward their optimal economies of scale.

As large-scale industries grow, they bring new risks to society. For example, before large-scale chemical plants, no one feared a chemical spill. New problems are dealt with either by reducing risk or by accepting it because the benefits outweigh the risks. For example, although no one wants to be in an elevator when a power outage occurs, most people are willing to live with the risk, being confident that safety regulations are in place and the technology is reliable enough. Sometimes the level of aversion to systemic risk changes over time, and citizens and governments make changes that push industries into different directions. That is the case with Japan and Germany, which used nuclear power for electricity generation for decades before the terrible Fukushima nuclear accident in 2011, at which point both countries decided to phase out use of this technology.

The main question regarding the cloud computing debate is whether we are at the dawn of computing as a utility industry. To answer the question, we have to break it down as follows:

Q: Is the cloud computing phenomenon a scaling process that produces cheaper access to goods and services?

A: Yes; witness the 99-cent songs and smartphone apps and the $5/user/month fee at Salesforce.com or Google Apps.

Q: Are consumers acting on the value proposition regardless of the governance structure?

A: Yes. In October 2011, Google reported 4 million business users and 40 million personal users of Google Apps.[1] The CEO of a Mexico City–based company that provides ICT services for SMEs told me, "Since 2010 most of my clients demand cloud computing solutions as their default option."

Q: Are there systemic risk factors?

A: Yes. There are many examples, some recent, such as the 2011 LinkedIn password leak and Amazon's cloud crash.

Another clue is that the self-regulation/government regulation debate only occurs when a market has reached the large-scale tipping point. We wouldn't consider having this discussion if the scale were small or if we all referred to cloud computing vendors as "small cloud computing companies." I will go even further and say that what we call "consumerization" is another symptom of how newcomers are shaping new markets.

New problems are dealt with either by reducing risk or by accepting it because the benefits outweigh the risks.

So it seems safe to say that regardless of the technical nitty-gritty — such as IaaS, PaaS, SaaS, and the complications to establish their boundaries — what we call cloud computing is the emergence of a new utility industry that, in its early stages, has to cope with the advantages and troubles of becoming a systemic part of the economy. Utility computing will be clearly identified by services that are:[2]

- Shared
- Charged by metered service
- Self-service
- Elastic
- Easily accessible to anyone

This being the case, it is plausible to apply mainstream tools used in other utility industries to propose solutions that can effectively shrink cloud computing's maturation stage.

INFORMATION ASYMMETRY AND CLOUD COMPUTING

Early in the development of economics, social and political scientists proposed analytical frameworks to deal with asymmetry problems and suboptimal economic outcomes such as monopoly formation and market failure. The broad field of economics that deals with agent behavior is called microeconomics, in which economists use an oversimplified analytical starting point of the "perfect market" or "perfect competition" conditions. One of the main assumptions is that for perfect competition to exist, there must be perfect information flows among agents. The perfect information assumption means that all agents know about all products, their characteristics, and their prices at all times. This allows for rational decision-making processes that maximize the utility of economic transactions.

> **It is interesting to note that the value proposition of cloud computing relates deeply to information asymmetry.**

Economists soon realized that markets lack perfect competition conditions, so they dug deeper into information asymmetry problems. Out of this inquiry came agency theory, or the description of the agent-principal problem, which ultimately saw major advances with the use of game theory. This, in turn, explained observed interactions among agents that produced suboptimal results, and it became a highly regarded analytical framework for market development and formation decisions. Here the key concepts that interest us are the problems that arise when information is asymmetric:

- **Conflict of interest** occurs when an agent invested in different interests can act in favor of one or more of its own interests and against the interests of others or the law (e.g., the consultant/auditor dilemma that fueled SOX compliance rules after the Enron/Andersen Consulting debacle).
- **Moral hazard** is a predisposition to take on excessive risk given the possibility to transfer negative outcomes to another party (e.g., local banks reselling subprime mortgages to other financial institutions such that risk was bundled up and transferred elsewhere).
- **Adverse selection** describes a situation where agents make uneconomic decisions given poor information (e.g., a trapeze artist who buys a cheap safety net not knowing that it's substandard).

It is interesting to note that the value proposition of cloud computing relates deeply to information asymmetry. On the one hand, building and operating information and communications technology (ICT) infrastructures and complex information systems is a tough and specialized job that only a few master. On the other, people wish to use their time and resources in activities that don't relate to learning how to build and manage expensive complex systems such as electrical grids or data centers. Technical sophistication has thus created an understanding gap between technicians and business owners, and this means that imperfect information has become inherent to ICT adoption and use.

Indeed, there has been a gap between technical people and businesspeople inside organizations ever since the first ICT departments were organized. Moreover, for many decision makers, the information asymmetry problem is just as bad with an inhouse ICT model as with a cloud computing model. The difference between the models has to do with the leverage the principal (the business owner) has over the agent (the ICT provider) in dealing with the consequences of information asymmetry and the total cost of ownership (TCO) each model represents.

Table 1 presents a generic agent-principal analysis that assesses the differences between inhouse ICT and the cloud computing model. It is important to say that agent-principal analysis is a tool that must be applied on a case-specific basis — it isn't a one-size-fits-all model. Like other tools, such as knives, it can allow us to quickly cut through complex client-provider interactions. So the reader must take into account that the descriptions offered in Table 1 address the current perceived conditions of the cloud computing environment and will vary as business models change, regulations are established, and technologies evolve.

The reader can use Table 1 to evaluate the ever-growing list of cloud computing issues and start thinking of solutions to overcome information asymmetry problems. In this way, agent-principal analysis can be a useful tool for paving cloud computing's road ahead.

CURRENT PRACTICES AND THE ROAD AHEAD

> To solve the problems of today, we must focus on tomorrow.
>
> — Erik Nupponen

A few months ago, Cutter Fellow Lou Mazzucchelli gave a conference address in Mexico City titled "Cloud Computing Promises and Pitfalls," in which he presented his "fine print" analysis of the Amazon Web Services (AWS) contract. He highlighted two specific provisions of the contract:

> **3.1 AWS Security.** Without limiting Section 10 or your obligations under Section 4.2, we will implement reasonable and appropriate measures designed to help you secure Your Content against accidental or unlawful loss, access or disclosure.
>
> **7.2. Security.** We strive to keep Your Content secure, but cannot guarantee that we will be successful at doing so, given the nature of the Internet….

Table 1 — Inhouse vs. Public Cloud Computing Characteristics

Characteristic	Description	Inhouse	Cloud Computing
Scale	Even large enterprises with inhouse infrastructure can be dwarfed by the size that cloud computing data centers and networks have achieved.	Small	Large
Systemic risk potential	Even basic cloud computing functions such as search are crucial for normal, everyday functioning of the digital environment.	Low	High
Cost	TCO is a big debate issue. The implicit assumption is that, for cloud computing to be real, it must demonstrate large price reductions.	Higher	Lower
Access	Although many discuss the conversion rate for current customers that move out of traditional or inhouse ICT solutions into cloud computing solutions, the cloud's lower prices mean the deepening and widening of the ICT industry toward a new customer base. This means that analyses of cloud computing must be centered on new users and not so much on current users.	Low	Very high
Principal	In both cases, the principal is the business owner.	Business owner	Business owner
Principal risk evaluation leverage	Even though information asymmetry exists in both cases, in the inhouse model, the principal has leverage to conduct independent risk assessment, whereas it is powerless in the cloud computing model.	High	Low
Agent	This is the main difference between the models. Access to the benefits of cloud computing does mean changing agents. In some cases, the designated agent (the ICT department) will hire another agent (cloud computing provider) — a more complex situation where information asymmetry grows from the principal's standpoint.	ICT department	Cloud provider
Information asymmetry	Any noncore, technically complex function creates information asymmetry between agent and principal. ICT has always been, information-wise, an asymmetric business.	High	Very high
Moral hazard risk	In the inhouse model, any problem has immediate consequences, as it affects the business, and the ICT department is accountable for problems. Cloud computing interactions are more complex and formally established through contracts that can be opaque (e.g., "To continue, please sign our terms and conditions").	Low	High
Conflict of interest risk	An inhouse ICT department is constrained by the business it works for, so it has few conflict-of-interest problems. On the other hand, cloud providers work and deal with many players with whom they establish deals that are not transparent across their client base. Issues relating to information privacy and trade rules encompass only a few of the conflicts of interest that can occur using the cloud computing model.	Low	High
Adverse selection risk	Cost containment/reduction is an important driver for ICT adoption and use. The inherent ICT information asymmetry problem makes adverse selection a recurrent industry-wide problem, where millions are lost on one side as a result of a few dollars being saved on the other.	High	Very high
Solution/ vendor lock-in	Every time an enterprise finishes implementing a solution, it locks itself into the solution's lifecycle, sunken costs, and vendor environment. Cloud computing solutions promise to lower an organization's exposure to solution/vendor lock-in. This will happen only if technical transparency becomes the rule, and industry-wide standards allow for easy portability.	High	Uncertain, can be low

One does not have to look far to find many other examples that show how information asymmetry is shaping the cloud computing industry. This phenomenon is fueled by the high adoption rates, mainly by newcomers at the base of the pyramid who are taking the opportunity to gain access to services that were formerly out of reach, such as CRM and ERP services. It is further inflamed by big companies that have fired up the contract-making machine to avoid any liability and are free to establish closed standards to effectively lock in their customers. This aggregates systemic risk to the Internet as payrolls, customer information, and a myriad of functions start to depend on a few large players. Considering the accumulating continuity and security problems that cloud computing companies have reported during 2012, and extrapolating these problems to the largest players (such as Facebook, which will soon have a billion registered users), it is easy to conclude that if things keep on their current course, the problems of cloud computing will soon outweigh the benefits. Something as simple as a key payroll processing business going bankrupt would bring millions of people out in the streets to claim their monthly payments, while no one could really do anything about it.

The best-case scenario for cloud computing is one in which an independent authority works with industry and consumer advocacy groups to establish sound regulation as a way of mitigating information asymmetry problems, using good governance and a level playing field as tools for growth. Counter to what many think, good governance practices and consumer confidence are scarce resources that are essential building blocks in the foundations of many advanced economies. Just imagine the slow pace industrialization would have taken if every electric company provided electricity with different voltage and amperage ratings, requiring different types of plugs in each region.

Understanding these three basic principles of economics — moral hazard, conflict of interest, and adverse selection — can keep us pointed in the right direction, away from the technical private-versus-public cloud debate and focused on industry growth and development. Technical people must also acknowledge that politicians and regulators already know agent-principal problems; it will take only a little effort to connect the dots before the government regulation pressure starts to grow. And when it does, building a common understanding around agent-principal problems will foster a productive dialogue among all stakeholders.

ENDNOTES

[1]Crandall, Marc. "Google Apps and Government Requests: Enhanced Transparency." *Official Google Enterprise Blog*, 27 October 2011.

[2]Croll, Alistair. Presentation to *Cloud Connect 2011*, Santa Clara, California, USA, 9 March 2011.

Carlos Viniegra is Head of the Digital Government Unit (Federal CIO) at the Ministry of Public Administration (Mexico) and has taught economics at the Law School of Universidad Panamericana. Previously, he was Director General of Information and Energy Studies at the Ministry of Energy and is the former CIO at both the Ministry of Agriculture, Livestock, Rural Development, Fisheries, and Food and the Ministry of Environment and Natural Resources. Mr. Viniegra holds a BA in economics from Universidad Autonoma Metropolitana and an MBA from Instituto Panamericano (IPADE). He also holds diplomas in enterprise technology from Instituto Tecnológico y de Estudios Superiores de Monterrey (ITESM) as well as parliamentary theory and practice from Universidad Autonoma Metropolitana. He can be reached at carlos@viniegra.org.

Seeing the Future of Cloud Computing Standards

by Duff Bailey and Jeffrey Wu

Standards in the cloud computing space are one of the major drivers affecting ongoing adoption. Attractive financial and service delivery models can only deliver business results if the underlying technology portfolio can interoperate effectively. This article frames cloud computing as the latest generation of extended value chains that require standards to enable providers to interact with their clients.

HISTORICAL CONTEXT

While enterprise concerns about consistent standards with respect to cloud computing are significant, they are not new. Consumers of data and application services have grappled with these issues since the first computerized service bureaus were established in the early 1960s. As that industry evolved from remote job entry and timeshare services to application service providers (ASPs), and then to today's cloud services, market forces have shaped a set of standards and practices that have allowed customers to evaluate and engage service providers while managing their operational and strategic risks.

As the latest phase of an evolving market for data and application services, cloud computing will drive further evolution of standards. We can understand the future shape of standards in the cloud computing world by examining the standards that have already emerged for data and application services. Subsequently, we can review how the market forces that are particular to cloud computing are likely to further shape those standards and practices.

The application and data services market comprises a broad range of activities that may require interaction in many fashions. At the highest level, organizations are connecting value chains with outside providers to enable their businesses to operate. These functions vary from simple tasks to the most complex operations in a global supply chain. Providers must have standard processes to deliver services with the needed quality and at an effective price point. Enabling this deep linkage of process requires a technology fabric that can connect many different parties without costly and brittle custom integration for every link.

Standards must exist to facilitate these interactions. The following are some illustrative examples spanning almost eight decades, during which services gained in complexity and depth (see Figure 1). As the services became more real-time in nature, and integral to the customer's business operation, market forces and legal mandates have driven the development of standards.

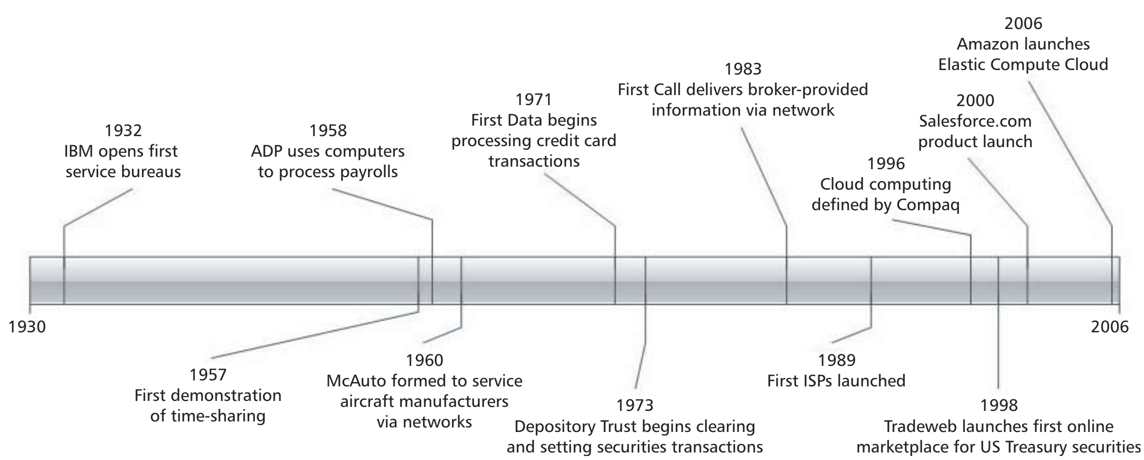

Figure 1 — Nearly eight decades of data and application services.

