Systems Engineering at MITRE

Cloud Computing Series

Ahead in the Clouds: *An Online Forum for Government*

Archival Content from January 2010 to January 2011

MITRE

MITRE

www.mitre.org

Ahead in the Clouds: Foreword

Cloud computing is an approach for deploying capabilities on private, community- shared, and public- shared information technology (IT) environments. Enabled by new technologies, it holds promise for more efficient use of IT resources, yielding significant cost and capability advantages in some circumstances. However, a cloud computing approach presents new acquisition, management, and security challenges as well as financial trade- offs for those responsible for IT decision making. Private clouds can offer some advantages over traditional IT while maintaining maximum control, while community and public clouds can provide the ability to leverage the provider's capital investment in return for "pay as you go" operational expense.

Special considerations arise when employing cloud computing in the DoD and Federal IT domain. DoD and Federal IT leaders are faced with a diversity of constraints that are unique to Government, including acquisition and funding models, particular cultural characteristics, and sophisticated adversaries. These constraints are coupled with the need for significant scale, assured performance, minimal downtime, and security needs that frequently exceed standard commercial practices. Many DoD and Federal programs are already wrestling with these considerations.

In the MITRE cloud computing series of white papers, we have addressed many of these issues to help Federal leaders understand the benefits along with the challenges. We offer insight to help senior leaders put their organizations on a course to realize the benefits of cloud computing while minimizing the risks. The series includes the following papers:

- A Decision Process for Applying Cloud Computing in Federal Environments provides an overview of cloud computing and a roadmap for a structured decision process.
- Cost and Business Case Considerations comprehensively explores the costs to examine when employing a cloud- based approach in Government.
- Cloud SLA Considerations for the Government Consumer describes the elements that should be considered by consumers as they examine service- level agreements (SLAs) to meet their requirements.
- Database as a Service: A Marketplace Assessment compares four public database as a service offerings from Amazon (2), Google, and Microsoft.
- Information Security in the Clouds addresses the security challenges presented by cloud computing with a focus on public and community deployment models.

- Leveraging Public Clouds to Ensure Data Availability addresses potential uses of cloud computing approaches to facilitate data access in the event of a failure in a primary location.
- Platform as a Service (PaaS): A 2010 Marketplace Analysis surveys many exemplar offerings and identifies considerations for using unique PaaS capabilities.
- Products to Build a Private Cloud identifies market segments and categories of products available on the marketplace for contemporary private clouds.

The efforts to address the challenges of cloud computing- based systems are significant. MITRE is collaborating with industry and other organizations supporting the Federal Government, along with leveraging lessons learned and our own research, to provide the best advice possible. We hope that you find the materials helpful, and if you have feedback or suggestions, please contact us or visit For more information about MITRE's cloud computing program, including all the papers referenced above, we invite you to visit www.mitre.org/capabilities/advanced-technologies/information-systems/cloud-computing.

Editor's Note: The "Ahead in the Clouds" forum ran from January 2010 through January 2011. This material is being offered in archive format, and no updates are planned. The posts appear in reverse chronological order (newest first), as they originally appeared on www.mitre.org. For more information about MITRE's cloud computing program, we invite you to visit http://www.mitre.org/capabilities/advanced-technologies/information-systems/cloud-computing.

Archives
Question for January 2011
Question for November/December 2010
Question for October 2010
Question for September 2010
Question for August 2010
Question for July 2010
Question for June 2010
Question for May 2010
Question for April 2010
Question for March 2010
Question for February 2010
Question for January 2010

Question for November/December 2010Technical Papers
The Cloud Computing Series

Question for January 2011

The Office of Management and Budget's 25 point plan describes a "cloud first" policy for the Federal Government. Is the approach described in Part I, Achieving Operational Efficiency, sufficient to deliver more value to the American taxpayer? What are the strengths or gaps in the plan regarding the use of cloud computing and what types of capabilities should be moved to the cloud first (e.g., within the first 12 months)?

Harry J. Foxwell, Ph.D.
Principal Consultant
Oracle Public Sector

US CIO Vivek Kundra's "25 Point Implementation Plan to Reform Federal Information Technology Management" is ambitious in its scope and timeline. Although cloud technologies are maturing rapidly, understanding of the benefits, risks, and costs of this approach to IT is evolving slowly. Clearly there are cost-saving efficiencies already being delivered through data center consolidation, virtualization, and massively-parallel, energy efficient, multi-core servers and integrated systems. Further exploiting these technologies to fully implement the NIST model of public and private cloud infrastructures will require not only significant technology changes but acquisition and management policy changes as well. The 25-Point plan's focus on identifying and developing government expertise and developing industry partnerships are essential first steps.

Efforts are currently underway within multiple government agencies to fulfill some of the 25 Point goals related to "commodity IT services" such as government-wide "cloud email". While even this project is a major undertaking, it is among the most feasible of the many candidates for initial cloud deployment. Converting special purpose and highly customized agency software to the cloud model will be much harder and will take significantly more time.

Cloud computing as a technology delivery model and as a business model do have the potential to provide significant cost savings and taxpayer value. But, as attractive as this may seem, its benefits should not be oversold nor its costs and risks underestimated. Additionally, although the "cloud first" policy is well intentioned, other data center consolidation technologies should not be overlooked.

For more information http://blogs.oracle.com/drcloud/.

Posted: January 28, 2011

Ron Knode
Director, GSS, LEF Research Associate
CSC

New Wine in Old Wineskins?

The "cloud first" policy declaration in OMB's 25-point plan of 9 December 2010 is aggressive thinking and terrific branding. The triple-play promises of economy, flexibility, and speed are precisely the kind of IT payoffs that any enterprise would want.

However, these promises are themselves based on another promise in the same plan, i.e., the promise of a cloud strategy that can deliver safe and secure cloud adoption across the U.S. government. While there is much to like about the ambitious vision and the no nonsense "let's get going now" message for cloud processing in the plan, real success hinges on making the underlying promise of a practical cloud strategy come true. That promise is the more difficult one. It must respond not only to the needs and realities expressed by (government) cloud consumers, but also to the needs and realities of cloud service providers who can actually deliver these payoffs. Only when both constituencies are accommodated in strategy and mechanics can we move from a hit or miss "Ready, Fire, Aim" process to a reliable "Ready, Aim, Fire" process for cloud adoption and payoff.

And, there's the rub. According to OMB plan, the promise for a practical cloud strategy is rooted in the development of standards for cloud service security, interoperability, and portability. The initial public draft of the Proposed Security Assessment & Authorization for U.S. Government Cloud Computing took a healthy first swing at such standards, but does not yet tend to the needs of all the constituencies involved. Continuing ambiguity about overall risk governance and accountability, a monitoring framework that excludes the cloud consumer, and a complicated scheme for trying to shape Spec Pub 800 53 for cloud services all present high hurdles to overcome.

One cannot but wonder if the biblical admonition against "pouring new wine into old wineskins"[1] must be observed here. Trying to bend the conventional machinery for C&A into a community process for A&A" without clarifying who is accountable for risk acceptance in cloud services only slows cloud adoption. The attempt to fashion existing Spec Pub 800 53 controls into a set of requirements suitable for cloud processing is laudable, but does not suit the consumption model of the cloud. In other words, the old wineskins of traditional C&A models and Spec Pub 800 53 cannot yet handle the new wine of cloud processing.

Until we fulfill the promises made in the OMB plan, we will be constrained to applications that satisfy the compensating techniques first introduced in Digital Trust in the Cloud and subsequently amplified in other places. We can gain some benefit from "safe" applications like non-sensitive email, development and test, backup and restore, and even a bit of collaboration and social networking. But, such applications do not deliver the kinds of payoff we need and expect from cloud processing.

In his earlier blog on this matter Chris Hoff declared we're gonna need a bigger boat." Simply enlarging the vessel may not be enough. The biblical warning declares that “both the wine and the skins will be ruined”1 if we try to pour new wine into old wineskins. The new wine of cloud processing may well need completely new wineskins (standards and practices) for us to enjoy the rich bouquet of enterprise payoffs.

See the full blog response at http://www.csc.com/cloud/blog.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: February 1, 2011

Peter Coffee
Director of Platform Research
salesforce.com inc.

When we answer the call for greater operational efficiency in IT operations, we should heed the warning ascribed to Peter Drucker: "There is nothing so useless as doing efficiently that which should not be done at all." Improved execution of current task portfolios is not enough: we should further strive to eliminate, or at a minimum delegate, any activity that does not directly contribute to mission performance. Tens of thousands of organizations use massively scalable multi- tenant services ("public clouds") to pursue that course successfully today.

U.S. CIO Vivek Kundra quickens the pace with his vigorous mandate to consolidate at least 800 data centers by 2015. This goal has the crucial merits of being countable, achievable, and uncomfortable. This goal will not be achieved by picking the low- hanging fruit of redundant or obsolete systems that are readily and painlessly retired as soon as someone decides to do so. Meeting Kundra's challenge will require fresh thinking about who performs what functions, and who needs to own what capabilities but it will not require lowering our standards for what constitutes satisfactory performance.

Indeed, the National Institute of Standards and Technology has urged us all to treat the move to the cloud as an opportunity for substantial improvements in IT reliability and governance. In its newly released draft document, Guidelines on Security and Privacy in Public Cloud Computing," NIST correctly asserts that.

Potential areas of improvement where organizations may derive security benefits from transitioning to a public cloud computing environment include the following:

- Staff Specialization: opportunity for staff to specialize in security, privacy, and other areas Platform Strength: Greater uniformity and homogeneity facilitate platform hardening and enable better automation of security management
- Resource Availability: Redundancy and disaster recovery capabilities are built into cloud computing environments
- Backup and Recovery: Data maintained within a cloud can be more available, faster to restore, and more reliable

- Mobile Endpoints: clients are generally lightweight computationally and easily supported
- Data Concentration: less of a risk than having data dispersed on portable computers or removable media

This list can aid us in choosing our targets for rapid cloud adoption. We should look for tasks requiring maximum speed and flexibility in deployment to mobile personnel or to frequently relocated sites. We should look for tasks requiring access to large collections of data, but using focused subsets of that data in typical situations. We should look for tasks requiring precise grants of privilege, and rigorous accountability for who has done what with sensitive information. All of these are criteria for which the cloud does not merely meet expectations, but rather elevates the standard of practice as widely demonstrated by enterprise customers today.

CIO Kundra's challenge comes at a time when technical transformation coincides with cultural readiness to consider dramatic change. Tightening resource constraints, combined with broad and growing public adoption of cloud services in both workplace and personal activities, create a powerful push- pull incentive to act and a basis for confidence in the outcome.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: January 4, 2011

Kevin Paschuck
VP, Public Sector
RightNow

'Cloud First'—An Important Move in the Right Direction

Federal CIO Vivek Kundra's 25- Point Implementation Plan to Reform Federal IT Management, is an important move in the right direction. With cloud technology positioned prominently at the center of the initiative, we are beginning to see a real shift toward recognizing the major benefits, including significant cost savings and decreased implementation times, that government can realize from cloud- based solutions.

The plan outlines a 'Cloud First' policy, which mandates that each agency identify, within three months, three 'must move' IT services and move one of those services to the cloud within 12 months. The remaining services should transfer to the cloud within the next 18 months.

Additionally, approval is reserved for major IT programs that utilize a modular approach, with new customer facing functionality provided every 6 months.

This is an important component and also addresses President Obama's Memorandum on Transparency and Open Government , issued on January 21, 2009. In this memo, the President outlined the Administration's commitment to creating an unprecedented level of openness in Government and instructed the heads of executive departments and agencies to work together to

ensure the public trust and establish a system of transparency, public participation, and collaboration. Cloud technology can help federal agencies comply with this mandate.

To deliver on the promise of open government and the plan to reform federal IT, agencies must identify services to transfer to the cloud. Specifically, Web Self Service applications and Pilot Programs are a good starting point to identify the best solutions for specific agency needs.

Coupling cloud solutions with Web self- service applications is an effective means to simultaneously improve constituent services and reduce overhead costs. With Web self- service, constituents can find information that they need on an Agency website quickly, without having to contact a live person. Additionally, the cloud provides Federal agencies with several benefits:

- Lower total cost of ownership
- Benefits from frequent solution innovation
- Increased reliability
- Speedy, measureable results on open government initiatives

Whether in the public or the private sector, identifying the appropriate IT solutions can be a daunting task. For this reason, working with vendors that provide pilot programs is a critical component in the decision making process. One of the unique things about cloud computing is the ability to test the solution first— before signing a contract. Identifying proof points and results up front, prior to making a large investment, is critical to ensuring success.

Cloud solutions provide the scalability that government agencies require to meet constituent needs—eliminating digital capacity limitation worries. By tra nsitioning to the cloud, agencies tap into an infrastructure that is as flexible as their needs are varied. Undoubtedly, these are some of the primary reasons why cloud is positioned as the center stone of the Administration's plan.

Posted: February 10, 2011

Question for November/December 2010

Given the rapid expansion of mobile computing devices such as tablets and smart phones, how do you see cloud computing technology enabling capabilities, such as location independent access for users, on these devices? Please identify the best uses for this technology and approaches for the government, taking into consideration security and privacy concerns.

Scinivas Krishnamurti
Sr. Director, Mobile Solutions, CTO Office
VMware

Enterprises have traditionally favored homogeneity since it enabled them to easily manage a huge fleet of devices deployed to their users. I'll call this stage Client Computing 1.0. This meant that enterprises typically standardized on as many aspects of their client strategy as possible including the hardware, OS, applications, application development frameworks, and management frameworks. The management paradigm was around managing the device and its contents. Unfortunately, the homogeneity that enterprises crave is slowly but surely disappearing.

Enter Client Computing 2.0 where almost every single vector in client computing is changing. Gone are the days of enterprises building Windows applications in Visual Basic or .Net. Many enterprises are embracing web applications, either hosted internally or in a SaaS delivery model. This trend is expected to continue at the expense of local thick- client applications. However, for the foreseeable future, Windows applications will coexist with web applications in many enterprises, so the underlying infrastructure needs to be flexible enough to support both traditional as well as emerging devices.