- **1932.** IBM opened its first service bureaus, providing data processing equipment that included punch cards, sorters, and tabulators.[1] The notion of elastic pricing of shared infrastructure was ground-breaking. It required a new service delivery model and a compelling economic model.
- **1957.** The first demonstration of computer time-sharing occurred on modified IBM 704s and 7090s.[2] This technical feat changed the approach to the design and use of computing machines and spawned new businesses and new business models.
- **1958.** Payroll service bureau Automatic Data Processing, Inc. (ADP) started using punch card machines, check printing machines, and mainframe computers.[3] Standards for submitting data and reporting were key enablers of growth and are ongoing technology concerns at ADP, as is compliance with financial reporting regulations.
- **1960.** McDonnell Douglas Automation (McAuto) formed to provide computer services, including computer-aided design and drafting, for other aircraft manufacturers.[4] Standards were essential for the exchange of data and metadata, including drawings and specifications.
- **1971.** First Data formed to provide credit card processing.[5] The company processed data for both Visa and Mastercard. Standards are critical to serve the many competing entities in this ecosystem, and compliance with regulations is a significant concern.
- **1973.** Depository Trust Company (DTC) formed to provide clearing and settlement of securities.[6] Major financial institutions entrusted DTC with critical client and trading data to facilitate the efficient operation of capital market transactions. Multiple regulatory bodies govern activities in this space.
- **1983.** First Call was founded to provide real-time, broker-sourced information to the financial community.[7] It tied together research, earnings estimates, equity and fixed income ownership information, insider trading information, and corporate news releases across an electronic network.
- **1989.** The first ISPs (PSINet and The World) were launched.[8] These services provided inexpensive Internet access to the mass market and required standards for interoperability of hardware and software.
- **1996.** The term "cloud computing" was defined by Compaq.[9] The concepts outlined then remain consistent with today's notion of cloud computing.
- **1998.** Tradeweb launched the first multidealer online marketplace for US Treasuries.[10]
- **2000.** Salesforce.com launched its first product.[11] At the time, SaaS providers were called application service providers. There is a thriving market of add-on products for Salesforce.com, made possible by a published API.
- **2006.** Amazon launched Elastic Compute Cloud (EC2).[12] Amazon made its celebrated platform available for use by enterprises of all sizes.

THE CURRENT STATE OF STANDARDS IN APPLICATION AND DATA SERVICES

As application and data services have become more strategic to the customer's business, the market has responded by holding service providers to higher standards of accountability in all dimensions. There is a need to ensure that services meet customers' requirements for security and reliability. Providers must also allow customers the flexibility needed to employ the service offerings that meet their needs best, at a competitive price, and in turn meet the business requirements of *their* customers. In industries where regulations exist, standards for compliance must also be met.

Data Security

For data security, there are well-defined standards that are enforced through contracts and audits on a wide range of application and data services providers. The Payment Card Industry (PCI) Data Security Standard, maintained by the PCI Security Standards Council,[13] governs the operations of every user of payment card data. Both affiliated and independent card issuers, such as AmEx, enforce this standard on their merchant banks, processors, and vendors. The standard is updated regularly to reflect new technologies and threats. It is accepted by proxy by the financial community as a whole — and applied by reference to most vendors that handle financial transactions. While the PCI standard is self-enforced, card issuers are required to engage in active due diligence to ensure that their business partners apply the PCI standard as a control objective and in turn provide an independent, SAS 70 audit as verification.

Personally identifiable health data is governed by US federal law under HIPAA, which mandates that covered entities implement written privacy procedures and employ the processes needed to ensure compliance. It also sets forth basic security practices, such as end-to-end encryption, role-based authorization, and strong user authentication for all systems that handle personal, identifiable health data. Covered entities must ensure that their ASPs comply with these standards as well.[14]

Reliability

Fully 74% of respondents to *InformationWeek's* 2012 State of the Data Center survey ranked reliability/availability as the most important requirement for application infrastructure.[15] Even so, a consistent standard for reliability that is both relevant to the business and practical to measure has not emerged. While data services contracts typically have defined hours of service and availability, it remains a challenge to contractually define reliability in meaningful terms, especially for complex services involving interconnected business processes such as the processing of securities trades. While it is easy to determine when a clear-cut outage has occurred, other reliability issues, such as processing errors or degraded service levels, are harder to define — especially when the process entails interaction with processes owned by the customer or other vendors. In addition, the data services providers have been largely successful at avoiding contractual penalties for the entire risk of a failed operation. Financial recourse is typically limited to the monthly cost of the service itself.

In spite of weak contractual terms, data services customers often find that they have more ability to drive service providers to deliver reliable services (in their terms) than they do with their own IT department. First, even limited financial recourse is better than none. Second, reliability issues have a significant impact on the business reputation of the service provider — who is thus strongly incentivized to ensure that such issues happen rarely if ever. Finally, when a service provider has a problem, it affects all of its customers; thus, the reputational damage each customer experiences is shared by any competitors that also use that provider. Our experience with vendors and clients is that vendors provide dramatically better levels of service than many internal IT organizations.

With cloud computing, a critical component of reliability is the availability of capacity on demand. To sustain guarantees of that sort, a service provider must either have large amounts of unused capacity or a large number of applications that can be shut down or operated in degraded mode when demand for critical applications soars. The data services field has many analogs to this concept. Securities processing, in particular, faces tight deadlines and extreme fluctuations in volume. Transaction-based pricing, with premiums charged when demand peaks, is one way to ensure that service providers make capacity available when needed.

Business Continuity

Business technology professionals rank business continuity as the most important organizational driver for using virtualization.[16] Standards for business continuity can be clearly defined in terms of elapsed time between a service outage at one location and the restoration of a given level of service at the same or an alternate location. Vendor contracts for ASPs typically specify these recovery windows, which we have found to be at parity with or better than most enterprise timelines.

In spite of weak contractual terms, data services customers often find that they have more ability to drive service providers to deliver reliable services (in their terms) than they do with their own IT department.

Application Integrity

Section 404 of the US Sarbanes-Oxley Act (SOX) defines legal compliance assurance for companies.[17] The independent auditors for a company must attest to the internal controls on financial reporting that management has implemented. As a result, controls over IT must be in place and verified. The implication is that any data that feeds into financial reporting systems, which includes numerous transactional systems, must have sufficient controls in place. Cloud-based systems are likewise subject to this scrutiny.

Auditors will demand documented evidence of controls in practice and also the results of tests to ensure the controls are functioning. In our experience, internal auditors typically review controls and results prior to independent auditing.

Interoperability

To date, there has been little demand for an overall standard for high-level interoperability, because customer decision making is primarily driven by whether a particular service will work effectively with the specific applications and services the customer currently uses — which often includes proprietary or customized applications with limited interoperability in the first place. An RFP for application or data services typically states the customer's specific interoperability requirements, which the vendor must demonstrate during product evaluation. Most providers have "fit for use" adapters that allow their services to work with the products and services commonly found in their customers' technology stack, which they will customize, for a fee, to meet client-specific needs.

Providers gain a market advantage with new customers when they make their services more interoperable, even if that reduces their leverage over existing customers by making it easier for them to add in competitors' services. Examples of this effect can be seen among CRM and accounting systems, where providers enable customers to easily upload data and transactions to foster the progressive adoption of their platforms.

Cloud services are notably more global than conventional application services, and this will result in increased demand for regulatory oversight.

Portability

Rather than seek a portability standard, data services customers have included the cost of migrating to a new application services platform in the business case for adopting a new provider. That practice puts pressure on service providers to minimize the cost of migrating new customers — whether by building inward portability features into their product set or by absorbing some of the new customer's migration cost. This approach leaves customers with the risk that their data may not be readily recovered from a provider at a later date.

While customers usually retain ownership of the data they provide and any data that they purchase to enrich it, there may be limitations as to the format in which the data is provided back to them as well as the amount of provider-owned metadata that is returned. More sophisticated organizations consider the "exit" cost of leaving a provider and ensure their contracts allow them to recover their data in a timely fashion. Our experience with migrations is that data and technology rarely cause business issues; problems arise from process changes and functional gaps between systems.

Pricing

Pricing is the least standardized and most opaque aspect of the data services market, and that status reflects the interests of both customers and vendors. For customers, a great attraction of data and application services is the ability to control their IT spending so that it is either a flat, predictable expense or one that varies with business activity and, ideally, revenue. Providers, of course, don't want to be measured against competitors on cost alone and thus seek to differentiate their service offerings in ways that limit direct cost comparisons. Both factors favor pricing schemes that reflect transaction volumes, number of users, or other customer business metrics for specific bundles of services. This pricing opacity does have benefits for savvy customers. In competitive situations, customers can extract significant discounts, concessions, and future price protection. We have seen discounting of over 90% in highly competitive deals.

CLOUD COMPUTING MARKET FORCES AND THEIR IMPACT ON STANDARDS

While cloud computing can be viewed as a logical outgrowth of data and application services, it nevertheless has some unique dynamics that will further affect the balance between providers and purchasers of cloud services.

Cloud services are notably more global than conventional application services, and this will result in increased demand for regulatory oversight by governments, supranational bodies, or self-imposed industry bodies. In addition to the security standards of the PCI Security Standards Council and mandates for privacy in HIPAA, additional bodies driving standards adoption include the National Institute of Standards and Technology (NIST), Desktop Management Taskforce (DMTF), Storage Networking Industry Association (SNIA), and the Cloud Security Alliance (CSA). Technical standards exist throughout the technology stack and continue to evolve in areas such as storage, management, security, and virtual machines.

There does not appear to be broad-based support for more standards in the customer community among organizations that have already deployed clouds; only 31% cited lack of standards as an issue.[18] Even so, there are strong economic incentives for cloud providers to adopt such standards. The economic model for cloud services is particularly dependent on economies of scale across many business and geographic markets. Providers that achieve massive scale will attain much lower costs and higher on-demand capacity, and the increased profit margins will justify increased investment to meet standards required by only a fraction of their customers.

The market pressure to adhere to standards that enable broader market penetration will be particularly strong for the IaaS providers, whose largest customers are likely to be the SaaS providers. The SaaS providers can be expected to push the IaaS providers to meet the highest standards demanded by their customers — particularly for data security and reliability. In doing so, the IaaS providers will find that it is most cost-effective for their entire infrastructure to comply with the highest commonly applicable standards for security and operational integrity, including PCI, HIPAA, COBIT, and the ISO 27001/27002 Information Security Management Standard.[19]

While reliability will remain hard to define precisely, cloud providers will be strongly motivated to maintain an extremely strong reputation for reliability, in the same way that application and data services providers do currently. In this respect, the fine-grained nature of cloud services provides some special challenges. A customer that obtains data storage from one provider, computing power from a second provider, and specialized application services from others can encounter unpredicted adverse interactions between cloud services — especially when one or more services experience peak usage or undergo failover. To shield themselves from the risk of reputational damage, cloud providers may require some kind of certification of the other services and applications in a customer's technology stack to ensure reliable interoperability of a customer's applications under all conditions, including peak usage, system stress, and the failure of individual components.

While concern for reliability may place some limits on providers' willingness to support full interoperability, this will be balanced by the pressure on cloud providers to maximize their economies of scale by supporting broad adoption. Hence, certification of applications is likely to be limited to screening out those that clearly threaten overall system stability. Customers will continue to bear the onus for verifying interoperability of their configuration with respect to supporting their overall business requirements. In addition, market forces will still drive acceptance of and accommodation for dominant technologies. Currently, 53% of virtualization stacks rely on VMWare vCenter or vCloud Director, with all other competitors having a 10% share or less.[20] That fact pressures providers to ensure that the dominant products work on their platforms.

Certification requirements by IaaS providers would also drive adoption of standards upstream to SaaS and application services as well as to customers' proprietary software. Once certified, applications could be defined as services and fully defined in a service directory, with interoperability defined in terms of which of these services' interfaces were supported, and certified, by other services.

Such standardization would drive portability by significantly reducing the cost of switching providers for interchangeable services. Vendors in the cloud market appear to be following the example of conventional data service providers by defining their role as custodians rather than owners of the customers' data, giving customers the ability to recover their data as needed. Service users have confidence that their vendors do not control them; only 9% of companies that have deployed clouds are concerned about vendor lock-in.[21]

For those services that become commoditized, a reputation for reliability and responsiveness to customer needs will become the greatest differentiating factors between providers. In fact, the underlying technologies are of little concern to business consumers, as 99% of the time the business achieves what is required without violating its internal IT standards. An *InformationWeek* survey found that 92% of business users only occasionally, rarely, or never find conflict between cloud solutions and their organization's application and hardware philosophies. When conflict arises, only 17% of the time are IT standards more important than business functionality.[22]

Vendors have significant motivations to standardize their offerings and compete on service, innovation, and other realizable customer benefits. Providing customized service to every cloud client would destroy the economic model of the cloud. Service bureau providers realized this early on. They homogenized offerings and developed well-defined service models while still driving innovation and improving what they delivered to the market.

Standardization affects service levels in many ways. For example, cloud vendors have a strong incentive to standardize on the highest level of security they can reasonably deploy, since the network infrastructure is only as strong as its most vulnerable parts. Standard interfaces foster adoption and lower operating costs for both customers and vendors.

Pricing for high-value services will remain tied to business activity metrics, but those for simpler services such as IaaS will be commoditized. As these services become commoditized, market pressure will force greater transparency and simplicity as to how these services are priced. To offset commoditization in areas such as IaaS, providers will work hard to differentiate their offerings from competitors in terms of reliability as well as interoperability and portability.

CONCLUSION

As long as the market for cloud services remains competitive and open, competitive pressures and market forces will continue to support the organic evolution of the de facto standards for service security, quality, and pricing. These evolving standards will enable service customers to make informed and appropriate choices between service providers and service offerings. Market-driven standards will also enable service providers to efficiently support a broad range of customers and achieve the economies of scale needed to economically support the cloud model. The evolution of these standards will increase the interoperability of different services as well as portability across providers and platforms. As this happens, service providers will increasingly need to compete on their reputation for reliability as well as the value that they add to their customers' business model.

ENDNOTES

[1]Sipples, Timothy. "Did the US Government Kill Cloud Computing Over 50 Years Ago?" *The Mainframe Blog*, 5 September 2011 (http://mainframe.typepad.com/blog/2011/09/did-the-us-government-kill-cloud-computing-over-50-years-ago.html).

[2]Woo, Elaine. "John McCarthy Dies at 84; The Father of Artificial Intelligence." *Los Angeles Times*, 28 October 2011.

[3]"ADP History — Stable, Reliable and Profitable." Automatic Data Processing, Inc. (www.adp.com/careers/uscareers/who-is-adp/history.aspx).

[4]Weisberg, David E. "Siemens PLM Software (Unigraphics)." In *The Engineering Design Revolution*. David E. Weisberg, 2008 (www.cadhistory.net/19%20Siemens%20PLM%20Software.pdf).

[5]"First Data Corporation History." Funding Universe (www.fundinguniverse.com/company-histories/first-data-corporation-history).

[6]"SEC Rule Filings." The Depository Trust & Clearing Corporation, 2012 (www.dtcc.com/legal/rule_filings).

[7]"Jeffrey P. Parker: Venture Partner." Grand Banks Capital, 2012 (www.grandbankscapital.com/our_team/bio_detail.cfm?id=129).

[8]"Who Was the First Internet Service Provider." *The Internet Business Man*, 8 August 2009 (www.theinternetbusinessman.com/uncategorized/who-was-the-first-internet-service-provider).

[9]Regalado, Antonio. "Who Coined 'Cloud Computing'?" 31 October 2011 (www.technologyreview.com/news/425970/who-coined-cloud-computing).

[10]"The Building of a Global Leader." Tradeweb, 2012 (www.tradeweb.com/about-us/history).

[11]"Company Milestones." Salesforce.com (www.salesforce.com/company/milestones).

[12]MacManus, Richard. "Amazon Launches Elastic Compute Cloud." *ReadWriteWeb,* 23 August 2006 (www.readwriteweb.com/archives/amazon_ec2.php).

[13]PCI Security Standards Council (www.pcisecuritystandards.org/index.php).