On the mobile phone application use is changing dramatically. In the past corporate mobile phones were synonymous with accessing email and calendar anytime anywhere. While this paradigm led to productivity increases, enterprises are realizing that more applications could and must, be mobilized to give employees the freedom to realize the true potential of mobile devices. Employees are buying cool and capable PCs, phones and other emerging devices for personal use and actually preferring to use them instead of the corporate issued device. Consequently, Macs, iPhones, Android phones and iPads are now entering the enterprise and the demand to support them is increasing.

As multiple devices enter the enterprise, each one brings yet another operating system into the enterprise. In Client Computing 1.0 era, most employees were allocated just one device either a desktop or a laptop. Today many employees are also given a mobile phone and tomorrow, employees may also carry a tablet. The expectation is that they will use these devices interchangeably to access the applications they need at any given point in the day. They may use

the desktop at work, cell phone during long, boring meetings, laptop/ tablet during the train ride to/from work.

Consequently, CIOs are faced with having to manage a very complex and heterogeneous environment that includes different device types, operating systems, processor architectures, sizes and shapes. As if the above matrix is not complicated enough, the additional dimension is the ownership of the device corporate owned vs. employee- owned. The one thing that has not changed is the need to comply with various government requirements (FISMA, SOX, HIPAA, etc.). The old management paradigms of managing the device just do not work in such a diverse environment. Going forward, solutions must offer Application Access providing the right set of applications to users irrespective of the device; Security securing enterprise applications, services and data; and, Scalability heterogeneous management of multiple different end points.

The challenges are many but also is the potential for increases in effeciency. The opportunity is how to merge the security and reliability of the 1.0 generation with the flexibility, mobility and choice of the 2.0 generation. You can read more about VMware's approach to these Emerging Devices at my blog.

Posted: December 14, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

As 2010 draws to a close, it's being suggested that the "Wintel axis" of the 1990s and 2000s (Windows operating system, Intel processors) is being overtaken by the "Quadroid" alliance (Qualcomm chip sets in Android- based devices: smartphones, tablets and other form factors yet to emerge). The dominant means of information delivery and of business and personal interaction will be, not the thick- client desktop or laptop PC, but a panoply of network- edge devices that rely primarily on the cloud to store data and run applications as dramatized by a Google video showing repeated (and perhaps improbably catastrophic) destruction of a series of Chrome OS netbooks, with no impact on a user's work in progress.

Going into 2011, it's clear that two of the most important use cases in public- sector computing are ideally served by the strengths of the cloud. First, there are tasks that fall to the government requiring rapid deployment in response to legislative initiatives. Economic stimulus programs, for example in the area of health information technology, have demanded implementation schedules far more rapid than those traditionally associated with massive government programs and the cloud has delivered. Second, there are scenarios such as disaster relief in which the government must be among (or even in charge of) first responders, with unpredictable but absolutely urgent timing of service delivery into inconvenient locations and again, the cloud has delivered, with multi- agency coordination in the field and with peak- load capacity for financial and logistic support in situations such as the devastating earthquake in Haiti.

Regardless of urgency, whether driven by man or by nature, there can be no slackening of attention to security or robustness in such situations: indeed, it is unfortunately true that low

barriers to entry into the cloud have facilitated fraud as well as enabling real aid. Fortunately, the apparatus of the cloud is increasingly being recognized as enabling rigorous security practices, and affording access to top- tier security and governance tools, that once were beyond the reach of resource- starved agencies in public- sector and non- profit domains.

With recent high- profile commitments to a "cloud first" strategy in the U.S. federal sector, we may hope that 2011 will bring increasingly confident use of the power of the cloud to serve the public interest.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: December 17, 2010

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP

Mobile Computing and the Cloud

Mobile computing and the Cloud can form the perfect storm for revolutionary change in the way that agencies do business. Individually each of these technology advancements shows great potential, but together they are creating possibilities for tremendous opportunity. The benefits are only limited by what we can think of.

Now, we don't want to get ahead of ourselves and must recognize that there are important and significant problems that must be solved. Privacy and security are certainly among these problems. Security is at the top of everyone's list in both of these technologies. But in one way, using the Cloud can help a mobile workforce with their security profile. For example, sensitive information can be stored on the Cloud and not on the device, thus reducing the risk of loss of data. Such an approach can be used even if the mobile device has a remote kill pill capability. The reason this is helpful is because there may be a lag in time before one discovers that their mobile device is missing.

Think of the ability to have any information you need, at any time, in any place. Examples of the potential range from the more mundane like being able to read reports on the go in a secure, ubiquitous and fast manner or being able to have your favorite music list follow you around while you change devices from home to the car to your office.

At the other end of the spectrum are examples of law enforcement situational awareness. It is possible to provide details that can help tactical operations come to a safe and successful conclusion. The stuff you see on the television series *24* is beginning to be possible. We can now use smart phones to collect and/or show surveillance video in a very unobtrusive way. Such uses do not require specialized equipment, but can take advantage of relatively inexpensive off the shelf mobile devices.

Think of the situation where an undercover operation is taking place and the undercover operative is dealing with some unknown people. Video or still photographs can be taken

unobtrusively and this information can be sent to the office for analysis against a known database, or even better, can be used with crawlers to check social media sites in order to identify the individuals. This gathered information can then be sent back to the operative. Again, the device in the hands of the operative is just an off-the-shelf mobile device.

Mobility and the Cloud go nicely together. They can very much complement each other and extend their capabilities. Many of the examples I have used could be achieved without using The Cloud or off the shelf mobile devices, however, these two technologies make the possibilities even greater and should be explored.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: December 22, 2010

Seth Landsman, Ph.D.
Lead Software Systems Engineer
MITRE

As mobile networks approach desktop quality network speeds there will be a compelling argument for a marriage between cloud computing and mobile computing. Cloud computing can be an enabler for mobile devices to have access to vast amounts of storage, processing capacity, and information *when needed and where needed*, enabling the *truly mobile, always available* promise for government users, as well as citizen consumers of government services.

Cloud computing is not new to mobile. Aspects of it have existed in the mobile world for many years, to delegate processing or storage to backend servers that are more capable than the mobile device. As an example, Research in Motion (RIM), makers of the Blackberry® series of devices, delegates the management and control of devices to data center servers. Opera Software™, the web browser company, provides a similar delegation example for web requests.

The integration of new cloud-based mobile capabilities will explode in the next several years. As mobile devices both smart phones and tablets replace the laptop as the tool of choice, the need for government to leverage the cloud for applications, storage and processing is going to become essential. As an example, Google Apps and Microsoft Office® 365 (Beta) provide cloud-based office applications that can be accessed from a browser over a network. But, this is just a start. "In the past corporate mobile phones were synonymous with accessing email and calendar anytime anywhere," writes Srinivas Krishnamurti. He adds, "While this paradigm led to productivity increases, enterprises are realizing that more applications could and must, be mobilized to give employees the freedom to realize the true potential of mobile devices."

Mobile applications can enable significant operations and business process advantages if the new capabilities are aligned with the goals of the organization and mission assurance is considered. Before government can fully realize the value of cloud-based mobile computing, network availability, service reliability and security are non-trivial factors that must be addressed. As Peter Coffee cautions, "There can be no slackening of attention to security or robustness…" Further, as these devices are expected to be more capable, their complexity will

continue to increase. The reliability and availability of the mobile platform, coupled with the network and cloud services will determine the ability of citizen or government users to interact with government systems, especially when the need is time critical. If the network, device, or cloud service is unavailable, due to lack of security, robustness, or other factors, the user may not be able to access essential capabilities. Given these risks, Federal IT leadership should consider all of these implications in planning their enterprise, and develop an appropriate IT architecture, policies and procedures to mitigate these factors, depending on the intended use of these technologies.

By the end of 2011, 1 in 2 Americans will have a smartphone . As these devices become pervasive and more applications become embedded within the cloud, more government business and more capabilities will be accomplished with the combination of these technologies. As such, cloud computing and mobile technologies coupled with appropriate risk mitigations will provide government and citizen users with more capability, increased mobility and, ultimately, enable them to leverage government services more effectively, regardless of where they are and what means of access they may have at their disposal.

Posted: January 7, 2011

Question for October 2010

For Federal IT leaders considering building a business case for a cloud computing investment, please identify the general cost categories/drivers to include in a business case, and if possible, suggestions on approaches for attributing value to new cloud features.

Douglas Bourgeois
VP, Federal Chief Cloud Executive
VMware

This is a really good question because it considers the overall value of the cloud beyond simply cost efficiency which is an important part of the value equation. As most are now aware, virtualization has become widely accepted as a key enabler for cloud computing. Infrastructure virtualization provides a significant means of achieving cost efficiency through increased asset utilization. So, the key driver there is the consolidation ratio. In my experience, another key driver of the business case is the VM density. As you know, not all servers are created equal and so it follows that not all virtualized servers are created equal either. In my experience, from a financial modeling perspective, VM density can be a major variable in a cloud cost model. The license cost of the software included within the cloud service offering can be another major driver. Some software products are more affordable than others and some software licensing models are more compatible with cloud computing than others. These structures can make it very difficult to get started in the cloud especially if software acquisition costs are allocated over a small, initial cloud customer base. In effect, the cloud economies of scale can work against you until sufficient scale is achieved.

The broader question of value to be derived from the cloud, of course, includes cost efficiency but does not stop there. In addition, the cloud service offering should be carefully considered and specifically selected to provide the most value to end users. One of the challenges is that the services of most value to an organization will vary depending upon the mission and capabilities of that organization. The best practice is to identify those widely utilized and common services that would be good candidates for migration to a cloud model and would therefore draw high usage throughout the organization. This widespread potential for adoption will accelerate the efficiencies as usage increases. The final "piece" of the broader value proposition from the cloud can be associated with service levels perhaps the most important of these is speed. One of the key features behind the end user interest in the cloud is customer self service. This capability, in itself, is not the appealing factor. Rather it is the underlying use of standards and technology to automate the processes for service provisioning that is appealing. The considerable potential to radically reduce cycle times for the provisioning of services is a major component of the overall cloud value proposition.

Since cost efficiency is the easiest to measure and budgets are tight and getting tighter, there is considerable attention given to this key driver. Don't lose sight of the fact that there is also

considerable value to be derived from the selection of cloud services as well as the speed in which cloud services are delivered. This latter component speed is the one that will "wow" your end users the most and perhaps have the biggest impact on changing the perception of the IT organization on the journey to becoming a service provider.

Posted: October 14, 2010

Nathan Rushfinn
Certified Enterprise Architect
CA Technologies

The promises of cloud computing can be nebulous. To build a business case, federal IT leaders need to balance costs of new capital expenditures with reduced operating expenses. They must also be able to measure the success of cloud computing from the viewpoint of the customer.

To realize the benefits of cloud computing, the cost of capital expenditures should be offset by reduced operating expenditures over time. Cost categories for capital expenses should include all of the hardware, software and installation costs to implement new cloud technologies.

Cloud computing will drive the adoption of open source software, reducing costs for operating systems, software development stacks, and applications like Appistry and CA 3Tera AppLogic. IT leaders should also carefully track capital expenditures for systems integration costs related to installation, configuration, and training.

The best way to track operating expenses is to use project portfolio management software (PPM) and track all expenses as services. Projects should be clearly defined so that cost codes can be assigned and broken out by specific tasks in the work break- down structure (WBS). Labor costs must be tracked for both employees and contractors and broken out for each FTE (full time equivalent). Operating expense categories should be tracked by service and should include- time to deliver, support costs, infrastructure, and electricity. While some operating expenses like electricity can be tracked against specific servers, many expenses like HVAC and floor space will have to be calculated.

When building a business case for cloud computing, it is especially important to quantify success from the customer's perspective. A short survey taking no more than two minutes can accomplish this. For example, a customer might be asked to rate a statement such as "I find it much easier to order IT services through the new self- service cloud computing portal" using a five- point Likert scale consisting of strongly agree, agree, neutral, disagree, and strongly disagree. Response rates of 50% or more can be interpreted with confidence. Follow- up reminders and incentives, such as a random drawing for a gift certificate, are good ways to increase response rates. Sample categories to include in any survey on cloud computing should include: reduction in time to deliver services; ease of use of ordering; improved confidence in IT; and reliability of delivered services.

There are many drivers for implementing cloud computing, and while initiatives or mandates do not require a ROI, a business case does. By clearly defining costs categories for both capital and operating expenses, and by using well- defined customer surveys, federal IT leaders can estimate

the success of cloud computing projects from both an ROI perspective and from a customer's vantage point.

Posted: October 25, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

There's no question that cloud computing can be amply justified on grounds of reduced IT cost. That doesn't mean that cost- based justification is the best way to drive a cloud initiative.

Cloud computing both reduces and re allocates the cost of managing data and supporting processes. In one widely cited study, Bechtel Corporation benchmarked its internal costs against its best estimates of the costs of cloud service providers. Across the board—storage, network, server administration, and application portfolio maintenance—the Bechtel estimates favored large- scale cloud providers by ratios on the order of 40 to 1. Economies on this scale are not merely attractive, but compelling.

Other IT costs are less readily measured, but perhaps even more vital to contain. Workers in many organizations report continual distraction and loss of time due to every individual user of a thick- client PC needing to serve, in varying degrees, as his or own system administrator. Tasks of accepting and activating software updates, managing mailbox quotas, and protecting thick-client systems from increasingly aggressive security threats demand effort from individual users but do not serve organizational missions.

Further: the acquisition and deployment costs of conventional IT will almost always precede, often by months or years, the realization of value from those investments. Hardware purchase, facility construction, software licensing, and labor- intensive custom application development divert present resources to deliver future value that is, in the best case scenario that a project achieves its goals on specification, on schedule, and on budget. Many projects are placed at grave risk by the growing complexity of the technology and the dynamic nature of the problems being solved: a recent federal analysis found 72% of major projects to be considered at serious risk.

Cloud systems align the cost incurred with the value received, sometimes on the scale of yearly or monthly subscriptions ; sometimes at the scale of hours (or fractions of hours) of service received , or bytes of data transferred or stored. Services that are evaluated on a discount or free trial basis are the services that will be used in production, not approximations of a future on premise configuration. Cloud- delivered applications, including custom applications as well as configured versions of packaged applications, are frequently developed and deployed in days or weeks.