[14]"Summary of the HIPAA Security Rule." US Department of Health and Human Services, 2009 (www.hhs.gov/ocr/privacy/hipaa/understanding/srsummary.html).

[15]Marko, Kurt. "Research: 2012 State of the Data Center." *InformationWeek Reports*, June 2012.

[16]Marko. See 15.

[17]Marks, Norman. *Sarbanes-Oxley Section 404: A Guide for Management by Internal Controls Practitioners.* 2nd edition. The Institute of Internal Auditors, January 2008 (www.theiia.org/download.cfm?file=31866).

[18]Fratto, Mike. "Private Clouds: Vision vs. Reality." *InformationWeek Reports*, June 2012.

[19]"ISO/IEC 27002:2005 Information technology — Security techniques — Code of practice for information security management." 2012 IsecT Ltd., 2012 (www.iso27001security.com/html/27002.html).

[20]Marko. See 15.

[21]Fratto. See 18.

[22]Marko. See 15.

Duff Bailey, PMP, is Practice Director, Master Data Management and Enterprise Integration, of Actionable Strategies, a boutique consulting firm with three focus areas: strategy and organization, business process excellence, and enabling technologies. He is a veteran technology manager who has managed and advised clients in investment banking, media, and telecommunications on the acquisition and integration of a wide range of data and application services, as well as overall strategies for master data management and application integration. He can be reached at dbailey@actionablestrategies.com.

Jeffrey Wu is the CEO of Actionable Strategies. He works with clients on cloud strategies as well as developing new services delivered via the cloud, and he is often invited to speak at industry events. Mr. Wu's experience with cloud computing began in the early 1990s when he built and delivered new products globally over the networks at Thomson Financial. As the CEO of Claritas Consulting, he worked with a number of high-growth firms to launch as application service providers and was a frequent speaker at industry events. As a CTO and VP of Product Development for ADP, he delivered a range of services across the Internet. He graduated from the Wharton School of Finance in 1985 with a double major in finance and management. He can be reached at jwu@actionablestrategies.com.

A Standard Isn't a Document — It's a Process

by David R. Bernstein

STANDARDS IN TECHNOLOGY: THAT'S THE OLD WAY TO MAKE PROGRESS, RIGHT?

In the last several hundred years of humankind, technology advancement has played an important role in the overall advancement of civilization. Electricity, communications, and computing are all 100% technology-based innovations. And in these fields, there are implicit requirements for standardization. Electric devices have to be able to plug in and use the type of power that is produced by the power plant. Communications involves multiple parties exchanging information using common protocols and formats, even over several service providers. And, of course, computing is filled with standards from the binary representation scheme of information all the way up to modern programming languages.

But now with the Internet, we see "new ways of doing things" exploding into common practice, along with open source software, which embodies the "how-to code" for that as well. Successful products and services take off overnight, with no apparent standards organization in sight. This leads one to wonder, are we "beyond" standards as a way to make progress? Are the people who criticize the formal standards approach as "too old and too slow" correct?

This article argues that the modern standards-making process itself can create the proper mix of technology portability to ensure freedom of choice for consumers and free market innovation and reward for entrepreneurs. We will see that standards, "common practice," and open source are complementary dimensions of technology advancement. It turns out that the best benefits to manufacturers, consumers, and the overall ecosystem around a particular technology are realized when we leverage all these dimensions.

What is true is that the old way of making standards — in isolation, without real-world trials, and without a process involving evolution of the standards — is just that: the old way of making standards. This method indeed does not work anymore. However, standards-making organizations have long jettisoned this historical approach, and for some time they have been creating standards in a contemporary and modern way, incorporating multiple dimensions of evolution. In other words, a standard isn't a document — it's a process.

A GREAT EXAMPLE OF A STANDARDS PROCESS: ETHERNET

To fully illustrate this process, let's take a look at one of the world's most successful standards-based technologies — Ethernet. Ethernet was originally developed by Xerox in the mid-1970s. It became immediately apparent that a technology whose essential purpose was interoperability needed to have a wider support base than one company. Xerox convinced Digital Equipment Corporation and Intel to work with them to promote a "standard" Ethernet, and together they published the so-called "DIX" specification for Ethernet. Although this initial specification was tremendously important, it was not comprehensive enough to cover the rapidly evolving marketplace around this new technology.

First of all, Ethernet allowed multiple physical implementations. At the time of the DIX specification, coaxial cables with cable "taps" were the only practical implementation. Soon, hardware vendors figured out how to drive twisted-pair Ethernet, leading to a variety of incompatible physical implementations, both in form factor as well as in rate capability (network speed). Ethernet also allowed for flexibility in higher-level network protocols. Several variations emerged, including token ring and token bus, which were similar but incompatible. Finally, while the specification called for every Ethernet node to have a unique identifier, known as the media access control (MAC) address, the specification did not answer the obvious question regarding how the world was to ensure that no two devices were manufactured with the same MAC address.

On the one hand, the late 1970s saw great success in Ethernet, with many breakthrough implementations and many alliances around particular "interoperable" implementations. However, these implementation factions began to compete with each other, complicating the overall marketplace and slowing down overall market uptake.

In 1980, the IEEE started the IEEE 802 project to standardize local area networks, and the "DIX Group" submitted their specification as a candidate for the specification. This allowed for larger industry participation in the discussion in an open forum, enabled feedback from real-world implementation experience to be included as practical input to the standard, and also provided an elegant solution to the MAC address numbering governance problem.

> **If the standards process hadn't yielded several important technology ecosystems, cloud computing would never have emerged.**

The IEEE 802 project (together with the IEEE MAC Address Numbering Authority) is still, over 30 years later, one of the most active and prolific projects in the standards world. Ethernet is still a broad and varied technology, spanning the wired and wireless worlds, and most recently the physical and virtual worlds.

The transition to the wireless world was a perfect example of "standard as a process." As the technology of wireless LAN (WLAN) was being developed, the IEEE formed the IEEE 802.11 project to address the emerging technology needs. Real-world experience with spectrum use, error-correcting protocols, and other technical details led to the evolution of the IEEE 802.11 project through IEEE 802.11a, b, c, g, and n amendments to the standard. Similarly, real-world experiences in the robustness of security schemes have led to the improvement of WLAN security schemes as specified in 802.11 as WEP, WPA, WPA2, and so on.

The transition to the virtual world is an example of "standard as process" that's still very much in flux. The original physical Ethernet adapters and switches are now to be found in virtual forms in almost every computing virtualization system, bringing new challenges to the technology. For example, the original notion of a MAC address representing a unique piece of hardware no longer applies. Security and identity products, which made that assumption, have now had the rug pulled out from under them. Software systems that "manufacture" random MAC addresses in the virtual environment are sure to cause interesting problems, as they are working outside of the traditional IEEE MAC Address Numbering Authority. This is a current and vibrant subject of discussion in this new world of virtualization and cloud computing.

A BRAVE NEW WORLD: CLOUD COMPUTING

Speaking of cloud computing, one can hardly listen to the radio or watch TV without hearing about this important new trend. In fact, if you are listening to Internet radio like Pandora, Shoutcast, Jango, Live365, or one of many other such "stations," you are using cloud computing. Likewise, if you are watching television using Netflix, Roku, Hulu, Xfinity, Joost, or one of the many other such "channels," you are also using cloud computing.

Cloud computing has swept the planet as a major new computing paradigm, making possible not only the aforementioned low-cost and ubiquitous radio and television, but also the features we know of best as cloud services, including search, email, and storage, as well as videoconferencing, voice chat (which your parents would have called "the telephone"), limitless free and open libraries, encyclopedias, and museums and virtual communities and games. Somehow, clouds have been able to deliver a computing platform capability for these services at such efficiency that many of these services can be delivered to consumers for free. Not only are the services free, but they are delivered with apparently infinite computing, storage, and bandwidth capabilities, and added to this, seemingly global coverage.

It is interesting to examine the dynamics of cloud computing and to understand how standards have played a key role in the development of this technology. What we will see is that if the standards process hadn't yielded several important technology ecosystems, cloud computing would never have emerged. We will also see that in the process, cloud computing has also created a tremendous new gap in standards and interoperability as a side effect.

To understand clouds in this manner, we must first understand how a cloud is constructed and what the "secret sauce" is that makes this type of data center fundamentally different. First off, physically, clouds are constructed out of servers, networking, and storage — but not just any old servers, networking, and storage. Cloud computing is based on the notion that automation and virtualization are applied to the physical data center in such a way as to let software figure out where to place application servers or run application code without human interaction, in a totally automated and dynamic manner.

This level of virtualization and automation extends to the way cloud computing approaches failure response as well. In a large cloud — for example, with a million servers — machines, network devices, or storage drives will fail regularly and often, albeit still randomly. Building hardware to prevent failures is cost prohibitive and also really impossible; because of the scale of the cloud, the statistics of large numbers will always prevail, and failures will appear. Therefore, clouds use the lowest-cost devices, and typically everything is made redundant in software or network configuration. Because the application servers and application code have no connection to the actual underlying physical environment, failures can be accommodated simply by no longer counting on that resource and placing servers or running code elsewhere. Application software is architected to be redundant and restartable, and the failed devices are simply powered down and ignored until they are replaced or removed.

The philosophy, then, is lowest cost per device and utilizing the most practical and simplest of technology. It is analogous to low-cost airlines purchasing all of one type of airplane to maximize efficiency. They will usually select a modern, very successful airplane model. While it may be "no frills," it will not be inferior; the cost reductions to be had will not come from lower cost or quality in the planes themselves. Rather the lower costs come from standardizing on a single model, which means that an airline can:

- Maintain parts inventories for only one type of plane
- Purchase gates and gate equipment for only one type of plane
- Train the workforce (technicians, flight attendants, and pilots) for only one type of plane, thus yielding lower staff costs and greater flexibility in work assignments

At one airline, eliminating seats in reservation systems saved millions of dollars. It's as though they virtualized the airplane — you have a seat, you just don't know where it will be. They know that each plane carries the same fixed number of people, so there are no issues with swapping out whole planes for any given route. Mechanics, pilots, flight attendants, gates, catering containers, and hangar slots are all swappable and interchangeable! This is the same philosophy for cloud computing — driving out costs through standardization, on top of which virtualization and automation are easier to apply.

Thus, in clouds, one uses standards-based servers, switches, and storage drives in a very basic and straightforward architecture, and the cloud OS software above it adds the automation and virtualization. If it weren't for the underlying standards in servers (bus, memory, instruction sets, sockets, etc.), network (Ethernet, TCP/IP, etc.), and storage (SAS, SATA, etc.), none of this would be possible.

So out of these smaller devices, we now have a cloud, which functions like a "datacenter-size networked computer."[1] Just as with any other computer, users are most familiar with the user-facing interfaces "on top." In this case, the on-the-top interfaces are the ones for automatically placing application servers or running code. Or you might think of these interfaces as the "sockets" for the network access. Standardized user-facing interfaces allow for application server or code portability across clouds.

The philosophy, then, is lowest cost per device and utilizing the most practical and simplest of technology.

Also just like any other computer, the cloud has system internal interfaces for virtualization and automation and for making one cloud out of multiple data centers. These are somewhat analogous to device driver interfaces in a typical computer. In the network case, it is rather precisely analogous to the routing protocols or Domain Name System (DNS) and the like. Standardized system interfaces allow for substitution of internal components locally and remotely, creating the ability to distribute or federate a cloud in a pluggable way. Routing protocols did exactly this for the Internet.

While most people think of "cloud standards" as the user-facing interface, when we think of the Internet, we think of both the user-facing interface and the system-facing interface. In fact, the creation of standards around the system-facing interface was the crucial development that allowed the Internet's explosive growth.

FROM ETHERNET TO CLOUDS: DIDN'T WE SKIP THE INTERNET?

Now we are prepared to really understand something interesting, this lesson learned called "the Internet." Indeed as we have seen above, the Internet really leveraged two key areas of standardization: the user-facing interface of the IP socket, on top of which HTTP and all other Internet user programs run, and the system-facing interface on which the network itself runs. The latter includes things like DNS, Autonomous System (AS), Internet Control Message Protocol (ICMP), Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and so on. The average Internet user never hears about these, but they are the key protocols that make the Internet interoperable.

> **While the interplay of new ideas may produce exciting new technologies that can compete and cross-pollinate, markets grow on the basis of confidence and consensus.**

How did these protocols get developed into standards? They were written in code, and in documents, and then placed into test beds, roughly in parallel. Pioneering university research groups and networking companies created these protocols and made equipment that embodied them. They formed the Internet Engineering Task Force (IETF), where they wrote up the protocol definitions in a collaborative request for comments (RFCs) form. Even calling the documents "RFCs" instead of "draft standards" showed how important the process of developing the specification was — the process of gaining consensus was more important than the specification itself! At the same time, the National Science Foundation (NSF), other research entities, and even the vendors launched test beds to try out interoperability. One example of this on the Internet was the Commercial Internet Exchange (CIX), in which companies and research entities connected for one of the first times to form the commercial Internet as we know it today.

The point is, the multiple dimensions of the standards process — the documents, the comments, the test bed, the equipment, the standards organization itself as a collaboration vehicle — were all part of the "standardization process" for the Internet. I think we can all agree that this worked out pretty well.

FROM ETHERNET TO INTERNET TO CLOUDS: AREN'T CLOUDS THE SAME THEN?

Cloud computing brings some of these same challenges and some new challenges. Standards processes are adapting because the industry has learned from previous efforts like those for Ethernet and the Internet:

- Ethernet had the standards specifications part as the leading part of the process. The multivendor test bed component took some time to add to the mix. Implementations were not open source.
- The Internet utilized standards specifications while also having a vibrant, open, multivendor test bed. Although there were some open source user-facing components, the system-facing parts of the Internet (routing protocols) were not open source, nor were the core operating systems (of the servers and routers). Server operating systems became open source, but even today open source routing software is relatively new.
- Cloud computing will need to utilize standards specifications, multivendor test beds, *and* open source. It will require a standards-based approach to source user-facing components and the system-facing parts. In fact, this is precisely how cloud computing standards are evolving today, in all of these dimensions.

CLOUD COMPUTING WILL BE SERVED WELL BY A CONTEMPORARY STANDARDS PROCESS

When a visionary market finds itself stalled in the early-adopter stage and unable to move into the mainstream, the problem is usually a lack of consumer confidence. Cloud computing appears to be caught in this situation today with the two big problems being consumer fear of lock-in and cyber security unknowns. While some businesses are used to single-source solutions for some IT elements, large telecommunications companies, government agencies, and many large enterprises have strict "standards based" or "multiple sourcing" procurement policies. Until there are cloud computing platforms from several different vendors, all compatible with certain standards, these organizations will have issues with rapid cloud adoption. In addition, industries that require security and compliance standards for credit card processing, healthcare information, or government/defense information handling will want to see that cloud platforms comply with these standards before they acquire them. While the interplay of new ideas may produce exciting new technologies that can compete and cross-pollinate, markets grow on the basis of confidence and consensus. Standards give the pragmatists that control entry into the mainstream a crystal ball that allows them a glimpse of the future.

Standards making is a dynamic process that takes into account the needs of all stakeholders over extended periods of time. It brings them to the table, creates a forum for negotiating tricky issues such as intellectual property compensation, and, perhaps most important, becomes a vehicle to move the technology forward over time. A technology foundation is established on consensus issues so that competition can focus on performance and fulfilling customer needs.

Although it has its own characteristics and eccentricities, cloud computing is not so different when observed at a higher level of abstraction in what might be called the "technology development" domain. In a technology as broad and complex as cloud, solutions cannot be predetermined. The process becomes the solution. Of course, the most contemporary of standards organizations — the ones that are successful and active in cloud standards — already know this.