But even this is an argument based on costs, when often the far more powerful justification is in the value to be gained by pursuing projects that today are deferred due to excessive cost or delay of any realistic availability date. The Bureau of the Census, for example, did not use a cloud

database to save money, but to meet a deadline that's in the Constitution and was not looking likely to be met by on premise technology.

Justification of cloud projects should therefore begin with expected improvements in cost, reductions of risk, and accelerations of service availability, but should not stop there: they should also make reasonable projections, based on growing collections of relevant examples, of the value of improved mission performance.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: October 29, 2010

Teresa Carlson
Vice President
Microsoft Federal

This is a question that every government technology leader must deal with when evaluating cloud computing options. What's the ROI Is this going to save us money? The short answer is unfortunately "maybe". In general, cloud computing offers cost benefits through increased efficiencies, pooled IT resources and "pay- as you- go" models. But when making the business case it's important to distinguish between different types of cloud offerings, because matching the unique needs of an organization to the right type of solution is the best way to maximize ROI.

The first step is identifying the right cloud level to implement at whether it's at the infrastructure level, the platform level or the software/application level. For example, the GSA recently announced that government agencies would be able to access Infrastructure- as a- service (IaaS) offerings through Apps.gov. IaaS options are great for agencies that want to get out of the business of buying servers, data center space or network equipment. It's an entire IT infrastructure in a pay- as you- go model, but it still requires general administration and maintenance.

For agencies that want to remove IT maintenance completely, SaaS is the way to go. SaaS allows organizations to consume finished applications on demand, and is typically far less expensive than software that includes a licensed application fee, installation and upgrade costs. Now if an organization has internal developers with the skills to build customized applications, Platform as a- Service (PaaS) becomes the best option. Government is seeing an explosion of Gov 2.0 application development for improving citizen services, and PaaS provides developers with the tools they need to test, deploy, host and maintain applications in the same environment.

Organizations have options and each model follows the same basic ROI principle you only pay for what you use. A pay- as you- go model combined with very limited upfront costs creates a low risk environment where organizations have the freedom to innovate. If an application or program is successful, cloud offers the scalability and elasticity to incrementally grow as needed. If a program or application doesn't catch on, the upfront investment was already extremely low.

For example, it's interesting to think about how a program like Cash for Clunkers may have been different in a cloud- based model.

Every organization has to crunch its own numbers to evaluate the cloud solution that makes the most business sense, but the number of cloud options and reduced implementation risk make the current IT environment ripe for innovation. That freedom should be factored into any ROI discussion.

For more information, please see Teresa Carlson's FutureFed blog at http://blogs.msdn.com/USPUBLICSector/.

Posted: November 1, 2010

David Mihalchik, Jim Young (pictured)
Google

Why the Cloud Makes Good Business Sense

Cloud computing offers the federal government an unprecedented opportunity to access more powerful, modern technology with constant innovation at a substantially lower cost. Similar to the existing practices of many businesses and government agencies who outsource functions like payroll, shipping, and helpdesk support it makes good business sense to use a cloud provider who offers better applications with government FISMA compliant security at a lower cost than an organization can provide on its own.

By taking advantage of the scale at which cloud providers operate, organizations using cloud-based applications drive down their own costs substantially. In fact, a recent Brookings Institution study found that agencies moving to the cloud can cut costs as much as 50%. The three main areas in which the cloud offers cost savings are labor, hardware and software. The primary driver of cost savings is the reduced amount of employee time spent patching and maintaining servers and software applications. This labor can instead be applied to the government's more mission critical systems. By using systems operated by cloud providers, agencies can decrease hardware costs and the associated costs of real estate, electricity, and more required to operate servers in an organization's own data centers. Additionally, instead of the traditional model of an upfront software licensing cost plus a recurring annual maintenance fee, cloud computing applications are paid for via an annual subscription fee. In addition to providing cost savings, this model offers both predictability and flexibility, as organizations evolve or change in size.

Harder to measure are the soft cost savings associated with cloud computing. Ubiquitous access, increased productivity and better security are all worth something to cloud users, but are not always easy to value. With cloud computing, employees can access their information anywhere they have access to an Internet connection, whether at work, home, in the field, or on travel. The cloud also makes people more productive by making it easier to collaborate with fellow employees and locate an organization's historical information, lessons learned, and improve organizational knowledge management. And if users ever lose a laptop or mobile device, with

their data stored in the cloud they can be back up and running in no time; not to mention the benefit of limiting the organization's risk of such a lost device. On the security front, in many cases cloud providers offer security capabilities such as redundant data centers for failover that would be prohibitively expensive for organizations to build on their own. All of this must be considered when building a business case for moving to the cloud.

Government agencies are already benefiting from moving to the cloud. Take for example Lawrence Berkeley National Laboratory. By moving to Google's cloud- based email and collaboration tools, Berkeley Lab expects to save in hardware, software and labor costs, while increasing email storage and improving collaboration tools. (See Government Computer News article for details.) With these results, agencies should take a serious look and independently assess the business case including mission, operational, and financial, plus workforce trends for user expectations in the workplace, for moving some of their applications to the cloud.

For more business case ROI information, see http://googleenterprise.blogspot.com/2010/11/how-much- is faster collaboration- worth.html.

Posted: November 7, 2010

Larry Pizette
Principal Software Systems Engineer
Formerly MITRE

The value that an organization obtains from well-publicized cloud computing benefits such as increased utilization of hardware, location independent access for users, and scalable computing environments, will vary based upon their unique goals and circumstances. "Every organization has to crunch its own numbers to evaluate the cloud solution that makes the most business sense, but the number of cloud options and reduced implementation risk make the current IT environment ripe for innovation" writes Teresa Carlson.

Government is both providing cloud environments and using them. In order to establish a business case for being a cloud provider, whether private or community, cloud-specific benefits and costs need to be estimated and analyzed. The owning organization invests their resources into their own hardware and software and operates and controls their own infrastructure. Through more efficient use of physical servers, reductions in cost categories such as capital investment and ongoing operating expense can be realized. The value of new capabilities for users and costs for delivering the capabilities should be included along with the costs for meeting rigorous requirements for COOP, location independent access for users, security, "up time" and help desk support. Nathanial Rushfinn notes, "By clearly defining cost categories for both capital and operating expenses, and by using well-defined customer surveys, federal IT leaders can estimate the success of cloud computing projects from both an ROI perspective and from a customer's vantage point."

When using a public or community cloud service, the acquiring government organization no longer needs to invest significant capital for building their own data center capability, which can include cost drivers such as buildings, storage hardware, servers, and HVAC systems. Associated cost drivers such as electricity, maintenance contracts, software license costs and support personnel for data center infrastructure are reduced. In addition to the cost reductions, there can be value from increased agility. Douglas Bourgeois states: "Don't lose sight of the fact that there is also considerable value to be derived from the selection of cloud services as well as the speed in which cloud services are delivered." These cost reductions driven by using public and community clouds need to be compared against the cost areas that will increase. In addition to monthly usage costs, there are on going costs to manage the relationship with the provider. These cost categories can include porting, integration, data migration, testing, security analysis and certification and accreditation (C&A) costs that impact the business case.

In addition to the above factors, there are many considerations relevant to organization- specific business cases that can drive costs, such as schedule demands, network dependency, security requirements, and risk analysis. The value can be more than cost savings, notes Peter Coffee. "Justification of cloud projects should therefore begin with expected improvements in cost, reductions of risk, and accelerations of service availability, but should not stop there: they should also make reasonable projections, based on growing collections of relevant examples, of the value of improved mission performance."

Posted: November 12, 2010

Question for September 2010

Often service level agreements (SLAs), contracts, or memorandums of understanding (MOUs) are used between organizations to define the relationship between the service provider and consumer. For a Federal Government or DoD context, please describe or suggest important attributes of SLAs, contracts, MOUs, or other status information that are needed to enable successful operational cloud deployments.

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP

The relationship between the provider and the consumer (or subscriber) is critical to success with Cloud Computing, as it is with any service. One piece of the relationship is to fully understanding what you are buying. For an Internal Cloud, the provider and consumer may be in the same organization. In the case of a Public Cloud or Virtual Private Cloud, the need for a good relationship cannot be over stressed. It has been said that good fences make good neighbors. Creating and maintaining a good set of SLAs are the fences. Accordingly, a clear and healthy relationship of mutual understanding and alignment is a critical success factor. For the IT shop providing or brokering Cloud Services to internal clients, getting the right SLAs is critical as they ultimately are responsible to the client regardless of the downstream agreements.

First, you should make sure that you are clear about what is most important to the consumer in terms of performance. For some it may be availability or system responsiveness, and others a timely back up schedule. I recommend listing the attributes of the service that are most important, and then figuring out how to measure those attributes in the most systemic and meaningful way possible. I had one agency tell me that a major problem came up because two vendors were using different sources of time. So you may even want SLAs on how time is used and what source it comes from.

Next, as it turns out, coming up with good metrics is one of the most difficult steps to establishing an SLA. Let's takes something as seemingly strait forward as availability. How do you measure it Is it the availability of the infrastructure and network or the availability of the application? If it is the application, which applications are critical to the end- users? How do you handle scheduled down time There is a twist to choosing metrics. In some cases metrics can drive unintended behavior. For example, we used to measure programmers by lines of code delivered (or bugs fixed). The natural result: very long, monolithic, inefficient code. Many of these unintended behaviors can be hard to predict. Make sure you have a way to adjust the SLA if it is creating unintended behavior or just not working for you.

Finally I would recommend that your SLAs have real teeth in them, aligning risk and reward appropriately. If the risks are outside of your immediate control, the SLA should address the risk and the consequences. If the SLA does make the provider a little uncomfortable, they may be more responsive and deliver quicker solutions. In either case, both parties must be involved in creating, monitoring, and fixing SLAs.

In summary, first focus on the attributes most important to you, remembering relationship is important. Next, build good metrics for those attributes. Monitor the SLAs to make sure they are working for you and change if necessary. Finally, make the consequences of failed SLAs painful enough to promote quick response.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: September 23, 2010

Erik Hille
Director, Cloud Business at CA Technologies
CA Technologies

Pressured to improve operational performance and accountability, many federal agencies have increased scrutiny over their outsourcing strategies. Ironically, as the outsourcing market has evolved to include cloud- based services, this level of scrutiny has not been applied to these emerging delivery methods. Cloud providers excel at communicating the business benefits of their service, but from an accountability perspective, many could stand to take a more proactive stance. Here are 5 things you should think about when establishing service level agreements (SLAs), memoranda of understanding (MOU), and performance measures with cloud providers:

1) An MOU won't cut it: Although it can express an obligation between the service provider and the agency, an MOU is not a strong enough document to govern the relationship. MOUs fall short of being an enforceable contract. Because the outsourced services may be mission critical and involve fundamental building blocks such as infrastructure or platforms, it is far more effective to leverage a contract that describes what services are outsourced, what the responsibilities of both parties are and what the performance characteristics will be.

2) SLAs and contracts are the same thing: Outsourcers have used a specialized version of an SLA called an underpinning contract (UPC) for years. This contract outlines the services the provider will deliver, the penalties and credits associated with under and over- performance, and the metrics that will be used to describe how the contract will be delivered.

3) Measure performance, not just provisioning: Note that the SLA is a contract, not an operational performance characteristic such as "% uptime." Instead of SLAs or UPCs, these measures are "metrics," indicators of the cloud service's operational parameters. Because much of the external cloud market grew out of the virtualization space, many providers offer provisioning metrics, but steadfastly avoid performance metrics (% uptime, throughput,

capacity, etc.). These cloud services are still fundamental building blocks for running your agency and must be protected operationally.

4) It is your data, and you are not doing the provider a favor: It's not uncommon for providers to keep performance metrics close to the vest. Many are trying to avoid obvious penalties and some go to great lengths to report performance only to the minimum required level. If a provider claims the performance data is proprietary, they are wrong. The agency needs to ensure that it has access to these metrics for a level of assurance against the obligations agreed upon in the contract itself.

5) Active not passive monitoring is key: Another way some cloud service providers might fail to report performance is to expect it of the agency). They leave it up to the customer to find the outage, report it, and collect the penalty. Instead agencies need to be proactive, either requiring the provider to implement a Service Level Management solution to actively monitor the agreement, or they need to do so remotely. In this way, both parties are able to agree to the performance of the contract.

For further information, please contact Erik Hille at Erik.Hille@ca.com.

Posted: September 24, 2010

Ron Knode, Director,
GSS, LEF Research Associate
CSC

In the Cloud, Security Begins with a 'T'

We've all seen clouds work. We've all read case studies of productive use of the cloud in both government and industry. We've all been inundated with a seemingly endless cascade of cloud technology announcements, offerings and alternatives. And, we're probably all near to some cloud technology testbed of one variety or another. In the face of such single-minded devotion to the "technology of cloud" we might conclude that all we need for a trusted cloud operation is the right technology arranged and configured in the right way. Clouds are technology, right ?!

Wrong! As much as we are (rightfully) intrigued by the technology of cloud, and as much as we impressed by the snappy and snazzy way cloud technology seems to respond to our needs, the real power of cloud is what happens *around* the technology. Clouds need technology for sure. But, trusted clouds need people, and process, and rules for operation, and governance, and accountability for outcomes even more. The technology of clouds is evolutionary. The consumption model for clouds is revolutionary. When those people and process and rules and accountabilities that are needed span organizations (internally or externally), then we inevitably must include some sort of agreed mechanisms for cloud service delivery, e.g., Service Level Agreements (SLAs), Memoranda of Understanding (MOUs), or contract terms and conditions (T&C's).

Okay, we're making progress. We know we need the important (revolutionary) mechanisms around sound (evolutionary) technology in order to generate and sustain trust in cloud operations and reap the payoffs that are promised to us. But, what should those mechanisms emphasize in order to capture the best payoff situation?

That's the thrust of the question for September. And, as usual, the Government's desire to find the important characteristics of such mechanisms is not much different from that of industry.