ENDNOTE

[1]Barroso, Luiz, and Urs Hölzle. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool, 2009 (http://books.google.com/books/about/The_Datacenter_as_a_Computer.html?id=Y3hhzdSSK58C).

David Bernstein is founder and Working Group Chairman of the IEEE P2300 Cloud Standards, in addition to serving the European Commission on the Expert Group/Roadmap Editorial Board on the SIENA/FP-7 Standards and Interoperability for eInfrastructure Implementation Initiative. He is also Managing Director of Cloud Strategy Partners, LLC, serving many clients in communications, networking, and security. A representative client list includes Cloudscaling, the US Department of Commerce/NIST, Barnes & Noble, Huawei Technologies, Aria Systems, Catbird Networks, Xtium, Finisar, Obopay, Veloxum, and Intelliden, Inc. Prior to joining Cloud Strategy Partners, Mr. Bernstein was at Cisco as VP/GM of the Network Applications Infrastructure Business Unit; he then started cloud computing for Cisco in the Office of the CTO. His previous experience includes executive positions at AT&T, Siebel Systems, Pluris, Intertrust, and Santa Cruz Operation. Mr. Bernstein holds nearly a dozen patents in software and communications and speaks regularly at IEEE, ACM, IARIA, and industry conferences. He has been a key author/contributor to many other industry standards such as OpenSOA.org, OASIS SCA, WS-I, JCP/J2EE, and IEEE POSIX. Mr. Bernstein holds degrees in physics and mathematics from the University of California, where he was awarded the UC Regents Scholar designation for his Office of Naval Research work. He stays active with webcasts, papers, blog posts, conference schedules, board activities, and success stories at www.CloudStrategyPartners.com. He can be reached at david@CloudStrategyPartners.com.

Cloud Standards? It's the Users, Stupid!

by Claude R. Baudoin

Many organizations have developed models to explain the introduction of a new technology in the marketplace. Few of these models, if any, identify the point at which standards become necessary and get developed, or how they influence — positively or negatively — the subsequent development of the products they affect. In the case of cloud computing, it is particularly worthwhile to fill this gap, not only because of the importance of the questions posed by Guest Editor Mitchell Ummel in his call for papers, but also because cloud computing has been developing very quickly compared to previous computing models. While the new model brings about a new set of problems, there are actually some promising standardization activities underway, including a user-driven advocacy effort worth mentioning.

FROM ZERO TO 60 IN JUST A COUPLE OF YEARS

It has become quite a cliché to talk about the "exponentially accelerating pace of technology." It is certainly true that the introduction of mainframes, minicomputers, personal computers, smartphones, and tablet computers, successively, has happened in progressively shorter time frames. As a result, adopters are faced with a large array of potential choices within shorter and shorter periods of time.

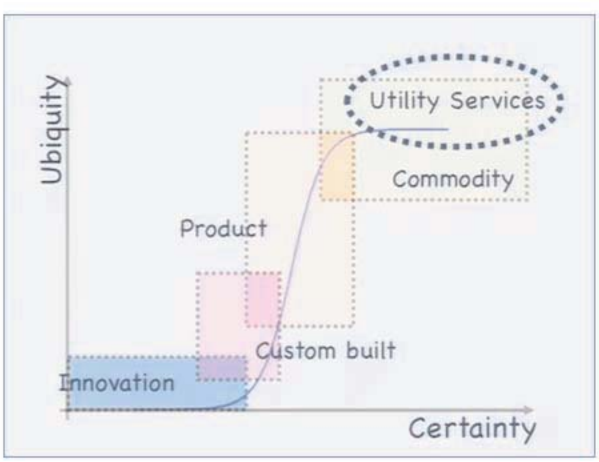

Figure 1 — Technology Diffusion Model. (Source: Wardley.)

The generally accepted model of innovation diffusion is an S-shaped curve, which dates back to the early 1960s work of Everett Rogers[1] and was further developed at Carnegie-Mellon University in the 1980s. Coincidentally, Rogers's work appeared at about the same time as the first mentions of a utility model for computing. Simon Wardley updated this model during his OSCON'10 keynote talk (see Figure 1).[2]

While it is well understood that an innovative company, with good thinkers and access to venture capital, may bring a product to market in a relatively short time span, the emergence of a new *model* of computing would seem to require more time, because it demands a change in the mindset of the users. So why did cloud computing come out of relative obscurity just a few years ago (in 2006, many of us were still preoccupied with "grid computing," which we would probably call nowadays a private IaaS cloud) to being what everyone was talking about all the time as soon as early 2009?

As the sidebar "A Cloud Computing Timeline" implies, a key reason is that cloud computing is not really new. It did not start at the bottom left of the S-curve, but somewhere in the middle, because we have a 45-year history of flirting with some form of shared hosted resources. So when the new phrase "cloud computing" emerged, we were in fact already in the middle of the curve, which is its steepest part.

With older shifts in technology, CIOs and business managers had the time to do their research, with or in spite of vendors and consultants, in order to decide on an adoption roadmap. In the case of cloud computing, IT in particular is being forced to adopt the model much faster than they probably wanted. If they don't, they are making themselves irrelevant in a world where a project manager can share files with a contractor using Google Docs and a sales manager can set up a CRM application in a few days by charging a few hundred dollars a month to his business credit card. The technology has gotten ready for prime time faster than the guidelines needed to use it wisely.

CLOUD STANDARDS: THE USERS' VIEW

There is an irony to some of the arguments that have been used to sell the cloud model. The most frequent argument goes like this: instead of investing in hardware and software that you will have to keep for five years in order to fully depreciate, you can rent the capability you need, with elasticity, paying only for what you use. If you change your business model or your enterprise architecture later, such that you no longer need this capability, you can cancel and stop paying. You're not "locked in" by your hardware and software investment.

But here's the rub: you may be trading in that form of dependency for another one. Now you may get locked into a specific cloud because of proprietary formats, protocols, and APIs. Ideally, as a customer, you would want the following areas to be standardized:

- The definition of a "virtual machine," so that you can decide to reassign a workload from cloud A to cloud B (supplied by different companies), knowing that the processes you ran on A's virtual machines are going to run exactly the same way on B's machines. In spite of some notable advances such as the Open Virtualization Format (OVF), there are differences in virtual disk formats that provide less than full migration possibilities.
- The ability to federate services from multiple providers, as well as federate a private cloud with a public cloud, in particular to be able to handle peak demand.

The current tension between these technical requirements and the state of the industry is twofold:

1. There are indeed emerging standards in these areas, but they form an alphabet soup of complex, overlapping specifications from too many organizations (see next section).
2. The end users are not confident (and who could blame them?) that the vendors are truly motivated to create standards that benefit them.

And this is just the technical view of what standards customers would benefit from. In addition, users need a number of other common specifications.

First, they have to be able to compare offerings from different suppliers. This requires a standard vocabulary, with standardized semantics for the terms and units that describe various infrastructure capabilities (processing resources, storage resources, transaction rates, network sockets, etc.).

A CLOUD COMPUTING TIMELINE

1961	John McCarthy forecasts computing as a utility
1971	IBM introduces the Time Sharing Option (TSO)
1996	The term "application service providers" appears
1999	Salesforce.com launch
2000	"Utility computing" buzz
2002	"Organic computing" and "on-demand computing"
	Amazon Web Services (AWS) launch
2005	Open Grid Services Architecture v1.0 published
2006	Amazon Elastic Compute Cloud (EC2) launch
2009	Oracle flip-flop on the meaning of cloud computing
	Google Apps launch
	Eruption of cloud computing articles and papers
2010	Microsoft's Windows Azure launch
2011	Cloud Standards Customer Council created

Second, they need a methodology to understand which applications and services are "cloud-ready" and, if they are not, what it may take to make them so. This will clearly not come from the cloud vendors, who will never say, "This application works better in your data center than it would in the cloud," or "The integration effort to make this work in the cloud is very large," even though this is sometimes the case.

Third, users need to be able to get control of the service-level agreement (SLA) discussion back from the vendors. SLAs have become a bit like end-user license agreements (EULAs) in personal-use software: you check the box and click "continue" because you don't care that much, you assume the terms of the EULA are fairly anodyne, and you don't want to spend the time needed to read about indemnification, jurisdiction, and other lawyer inventions. But when you do read some of the early (and some current) SLAs from cloud providers, there are reasons for alarm: the provider guarantees almost nothing, gives you minimal reductions in fees if its service is down (even though the outage may have a much larger impact on your business), and invokes the inherent unreliability of computers and networks as an excuse for potential service interruptions as if this were still 1975. Clearly, the issue is not to have a standard SLA, because that's not how a free market works, but it is to have a standard way to *describe* an SLA so that two vendors can be compared.

Fourth, there needs to be a way to measure and rate the things that can be. In other words, users need benchmarks. When computers started to be serious business in scientific research centers, each supplier had a different, self-serving way to describe how fast its machines ran. So in 1972, the National Physical Laboratory in the UK (essentially the equivalent of today's NIST in the US) created the Whetstone benchmark to provide a standardized comparison of how fast scientific programs would run on different machines. Today, users need a similar benchmark for cloud offerings, especially for IaaS, where numerical measurements are easier to specify than for PaaS or SaaS.

> **Clearly, the issue is not to have a standard SLA, because that's not how a free market works, but it is to have a standard way to *describe* an SLA so that two vendors can be compared.**

These four requirements may seem a tall order given the current state of disparate cloud offerings, but you only need to go back to the utility metaphor used to define cloud computing to find the equivalent. When I connect an appliance to an electrical socket in my house:

- The plug has the same shape on all appliances, so I do not need to have a special socket designed for that appliance.
- I don't need to call the utility or an electrician to deliver a special kind of electricity with a different voltage or frequency to that socket.
- There is a label on the appliance that says how much power it will consume, and there is a sticker on the circuit breakers in my garage that says how many amps each circuit will bear.
- In places where electricity is deregulated, and I can buy it from different companies, it comes on the same wire, and I can compare prices based on a standard unit of service (kilowatt-hours).

It should be realistic to expect the same capability for cloud services in the relatively near future. One might observe that one area where the utility companies are still behind user expectations is SLAs, because they come from a recent past of having monopoly positions where they did not need to compete on percentage of uptime or how they compensated users for downtime. While cloud providers do not have that comfortable past history, the fact that utilities are not yet offering comprehensive or standard SLAs should serve as a caveat: this may well be the issue that takes the longest to resolve.

CLOUD STANDARDS: THE SUPPLIERS' VIEW

In fact, there has been notable standardization activity for the cloud in the past few years. Most of it came from cloud infrastructure suppliers and therefore has been focused on IaaS. In addition, cloud standardization efforts became a way for existing industry consortia to continue to justify their existence by repurposing prior efforts. For example, the OVF was first proposed in 2007 and was not explicitly related to cloud computing, but now it is presented as one of the efforts that contributes to the standardization of cloud interfaces.

Why do vendors contribute to standards? For small vendors, the motivation is clear: if they do not provide "plug-and-play substitutability" for the larger vendors, they have little chance to gain entrée to the accounts that are already occupied by those larger companies. But what's in it for larger vendors?

In fact, in the history of computing, we have seen many examples of larger vendors who disregarded standards for a long time. For example, IBM was very late in adopting almost every standard during the past century: it used EBCDIC coding instead of ASCII, it used SNA long after TCP/IP had taken off, its email system did not interoperate with SMTP until the mid-1990s, and so on. Microsoft decided to join the OMG and participate in CORBA discussions in 1995, but then vanished from one meeting in order to block the progress of work that threatened the position of its proprietary COM protocol.

Eventually, several things happen even to the larger vendors:

- Their customers learn about the standardization efforts in progress and ask the vendors why they are not participating. The customers care because they want to avoid lock-in. The vendors initially say, "You don't need us to adopt those standards; our proprietary system is better." Eventually, though, enough customers say, "Sorry; if you don't converge on standards, we won't buy from you anymore" to reach a tipping point. So the realization sinks in that the lack of standards compliance is hindering the development of the market.
- They have trouble recruiting and training engineers to work on their proprietary systems. Universities train

students using the most prevalent, multivendor systems because they can get them cheap and because it makes the students more valuable to more employers.

- Especially since the open source movement started, using a standard means that you can reuse components developed and improved by a community, and this has value even for a large company.

As with all tipping points, the situation can sometimes shift quite rapidly. The Unified Modeling Language (UML) is a good example. In August 1993, a gaggle of methodologists led by Steve Mellor published a diatribe in the *Journal of Object-Oriented Programming* (JOOP) called "Premature Methods Standardization Considered Harmful" in response to a call for unification by Ivar Jacobson. Two years later, most of the people involved in this conflict were writing *together* the requirements for what would become the standard method for software modeling, UML.

Today, the field of vendor-led cloud standards resembles the early 1993 situation. There are several offerings, they seem to address many of the same issues, they are hard to compare, they all seem to keep progressing in parallel without much coordination, they are complex, and they mostly affect the people who manage the infrastructure, not the end users. A detailed review of each of these standards would take an entire issue of this journal, so I present a simple summary here:

- **Open Virtualization Format (OVF).** Developed by VMware and other suppliers, this is a format to package and distribute a virtual machine.
- **Cloud Server API.** Developed by Rackspace, this specification is published under a Creative Commons license. Rackspace claimed that "while OVF provides a standard packaging, it doesn't address the problem of different virtual disk formats."[3]
- **Cloud Data Management Interface (CDMI).** This specification from the Storage Networking Industry Association (SNIA) "defines the functional interface that applications will use to create, retrieve, update, and delete data elements from the cloud."[4] We're talking about management and administrative applications here, not business applications.
- **Open Cloud Computing Interface (OCCI).** This is "a protocol and API for all kinds of management tasks" from the Open Grid Forum (OGF).[5] This is a rather grandiose statement, so we would do well to remember that before the definition was expanded to its current fuzzy state, OCCI was targeted at deployment, "autonomic" scaling, and monitoring of IaaS.
- **Cloud Infrastructure Management Interface (CIMI).** CIMI is a protocol for the creation and management of virtual machines and volumes, developed by the Cloud Management Working Group of the Distributed Management Task Force (DMTF), a not-for-profit association of industry members "dedicated to promoting enterprise and systems management and interoperability."[6]
- **Open Systems Group (OSG).** The OSG, the oldest group within the Standard Performance Evaluation Corporation (SPEC), is now working on justifying and developing benchmarks for IaaS.
- **Global Inter-Cloud Technology Forum (GICTF).** Started and led by Japanese computer manufacturers (Hitachi, NEC, etc.), the aim of the GICTF is "to develop an inter-cloud world in which multiple cloud systems, running with different policies, interwork with each other to share resources."[7] As of July 2012, the GICTF has started collaborating with the DMTF.
- **Zend Simple Cloud API.** This is an open source library of interfaces to access multiple cloud services in a uniform way using PHP, one of the key technologies used (on the server side) to build Web pages.

Suppliers will tell you that the main reason why they drive standards is that the customers don't come to the table even when they are invited.

GIVING THE CUSTOMERS A STRONGER VOICE

Suppliers will tell you that the main reason why they drive standards is that the customers don't come to the table even when they are invited. Even though this sounds like an easy excuse, and it sometimes is, it is often true.

Taking part in standards work takes time and money. Developing the "meat" of a standard takes weeks of hard work; it demands painstaking review and verification; it often requires equipment and software to test that the standard can actually be met; and there is a substantial need for face-to-face working sessions that involve travel costs. Most end-user organizations just do not see why they should invest so much. For the suppliers, there is the prospect of larger sales, but for the user, it's much cheaper to wait for the standards to emerge and then see how soon the products evolve to comply with them.

Once vendor-driven standards do emerge, however, users are quick to point out that the standards are hard to understand, are biased toward making the vendors' — not the users' — lives easier, and do not reflect the true interests of the users. No one has found a real solution to this problem. There are abundant calls for end-user involvement in the standards process in various reports (see the sidebar "Wishing for End-User Involvement"), but they sound more like wishful thinking than reality.

There are exceptions. In the cloud arena, the nonprofit Open Cloud Consortium (OCC) acts both as a supplier of cloud resources to scientists and as a think tank about the needs of its members. One of its working groups manages a test bed for open cloud technology and "intercloud" computing. Another one develops benchmarks. The Large Data Cloud Working Group is developing APIs for cloud storage services.

While it is dominated by hardware, software, and cloud suppliers, the Cloud Security Alliance (CSA) has a larger member base than traditional vendor-led standards organizations and includes some very large consumers of computer resources (eBay, Orange, Raytheon, etc.). Its products, however, are more oriented toward training and certification than toward true standards that would improve security.

When it comes to advocating for actual standards from the end user's perspective, the most promising effort to date appears to be the Cloud Standards Customer Council (CSCC). It was formed in 2011 under the auspices of the OMG and holds its meeting in conjunction with the OMG, but the CSCC is an independent organization with over 300 members, many of whom are not OMG members. Its charter is to be "an end-user advocacy group dedicated to accelerating cloud's successful adoption, and drilling down into the standards, security, and interoperability issues surrounding the transition to the cloud."[8] The wording is important: the CSCC does not write standards, but it promotes cloud computing and identifies issues that would benefit from standardization.

Not being a standards organization has its advantages. The CSCC has been able to develop two useful documents in about six months each:

- The "Practical Guide to Cloud Computing"[9] proposes a methodology to prepare for transition to the cloud, with nine specific steps.
- The "Practical Guide to Cloud Service Level Agreements"[10] helps end users make sense of what vendor SLAs do or do not contain and decide what to ask for.

After completing the second deliverable, the CSCC considered creating a list of requirements for cloud benchmarks, but this is an area that the OSG consortium has already started addressing. What the CSCC is now planning is a "deeper dive" into the specific content and recommended language of SLAs, complete with examples and comparisons of existing agreements.

WISHING FOR END-USER INVOLVEMENT

> It is generally accepted in Australia that user involvement in standards setting is an integral part of the telecommunications regulatory environment ... the Minister may … make a grant of financial assistance to a consumer body for purposes in connection with the representation of the interests of consumers in relation to telecommunications issues.
>
> — Australian Communications and Media Authority[1]

> Standards are intrinsically difficult to implement and adopt. In the industrial automation business, OPC is a unifying standard that allows true interoperability. It needs more end-user support and involvement.
>
> — Automation.com[2]

> The emergence of consortia and other ad hoc mechanisms replacing the traditional Standards Development Organization (SDO) processes has raised questions about the importance of due process and end-user involvement in standardization.
>
> — Researchers at the University of Pittsburgh[3]

[1]"User Involvement in Standards Setting." Australian Communications and Media Authority (www.acma.gov.au/webwr/_assets/main/lib310131/9a-user-involvement-in-standards-setting.doc).

[2]Pinto, Jim. "OPC — The Standard That Makes Other Standards Interoperable." Automation.com, 17 October 2006.

[3]Spring, Michael B. et al. *Improving the Standardization Process: From Courtship Dance to Lawyering; Working with Bulldogs and Turtles*. University of Pittsburgh, 1995 (www2.sis.pitt.edu/~spring/papers/improve.pdf).

GOVERNMENT ROLES

So far, we have talked about vendors and users, but what is the role of government entities? Most writers of English-language publications take a US- or UK-centric view of this question, which should be avoided. As soon as you broaden your view of the world, you will realize that there are almost as many models of government involvement in standards as there are countries.

In the US, the National Institute of Standards and Technology (NIST) has been active in the cloud arena, but it operates within a free-market ideology in which interference with the market-based decisions of the suppliers would be frowned upon. NIST's cloud-related work so far has consisted of publishing a definition, which has been criticized for being much too long, and a reference architecture, which has the merit of defining certain important concepts such as the different types of clouds (IaaS, PaaS, SaaS) or the different modes of deployment (public, private, hybrid, community).

In the European Union, there was initially much reluctance to adopt cloud computing, largely because of concerns about the protection of individual information. Since privacy protections are so strict in Europe, placing "personally identifiable information" (PII) in the cloud seemed like courting disaster. This is changing, however, with the publication, in mid-2012, of European Commission guidelines that explicitly endorse the use of the cloud if accompanied by appropriate data protection measures.

The Cloudbook Community publishes on its website a list of notable government efforts that promote cloud computing in the US, Japan, Canada, the UK, and the EU. Most of these efforts are aimed at accelerating adoption or creating pilot clouds shared by several agencies, but they are not oriented toward the development of standards.

MIX INGREDIENTS AND STIR WELL

Where does this leave us? As is often the case, we need a little bit of everything, in the right proportions, to end up with an appealing dish. Everyone has a role to play, but so far history has repeated itself: the emerging cloud standards have been dominated by the suppliers, and we need more scrutiny by the users to make sure that the standards serve them well. The work the CSCC has started is promising in this regard.

Going back to the initial point of this article, the technology adoption curve has been steeper than what many of us were expecting just three years ago, so the standards have had trouble following. We do not need *more* standards efforts: the alphabet soup of the earlier "suppliers' view" section proves that point. We need several of these organizations to merge or to join forces and create a single, consistent body of cloud standards consistent with customer priorities and requirements.

ENDNOTES

[1]Rogers, Everett M. *Diffusion of Innovations*. 5th edition. Free Press, 2003.

[2]Wardley, Simon. "Situation Normal, Everything Must Change." Presentation to the *Open Source Convention*, Portland, Oregon, USA, 23 July 2010 (www.youtube.com/watch?v=5Oyf4vvJyy4).

[3]Williams, Jorge L. "An Implementor's Perspective on Interoperable Cloud APIs." Presentation to *Cloud Interoperability Roadmaps Session*, Long Beach, California, USA, 10 December 2009 (www.omg.org/news/meetings/tc/ca/special-events/ci/Rackspace.pdf).

[4]"Cloud Data Management Interface (CDMI)." Storage Networking Industry Association (SNIA) (www.snia.org/cdmi).

[5]"An Open Community Leading Cloud Standards." Open Cloud Computing Interface (OCCI) (http://occi-wg.org).

[6]"Cloud Infrastructure Management Interface (CIMI) Model and RESTful HTTP-Based Protocol: An Interface for Managing Cloud Infrastructure." Distributed Management Task Force (DMTF), 2012 (http://dmtf.org/sites/default/files/standards/documents/DSP0263_1.0.0e.pdf).

[7]"Background to the Establishment of the Forum." Global Inter-Cloud Technology Forum, 2009-2010.

[8]"OMG Announces Cloud Standards Customer Council." Press release, OMG, 7 April 2011 (www.omg.org/news/releases/pr2011/04-07-11.htm).

[9]"Practical Guide to Cloud Computing." Cloud Standards Customer Council (CSCC), October 2011 (www.cloudstandardscustomercouncil.org/10052011.htm).

[10]"Practical Guide to Cloud Service Level Agreements." Cloud Standards Customer Council (CSCC), April 2012 (www.cloudstandardscustomercouncil.org/04102012.htm).

Claude R. Baudoin is a Senior Consultant with Cutter's Data Insight & Social BI *and* Business & Enterprise Architecture *practices. Prior to becoming an independent consultant, he spent 35 years in the corporate world, occupying various positions in software engineering and IT management, in France and the US, mostly for the global oilfield services company Schlumberger. Over his career, Mr. Baudoin developed expertise in enterprise and information architecture, business process modeling, enterprise security strategy, building communities of practice, collaboration methods and environments, and the use of social media in business. He has given several talks and led workshops for Cutter, advocating a balanced and pragmatic approach to the adoption of cloud computing.*

Mr. Baudoin is the coauthor of two books: Realizing the Object-Oriented Lifecycle *and* Méthodes de Programmation. *He has served as Guest Editor for* Cutter IT Journal, *presented at international IT events, published many papers, and holds two patents in networking and information security. Mr. Baudoin has an MS in computer science from Stanford University and an undergraduate degree in engineering from Ecole Polytechnique (Paris). He can be reached at cbaudoin@cutter.com.*

An Open Source Approach to IaaS Cloud Standards

by Beth Cohen

"Cloud" as a popular term has been around for only a few years, but the concepts of pay-as-you-go services and IT as a utility go back to the first hosted Web servers starting in the late 1990s. After 15 years or so, you would think that cloud technology would have standardized. While basics such as the ASCII binary text standard that helped define the Internet have been widely accepted, cloud standards have been much slower to be adopted. As cloud infrastructure technology matures, increased interest by enterprises and emerging vendors is driving a renewed effort to create viable standards. The logical next step in the maturation process is the long-overdue development of real cloud computing infrastructure and interoperability standards. Because of its unique leadership role in the IT community, open source offers the best option for the development of these standards.

After years of unreliable infrastructure, questionable business models, and application service provider (ASP) false starts, the cloud has finally matured into a robust industry. Projections of US $14.5 billion in SaaS revenue in 2012 alone[1] demonstrate the proven success of the many products and solutions available today. In the IaaS segment, popular proprietary platforms such as Amazon's AWS, Rackspace, VMware's vCloud-based Terremark, and rising star Bluelock[2] have demonstrated the viability of a cloud infrastructure utility, pay-as-you-go business model.

Everyone benefits from technology standards in the long term; you would think, therefore, that everyone would agree to give cloud infrastructure technology standards high priority. However, creating compelling proprietary systems and discouraging standards gives early adopter companies competitive advantage in the short term. Thus, until recently, the high opportunity costs, the difficulty of building data centers, and the lack of a perceived need for enterprises to build private cloud infrastructures translated into little incentive to develop standards. All that is changing now that enterprises are seeing the value of building private clouds that can take advantage of the same cost efficiencies and self-service capabilities that incumbent commercial services offer. Now, the lack of real standards is a major factor holding them back.

With the emergence over the last few years of a number of cloud standards bodies, such as the Cloud Security Alliance (CSA), OpenStack Foundation, Distributed Management Task Force (DMTF), and Cloud Computing Interoperability Forum (CCIF), just to name a few, the open source community is poised to address the standards gap. By leveraging the proven model that created such widely adopted standards and tools as Mozilla, Linux, and Apache, there is an opportunity to create powerful, long-lasting IaaS cloud technology standards more quickly than by relying on commercial interests alone. Some of the issues discussed in this article include:

- The real benefits of an IaaS cloud standard for businesses and consumers
- The difference between cloud reference architectures and cloud standards and why there is a need for both
- The available open source cloud standards and architectures
- Why open source is the best approach to creating a workable IaaS standard

CLOUD REFERENCE ARCHITECTURES AND STANDARDS

Many people assume that reference architectures and standards are one and the same. Yes, they are clearly related concepts, but a reference architecture is a far more abstract construct than a standard. Standards are more concrete because they need to define how disparate systems will interact. Think of a reference architecture as the starting point for framing the conversation about what standards need to be defined and how they will interface with other components inside and outside the system. For standards to be successfully adopted, vendors and users need to embrace them. Users must encourage vendors to build systems that meet standards so that different clouds will interoperate properly.

There are several widely published cloud reference architectures, but none of them alone is enough to translate to architectures that businesses and consumers can use to build more robust public or private clouds. Like the Open Systems Interconnection (OSI) and cloud stack models, they are artificial constructs or frameworks that can be used to create real systems. By themselves, cloud reference architectures tend to be so generic and broad that using them to build an enterprise cloud without considerable work defining business objectives and technical system requirements would only be an exercise in frustration. To build clouds, reference architectures need to be combined with a rich set of standards to take advantage of the cloud's promise of utility computing and full interoperability.

For the most part, there is consensus that cloud architectures roughly divide into three logical abstraction layers: SaaS, PaaS, and IaaS. This is often referred to as the "cloud service stack model." Starting from the most abstract or top layer, SaaS is the public face of the cloud, encompassing applications such as Salesforce.com, Workday, Mozy, and other tools used by millions of businesses and consumers. When applications are built in a private cloud within the enterprise, they are often referred to as cloud-ready applications.

Moving down the stack, PaaS was originally envisioned as strictly the virtual machine image of the operating system, but increasingly it includes the development environment and the operations management tools as part of the package. The logic is that an operating system image file is by itself unusable; you need tools to make it useful. Microsoft's Azure and VMware's Cloud Foundry are examples of pure PaaS platforms. PaaS is still an emerging concept, so expect changes in the definition as cloud architecture models mature.

Finally, the most conceptually concrete layer is IaaS. By definition it generally covers any part of the system below the virtual machine. This includes the data center, physical hardware, network, hypervisors, orchestration tools, and operations systems. The IaaS layer is by far the most complex in terms of the number of components and how they integrate.

A quick review of a few of the published models reveals that a reference architecture means very different things depending on the agenda of the creator. A useful way to compare the perspectives is to map them to the cloud service stack model. For example, both the NIST Cloud Standards[3] and the closely related IBM architecture[4] are relatively generic and high level, but they both have a strong operational focus that closely matches the IaaS layer concerns. The Rackspace Private Cloud Reference Architecture,[5] while specific to OpenStack, also primarily has an operational/IaaS bias. Microsoft has defined the cloud, not surprisingly, more from a development platform/PaaS view, but it also has an IaaS-flavored version based on its proprietary hypervisor, Hyper-V. The HP and VMware varieties are more appropriate for companies building end-to-end applications, SaaS or otherwise. It should be noted that the VMware and Microsoft architectures are both based on systems that did not originate as clouds per se, but rather started as hypervisors with cloud tools wrapped around them. The hypervisor-centric approach remains with their focus on the virtual image layers of the cloud rather than the orchestration component.

A reference architecture means very different things depending on the agenda of the creator.

While I would not go so far as to say that the boundaries between each of the three layers have been completely standardized, for the most part there is a sense of agreement on these boundaries within the cloud community. However, standards beyond the three simple divisions are fuzzy at best. There is much room for the creation of viable standards within each of the layers. For example, data exchange standards would allow easier sharing of SaaS data between applications and simpler integration across SaaS providers. In an ideal world, not only would this enable consumers to migrate between SaaS product offerings as their requirements change, it would also allow businesses to create mixed portfolios of SaaS and internal cloud-ready applications. No SaaS vendor should be so arrogant as to think that its solution is the only one its customers will be using. Cloud data transfer standards would also reduce customer worries about the very real threat of vendor lock-in.

At the PaaS layer, establishing cloud standards would let companies create applications that work on a variety of different platforms without the need for costly ports or application refactoring. PaaS also offers the opportunity for companies to enforce tool and platform standardization without risking the wrath of developers who demand specialized toolsets. One enterprise used PaaS tools as a way to rein in rising development costs and complexity. It offered developers the choice of using standard tool templates, which were immediately available through a self-service portal, or custom tools, which were available if the developers were willing to wait for the typically overworked IT department to hand-build them.

Finally, IaaS layer standards open up the possibility of wider adoption of cloud architectures in the enterprise setting by reducing complexity and enabling companies to take advantage of operational best practices and widely understood use cases. Anyone who has attempted to build even a simple enterprise cloud knows just how difficult it can be. The growing popularity of cloud-in-a-box hardware offerings such as VCE's Vblock and NetApp's FlexPod attests to the need for good IaaS cloud standards. There are even some OpenStack packages available from StackOps and Cloudscaling that are designed to simplify implementation projects.