The answer to that question lies in the recognition that:

In the cloud, 'security' begins with a 'T'.

Transparency is the single most important characteristic required to generate trust and capture payoffs. Beyond the standard characteristics of availability and incident response timeliness (for which SLAs are well known), additional SLAs, MOUs and/or T&C's should reinforce the characteristic of transparency in cloud service delivery. While the technology must support the delivery of transparency of service, it is the accompanying mechanisms of service definition that provide the real payoffs.

The Precis for the CloudTrust Protocol includes a list of the *elements of transparency* that can be the basis for an SLA that requires the measurement and reporting of such metrics. The CloudTrust Protocol is intended to provide a vehicle for such reporting. In addition, that same reference also describes a recommended SLA for self-reporting by cloud service providers (to reduce the chances of 'gaming' results). Whether in SLAs or MOUs or T&Cs, or even in standards and 'best practices' themselves, attention to the transparency of service is essential.

Just remember your spelling lesson for the cloud and payoffs can come your way. See the full blog response at www.trustedcloudservices.com.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: September 27, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

IT-using organizations want service today, not a credit for service tomorrow or any other compensation for a service provider's failure to provide what was promised. Proven cloud providers like salesforce.com, Amazon Web Services, and Google are meeting the need for true service by giving customers prompt and detailed information via Web sites like trust.salesforce.com, status.aws.amazon.com, and www.google.com/appsstatus to provide the record of reliability, and disclosure of even slight departures from normal operation, that let customers plan with confidence.

Organizations should make a cloud comparison, not against an ideal, but against legacy IT's reality. When organizations own and operate their own IT infrastructure, they assume the entire burden of providing reserve capacity and protection against interruption of capability.

Backup storage, backup power, even entire backup data centers are routine expenses in the world of traditional IT. If these protections turn out to be inadequate, the user organization bears all costs of mission failure.

In contrast, cloud providers provide reserve capacity and redundant capabilities, such as backup power and backup connectivity, at a lower cost per customer due to massive economies of scale. Moreover, cloud providers have enormous incentive to assure that their customers do not experience service degradations that lead to unfavorable publicity. Further, each customer of a true multi-tenant cloud enjoys "sum of all fears" protection: the provider must satisfy all concerns of all customers, and in the process will generally address a superset of the concerns of any single customer.

If cloud providers price their services to cover worst-case consequential damages of any service limitation, as felt by their most demanding customer, the result will not be economically attractive to the vast majority of customers. Those customers will be better served when they make their own assessments of risk and cost, and mitigate those risks in a way that meets their own requirements.

Cloud data centers will be, for the first time, statistically similar enough to enable accurate pricing of risk: a vast improvement over the unique, unpredictable risk of a myriad of individual and uniquely configured onpremise data centers. We may therefore expect to see an efficient marketplace of risk that the IT professional has previously lacked.

Compare the cloud's risks to the risks that are experienced by the customers of a shipping service. Those customers make their own judgment of the reliability of that service, and of the consequences of any delay or damage or loss of a shipment. They purchase insurance, or keep extra inventory on hand, or take other measures to limit their exposure. In a similar way, the most risk-sensitive customers of cloud services will choose their own measures that are suited to their own particular circumstances.

A "service level agreement" does not give the customer what's really needed which is reliable, secure service that gracefully handles peak loads without the customer needing to own peak-load capacity. That's what the cloud is all about, and that's what cloud service customers are quickly learning to expect.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: September 29, 2010

Teresa Carlson
Vice President
Microsoft Federal

The same terms always pop up when discussing cloud SLAs uptime, availability, reliability. These words speak to the really innovative quality of cloud computing—how computing resources are accessed. You're not buying a product with a set of agreed upon features, you're

buying a new way to house and tap into your IT assets. Customers want assurance that they will have access to their data and applications, and it's up to vendors to guarantee this access. When reliability is combined with security, cloud computing becomes a no brainer, and SLAs are absolutely necessary to outline agreed upon service expectations that meet customer needs.

But as cloud infrastructures have improved, *access* seems like a pretty low bar. If I'm a Public Sector CIO evaluating cloud computing options, I'm not willing to accept a significant decrease in access (uptime, availability, reliability) in order to gain the other benefits cloud offers (efficiency, scalability, cost reductions). A large part of my decision will be based on a cloud solution's ability to be there when I need it, and it shouldn't be much different the reliability of traditional IT infrastructures. Federal agencies can't afford regular, unexpected service interruptions. The data and the mission is too important. This is why data portability is essential. It gives agencies the ultimate option to immediately relocate to another cloud provider if their service needs aren't being met. Agencies need the freedom to move their data to an environment they trust, and SLAs that include data portability language protect customers more effectively than any other metric or clause.

It's common for SLAs to include financial compensation for service outages, and that's an important start. Customers should be compensated for lost access, but if there are repeated, unscheduled breaks in service, that policy is failing to provide value. All enterprise organizations require consistent access to their computing resources, and when service needs aren't being met, data portability adds another layer of assurance beyond financial return.

It's true that service interruptions often occur because of network outages rather than issues with the cloud solution itself. Unfortunately, the result is the same for customers lack of access. To limit these breaks in service, vendors should address minimum network connectivity requirements in the SLA. Network monitoring is a key component of a holistic cloud implementation, and vendors should continually and proactively work with network providers to ensure connectivity needs are being met. SLAs can address these issues at the outset, and can even outline network backup options like leveraging satellite connectivity.

Overall, SLAs are extremely important, but they are evolving as cloud offerings improve. Customers are rightly expecting more, and vendors must step up their game to deliver. Ensuring data portability in SLAs avoids vendor lock-in, promotes choice, increases competition and allows government enterprises to freely choose the best available solutions.

Posted: October 4, 2010

Lynn McPherson
Lead Software Systems Engineer
MITRE

An SLA is an agreement between two parties, the service provider and the service consumer, that defines a contractual relationship. As Skip Bailey stated in his October response above "The relationship between the provider and the consumer (or subscriber) is critical to success

with Cloud Computing, as it is with any service." As is true in any successful relationship, both parties must understand and accept certain **responsibilities** —successful relationships are rarely one-sided. Among other things, the responsibilities of the service provider include providing the described service within defined constraints, collection of agreed upon metrics, timely production of predefined reports, and adherence to an agreed upon incident management and resolution process. Likewise, the consumer bears certain responsibilities which include, but are not limited to, ensuring that they don't exceed the agreed upon workload as well as validation that the provider is collecting and reporting metrics properly through a quality assurance surveillance plan. Other necessary and fundamental aspects of the relationship include:

- **Control:** Delineates the aspects of the service which are and are not under the control of the provider and is critical to writing an effective, enforceable SLA. The provider is held accountable for delivering a particular level of service which is agreed upon in the SLA; however, the provider should not be held accountable for failures that occur which are outside their control. The complexity of today's computing environments necessitate that the SLA clearly describe those aspects of the service which are and are not under their control. In general, descriptions such as this necessitate the inclusion of an architecture diagram to supplement the verbiage provided.
- **Measurement:** Ensures and demonstrates that the agreed upon level of service is delivered. Measurement encompasses measures and metrics. A measure is a value that is recorded as a result of a physical measurement such as a single instance of a response time. A metric is a quantitative measure of the degree to which a system, component, or process possess a given attribute. Metrics are the foundation of a well defined SLA; they must be objectively measureable or calculated from objectively defined measures. The lack of objectively measureable metrics may result in an SLA that is unenforceable.
- **Transparency:** Implies openness, communication, and accountability. A successful relationship is always based, in part, on trust and transparency is fundamental to trust. Transparency applies to many aspects of the SLA including the definition of unambiguous responsibilities and metrics as well as a clear understanding of the provider's span of control. In addition, it is extremely important that the reporting process, scheduled reviews, and methods for computing incentives and penalties be completely transparent to all involved parties.

Taken together, responsibilities, control, measurement and transparency in an SLA can help to establish trust and facilitate a successful cloud computing relationship.

Posted: October 5, 2010

Question for August 2010

Mr. Wes Schooley, US TRANSCOM J6, asks: "Can industry share some examples, including high level architecture details, regarding successful implementation of cloud computing-based 'Data- as-a-Service'?"

Alternate question: Please share considerations and/or architectural approaches Federal IT leaders should examine for providing data services with a cloud based model.

Peter Coffee
Director of Platform Research
salesforce.com inc.

Data services are evolving from a commodity, priced on simple measures of data volume and transfer rate, toward a more differentiated market with a lengthening list of measures of quality including accuracy, timeliness, integration/conversion flexibility and disaster preparedness. Data services are thus an excellent example of the need to think of the cloud as something more than a relocation of familiar infrastructure. The value is in the services that are added, at least as much as in the core costs that are (one hopes) substantially reduced.

One example of this transition is incorporation of the Jigsaw "crowdsourced" database into salesforce.com cloud applications such as sales force automation.

When a salesperson touches a contact or company record, real time triggers can highlight possible conflicts between that record and the latest available information. It's clear that similar mechanisms of real time update and discrepancy detection could have great value in many Federal task domains, and that these should become part of the basis for data service selection.

One can't buy a hard drive that comes with built-in data quality maintenance: such a feature is only conceivable in a cloud-based service, and buyers will benefit by expanding familiar frames of reference for data storage procurement to think in these new terms.

Agencies will doubtless be concerned about the issue of data residency, both in terms of physical location and operational control of data storage devices. It's essential to understand the flexibility achievable with cloud integration points and partner services. Agencies will increasingly come to appreciate the rigor and cost-effectiveness of cloud service providers' protections, but may also choose to take advantage of selective masking capabilities (including encryption) for cloud-resident data or even to partition their applications, keeping the most sensitive data fields in local storage while associating them with cloud-based records only by means of anonymous identifiers (e.g., "Case Number").

To be sure, the fundamental challenges of volume and throughput continue to grow. One analytics company recently posted a job description headed "Petabyte Platform Engineer"; Amazon Web Services invites cloud customers with large data volumes to deliver physical storage devices directly to an AWS data center for faster transfer of volumes ranging from 100 GBytes (82 days on a T1 line) to several TBytes (many days on even a T3 connection). At the same time, however, Federal agencies with critical data dependencies will do well to consider the next-generation services of cloud-focused systems partners such as Model Metrics; Informatica and Cast Iron Systems—with the acquisition of the latter by IBM giving emphasis to the growing importance of services, added to core capabilities, to create compelling cloud options.

An IT buyer's conversation with a data services provider will still begin, most likely, with questions of gross capacity and speed, but that's just the introduction to a discussion of more interesting value-adds. Consideration of cloud services should focus on the need to be met, not on the hardware attributes to be virtualized.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: September 1, 2010 1:20pm.

Nicklous Combs
Chief Technology Officer
EMC Federal

There are many specific examples of implementing cloud based data as a service and I would be happy to follow up with specific customer references in person but I will talk to a couple examples where we have implemented such services to support our customers both in and out of DoD.

A specific DoD customer was interested in moving 24K user exchange and SharePoint into cloud services but was challenged by the distributed nature of the environment, need for continuous operations in multiple locations, a frequently moving user base and the need to rapidly scale to 45K users on demand. This was becoming more and more challenging as demands for flexibility and resiliency were seen as a huge cost hurdle that they were not able to overcome. By implementing a 100% virtualized cloud environment we were able to provide the flexibility needed to support their needs. The next challenge was how do we make data available wherever the customer traveled and make it transparent to the user. We addressed this by implementing object based Cloud Optimized Storage to present data to users and distributed applications using simple SOAP and REST protocols. By presenting the data as objects to the user and applications, data mobility became a by-product of implementing a cloud based data as a service architecture and existing scalability and replication challenges within the SharePoint environment were eliminated.

A similar example in the commercial space was implemented when a customer needed to support five geographically distributed data stores which exceeded three petabytes of

information. The data was changed frequently by many millions of users and the need to avoid a centralized data management application, support continuous changing IO requirements and to do this in a cost effective centrally managed infrastructure required us to build a data as a service architecture that was based around user defined policy management of information. This approach takes many decisions typically made in the application and pushes them down the technology stack where data could be managed by policy at the storage layer. This approach dramatically reduced the operations and maintenance costs because O&M personnel were not required at each site to do data management and replication tasks. This is an example of how automation is beginning to and will eventually replace many of the lower level IT functions currently being performed and will free the IT specialist's time to focus more on how IT is supporting the business.

For more information: http://www.emc.com/ ?fromGlobalSiteSelect.

Posted: September 10, 2010

Larry Pizette
Principal Software Systems Engineer
Formerly MITRE

Many organizations can reap significant benefit from the features associated with cloud computing, such as location independent access to information and the ability to access data services with up- times guaranteed by service- level agreements (SLAs). Additionally, many large cloud- based offerings can provide extensive "on demand" scalability that can help an organization to increase their data service usage without a large, planned capital investment in storage hardware and infrastructure.

A cloud-based data service approach can be used for a variety of purposes in Government, ranging from highly secure private cloud purposes to very visible public capabilities supporting the White House's Open Government initiative. For example, the NASA Open Government Plan describes the NASA Nebula cloud computing environment as "an open source cloud computing platform" which provides "an easier way for NASA scientists and researchers to share large, complex data sets with external partners and the public." USAspending.gov 2.0, a website for government budget information, is hosted on the Nebula cloud.

While the potential benefits are numerous, Federal IT leadership should consider the business case for data services. For private cloud deployments, there will be an investment in hardware and software, and costs for implementing, integrating and testing data services. For community and public deployments, there can be significant savings in capital investment; however, leadership should ensure that the cost/benefit analysis includes the costs for applications that have to be ported to take advantage of data services. For example, some public database cloud computing offerings can only be accessed with proprietary application programming interfaces (APIs). Usage of these APIs can require significant porting, integration and testing of legacy applications.

Network latency and throughput should be considered in determining the approach to using a cloud-based data service. A degraded network could slow the upload and retrieval of information, and in the event of a network failure, access could be totally severed. Therefore, data services that are placed in off-premise, cloud-based environments will not be available for "disconnected operations."