As we can see, cloud reference architectures are certainly a good starting point for building cloud standards, but they are just not concrete enough to build an entire corporate cloud strategy around, and it is certainly not appropriate to equate them with cloud standards. Cloud standards need to be created by collective bodies with representatives from the various constituent communities. Traditionally, standards organizations such as the Internet Engineering Task Force (IETF) and the IEEE are hybrids of trade groups and nonprofit foundations. The open source community has a long history of creating and supporting major components of what makes up the Internet, so it is a fitting choice for taking a leading role in creating cloud IaaS standards.

As MINIX creator Andy Tanenbaum put it so succinctly, "The nice thing about standards is that you have so many to choose from."

CURRENT OPEN SOURCE IAAS CLOUD STANDARDS

To put the discussion of IaaS open source standards in perspective, a short summary of the currently available IaaS stacks (courtesy of John Treadway's excellent blogs[6] on the subject) is in order. The IaaS cloud stacks break neatly into propriety and open source efforts. So far the only open source projects making any serious community headway toward becoming the definitive cloud standard are Eucalyptus, OpenStack, and CloudStack. As is typical in the open source world, each of these projects is approaching the creation of an IaaS cloud standard differently. As MINIX creator Andy Tanenbaum put it so succinctly, "The nice thing about standards is that you have so many to choose from."

Eucalyptus

Eucalyptus, first conceived in 2008 as a high-performance computing research project at the University of California at Santa Barbara, is an open source emulation of Amazon's AWS. Both the company, founded in 2009, and the project have had problems with funding and lack of strong direction. The company's attempt to play both sides of the open source fence by maintaining a dual commercial and open source code base, combined with its weak business model (who really needs a pale imitation of Amazon anyway?) has made for rocky going. That said, Eucalyptus's fortunes and direction have changed radically with the March 2012 Amazon announcement of full support and partnership.[7] Rampant speculation about Amazon's motivations aside, the move provides more options for companies that already have substantial investments in applications built on the Amazon platform, while essentially ducking the standards issue.

OpenStack

With over 3,300 contributors and 183 companies backing the project, including such giants such as Dell, Rackspace, IBM, HP, and NetApp, OpenStack is by far the largest IaaS standard effort to date and one of the fastest-growing open source initiatives ever undertaken.[8] The project began in July 2010 and released its first prototype the following October. Rather than just aligning itself with existing open source foundations, OpenStack is now in the process of establishing an official foundation modeled along the lines of the Linux Foundation.

It is a positive sign that the OpenStack community is serious about keeping the technology available to any and all who want to create software and tools that support real cross-platform IaaS cloud integration. The Apache 2.0 licensing model means companies using this standard to provision their cloud services will not incur any licensing costs and are allowed to freely interoperate across OpenStack implementations.[9] There are already many variations, including Ubuntu Enterprise Cloud (Canonical), Piston, and Nebula, among others. The technology has gotten the attention of big Amazon rivals such as Rackspace, HP, AT&T, and KT (Korean Telecom), each of which has built commercial cloud services based on the OpenStack core. While technologically OpenStack is still quite immature, it is far more ambitious than any of its rivals, with modules for virtual networking (Quantum), encryption (Keystone), object storage (Swift), and block storage (Ceph), in addition to the basic cloud computing platform (Nova). Of all the current contenders for the IaaS crown, this one has the highest chance of success due to its broad community support.

Cloud.com/CloudStack

CloudStack has always been a bit of a mixed bag. Started originally as a proprietary project by Cloud.com (formerly VMOps), the first version of the product was released in April 2010, giving it a jump on its chief rival, OpenStack. The company's purchase by Citrix (which also originally backed OpenStack) in July 2011 has further muddied the waters. Citrix has since dropped support of OpenStack and seems to be betting on its investment in Cloud.com. In April 2012, Citrix announced that it is donating the code base to the Apache Foundation as an incubator project. Strategically, Citrix is trying to drum up more open source community interest, which has always been somewhat anemic. The switch to the more open Apache licensing model over GPL might encourage more companies to participate. On the positive side, CloudStack has an established base of enterprise private cloud deployments. Its support tools and feature set are more mature — not surprising given that it had a year's technology head start. At this point, though, CloudStack's future as an open source IaaS standard seems, well, cloudy.

Other Contenders

Rounding out the list are several older, smaller projects, including OpenNebula and Nimbus, which were both begun as research projects. Neither of them is gaining much interest in the commercial software communities. Finally, a recent Japanese project, Wakame-VDC, has some deployments at Kyushu Electric Power Co., Inc., and the Japan National Institute of Informatics.

On the proprietary side, the major stakeholders are Amazon, VMware, Microsoft, and IBM. Amazon's AWS, the 800-pound gorilla, is strictly offered as a service so far. VMware's vCloud is for the most part sold as an enterprise cloud solution. The service providers who have packaged it are forced to charge a premium due to VMware's licensing model. Microsoft and IBM also have commercial products with Hyper-V and SmartCloud, respectively. None of these companies have a strong history of supporting interoperability standards. Judging from their behavior so far, this will not be changing anytime soon. Oddly, both Microsoft and IBM are supporting OpenStack. Industry opinion is that this is mostly a way to confront the Amazon and VMware monopolies in the commercial IaaS services and private cloud sectors, respectively.

For companies looking for interoperability and the capability to build hybrid clouds, various proprietary and open APIs have been proposed to provide interoperability among IaaS platforms at the virtual machine level. The first (and so far only) cloud-oriented standard that has been ratified is the Open Virtualization Format (OVF), which was approved in September 2010 after three years of processing by the DMTF. OVF's open packaging and distribution format offers some platform independence by allowing migration between IaaS platforms, but it does not provide all the tools needed for full cloud interoperability.[10]

APIs are all well and good, but they do not guarantee true interoperability.

It is quite clear that with all the chaos in the IaaS platform space, the standards are equally immature. The biggest needs remain for interoperability standards to allow virtual machines to be migrated between clouds transparently and for more robust hybrid cloud solutions. For the moment, companies that want to use multiple IaaS platforms or a mix of public and private options are stuck with complex architectures and emerging orchestration tools such as enStratus and Rightscale to bridge the gap. Despite the obvious need for more work, the open source community is still far more amenable to working on the IaaS cloud standards than the commercial interests are.

USING OPEN SOURCE TO CREATE STANDARDS

Users of the cloud are asking when cloud computing standards will be mature enough so that more companies will feel comfortable implementing cloud architectures and using cloud services without feeling locked in. Ironically, while the commercial cloud offerings have been growing, built on the very standards that created the Internet itself, Amazon and others have been reluctant to publish their architectures. APIs are all well and good, but they do not guarantee true interoperability. Downstream vendors quickly find that they need to build API interfaces for several services, adding significantly to the development and maintenance costs. To address this problem and transparently transfer workloads among the different vendors based on predefined business rules, there must be much more comprehensive standards.

One obvious question is whether there is an opportunity for the commercial cloud IaaS systems to become standards. After all, there are precedents for formerly proprietary formats, such as VMware's VMDK (Virtual Machine Disk), becoming de facto standards simply by being widely adopted by the industry. Cloud computing itself originated as proprietary services built on top of such long-time robust standards as TCP/IP, HTTPS, and SSL. One could even argue that IaaS platforms such as Amazon's AWS are already standards. However, as many companies have found to their chagrin, while AWS has a vast variety of services, easy-to-use tools, and a 15-year technological head start on the competition, it is something of a cloud world roach motel ("Enterprises check in … but they don't check out!"). Lots of companies have found it easy to get applications running quickly, but changing providers or taking the applications back in house as requirements change is fraught with unexpected perils. Amazon's backing of Eucalyptus does not address that problem directly, but it does offer a feasible option for companies that want to build what Amazon euphemistically calls "on-premises services."

> **Given its history, open source development is the best framework for creating effective IaaS cloud standards.**

Open source has been successfully used to create widely used standards such as Linux and the Apache suite of Web services. Given its history, open source development is the best framework for creating effective IaaS cloud standards. Other community-driven cloud standards projects, such as the CSA's Cloud Security Guidelines, have been widely adopted. After years of neglect, as it concentrated on polishing Linux, MySQL, and the LAMP stack, the computing open source community is finally starting to devote its considerable resources and energy to the long-recognized need to create IaaS cloud computing standards. Early projects, which mostly came out of the academic community, were either not commercial enough or were available before there was a sustainable market. But the wild success of the OpenStack project is enough to demonstrate the potential for leveraging the open source approach to drive standards. With strong financial and technical backing from several large commercial interests from the beginning, it doesn't have the baggage that the single-company product CloudStack has to contend with or Eucalyptus's association with Amazon's proprietary approach.

WHERE IaaS CLOUD STANDARDS ARE GOING

Now that the open source community recognizes the need for viable cloud IaaS standards, progress over the last two years has been impressive. With growing pressure from enterprises that do not want to be too dependent on proprietary cloud systems that were not built with their requirements in mind, the need for better integration and cross-platform standards has never been higher. In the end, a good cloud IaaS standard should be robust enough to allow it to be used from a variety of different perspectives: business, operations, development, and consumer. The financial, technical, and organizational frameworks are in place, and the open source community is eager to take up the challenge of making these much-needed standards a success. Expect to see major shifts in the technology with better integration tools and simpler cloud architectures as the open source community takes a greater leadership role in cloud standards in the next few years.

ENDNOTES

[1]Eddy, Nathan. "SaaS Revenue to Reach $14.5 Billion in 2012: Gartner." *eWeek*, 29 March 2012 (www.eweek.com/c/a/Cloud-Computing/SaaS-Revenue-to-Reach-145-Billion-in-2012-Gartner-567722).

[2]Burns, Christine. "10 Most Powerful IaaS Companies." *Network World*, 9 April 2012 (www.networkworld.com/supp/2012/enterprise2/040912-ecs-iaas-companies-257611.html).

[3]"Two New Publications Provide a Cloud Computing Standards Roadmap and Reference Architecture." *NIST Tech Beat*, 13 September 2011 (www.nist.gov/itl/csd/cloud-091311.cfm).

[4]Stifani, Raffaele, Stefan Pappe, Gerd Breiter, and Michael Behrendt. "IBM Cloud Computing Reference Architecture." *Academy TechNotes* (IBM Academy of Technology), Vol. 3, No. 1, 2012 (www-05.ibm.com/it/cloud/downloads/Cloud_Computing.pdf).

[5]"Rackspace Cloud: Private Edition Reference Architecture." Rackspace US, Inc., November 2011 (www.referencearchitecture.org).

[6]Treadway, John. "Cloud Stack Red Ocean Update — More Froth, but More Clarity Too." *Cloudbzz*, 22 October 2011 (www.cloudbzz.com/cloud-stack-red-ocean-update-more-froth-but-more-clarity-too).

[7]"Amazon Web Services (AWS) and Eucalyptus Partner to Bring Additional Compatibility Between AWS and On-Premises IT Environments." Press release, Eucalyptus Systems, 22 March 2012 (www.eucalyptus.com/news/amazon-web-services-and-eucalyptus-partner).

[8]OpenStack (www.openstack.org).

[9]Iyer, Bala, and John Henderson. *Wisdom of Clouds*. SIM Advanced Practices Council, 2012.

[10]Claybrook, Bill. "Weighing the Cloud Computing Standards Dilemma." Search Cloud Computing, February 2011 (http://searchcloudcomputing.techtarget.com/feature/Weighing-the-cloud-computing-standards-dilemma).

Beth Cohen is a Senior Cloud Architect for Cloud Technology Partners, Inc., a firm focused on delivering solutions to help enterprises leverage the efficiencies of cloud architectures and technologies. Some of Ms. Cohen's recent projects include: building the cloud infrastructure for a major international electronics company in support of consumer-facing applications that will scale to millions of users; developing a cloud-based collaborative community for the dental industry to allow for the secure sharing of images and records across suppliers; creating migration architectures and processes to move over 11,000 applications into the cloud for a major international telecom; building the operational systems to support a cloud-based portable medical record service designed to securely and transparently store millions of records; and creating a unique cloud-based data protection service that delivered cost-effective data protection, disaster recovery, and business continuity solutions for the SME market. Previously, Ms. Cohen was the Director of Engineering IT for BBN Corporation, where she was involved with the initial development of the Internet, working on some of the hottest networking and Web technology protocols in their infancy. She can be reached at beth.cohen@cloudtp.com.

CARMA: The Open Cloud Standard

by Aditya Watal

WHAT IS CLOUD REALLY?

Richard Stallman[1] is the founder of GNU Linux and the Free Software Foundation and one of the key experts in the field of open computing. In June 2012, I interviewed him for my blog, Watalon.com. The aim of the interview was to get Stallman's views on cloud computing and how open standards should be built around the cloud. In the interview, Stallman vehemently argued that the cloud does not exist and the term "cloud" as such is a complete misnomer.

He feels that most major vendors are trying to pass off anything they offer as a cloud solution. This may not be far from the truth, as the industry already had a term for such labeling: "cloud washing."[2] After I pushed him hard to give an alternate term to the cloud phenomenon, he chose "network service layers."

The networking world already has a well-evolved and de facto standard, the Internet protocol suite (or TCP/IP). The main design goal of TCP/IP was to build an interconnection of networks, referred to as an internetwork, or "Internet," that provided universal communication services over heterogeneous physical networks.[3]

Given that cloud is primarily about services delivered over a network, I believe that defining additional protocols and standards for these protocols within the layers of the TCP/IP stack would result in a more open standard for the cloud. However, this still does not answer the question, "What is cloud really?" In an attempt to answer this question, I have written a book entitled *Cloud Basics*, based on my experience and research. In the book, I define cloud as the convergence of two major trends in the last stage of the "First Epoch of Computing": layer elasticity and layer bundling (see Figure 1). A full convergence would mean the ability to stretch any layer elastically almost infinitely without breaking the bundle of which it is a part — that would be a "Real Cloud."[4]

WHAT WOULD BE A GOOD CLOUD STANDARD?

One can always measure a standard by applying the API test. First, this means that there are well-accepted or likely-to-be-accepted and open APIs that form the standard. Second, it means that the APIs meet the "A-P-I" test:

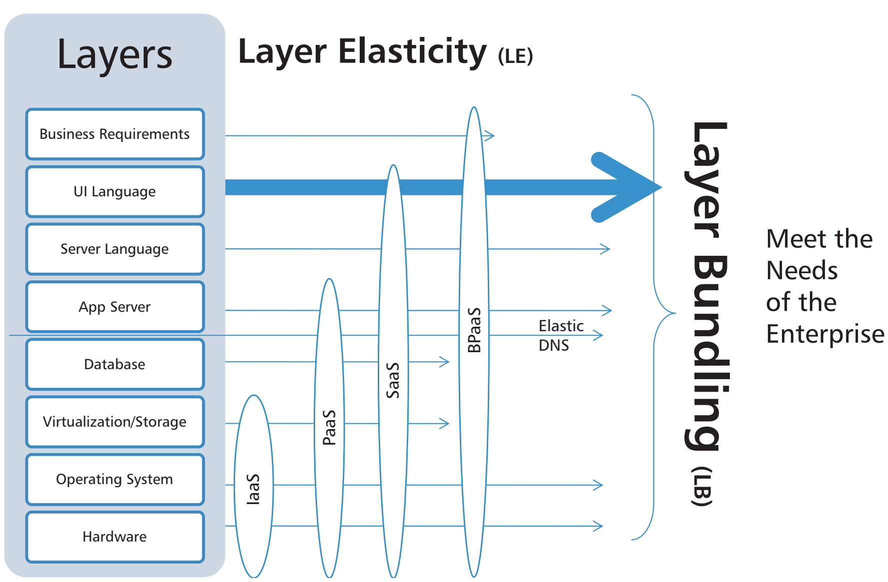

Figure 1 — The cloud is a convergence of layer elasticity and layer bundling.