Security solutions should be employed (e.g., encryption) and SLAs should be codified with providers to meet the organization's needs for protection of data at rest and in motion. "While data outsourcing relieves the owners of the burden of local data storage and maintenance, it also eliminates their physical control of storage dependability and security, which traditionally has been expected by both enterprises and individuals with high service-level requirements," writes Wang et al in IEEE. Similarly as Gartner has noted, Federal leadership should examine the location of data (e.g., CONUS) and controls for segregating it.

While there are many benefits and considerations for cloud-based data services, as Peter Coffee writes, the conversation will likely start with a discussion of "capacity and speed" before moving to the more interesting value-adds. It's this value coupled with meeting requirements that should be compared against the costs and considerations for each organization to determine their return on investment (ROI), benefits, and trade-offs.

Posted: September 8, 2010

Question for July 2010

The use of standards-based solutions can be an important risk reduction approach for Government. Please describe current standards that could help the Government in its adoption of cloud computing. Also, what cloud standards efforts would you like to see in the future?

Winston Bumpus
Director of Standards Architecture,
Office of the CTO
VMware

Over the last year much progress has been made on new standards for improved cloud interoperability and reduced vendor lockin. Standards development organizations (SDOs) have been applying the expertise of their constituencies to the problem and new organizations like the Cloud Security Alliance, have emerged to focus on unique challenges of cloud computing. Existing standards are being adapted as well to address cloud computing interoperability such as the Open Virtualization Format (OVF) from the Distributed Management Task Force (DMTF). OVF was originally developed to address portability concerns between various virtualization platforms. It consists of metadata about a virtual machine images or groups of images that can be deployed as a unit. It provides an easy way to package and deploy services as either a virtual appliance or used within an enterprise to prepackage known configurations of a virtual machine image or images.

For example, it may contain information regarding the number of CPUs, memory required to run effectively, and network configuration information. It also can contain digital signatures to ensure the integrity of the machine images being deployed along with licensing information in the form of a machine readable EULA (End User License Agreement) so that it can the terms can be understood before the image(s) is deployed.

OVF is currently being explored to validate if other metadata should be added to help improve the automation of intercloud workload deployment. Concepts such as standardized SLAs (Service Level Agreements), sophisticated inter virtual machine network configuration and switching information and software license information regarding all of the various components that make up the workload are possibilities.

Other standards are still emerging including a common cloud API (Application Progamming Interface). These higher level APIs will be important as the pendulum settles in towards a series of multiple cloud offerings, each have different underlying implementations, but a standard many in which to interact. Having a standard set of foundation APIs, similar to the POSIX standards of the past, will help to ensure that cloud management and compliance tools will not be overly complex in order to handle different cloud implementations.

Efforts such as NIST's Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) and Federal Risk and Authorization Management Program (FedRAMP) are two key initiatives to help both the government and industry to better define a robust, scalable and economically viable set of standards to accelerate the adoption of cloud computing.

For more information on what is going on in the various SDOs and the current and emerging industry standards for cloud computing please go to www.cloud standards.org.

Posted: July 13, 2010

Ron Knode, Director,
GSS, LEF Research Associate
CSC

Cloud Standards Now!?

Wouldn't it be wonderful if we could simply point to cloud standard(s) and claim that such standard(s) could reliably lubricate government adoption of safe, dependable, accreditable cloud computing?! Sadly, we cannot. At least, not yet. And, this fact is as true for commercial adoption of cloud computing as it is for government adoption.

However, what we do have is the collective sense that such standards are needed, and the energy to try to build them. Furthermore, while the "standards" we need do not yet exist, we are not without the likely precursors to such standards, e.g., guidelines, so called best practices, threat lists, special publications, and all manner of "advice- giving" items that try to aim us in the right direction (or at least aim us away from the very wrong direction). In fact, we have so many contributors working on cloud standards of one kind or another that we are in danger of suffering the "lesson of lists" for cloud computing.

Nevertheless, given our desire to reap some of the benefits of cloud computing, should we not try to accelerate the production, publication, and endorsement of cloud computing standards from the abundance of sources we see today?

Wait a minute! Standards can be a blessing or a curse. On the one hand, standards make possible reasonable expectations for such things as interoperability, reliability, and the assignment and recognition of authority and accountability. On the other hand, standards, especially those generated in haste and/or without widespread diligence and commentary, can bring unintended consequences that actually make things worse. Consider, for example, the Wired Equivalent Privacy (WEP) part of 802.11 or the flawed outcomes and constant revisions for the PCI DSS (remember Hannaford and Heartland !?).

What we seek are standards that lead us into trusted cloud computing not just "secure" cloud computing or even "compliant" cloud computing. Ultimately, any productive stack of standards must deliver transparency to cloud computing. Simply having cloud standards just to have standards does not bring any enterprise closer to the promised payoffs of the cloud. So, let's proceed with all deliberate speed through some of the worthy efforts ongoing, but not declare success merely for the sake of an artificial deadline or competitive advantage. The cloud definition and certification efforts sponsored by NIST and GSA, the security threat and

guidance documents authored by the Cloud Security Alliance, the cloud modeling work of the OMG, the cloud provider security assertions technique proposed by Cloudaudit.org, and the CloudTrust Protocol SCAP like extension to reclaim transparency for cloud computing, all of these efforts certainly hold promise for accelerating the adoption of cloud computing for government and industry.

Let's push and participate in the actions of these and other groups. Ask questions, experiment, build prototypes, and seek extensive and deliberate peer review. Standards that survive such a process can be endorsed. But, like fine wines, cheeses, (and even thunderstorms), "We will accept no cloud standard before its time."

See the full blog response at www.trustedcloudservices.com.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: July 21, 2010

James A. St. Clair
Sr. Manager, Global Public Sector
Grant Thornton LLP

While the Cloud does prompt consideration of unique standards, many of the "same old thing" still pertain and should be considered.

At a high level, existing compliance standards such as The Health Information Portability and Accountability Act (HIPAA), Sarbanes Oxley (SOX), and the Federal Information Security Act (FISMA), embodied by NIST guidance, all provide specific considerations for security that pertain to any computing environment. In basic terms, systems are expected to provide a level of confidentiality, integrity and availability commensurate with the sensitivity of the information and the identified level of risk, whether the system is a one-sever LAN or cloud-provisioned infrastructure.

However, how provisioned services and cloud computing manage information and limit risk is arguably different than traditional client/server architecture. As such, new standards are developing that help the "apples to apples" comparison of traditional security control objectives and new cloud provisioned services:

- The Cloud Security Alliance. The Cloud Security Alliance (CSA) is a non-profit organization formed to promote the use of best practices for providing security assurance within Cloud Computing, and provide education on the uses of Cloud Computing to help secure all other forms of computing. The CSA's "Security Guidance for Critical Areas of Focus in Cloud Computing" in one of the first security standards for cloud computing, based on the contribution of many industry participants.
- FedRAMP. The Federal Risk and Authorization Management Program (FedRAMP) is a unified government-wide risk management program focused on large

outsourced and multi-agency systems. The program will initially focus on cloud computing but will expand to other domains as the program matures. FedRAMP provides security authorizations and continuous monitoring of shared systems that can be leveraged by agencies to both reduce their security compliance burden. FedRAMP authorization standards are built upon the tested application of tailored security controls and enhancements of NIST Special Publication 800 53.

Additionally, other groups are fostering technical and architectural standards that provide important elements of a framework to develop virtualized environments and provision cloud computing:

- For several years, The Open Group has developed and promoted The Open Group Architecture Framework (TOGAF), now at version 9. TOGAF has been built as an industry consensus framework and a method for enterprise architecture that is available for use by any organization around the world, and available for download.
- The Open Web Application Security Project (OWASP) is a global, collaborative, not-for-profit association dedicated to advancing knowledge and awareness of security in application security. Perhaps most importantly, OWASP tools and initiatives directly benefit virtualization and cloud services which heavily leverage web application for service delivery.

It is envisioned that continuing adoption of these and other standards will also help drive broader adoption of more transparent and secure cloud computing services. Ultimately, whether an organization pursues creating their own cloud, or leveraging hybrid or public clouds, these standards will make adoption easily acceptable and promote cost-effective IT investments.

For further information, please contact James A. St.Clair at Jim.StClair@gt.com.

Posted: July 27, 2010

Lew Moorman
President, Cloud and Chief Strategy Officer
Rackspace Hosting

Many suggest that standards are the key to encouraging broader adoption of cloud computing. I disagree; I think the key is openness and a competitive market. What's the difference? In the standards approach, a cloud would look and work as described by the standard it is implementing. If only one commercial implementation of the standard exists, this limits choice and freedom. Open clouds, on the other hand, could come in many different flavors, but they would share one essential feature: all of the services they'd offer could be run by the enterprises or agencies themselves without requiring a service provider.

Why is openness so important? If the web has taught us anything, it is that open systems, portability, and choice drive innovation. The open Linux system brought us a mountain of software and tools to help accomplish almost any task. And, each open component, whether a database or a widget, could be moved in and out freely to get the job done. These components often followed standards but were not limited or locked- in by them; and as long as their changes were open and accessible, they saw adoption.

Last week, Rackspace announced a new open source cloud platform with the support of more than 25 technology industry leaders: OpenStack. With initial source code contributions from Rackspace and the NASA Nebula cloud platform, OpenStack forms a foundation of cloud technologies used at scale in production today, including a compute provisioning engine OpenStack Compute and a fully distributed storage engine OpenStack Object Storage.

We expect an open source cloud platform like OpenStack to enable several things. One, anyone will be able to run this cloud and do it anywhere. Enterprises and agencies will be able to build private clouds.

Workloads will be moved among these clouds easily from private to community to public clouds, or among different service providers, without having to rewrite the software to do it. Two, the entire tech ecosystem can build around this open source foundation. With wide adoption, there will be a more robust market of services around this flexible engine, from storage systems to monitoring tools to management systems. Three, the cloud will advance faster than ever.

It has been rewarding to see more formality around cloud standards development (from the DMTF, OGF, and others) as well as a coalescing of various standardization efforts (e.g. http://cloud- standards.org). Customers demand a faster pace. That's why we established OpenStack to accelerate innovation around open cloud software, and have also committed to adopt and support open and extensible standards as they emerge.

For further information, please visit http://www.rackspace.com.

Posted: July 29, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

In the cloud, if you're not interoperable, you're irrelevant. Any cloud service that can't interact with other services, and integrate with legacy IT assets, is too crippled to be competitive: it will never be important enough to make its proprietary nature a problem to any large community of users.

Even so, the world of the cloud still demands a role for standards but not the role that standards have played in the past.

With locally operated IT, a dominant vendor could exploit increasing returns to scale: the more of that vendor's technology a customer adopted, the greater the incentive to use the same

vendor for the next incremental need. Customers needed standards to protect them from dominant vendors' temptation to strengthen barriers to competition.

The cloud's best-paved path does not lead to natural monopoly: rather, the cloud invites any provider of even the narrowest specialty capability to offer that service through an appropriate API, accessible to other services through established Web service protocols. A vendor of risk management tools, for example, need not create its own sources of data or its own facilities for documenting results: other cloud services can perform those functions, while the vendor puts its resources into what it does best.

The cloud is an ecosystem that favors specialization and symbiosis, rather than generalization and top-predator dominance. The question that should be asked, therefore, is not, "are standards desirable?" The question should rather be, "what standards will benefit customers and users?"

In the old IT model, customers needed standards that made different vendors' products readily substitutable even if this led to a commoditized marketplace, such as the desert of "beige box" Windows/Intel desktop PCs that competed only on price.

The cloud invites providers to develop distinctive competence, while customers focus on standards that maximize interoperability. Already, these standards largely define the cloud marketplace: for example, more than half the workload borne by salesforce.com systems supports other services invoking salesforce.com APIs, rather than executing salesforce.com's own cloud applications.

Cloud service customers will still benefit from some degree of substitutability, both commercially (providers will not be able to raise prices without inviting competition) and technically (rare service interruptions may be inconvenient, but need not be catastrophic).

Customers will do well, though, to think on a higher level than mere substitution: should cloud storage services, for example, be purely interoperable, or should there be stratified storage offerings with varying reliability/speed/cost that can provide "defense in depth" against common-mode failures What protections matter most?

Any competitive IT platform should give application developers and customers the freedom to put some things where their productivity and capability will be greatest, while putting other things where substitutability is maximized. When writing an application with a useful life of months, developers want maximum leverage; when developing intellectual property with a lifetime of years, they may prefer a path that can lead to as many different environments as possible.

The proper role of standards is to provide, and preserve, a choice among varied paths. For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: July 30, 2010

Teresa Carlson
Vice President
Microsoft Federal

Standards can be extremely valuable in providing security and privacy assurances to organizations exploring cloud computing options, and they are also critical to laying a foundation of interoperability within the IT industry. Interoperability is really essential because it promotes competition, innovation, and customer

choice, which are all key to ensuring the government has access to the best solutions at the best prices. It's important to always think about standards as a means to this end, because creating standards for the sake of creating standards has the potential to hinder innovation.

History shows that standards tend to emerge as they are needed. Industries adopt standards when organizations demand them. They have the power to level the playing field and provide access to the best solutions. In the cloud, this means security standards, protocols, Internet standards and storage standards. Some of the best cloud standards that exist today are the same ones that advanced the Web, including HTTP, Simple Object Access Protocol (SOAP), Representational State Transfer (REST) and Extensible Markup Language (XML). These are all market-test standards carried over from Web 2.0 or grid computing and they all support interoperability. In some cases these existing standards aren't a perfect fit for the cloud, especially when it comes to connectivity, datacenter proximity and privacy, but new standards that address these issues specifically will continue to build upon Web services and REST-based approaches. Our developers always talk about how high level, semantic standards tend to work better than syntactic standards because they avoid the friction and overlap that can actually hinder progress.

But standards don't create interoperability on their own. If two providers implement the same standard within their cloud offering there is no guarantee that those products will be able to interoperate with each other. It still requires ongoing technical collaboration amongst companies, governments and standards-setting organizations. In the federal space, standardized terms, language and processes will go a long way to achieving this goal. On the security side, FedRAMP is a great example of collectively addressing standards and streamlining the process of evaluating cloud solutions. Making authorization government-wide will eliminate the time and resources that each agency needs to devote to risk management.

The best standards come from an open, collaborative process that is driven by market need. Worthwhile standards need to be tied to a specific "use case" that addresses a critical outcome like interoperability or security. There are some great standards that already exist today that we can build upon to address current cloud gaps, which will keep the playing field level and offer the best solutions for federal agencies.

Posted: August 5, 2010 6:00 pm.

Marie Francesca
Director Corporate Engineering Operations
MITRE

Many thanks to this month's submitters for sharing their insights and perspectives on cloud standards. As our submitters have noted, there are multiple on going activities by government and industry with many market- leading companies participating. Winston Bumpus states much recent progress has been made. However, more effort is needed to facilitate widespread government adoption. NIST is leading the way in government and there are industry- based organizations such as DMTF pursuing standards that can move the community to the next level. The history of technical standards has shown that they can be highly successful in facilitating interoperability and portability as well as lowering costs and enabling new products.

Successful standards development involves people, process and technology, and all three aspects must be addressed. Lew Moorman hinted at this in the discussion on openness in his response. Several sources, including IEEE, indicate that the better standards themselves are open (versus closed/proprietary) and are developed in an open, community process. Business opportunities need to outweigh the loss of competitive advantage attained through a closed, proprietary approach. In his Essential Guide to Standards, Andrew Updegrove places significant weight on the right membership in a standards development consortium, including a broad list of market leading vendors and government. Peter Mell and Tim Grance of NIST both indicate minimum cloud standards should ensure interoperability without inhibiting innovation. Similarly, several of our blog responders commented that preserving innovation (and thus competition and business opportunity) is a necessary standards attribute.

Peter Coffee proffers an important question that all of us should ask: "What standards will benefit customers and users?" This speaks to the question of relevance, which Peter J. Ashendon mentions in What Makes a Good Standard ?. There should be a clear business purpose that standards address. Ron Knode suggests that we should seek cloud standards that lead us to "trusted cloud computing" and Jim St.Clair points out a number of activities already underway.

In examining successful modern technical standards, the degree of adoption was a critical tipping point. HP CTO Paul Congdon coined a term called the "thud factor" referring to the amount of documentation needed to describe the standard. If it is difficult to understand, adoption is not likely, or, it will not be implemented fully and lead to the occurrence of non- standard subsets. The W3C also makes the case that "learnability" is important. An open process is likely to help in reducing complexity as noted by several sources.

Many technologies comprise cloud computing and there is likely a need for standards that address different aspects of cloud implementations. For example, portability and interoperability standards have very different implications for IaaS than PaaS or SaaS. Successful standards committees will need to narrow down the scope of the specific standards while looking for opportunities to have a positive impact for both consumers and providers.

In summary, cloud computing standards should be open, simple, interoperable, and relevant to business needs as well as targeted for a given application. They should enable innovation and the process by which they are developed should support community- based participation.

Posted: August 5, 2010 5:03 pm.

Question for June 2010

Major Manny Dominguez, USAF, Chief Information Officer (CIO) for the Medical Education and Training Campus (METC) asks: "In moving capabilities to the cloud, it will be important for Government/DoD organizations to have an understanding of continuity of operations, failover, and backup and recovery capabilities, with associated SLAs. Please describe the key elements of these capabilities and how you believe

Teresa Carlson
Vice President
Microsoft Federal

In a cloud environment, the principles of continuity of operations planning, failover and backup and recovery aren't much different from a traditional IT infrastructure. The big difference is that the potential scale of cloud computing ensures computing resources are available to agencies when they need them.

Large cloud providers offer environments that are worldwide in scale, with the ability to handle and route massive amounts of data. The data centers are enormous, and when there is spillover, or if a data center experiences a service interruption, traffic is automatically transferred to another datacenter with availability. The best cloud-based systems are redundant by design, with standardized processes for dealing with unexpected or unusual computing patterns. Not only does this provide greater flexibility and resources, but it allows providers to be completely transparent about where data is being stored or relocated to.

In terms of backup and recovery, large cloud environments provide capabilities that protect both the physical equipment and the applications themselves. Applications are replicated and stored in multiple data centers, so that if one location experiences a problem, the application can be accessed from a secondary data center. It's failover on steroids, and it's all because of scale. Major cloud providers build these capabilities from the ground up, and they add an incredible amount of resiliency to the entire operation.

Verifying these capabilities and ensuring effectiveness is a major issue not only for providers, but for legislators as well. Vendors have to be more accountable from a legal perspective especially when protecting sensitive government data and applications. Citizens and organizations need guaranteed access to secure data, and cloud vendors must be transparent about documentation and controls. To make this a reality, the U.S. needs to adapt its communications technology laws to reflect the modern computing environment. Microsoft's Brad Smith has called for the creation of a "Cloud Computing Advancement Act" to establish best practices and increase confidence in the privacy, security and resiliency of the cloud. Steps in the right direction include:

- Reforming the Electronic Communications Privacy Act to include stronger security

protections

- Updating the Computer Fraud and Abuse Act to provide law enforcement with the resources it needs to combat emerging forms of online crime
- Transparency provisions that ensure citizens and organizations have a right to know exactly how their information will be used, accessed and protected by service providers
- Initiate discussions with countries from around the world to establish global cloud standards, because it's not uncommon for data that originates in one country to be hosted in another Industry, government and consumer groups must work together to create legislation that encourages innovation while demanding security and protecting privacy. In the meantime, it's up to vendors to be completely transparent with government agencies about resiliency and continuity capabilities, and for agency IT leaders to demand adherence to industry best practices.

Posted: June 14, 2010

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP

It is interesting that the issues of continuity of operations, failover, and backup and recovery capabilities are great strengths of cloud computing. In fact, these may be areas to launch your cloud experience in. There are at least three possible ways to use the cloud in your continuity of operations plan (COOP) and backup plans. First, you can provision and use the cloud to be a backup site for a traditional data center application. Second, you can use a traditional data center to back up a cloud implementation. Third, you can use a cloud backup to a cloud implementation. With all three approaches, security is the place to begin. For the purposes of this discussion, I will assume that you are comfortable with your security platform (a discussion for another time). I have also heard of using the cloud to test the COOP process without using live data. This solves some of the security concerns.

Let us take the first scenario, which I believe is one of the most effective ways to get experience with the cloud. In this scenario, you would have your traditional data center in whatever state of virtualization you may be in, and you would use a cloud offering to provide COOP and/or backup. Backup may be the easiest to provide security for as you can keep data encrypted the entire time it is on the cloud. Either way, you would be able to procure the compute or storage capabilities on the fly and only pay for what you use. If one service or application is deemed mission critical today, then you can provide for it. If over time you remove that application from the mission critical list, you can ratchet back the appropriate cloud services. In this scenario, you could gain valuable experience without interfering with the day- to day operations of your information technology services.

In the second scenario, you could make use of your legacy data center to support your new cloud offering. This is not a likely scenario, but in some cases could be useful if the legacy environment were capable of backing up your new environment. In the third scenario, you would be moving to a total new environment and the issues would be very much like the ones discussed in the first scenario.

Now, a word about SLAs and the ability to determine if the uptime and service are acceptable. Obviously, the availability needs will depend on your mission and the needs that are required for your work. You should secure appropriate guarantees based on your needs. The good news is that most cloud providers have built in the precautions needed for such SLAs, such as power and communications redundancies and in some cases geographic diversity. The key is to know what the provider is providing and what you are responsible for. We are not to the point in maturity that you can treat the cloud as a black box. You should understand what the cloud is and how it is being provided to you.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: June 23, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

Public cloud services are clearing and illuminating the landscape of IT risk. Major cloud providers run homogeneous systems at nearly constant workload, with high degrees of automated or otherwise systematized management and fault mitigation. Further, the multi-tenant architecture of true clouds enables enormous reduction of points of failure and number of distinct failure modes improving reliability, and also enabling superior visibility into operational state (as demonstrated by public Web sites such as trust.salesforce.com/trust/status and status.aws.amazon.com).

The initial question of SLAs suggests a more important question. Wouldn't any customer organization, in public sector or private, prefer reliable service combined with ample warning of degradation or interruption rather than merely receiving after-the-fact credit for the price of any services not received Anything more than this would get into the realm of consequential damages, which is the domain of insurance companies rather than cloud service providers but statistically significant data from large cloud providers will enable, in the very near future, a far more efficient marketplace in such coverage. There's every reason to expect that insurers will reward cloud service providers who set high standards for transparency, operational excellence, and consistently high performance by giving the customers of those services better rates for service interruption policies.

We must compare service level confidence, not to a theoretical ideal, but to present day reality. At a conference in Singapore last year, a Red Hat executive observed that people ask all the time about the Service Level Agreement that they'll receive from a cloud service provider

while seeming not to notice that there are rarely any service level agreements to protect their on-premise IT assets.

"If the data center goes dark, or that server in the corner bursts into flame," Red Hat's Frank Feldmann rhetorically inquired, "do you have SLAs with the power company and the fire department?" Corporations don't have SLAs with such services, because those services respond all the time to similar incidents. This gives people a reasonable basis for estimating the risk they'll take by relying on those services, and lets them make a sound decision about any further steps they might need to follow (such as engaging private emergency-response firms, or employing their own local disaster teams). Cloud services' customers also have comparable options.

As more information services move to public clouds, agencies will be able to modernize their mitigation of IT risk. Today, every traditional data center is its own, uniquely configured, hand-built tower of toothpicks with failure modes that are not precisely the same as those of any other facility. The hardware in use, the software version and configuration, and the skills of the operators vary enormously from one such center to another and even from one work shift to the next. There's no statistically meaningful data on operational reliability, which means there's no sound basis for pricing risk.

The operational advantages of true clouds, and the relentless competitive pressure of a marketplace of transparent service measurement, will drive cloud-based IT to achieve new standards of assurance.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: June 25, 2010

Larry Pizette
Principal Software Systems Engineer
Formerly MITRE

Thank you, Major Dominguez, for sharing an insightful question that many civilian and Federal/DoD IT leaders are considering. Also, many thanks to the private sector leaders who shared their insight and perspective this month.

Continuity of operations, failover, and backup and recovery are capabilities that must be periodically tested and exercised, and this testing is not without service provider effort and expense. Consequently, from its initiation, a contract with a service provider should have language to account for the testing, expected performance levels, and expected measurement techniques for the COOP- related SLAs mentioned in the question. The contract will give the Government a means to exercise and verify the capabilities in an operational context. During the solicitation process it is useful for the Government to offer a document template of the service SLAs and their measurement techniques, along with a post- award operational verification plan. The content of these documents may be developed as part of the proposal by the potential vendors, as framed by the Government's requirements. A vendor's willingness to do special continuity testing for a government customer will likely depend on the business return the vendor anticipates, and not every procurement will provide that business return.

While COOP-focused requirements for all cloud deployment models flow from a systems engineering process driven by operational requirements, key elements include the specification of multiple distributed physical locations for data and processing capability, periodicity of backup activities, timing metrics, up- time expectations, failover recovery metrics, and network connectivity/throughput. It's important to include points of escalation and delineation of responsibility (e.g., responsibility for the network between a vendor and government organization). For community and public clouds, the agreement between the Government organization and the cloud provider should include mechanisms for verification and joint risk reduction.

Potential areas for Government risk reduction are audits, testing backup and recovery capabilities in concert with the provider, and review of periodic metrics provided by the provider to the Government to give visibility into their operational performance.

Please note that there are several efforts in the works across the Government in this general area. For example, the GSA is "currently undergoing a procurement to award multi- vendor BPAs for IaaS offerings available on Apps.gov." The request for quotation (RFQ) requires backup and recovery capabilities and an SLA with a minimum of 99.5% uptime. The RFQ has categories for Operational Management and Trouble Management. Similarly, for security certification and accreditation, government organizations may be able to leverage the upcoming Federal Risk Authorization and Management Program (FedRAMP).

In the influential white paper Above the Clouds: A Berkley View of Cloud Computing , the authors listed availability of a service as the number one obstacle to cloud computing. To many, the question is about trust: Will the capability be there all the time when I need it And will it be available under all circumstances, including crisis situations natural or man- made that may affect a broad area of infrastructure and other capabilities As noted by Major Dominguez's question, Government organizations must pay careful attention to COOP- related SLAs, as they are key to establishing trust and meeting operational needs.

Posted: July 2, 2010

Question for May 2010

Brian Shaw, DASN C4I/IO/Space, Director of Cyber Warfare asks: "How could a government system be more resilient to attack if hosted on a public cloud computing model vice a private one and what are the added vulnerabilities the government would need to consider ?"

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP

The public versus private cloud approach will be debated for some time. A key to understanding this debate is to distinguish between vulnerability and resiliency. With any public option that we are aware of, there are some inherent vulnerabilities that must be addressed. These vulnerabilities are beyond the security measures that take place in a private cloud. There are at least two big differences between public and private cloud offerings with regard to vulnerabilities.

First, the number of external groups that you are sharing space with and the level of trust you have with those groups. This is the reason that Community Cloud is so popular. In a community cloud you can get some economies of scale, yet limit access to external groups that you have a working relationship with as well as a level of trust. In a public cloud offering you have the whole internet in the transport portion of the cloud offering and you have the client base of your provider at the compute and storage part of the cloud transaction. Both of these are areas of concern that need to be addressed. In addition, you have to be concerned about insider threats from the provider themselves.

The second big difference in the vulnerability between public and private is that some of the security measures in your system are provided by someone else, namely the cloud provider. This means you have to trust your provider to do what they say they will in terms of security (although there are ways to verify it). You also have to understand what steps they will take to maintain security and make sure your security efforts mesh nicely with the provider's efforts. Recently, a number of cloud providers have been working to get their offerings Certified and Accredited (C&A) by the Federal Government. This effort should help alleviate some of this concern.