- ***A*bstraction.** A good cloud API would always abstract the cloud developer from the underlying vendor implementation. So no matter the underlying vendor platform, the coding to the API is generic and reusable. Think *virtualized*.
- ***P*ortability.** A good cloud API would always allow the developer to easily move or port the application or data without the need for a vendor-specific data transformation or a custom application container or wrapper. Think *passport*.
- ***I*nteroperability.** A good cloud API can talk to or be integrated with other APIs with ease. This means that the API uses a fairly standard set of nouns, verbs, plurals, and associations. Think *lingua franca*.

CLOUD STANDARDS TODAY

Collaboration between the US and European standards organizations to bring together the work of the Open Data Center Alliance (ODCA), the Distributed Management Task Force (DMTF), the Storage Networking Industry Association (SNIA), and the European Telecommunications Standards Institute (ETSI)[5] is already in progress. Further, organizations such as Green Grid and ODCA have formed an alliance to improve collaboration in standards.[6] IEEE has also launched projects P2301 for Cloud Portability Standards and P2302 for Intercloud Interoperability and Federation.[7]

IaaS Standards

Standards in the IaaS space have evolved mainly on account of various open source API initiatives, including the Apache Lib Cloud, Red Hat's Deltacloud, and Rackspace's OpenStack. Lib Cloud and Deltacloud have good abstraction as compared to OpenStack.[8] Further, Amazon's announcement of its partnership with Eucalyptus in 2010[9] paved the way for movement of workloads between private clouds and Amazon.

Deltacloud today provides support for all major APIs (see Table 1). The storage driver functionality of Deltacloud works across Amazon S3, Eucalyptus Walrus, Rackspace Cloud Files, Microsoft Azure, and Google Storage, allowing users to create new buckets, update/delete buckets, create new blobs, update/delete blobs, and read/write individual blob attributes.[10]

Table 1 — Deltacloud's Support for Major APIs

IAAS API	Create New Instances	Start Stopped Instances	Stop Running Instances	Reboot Running Instances	Destroy Instances	List All/ Get Details About Hardware Profiles	List All/ Get Details About Realms	List All/ Get Details About Images	List All/ Get Details About Instances
Amazon EC2	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Eucalyptus	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IBM SBC	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
GoGrid	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OpenNebula	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Rackspace	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEV-M	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
RimuHosting	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Terremark	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
vSphere	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OpenStack	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FGCP	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes

SaaS Standards

Over 26% of the programmable Web APIs today are SaaS APIs.[11] Salesforce.com API is the de facto standard within the SaaS CRM space, and Oracle and SAP are making a strong bid via acquisitions to establish their presence in the human capital management (HCM) space.[12]

PaaS Standards

Major PaaS platforms such as Google App Engine and Heroku fail the abstraction test, as developers need to write code specific to the platform.[13] Force.com from Salesforce.com is worse, since it expects the programmer to write code in a proprietary language called Apex.

THE CURRENT STATE OF PAAS STANDARDS

Today, code written on the Google App Engine or Force.com must be rewritten if a developer wishes to run it on Heroku. Clearly, some of the top vendors offer little or no portability across platforms. Code portability is something that an enterprise app developer takes for granted within the enterprise, where JVM and Java come to the rescue. Therefore it would be safe to conclude that while in the IaaS and SaaS areas there are a multitude of APIs and standards, the PaaS area lacks standards. This therefore brings up the question what an ideal PaaS standard would be and what would make it successful.

WHAT DO WE NEED FROM A PAAS STANDARD?

An ideal PaaS standard would be one that provides functionality and openness.

Functionality

This requires a common framework for all the features that an application developer would expect from a platform (see Figure 2). From a PaaS standpoint, this functionality would be whatever an enterprise developer is already used to:

- **SDLC tools.** These include browser-based tools to track functionality, changes, and bugs; a data modeling tool that helps the developer visualize the data model and optimize it for high performance; a good IDE or plug-ins for Eclipse or other IDEs the developer already uses; and code version control and code merging. Last but not least, there should be a good tool for code and API documentation.
- **Integration, build, and deployment.** This includes support for various continuous integration and deployment tools; ease of monitoring processes, jobs, logs, and resource usage across various machines or VMs; and micro VMs (needed for development and test environments).
- **Standardized and portable containers.** These include standard containers with interfaces to other vendor clouds and other layers within the cloud that support all popular operating systems, enterprise platforms/middleware, and other cloud brokers. These containers should not require any custom coding to make them run on specific IaaS platforms and should be easily portable across VMs or environments.
- **Infinite elasticity.** Inherent in almost all application components today is some degree of elasticity. Such elasticity should begin at the Domain Name System

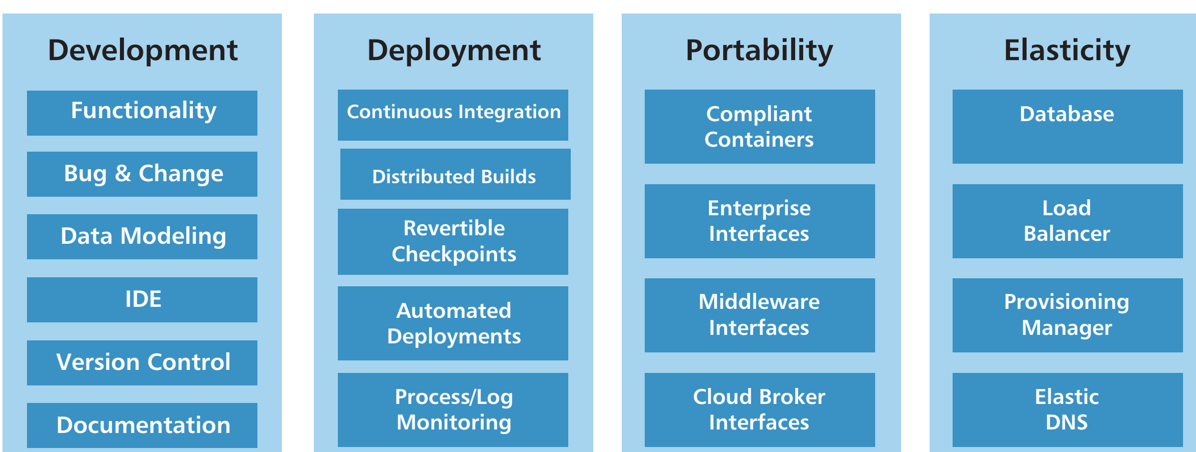

Figure 2 — A PaaS standard must offer a framework for all the functionality a developer expects from a platform.

(DNS) level, whereby an application can automatically failover from one location to another without the need for DNS remapping and associated downtime. The application should also be able to scale up or down automatically, without any limitation on the number of VMs or where these VMs are located. APIs that allow for automated billing and provisioning/deprovisioning are a key requirement for such elasticity. Without them, the obvious benefits of elasticity — such as optimized cost management and minimal downtime — would not be realized.

Openness

A truly open PaaS standard would be one that gives the developer a lot of open source components to work with and adequate control[14] over the computing:

- **Open source.** Various open source components already exist for almost all the functionality described above. These include Eclipse/NetBeans as an IDE, Jenkins for continuous integration and deployment, Subversion/Git/Mercurial for source code version control, nginx/Apache Web servers, Tomcat/Cherokee app servers, and MariaDB/MongoDB databases, to name a few applications that enable open development work today.

- **Control of computing.** A truly open standard would give developers adequate control over the computing. Some key questions that must be answered from a PaaS standpoint would be: Who owns the development and production environments? Who is given the access to the underlying infrastructure, and what is the role-based access mechanism around it?

The more control the PaaS provider gives the developer over the IaaS and the layers that form the core PaaS functionality, the more portable — and hence standards-compliant — that provider's offerings would be. So let's take a look at the current top PaaS providers and see how compliant they would be with a truly open PaaS standard, such as the CARMA standard I propose in this article.

PaaS PROVIDERS AND HOW THEY STACK UP

There are over 25 PaaS providers in the market today. In this section, I profile the top six providers — Google App Engine, Force.com, Azure, Heroku, Stackato, and Joyent — according to the PaaS standards framework outlined above (see Table 2). The purpose of this comparison is twofold: firstly, to show how these platforms stack up in relation to the standard and to test if the standards framework is adequate, and second, to give a sense of how these platforms compare, how likely they are to be adopted in the future, and thus whether there is a possibility of a de facto standard emerging.

Table 2 — How the PaaS Providers Stack Up Against the Proposed Open Standard

PaaS Layers	Required Component	Google App Engine	Force.com	Azure	Heroku	Stackato	Joyent
Development	Functionality			Yes	Yes	Yes	Yes
	Bug & Change			Yes	Yes	Yes	Yes
	Data Modeling	Yes	Yes	Yes	Yes	Yes	Yes
	IDE/IDE Plug-ins	Yes	Yes	Yes		Yes	
	Version Control			Yes	Yes	Yes	Yes
	Documentation			Yes	Yes	Yes	Yes
Deployment	Continuous Integration			Yes	Yes	Yes	Yes
	Distributed Builds			Yes	Yes	Yes	Yes
	Revertible Checkpoints			Yes	Yes	Yes	Yes
	Automated Deployments	Yes	Yes	Yes	Yes	Yes	Yes
	Process/Log Monitoring	Yes	Yes	Yes	Yes	Yes	Yes
Portability	Compliant Containers			Yes	Yes	Yes	Yes
	Enterprise Interfaces						
	Middleware Interfaces			Yes	Yes	Yes	Yes
	Cloud/Broker Interfaces			Yes	Yes	Yes	Yes
Elasticity	Database	Yes	Yes	Yes	Yes	Yes	Yes
	Load Balancer		Yes	Yes	Yes	Yes	Yes
	Provisioning Manager	Yes		Yes	Yes	Yes	Yes
	Elastic DNS	Yes	Yes	Yes	Yes		

Comment:

1. Google offers some of the PaaS elements as SaaS offerings, but those are not integrated into its PaaS offering.
2. Most vendors do not provide version control out of the box. Heroku uses Git for version control out of the box.
3. Heroku has the most add-ons and interfaces and provides good portability.
4. Other names considered but found significantly inadequate were Jelastic, CloudBees, Cloud Foundry, Apprenda, and AppFog.
5. Azure is highly compliant; however, it is not truly open, as some of the tools and underlying controls to computing are not open.

Google App Engine

The Google App Engine is highly scalable given that it is backed by Google's technology and infrastructure. However, it has many deficiencies when it comes to the proposed open PaaS standards framework. It is a highly secure PaaS, but this security comes at a cost. The platform offers little or no portability and does not allow the developer any lower-layer access. It also is not very efficient when it comes to ongoing development and recurring releases, as it has no continuous integration and deployment features.

Force.com

Much like the Google App Engine, Force.com also offers a high degree of scalability and reliability, as it uses the same infrastructure that powers Salesforce.com, the leading cloud SaaS CRM solution. It also has a metalayer that allows the developer's release cycle to be independent of the platform release cycles. This, while beneficial in some ways, also takes the control away from developers. Developers must program in Apex, which has a Java syntax, so it does not offer as much in terms of flexibility and openness. Force.com does not provide any meaningful portability, because it has no platform-level interfaces, nor does it offer any standardized application containers or IaaS portability.

Azure

Azure offers significant portability due to its private cloud offering and support for standard stacks and operating systems. Azure is great for developers because it provides support for almost all languages, works with a varied set of tools, and thus helps improve developer productivity. The major downside with Azure is that it is built using few or no open source components. Hence, if there is a platform bug, the developer is largely stuck and at the mercy of the vendor.

Heroku

Heroku is truly open, since it is based on open source components for all of its elements. It has a robust add-ons library, which makes it highly compliant with the proposed standards framework. Heroku has a worker thread–based pricing model, meaning that the developer is charged based on the number of processes running and the overall CPU time those processes use. If applications are not programmed efficiently, this model can turn out to be very expensive as compared to a VM-based pricing model (e.g., a runaway process that spawns hundreds of threads due to a software bug may cost much more). Heroku does not offer any IDE or enterprise interfaces, making it a less than ideal choice for large enterprise-class cloud implementations.

Stackato

Stackato is another truly open platform. It focuses mainly on the enterprise market and hence has good support for IDEs. While a good provisioning capability exists, it offers little or no granular control to the underlying infrastructure being provisioned. Stackato's major gap is inadequate support for elastic DNS, which makes it a problematic choice if the use case for your implementation is a federated or hybrid cloud.

Joyent

Joyent is written with a mix of open source components; however, the core PaaS Smart Machine platform is entirely proprietary. Like Stackato, it has a good provisioning capability, but it offers little or no granular control of the underlying infrastructure being provisioned. While Joyent offers only a moderate add-ons library, that is not such a big concern, as the native platform fulfills a lot of the proposed standard requirements out of the box. Lacking an IDE and elastic DNS, though, Joyent does not provide full joy for the developer.

A NEW WAY OF WRITING APPLICATIONS

The above definition of the PaaS standard and the subsequent discussion would be inadequate if I did not discuss what an ideal application sitting on such a standards-compliant platform would look like. This brings us to the thought behind the Cloud Architecture Reference Model for Applications, or CARMA (see Figure 3).

Any application must have the following core components in order for it to work efficiently with underlying open standards–based PaaS:

Functionality

This would encapsulate the core application-specific functions:

- **Logic —** business logic or business rules that are encoded as part of the application
- **Computing —** core computations that are highly process-intensive or need to be distributed over a grid-like environment and the provisioning rules to support these computations
- **Plumbing —** the functional components that help move the data around in order for efficient computing and efficient execution of the logic

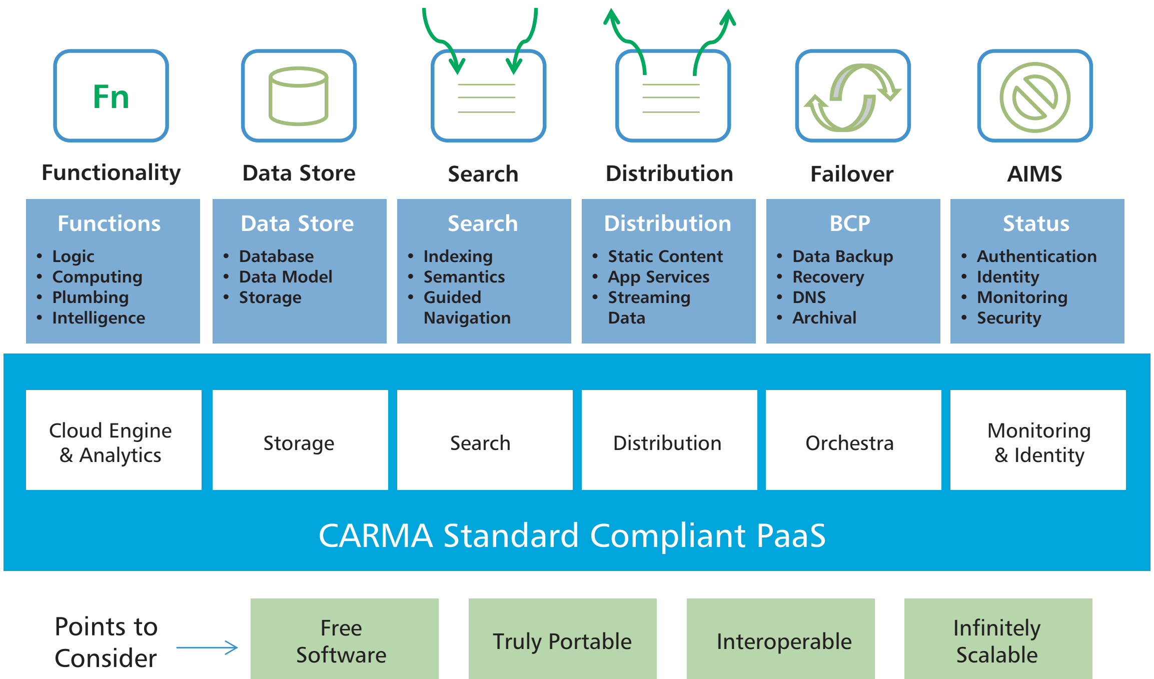

Figure 3 — A CARMA application reference model.