Interestingly, many of the situations that can cause extra risk in terms of vulnerability are the very things that can provide better resiliency. For example, there are many content management systems that can distribute your content across the Internet, making a Denial of Service (DOS) type attack much more difficult. When you have large systems that are virtualized, risk is reduced (if things are done right). In most public cloud offerings the scale is much greater than private cloud offerings. This scale (of public cloud offerings) can provide the resiliency that would be difficult in a smaller private cloud. As a general rule (if all other things are equal), when the size of the cloud increases the vulnerability and the resiliency both

increase. In this case, size is a function of how many players are involved in the cloud. It is our belief that the vulnerabilities will decrease in public offerings as the vendors become C&Aed and improve their security profiles. We also believe that public resiliency will continue to increase as economies of scale continue to grow and effective practices are implemented.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: May 12, 2010

Nicklous Combs
Chief Technology Officer
EMC Federal

Good Question, Although building private clouds will always provide a more resilient environment to attack than public clouds if built correctly, there are a few reasons that public clouds can be more resilient. The main reason is the

statement I made about "if built correctly". Public cloud providers will normally be subject matter experts in delivering resilient cloud solutions and therefore provide high availability environments at a great price point. As many organizations start to build private clouds, they may not have the expertise to build them correctly and some will no doubt create environments that are less than adequate to meet their security needs. This is why choosing a partner with experience is critical in moving towards a cloud environment.

Another reason public clouds could be more resilient is that most public cloud environments will be limited in the types of services that they can provide therefore decreasing the attack surface of the IT environment. Private clouds will no doubt have to support a much wider range of services which will make their attack surface much larger than public cloud offerings. As far as added vulnerabilities, there are many that need to be considered, there are way too many to cover in this short blog. Customers that need to provide non-repudiation, data erasure and cleansing procedures will be challenged when working with a vendor who provides multi-tenant environments through public clouds. Who is going to be doing the auditing in a public cloud? Do you really think companies like Coke and Pepsi will want to put their data in the same multi-tenant cloud controlled by someone else?

For more information: http://www.emc.com/ ?fromGlobalSiteSelect.

Posted: May 14, 2010

Ron Knode, Director,
GSS, LEF Research Associate
CSC

Fight Fire with Fire? in Clouds?

One of the most frequently used tools to fight forest fires is more fire! At first blush, this approach is counter- intuitive. But, the use of "back burns" to reduce the amount of flammable material and (ultimately) control the fire itself is a well- known and effective technique.

The irony of "fighting fire with fire" lies at the heart of this month's question. And, since the issue is equally relevant for both government and industry, let's restate the question as, "Can we use cloud processing to help solve the security and availability problems normally aggravated by cloud processing ?"

Once again, at first blush the answer would be "No." The security issues of cloud processing are well- advertised, and those issues continue to be the number one stumbling block to greater industry and government use of cloud processing of all types. The lack of transparency (especially in public clouds) is the root of most anxiety about cloud usage, and thus represents the biggest restraint on enterprise use of (public) cloud processing. Even if the contradiction of "fighting fire with fire" in the cloud can be successfully applied, this lack of transparency will still need to be overcome before industry and government are liberated to use (public) cloud processing for important mission functions.

Yet, there are some features of cloud processing that do suggest we can "fight cloud insecurity with cloud characteristics"! Consider, for example, the superior scalability, flexibility, adaptability, and redundancy of the public cloud. Then, imagine using those characteristics to deploy threat and vulnerability countermeasures in thousands of locations, many of which are (dynamically) placed closer to the threat source than any conventional static system. Such a dynamic operating characteristic would provide a new dimension for a classic "defense in depth" architecture, and could result in greatly improved resilience to attack. In particular, resistance to Distributed Denial of- Service (DDoS) attacks could be enhanced with use of a public cloud. This very same architectural model for public cloud usage has already demonstrated its effectiveness against one of the largest DDoS attacks against the U.S. government. The use of Akamai's EdgePlatform (public cloud) prevented a huge DDoS attack in July 2009 from disrupting operations of protected locations.

No doubt, we could imagine other examples of how certain cloud characteristics can be used to improve some security capabilities in some circumstances. Certainly, protection against DDoS is one good example. But, not every security need can be improved by such an ironic application of the cloud. (After all, we don't save drowning people by pouring more water on them!) And, every use of a cloud brings with it the issues of lost transparency for the cloud consumer (e.g., configurations unseen, vulnerabilities unmeasured,ccesses unreported, data and processing unanchored, ...).

So, the cloud can improve security in certain important ways. But, all fire, no matter how ironically used, is hot and dangerous. The cloud is no different.

See the full blog response at www.trustedcloudservices.com.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: May 17, 2010

Rick McEachern
EVP of Business Development
LongJump

Besides being affordable, cloud computing offers the opportunity to run within multiple distributed and replicated systems. For example, through Amazon EC2 and other IaaS (Infrastructure- as a- Service), your application servers are virtualized and replicated as you need them. This is likely the best approach to dealing with a physical attack because there is no physical box you can locate to attack. Instead, would- be attackers would instead be directly focused on cyber- warfare in the form of DOS or denial of service and intrusion.

The primary DOS attack happens through overloading existing services. Again, through an IaaS and in most hosting environments, when spikes happen, the IaaS provider generally increases capacity "elastically." You can then determine if you need to block out access from someone. Without an IaaS, monitors would indicate traffic increases and it's up to staff to respond by replicating servers to support the increased traffic or simply holding back access. You should keep in mind that DOS attacks can happen in any system where public networks are involved, not just in the Cloud.

For an intrusion attack where the goal is to access, steal, modify or delete the data itself, then it is all about the software and the software service provider.

Organizationally speaking, all the core system administration best practices must be adhered to. Have we set up proper roles, password policies, and limited IP addresses particularly for administrator access? Does stored data have the option of being encrypted Are there features to monitor changes in the data or to maintain access logs Have we limited views of sensitive data to only the select few? How secure are the APIs into the system? Does the system allow me to limit API calls?

And the service provider must be geared to handle vulnerabilities from their internal systems into their clients' systems. This means having an established set of policies in place and documented trust in their employees. In addition, the hosting facilities must be rock solid and trustworthy.

Even then, the one true option to circumvent many of these challenges is to consider working in a private cloud, which can offer the best of both elasticity and control. A growing number of PaaS (Platform as a- Service) vendors provide scalable, multi- tenant application platforms that work exactly like their public cloud offerings with the added flexibility to host all or some

of the tenants in private environments, including behind the firewall or on an IaaS with a protected URL. PaaS within a private cloud can be an ideal fit for data-sensitive, mission critical public sector computing, where you need to maintain public use of web-centric applications while maintaining controlled access to a secure backend. With a multi-tenant PaaS, government agencies can in a sense become cloud-based application vendors themselves, creating, managing, and charging back for web applications to their clients with complete control just like a SaaS provider.

For further information, please visit http://www.longjump.com.

Posted: May 17, 2010

Jeff Bergeron
Chief Technologist, U.S. Public Sector
HP

Resilience is, at least partly, one factor in a number of events necessary to take a system offline. A system housed on a single server instance in a single location on a single network no matter how secure the facility, is susceptible to any number of events that could take that server offline. Conversely, a system that can be rapidly instantiated on virtual servers across many providers, physical locations and networks is able to withstand events that would normally cause a system outage. In a distributed model, servers, facilities and networks are highly resilient and all cloud service providers would have to be compromised to fully disable the entire system.

The use of public cloud providers to increase resilience does introduce a number of potential new vulnerabilities, including the need to improve management and automation of provisioning in response to outage events or attacks, the protection of data and software spread across many locations that are not under the control of the government, the possibility of a Distributed Denial of Service (DDOS) attack due to an oversensitive migration response, and other attacks made possible when infrastructure control is delegated in a distributed fashion. The addition of smart, seamless, cloud management services capable of sensing vulnerabilities within the cloud and dynamically reallocating resources based on these threats could provide additional safeguards. Cloud resilience could be obtained through the use of geographically dispersed cloud service providers operating in a secure Inter-cloud service model for transparent service and data movement between trusted public cloud providers.

Posted: May 21, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

There are compelling reasons for government IT to adopt the public cloud. The GSA has estimated that Web site upgrades formerly requiring six months are done in the cloud in a day. The U.S. Census Bureau used a cloud platform to achieve 12- week deployment of a system to

manage its nationwide temporary labor force. The Family Service Agency of San Francisco estimates 50% reduction of administrative time, combined with improved outcomes tracking, thanks to cloud- based re engineering of mental health case management.

Public clouds can readily handle sudden bursts of activity. The government must be the responder of first resort to massive but infrequent workloads associated with natural or man-made disaster, but it's wasteful to provision static resources—regardless of their mode of operation that will be idle for most of their useful life. Public clouds provide scalable capacity on demand.

For all these reasons, agencies at every level of government are getting a green light from Washington D.C. to pursue cloud options. Casey Coleman, GSA's CIO, stated last year that "We will...work with industry to ensure cloud- based solutions are secure and compliant." Coleman's statement both acknowledges, and vows to address, the appropriate commitment of those who hold public trusts.

Public cloud providers must make the case for their reliability, security, and governability but they are already doing this in financial services, health care, education, and any number of other domains in which cloud services already enjoy wide acceptance. Public clouds are forced to provide "sum of all fears" protection that addresses the demands of the most demanding customer to the benefit of all customers.

The magnitude of threats to government IT will rise as agencies expand their use of public-facing Web sites to make services more available to citizens. Dean Turner, director of Symantec's Global Intelligence Network, puts it simply when he says that "[Attackers] aren't breaking into your network. They don't have to. You are going to them." Governments will be targets for a growing range of increasingly sophisticated attacks and these will arise, not only from the connected outside world, but also from within.

Every subscribing organization, therefore, will still need to assign privileges appropriately, audit actions effectively, and control access to information on its way in and out of the system. This is not a new problem arising from the cloud, but the agency using a public cloud can focus its resources on its own specific mission while common concerns are addressed by the public cloud service provider, with attendant massive economies of scale.

The issues for discussion in the public cloud are not qualitatively different from those already encountered as governments continue their turn toward use of public networks and Web resources, for all the reasons discussed above. What's different in the enterprise public cloud is the far more affordable cost of providing the rigor, accountability and transparency that the market demands to meet the needs of serious customers whether in the private sector, or in the pursuit of the people's business.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: May 23, 2010

Simon Crosby
CTO, Data Center and Cloud Division
Citrix

Picking a cloud is like picking pizza. Give it a try.

There is a misconception that public clouds are risky, but private clouds offer benefits with no downside. While there are legitimate concerns about the maturity of public clouds, where their service abstractions match application needs they can offer a superior service. For example:

You can encrypt your data. Combined with secure access control and opaque object name spaces you can make the likelihood of data leakage effectively zero. Fewer humans, simpler, infrastructure services, and secure isolation at multiple layers can offer better security.

By virtue of scale and geographical distribution, clouds can make applications available under conditions that would render your private cloud useless, including attacks and failures.

Finally, because of their rich connectivity, they are far better placed to deliver applications to end users who are geographically dispersed.

Consider three types of public clouds using an analogy of ordering pizza: "Margherita (with add-ons)", "Build your Own", and "Ready to eat". Compare each to the private cloud equivalent: build the oven, chop wood, light the fire, prepare dough and toppings and then create your dream pizza.

"Margherita (with add-ons)" This kind of cloud offers a set of simple, powerful service abstractions that are easy to understand, with additional services that can be added a- la- carte. None of the services can be changed, but you can combine them. A good example is Amazon Web Services (AWS).

The chief drawback is the relative simplicity of the service abstractions, which limits their applicability. Moreover you can't install an IDS, firewall or router in your virtual network, and they lack audit, role based access controls and other management abstractions. However, if the service model fits your app well, you benefit from lowest possible cost, and maximum scale. "Margherita" clouds are well suited to web apps, even if you retain the database tier in the enterprise, and they are increasingly useful for other enterprise apps.

"Build your Own" public clouds

These providers allow customers to build rich, isolated, private enterprise infrastructures that are verifiably secure. They are typified by hosters and carriers that add elastic resource consumption. Example: Carpathia which serves many agencies of the Federal Government today.

Such clouds let you choose your core infrastructure, layering on top security, compliance testing, auditing, granular access controls and rich management capabilities. They are fully certified to host highly regulated apps/data and use hardened facilities, with redundancy, replication, and high availability. To this they add elastic compute, network and storage, with granular access control. You can also dynamically instantiate network capabilities such as IDS, firewalling, load

balancing and PCI compliance. In summary, there are no enterprise apps that cannot run in this type of cloud, and all the benefits of cloud apply.

"Ready to Eat"

These are managed service providers who host their own services. Examples: hosted virtual desktops, Disaster Recovery, or managed email. An organization with specific competence in the service also runs it, combining economies of scale and automation with their service- specific operations skills to deliver a lower cost service with strong SLAs. Dedicated or elastic services are available, making these useful wherever it makes sense to outsource a traditional enterprise function.

For more information, please see Simon's blog.

Posted: May 27, 2010

Jim Young
DoD Manager
Google

Thank you Brian for your question as it raises many issues and questions that should be addressed by providers.

Skilled administrators can run Internet- based services in a highly controlled traditional environment in which certain security controls are assumed, but flexibility and innovation on the system are likely to be negatively impacted. Organizations responsible for Internet- facing networks can offer much more flexible services that dynamically scale more elastically, but they also must be particularly vigilant about ensuring security because the networks are exposed beyond the specific organization.

Even so, cloud computing can in many cases be as secure, if not more secure, than traditional on premise environments. To understand how, it's important to consider how most agencies and departments, and their respective networks, function today. Often, those that run client applications and have to manage a large heterogeneous environment encompassing different operating systems, platforms, and devices running multiple versions of applications, deal constantly with security patching issues and challenges that sometimes disrupt operations in the process. It is this variation that introduces complexity, increases attack avenues, opens larger windows of exposure, and leads to more security vulnerabilities in traditional networks.

NIST has done an excellent job of defining cloud computing definitions so that we have common vocabulary, plus they identified key security areas that need to be discussed with any provider. To help simplify for organizations like yours, the Federal Risk and Authorization Management Program (FedRAMP), is a unified government- wide risk management program focused on large outsourced and multi- agency systems. The program will initially focus on cloud computing but will expand to other domains as the program matures. FedRAMP provides security authorizations and continuous monitoring of shared systems that can be leveraged by agencies to

both reduce their security compliance burden and provide them highly effective security services.