- **Intelligence** — the services that are core to executing the logic, by finding the most optimal computation given the constraints and dependencies of the available plumbing (data search, distribution, business continuity programs, and others)

Data

This would encapsulate everything to do with making the data more efficient:

- **Database** — the ability of the application to provide transactional integrity, also known as atomicity
- **Data model** — relationships, graphs, and metadata that tie all the data together
- **Storage** — the core aspects related to data namespaces; chunking, writing, and reading the data over a large distributed PaaS platform
- **Archival** — records management functions that specify data archival and retention policies that must be applied and the keywords they must be applied to

Search

This would include any function related to user searches performed on the stored data:

- **Indexing** — organizing the data into buckets or sorting it on a dimension so it can be accessed faster; specification of the rules for such indexing
- **Semantics** — the ability of the application to parse queries and make sense out of them in order to make the search more contextually relevant; specification of the core semantic logic and semantic services to be called
- **Guided navigation** — drill-down into the search results and recommendation-based insights; specification of the recommendation rules or recommendation service credentials

Distribution

This would include any function related to the serving of data to a user/application response to a user request:

- **Static content** — rules related to handling of static content and its location
- **App services** — rules related to handling of application-specific services, which allow scalable distribution of this data through the use of appropriate load balancing or other services
- **Streaming services** — functions related to efficient and optimized streaming of live streaming data and the location of such content

Failover

This would include any function related to ensuring that there is zero downtime or, in the event of a disaster, instant healing and recovery of the data or the application:

- **Data back —** automated synchronization of data on remote nodes as per the defined business continuity rules within this section of the application
- **Recovery —** automated recovery rules in the event of any application service or data failures
- **DNS —** automated load balancing sequence and reinstitution of application services in the event of complete failure of a physical location or network

A truly open standard would be one that gives developers greater control over the underlying source code and the computing.

Status

This would include all other functions that are not otherwise included in the application but are required for the application performance:

- **Authentication —** the permission and role-based access for the various features within the application
- **Identity —** the single sign-on to be used within the application and associated identity models within other applications
- **Monitoring —** key metrics and parameters to be monitored within an application across features and functions
- **Security —** compliance with the 13-point CARMA security framework, which consists of hardware-specific security settings (if any), extensions to existing security protocols used by the PaaS, session authentication, audit policy, password policy, DOS Policy, multichannel authentication settings, data segregation policy, data masking policy, file and data scanning policy, application-specific firewall configurations (if any), encryption level, and the encryption algorithm to be used

CONCLUSION

As we survey the cloud standards landscape, it appears PaaS is the area that needs most help. In this article, I have proposed an ideal PaaS standard reference framework called CARMA. A truly open standard would be one that gives developers greater control over the underlying source code and the computing. Such access comes with the tradeoff that it could lead to greater security risks. Organizations like the Cloud Security Alliance (CSA) are currently working on standards that can help allay some of the security concerns about a more open PaaS.

Comparing all the existing PaaS vendor services against the CARMA reference standard raises some interesting questions about the portability and interoperability of such platforms as the Google App Engine and Force.com. There is a distinct possibility that a de facto standard could emerge as platforms such as Heroku expand to support languages beyond Ruby and make the underlying access to computing more open. In April 2012 the W3C, the Internet standards organization, proposed a concept called Web Application Cloud Interface (WACI).[15] My colleagues and I are now working closely with the W3C to help develop the community around WACI[16] and to see how CARMA can be reconciled with WACI and made into a real standard that people will adopt.

ENDNOTES

[1]Stallman, Richard (Wikipedia).

[2]Babcock, Charles. "5 Worst Cloud Washers of 2011." *InformationWeek*, 13 December 2011.

[3]Parziale, Lydia et al. *TCP/IP Tutorial and Technical Overview*. 8th edition. IBM, 2009 (www.redbooks.ibm.com/redbooks/pdfs/gg243376.pdf).

[4]Watal, Aditya. *Cloud Basics*. 2nd edition. Kavan Life and Computing, 2011.

[5]Kern, Justin. "The State of Cloud Standards." *Information Management*, 10 February 2012.

[6]"Green Grid, Open Data Center Alliance to Drive Efficient Cloud Standards." *The Whir*, 9 November 2011 (www.thewhir.com/web-hosting-news/green-grid-open-data-center-alliance-to-drive-efficient-cloud-standards).

[7]"IEEE Launches Pioneering Cloud Computing Initiative." Press release, IEEE, 4 April 2011 (http://standards.ieee.org/news/2011/cloud.html).

[8]Hirschfeld, Rob. "Cloud APIs: LibCloud, DeltaCloud. and OpenStack." *Dell TechCenter Blog*, 13 June 2011 (http://en.community.dell.com/techcenter/b/techcenter/archive/2011/06/13/cloud-apis-libcloud-deltacloud-and-openstack.aspx).

[9]Brockmeier, Joe. "Amazon Taps Eucalyptus as Private Cloud Partner." *ReadWriteWeb/Cloud*, 21 March 2012 (www.readwriteweb.com/cloud/2012/03/amazon-taps-eucalyptus-as-private-cloud-partner.php).

[10]"Driver Functionality and Credentials." Apache Software Foundation (http://deltacloud.apache.org/supported-providers.html).

[11]Willmott, Steven. "Surveying 1000 APIs — Infrastructure and SaaS APIs Power New API Growth." 3scale, inc., 28 June 2012 (www.3scale.net/2012/06/surveying-1000-apis-infrastructure-saas-api-currently-powering-api-growth).

[12]Kanaracus, Chris. "Gartner: Software as a Service Market to Grow 17.9 Percent to $14.5 Billion." *Network World*, 27 March 2012.

[13]Synytskyy, Ruslan. "PaaS 2.0 Add Standards and Greater Developer Control." *Cloud Computing Journal*, 1 December 2011 (http://cloudcomputing.sys-con.com/node/2083113).

[14]Che, Bryan. "The 1st Tenet of Open Cloud: Open Is About Control." *JBoss Planet*, 15 February 2012 (http://planet.jboss.org/post/the_1st_tenet_of_open_cloud_open_is_about_control).

[15]"W3C Headlights Project — WACI: Web Application Cloud Interface." W3C, 5 April 2012 (www.w3.org/wiki/CloudComputing).

[16]"W3C Headlights Program — Cloud Computing Report." W3C, 4 May 2012 (www.w3.org/wiki/CloudComputingReport).

Aditya (Ed) Watal is a Senior Investment Manager with Kavan Life. Mr. Watal is responsible for Kavan's investments in Life and Computing verticals. He is also the founder of Guzool.com, a cloud-based workspace and cloud platform company that is the main supporter of the proposed CARMA: Open Cloud Standard. Mr. Watal is an Adjunct Assistant Professor at New York University (NYU), where he is responsible for helping define the cloud computing curriculum for its School of Continuing and Professional Studies. In his role as a true cloud evangelist, Mr. Watal writes actively for various leading publications and is the author of Cloud Basics*, the first book in the series* Cloud: The Paradigm Shift*. His second book,* Cloud Standards*, will be released later in 2012. Mr. Watal's contributions extend beyond cloud computing as well. Some of these include Watal's 9 Elements of Corporate Strategy, which bring fresh thinking to Michael Porter's Five Forces of Competition; and Watal's Innovation Curve, a body of work that further develops the work of Nobel Prize winner Simon Kuznets. Mr. Watal shares his thoughts through his blog, Watalon.com. Prior to joining Kavan Life, Mr. Watal worked as an analyst at major financial institutions and as a consultant with global system integrators, experiences that provided him with a rock-solid foundation for his research and writing. He can be reached at awatal@gmail.com.*

THE CUTTER DIFFERENCE

Access to the Experts

The processes Cutter employs to generate research and to tackle client problems on the consulting/training front reflect our view that fostering collaboration among experts is the most effective way to create valuable new ideas and to fashion effective solutions. The composition of our community is also based on the belief that a group of cognitively diverse experts is hard to beat when it comes to addressing tough business challenges.

This philosophy is part of what differentiates Cutter Consortium from other IT analyst and consulting firms. Our more than 150 internationally recognized experts have come together to offer research, consulting, training, and executive education to our clients; they are bound together by a commitment to deliver top-level, critical, and objective advice. They have done, and are doing, ground-breaking work in organizations worldwide, helping companies deal with issues in the core areas of business-IT strategy, enterprise agility and innovation, agile software development and product and project management, enterprise architecture, business technology trends, risk management, metrics, team building and leadership development, business intelligence, social networking/Web 2.0, and sourcing.

Cutter's goal is to help organizations leverage technology for competitive advantage and business success. We accomplish our mission by furthering the thinking in business technology, serving as a catalyst for collaboration between thought leaders around the globe, then giving clients access to this think tank through our products and services. We offer a different value proposition — *Access to the Experts*. You get practitioners' points of view, derived from hands-on experience with the same critical issues you are facing on a daily basis, not the limited perspective of a desk-bound analyst.

The Cutter community not only draws upon expert practitioners, but also innovative academics, allowing us to blend the frameworks provided by the top professors in key universities with the hands-on experience of best-in-class consultants worldwide.

Cutter's clients are able to tap into this collaborative, cognitively diverse community in a variety of formats, including online and multimedia research, mentoring and inquiries, events, and training/consulting services. With Cutter Consortium, you get the best practices and lessons learned from the world's leading experts — experts who are implementing these techniques successfully in organizations like yours right now.

The Cutter Difference

Cutter's research, inquiry response, consulting, and training are qualitatively different from that of other analyst firms in many vital ways:

- **Cutter Research gives clients Access to the Experts — their latest writings and thinking.** All of Cutter's research, both written and multimedia, is provided exclusively by internationally recognized expert practitioners. Cutter has no desk-bound analysts. Cutter's research allows clients to tap into this brain trust and get the latest thinking on the business-IT issues challenging enterprises worldwide.

For More Information

For more information on Cutter Consortium, contact us by:

Tel: +1 781 648 8700
Fax: +1 781 648 8707
Email: sales@cutter.com
Web: www.cutter.com
Mail: 37 Broadway, Suite 1, Arlington, MA 02474-5552 USA


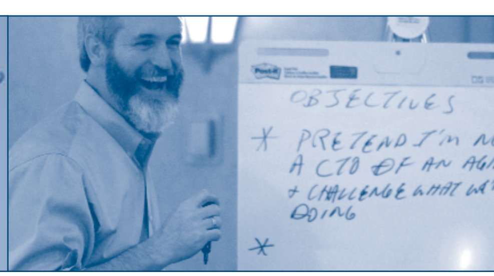

"Cutter's focus is on the user community — the people who use technology-related products and services to deliver business results. Vendor independence (particularly in terms of funding) means that Cutter's reports don't read like commercials. The advice given by Cutter is practical and pragmatic — it draws on a wealth of experience from a wide pool of 'hands-on' industry practitioners. Their breadth and depth of coverage is reflected in the frequency and format of reports — the underlying focus is on the supporting principles which enable success, not the latest fads."

— Paul Ramsay,
Service Delivery Manager,
Equinox,
Auckland, New Zealand

- **Cutter Inquiry Privileges — personalized guidance from the world's top practitioners.** Every inquiry is fielded by a Cutter Senior Consultant, Fellow, or Faculty member. Clients can purchase bundles of hours and may allot some of the inquiry time to senior executives, such as the CIO, pairing him or her with a top business-IT strategist so they can brainstorm every month. This allows the two of them to build a rapport and gives the Cutter expert a deep understanding of the issues the enterprise is facing. The Cutter expert quickly becomes a valuable advisor.
- **Cutter Consulting and Training — hands-on assistance from best-in-class consultants.** Unlike other consulting and training firms that use senior partners to sell the work but junior staff to execute, Cutter uses only its best-in-class experts for every assignment. The Consortium's great strength is that it can draw from its more than 150 internationally recognized experts to assemble the ideal team to help your enterprise tackle any business-technology challenge it faces.
- **Clients call Cutter "the thinking person's research firm"** because of Cutter's dedication to debate of the business-technology issues enterprises face and its success in attracting the thought leaders worldwide to conduct this debate. Clients get cutting-edge thinking plus multiple viewpoints so they can determine what's best for their individual enterprise/department/project.
- **Emphasis is on strategies and processes**, not on vendor/product detail.
- **Written content is likened to "consultancy in print"** since it provides hands-on, actionable solutions from expert practitioners who are successfully implementing these ideas, whether it be IT strategic planning, security strategies, or risk management. This relates back to the fact that the content is written by the people who are at the forefront of their fields and who are guiding companies daily.
- **Cutter Research — from reports to *Advisors* and webinars and podcasts — is prepared with both the IT and the business user in mind.**
- **Information and advice is truly objective.** Cutter is unique in the research space in having no relationships with vendors. It is well known that other analyst firms derive significant portions of their revenues from vendors and that the choice of which vendors to evaluate, and the slant of the research, can be influenced by this factor.
- **Emphasis is on agility across the enterprise.**
- **Focus is on sound project management, risk management, and change management** so IT can accomplish the changes business demands.
- **Emphasis is on IT as an enabler of innovation**, as the hyperdifferentiator that will give your business the competitive edge it needs to succeed.

About Cutter Consortium

Cutter Consortium is a truly unique IT advisory firm, comprising a group of more than 100 internationally recognized experts who have come together to offer content, consulting, and training to our clients. These experts are committed to delivering top-level, critical, and objective advice. They have done, and are doing, groundbreaking work in organizations worldwide, helping companies deal with issues in the core areas of software development and agile project management, enterprise architecture, business technology trends and strategies, enterprise risk management, metrics, and sourcing.

Cutter offers a different value proposition than other IT research firms: We give you Access to the Experts. You get practitioners' points of view, derived from hands-on experience with the same critical issues you are facing, not the perspective of a desk-bound analyst who can only make predictions and observations on what's happening in the marketplace. With Cutter Consortium, you get the best practices and lessons learned from the world's leading experts, experts who are implementing these techniques at companies like yours right now.

Cutter's clients are able to tap into its expertise in a variety of formats, including content via online advisory services and journals, mentoring, workshops, training, and consulting. And by customizing our information products and training/consulting services, you get the solutions you need, while staying within your budget.

Cutter Consortium's philosophy is that there is no single right solution for all enterprises, or all departments within one enterprise, or even all projects within a department. Cutter believes that the complexity of the business technology issues confronting corporations today demands multiple detailed perspectives from which a company can view its opportunities and risks in order to make the right strategic and tactical decisions. The simplistic pronouncements other analyst firms make do not take into account the unique situation of each organization. This is another reason to present the several sides to each issue: to enable clients to determine the course of action that best fits their unique situation.

For more information, contact Cutter Consortium at +1 781 648 8700 or sales@cutter.com.

The Cutter Business Technology Council

The Cutter Business Technology Council was established by Cutter Consortium to help spot emerging trends in IT, digital technology, and the marketplace. Its members are IT specialists whose ideas have become important building blocks of today's wide-band, digitally connected, global economy. This brain trust includes:

- Rob Austin
- Ron Blitstein
- Tom DeMarco
- Lynne Ellyn
- Israel Gat
- Tim Lister
- Lou Mazzucchelli
- Ken Orr
- Robert D. Scott