FISMA security controls and follow- on, continuous monitoring requirements can be directly applied to public cloud computing models often more effectively than for current on premise systems. In addition, providers can offer community clouds that only host US government data. For more information on this, see http://googleenterprise.blogspot.com/2009/09/google- apps and- government.html.

Posted: May 28, 2010

Teresa Carlson
Vice President
Microsoft Federal

The great part about cloud computing is that government organizations have choice. Some data makes sense in the cloud and some data may not. It's not an all or nothing discussion. Security and privacy are rightly the top concerns for most government leaders, and some are far more comfortable housing sensitive information on premise. That's OK. Agencies should move to the cloud as they're ready, and when they do, they have both public and private options to choose from.

In terms of resiliency, public clouds are extremely robust. They offer scale in terms of underlying architecture that private cloud infrastructures typically don't provide. The ability to move traffic and data throughout large (or multiple) data centers is a major advantage when it comes to performance, but it's also an advantage in terms of resiliency to attack. Public cloud providers understand the inner workings of large data center environments what data should be flowing, how it should be flowing, and when something doesn't seem quite right. This experience and knowledge enables public cloud administrators to spot anomalies quickly, and take action against a threat when it's required. Large public cloud offerings also allow application and data to be replicated and stored in other data center locations, ensuring critical information isn't lost during a targeted attack.

The risk lies less in the public vs. private argument than in the fact that it's a whole new approach to computing. Virtualized computing requires a different mindset, and it brings the potential for a whole new class of vulnerabilities. As cloud adoption increases additional threats may emerge, which is why government agencies need to be thoughtful about where their data is hosted. Hosting data on premise doesn't necessarily guarantee that it's more secure, which is why rigorous security methodologies need to be implemented at the outset of development in any cloud environment. Adhering to the best IT security standards that exist today like ISO 27001, FISMA, SAS 70 Type 1 and HIPPA not only ensures the highest levels of data protection, but also increases transparency by providing government leaders with specifics on how it's being protected.

Cloud is a major paradigm shift that government leaders are still wrapping their heads around. Data centers aren't just a room with a bunch of servers anymore. Computing has become a

utility a virtualized infrastructure that scales in accordance with need. The balance lies in maximizing these incredible efficiencies with robust privacy and security controls in place.

Posted: June 14, 2010

Emily Hawthorn
Principal Infosec Engineer/Scientist
MITRE

Thanks to our respondents for their very thoughtful remarks!

The European Network for Information Security Agency (ENISA) lists resiliency as market differentiator that will drive cloud service providers. ENISA states, "Security is a priority concern for many cloud customers; many of them will make buying choices on the basis of the reputation for confidentiality, integrity and resilience of, and the security services offered by, a provider." Public clouds offer the potential of resiliency through a number of means including the transparent use of multiple physical sites, redundant networking paths, and the automation of many administration tasks to backup data across physical boundaries. Private clouds can be engineered to provide the same benefits, though rather than leveraging the potentially large capital investment of a multi-tenant service provider, the private cloud provider must fund, secure, and manage the capabilities internally.

The following use cases illustrate the interplay of resiliency and vulnerability in cloud computing:

First, consider moving user desktops to a cloud service, where when users log into the network, they get a virtual desktop with the expected look and feel, however a fresh virtual machine image can be downloaded each time the user logs in. The first benefit is that no matter what malware a user encounters, it is flushed from the virtual desktop when the user logs out. The next time the user logs in, they can get another clean machine image. Additionally, when patches are required, the "gold disk image" need only be patched once, and the new image is delivered consistently at each user login.

For a second use case, please consider moving an enterprise service such as email into the cloud. The user's email client connects to one of several email servers that may be hosted in a virtual environment. If one virtual machine (VM) or server fails, the user session can automatically migrate to another VM. As defined by SLAs, service to the user could continue uninterrupted, as the provider manages the resiliency of their offering including security, configuration, performance, and patching.

Potential vulnerabilities of public cloud services can come in many forms. For example, multi-tenancy brings with it the risks of attack from within the infrastructure, by another customer of the same service, and the virtualization mechanisms can also expose an attack surface. However, as recently stated on the Navy CIO's website, "There is […] some good news. Since most of the underlying building blocks (e.g., servers, network and storage devices, and software—operating systems and applications) of cloud computing remain the same as those

used in traditional information technology systems, much of the existing security policies, practices and solutions can be readily repurposed to fit the new cloud computing paradigm." Still, given the number of newly initiated and planned projects employing cloud computing, many organizations (e.g., University of California, San Diego, Massachusetts Institute of Technology) are researching new attack vectors for a better understanding of cloud implementation vulnerabilities and mitigations.

Posted: May 28, 2010

Question for April 2010

> Private cloud computing implementations: What do you consider to be the essential components or capabilities, necessary to create a private cloud computing environment?

(Responses posted on an ongoing basis in April.)

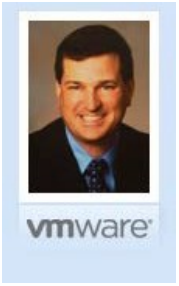

David P. Hunter
Chief Technology Officer
VMware Public Sector

There is no doubt; Cloud Computing is the talk of the town. A common perception is that cloud computing implies an external cloud, based on public cloud services. The fact is that cloud computing is how you approach IT and a way of doing computing differently. Most governmental agencies can benefit from adopting and evolving their existing infrastructure to a private cloud computing approach today within their own datacenters. As with any journey you need a starting point, a destination and a map detailing the path.

Like many disruptive events, the motivation for moving to a Private Cloud is the existence of an inflection point that necessitates a shift. Today's complex, often brittle, IT infrastructure and the desire to simplify IT is that inflection point. The current economic realities and budgetary climate in Washington and State Capitals, along with the desire to save energy, reduce capital and operational expenditures, share resources and provide more transparency to citizens on how tax dollars are spent are adding additional rationale for governments to implement Private Clouds to address this complexity.

The initial capability to begin the journey is a fully virtualized datacenter. Virtualized resources allow the pooling of compute, network and storage resources that are then shared across applications and users which enables on demand resource allocation. Once virtualized the following check-points along the route to a Private Cloud need to be considered:

- **Automation of management operations.** The result will be a zero-touch Infrastructure model that is driven by policies that automate routine tasks, minimize operational expenses and overhead.
- **Development of service delivery models.** Standard definitions for services and service levels allow the reduction of the number of variations supported and allow for the enforcement of standard methods and procedures that can then be easily automated.
- **Metering of Services.** A metering model for tiered services will allow the business

to have a transparent view of the cost associated with the various lines of business applications.

- **Providing user self-service capabilities.** The provisioning and deployment of services within the parameters of defined business and governance policies provides a distributed, time essential execution capability that is controlled by automated policies increases the ability of a agencies and departments to react to changing requirements while maintaining compliance with centralized policies and security models.
- **Open and interoperable standards.** Application mobility between clouds within a common management model, based on open standards, extending to other public or private clouds is a key condition to achieve flexibility.

The promise of reduced capital and operational expenses, higher customer satisfaction and greater control over security are some of the contributors today that are convincing governments to move to Private Clouds. Those organizations that chose to transform their computing model on a robust platform that provides core features such as high availability, the ability to optimize resource allocations to ensure service levels; built-in disaster recovery mechanisms to ensure business continuity; a security model that encompasses dynamic infrastructure and boundaries; and application-aware infrastructure to self-optimize application performance will be in the best position to achieve the promise of the Private Cloud.

For further information, please contact David Hunter at hunter@vmware.com .

Posted: April 15, 2010

Rich Wolski
Chief Technology Officer
Eucalyptus Systems, Inc.

Implementing a Private Cloud

Private clouds consist of several components, not all of which are technological. The cloud platform itself is deployed as one or more technologies, but in addition to the platform, the organization building the cloud must also define policies governing its usage, processes describing its maintenance, accounting schemes for determining its budget, and plans for managing its lifecycle. From a technical perspective, the cloud platform must be able to support these management activities in addition to the cloud abstractions it implements for its users and administrators.

One way to meet these requirements is to architect the cloud platform so that it can conform the infrastructure upon which it is deployed, particularly with respect to the mechanisms with

which policy is implemented. That is, the private cloud platform must be able to accept infrastructure governance defined for its environment rather than dictate governance requirements.

Open-source as a distribution style for the cloud platform software is particularly good at facilitating this form of policy malleability. Community contributions often take the form of modifications to specific configurations and environments. The source code is available so that customization is possible, and the interaction of the platform and the infrastructure is transparent.

From a more technological perspective, the "scale" of the private cloud platform is often a metric of great interest. There are two types of scale, however, that must be considered: request scale and resource scale. Request scale refers to the number of requests (usually from separate users) that the cloud can support per unit time. For IaaS-style clouds, these requests are transactional. That is, each request must either complete or fail unambiguously, usually within a specific timeout period (we at Eucalyptus use 60 seconds).

On the backend, the cloud platform must be able to use (efficiently) large collections of widely varying resources (machines, networks, storage devices, etc.) The key to achieving both user scale and resource scale reliably is to exploit eventual consistency in the internal state management of the cloud platform itself. As with user-facing cloud storage abstractions (e.g. "blob storage"), eventual consistency enables both reliable operation and vast resource scale. Managing eventual consistency, particularly to implement the platform, can be complex but it is the purpose of the cloud to hide that complexity in the cloud platform so that it is not exposed to the applications, the users, or the cloud operators.

Finally, private clouds must implement cloud provisioning abstractions. Virtual machines in a cloud, for example are similar to but not exactly like virtual machines in a data center. The same relationship exists between cloud Internet addresses, storage abstractions, firewall rules, etc. The cloud is a more dynamic usage model, and as a result, a more efficient model for managing IT resources. To exploit the maximum benefit it offers, it must support services that allow applications to take advantage of this dynamism.

For further information, please contact Rich Wolski at rich@eucalyptus.com or visit www.eucalyptus.com.

Posted: April 19, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

The question of "private cloud" versus "public cloud" arises when people think of cloud computing as a model of technology deployment. That's a path that leads to superficial economies, and leads away from the most transforming results of adopting the cloud computing model.

If an organization decides that it needs "a cloud computing strategy," it's likely to take its existing IT practices and look for a way to migrate those practices into a scalable environment with a high degree of resource virtualization. Any number of vendors will be happy to offer hardware and software to support those aims, and most of those vendors will today use the label of "private cloud" to describe the result.

What people actually want, when they talk about cloud computing, is a far more radical improvement in the way that they acquire and use information management and business process automation. Cloud computing is far less a technology model than it is a model of service delivery. It's a set of promises that service providers make to their customers: promises that may have been made in the past, but are only being truly fulfilled today.

- Cloud computing is a promise that a business process initiative can get off to a rapid start, focusing on the problem to be solved not on the limits and delays of a capital budgeting process.
- Cloud computing is a promise that the customer's scarce resources can be reserved for the creation of competitive advantage, with the service provider assuming the burdens of maintaining the security and performance of the software stack beneath the customer's applications.

Enterprise cloud computing is a distinct category of cloud computing, as opposed to consumer Web applications:

- "Enterprise cloud" implies a further promise of rigorous and audited security, high availability and robust capabilities for customization and integration.

If it's cloudy, the customer shouldn't need to purchase and support peak- load capacity that exceeds the everyday need. If it's cloudy, the workload of security patches and other updates should be the provider's problem; the benefits of continual upgrades should be an unmixed blessing for the customer, included in predictable subscription pricing.

Is it possible to deliver the cloud's distinctive advantages in an on premise installation, or in a reserved instance of hardware located off the customer's site It's possible, but most "private cloud" efforts let the adjective vastly overshadow the noun and wind up constructing a best practice data center, using technologies such as virtualization that substantially improve hardware utilization, but falling well short of the full potential of the cloud computing model as applied to appropriate tasks.

It's well and good to like the idea of "private," but that goal should not take precedence over the compelling economics and the benefits to business agility that come from being "cloudy."

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: April 25, 2010

Teresa Carlson
Vice President
Microsoft Federal

The idea of a "private cloud" really starts with how you define it, and there are many different definitions out there. NIST defines a private cloud as "cloud infrastructure operated solely for an organization." That's a good distinction, but others push it further and demand that data be hosted within a certain facility. Some define private clouds as a way to access services within an infrastructure that is closed by design without connection to the Internet. Along this line of thinking, you could say that traditional hosting providers and Federal Systems Integrators have been offering a type of private cloud for decades. For me it's really about changing the traditional hosting paradigm to allow efficient access to services on demand with a pay- as you- go consumption model. The procurement characteristics change, as does the concept of metered service, but the basic premise of the cloud is not a science fiction project.

The Internet has become part of cloud's evolution because of its prolific adoption worldwide and the ever growing consumerization of IT. Citizens are increasingly using Web tools to communicate, make purchases and access information blurring the line between enterprise and consumer based solutions. People are expecting more from IT both at home and at work, and it's forcing government's hand as it strives to connect with citizens, attract workforce talent, offer services more efficiently and become more open and transparent.

The cloud industry is currently building solutions within private data centers based on existing best practices in security, privacy and governance models. But leaders still have security concerns because they can't touch the servers and customize the solution to the granular level of detail they are accustomed to. To alleviate these concerns, there are great questions being asked by industry stakeholders: Should existing standards be modified to fit the cloud? How do government agencies know that C&A requirements are being met in public cloud solutions What if there is a data breach or data leakage?

The challenge involves highly specialized systems, perception and lack of maturity. The thought of hosting data in a non- government data center, on public servers owned by third party vendors, has always raised security and privacy concerns for government agencies. Traditional hosting solutions eased some of this concern by adhering to Federal C&A and allowing government customers to customize, audit and access facilities all of which drive up the cost and time to market of the solution. The promise of cloud computing takes traditional hosting to the next level, offering commodity based services that are cheaper and faster to market because they are not unique services tailored to individual agency needs.

The end state goal for the cloud has always been "dynamic IT" the ability to deliver computing services to people, devices and applications when and where they need them in a metered, only-pay- for- what you- consume procurement model. Business models will change, standards will emerge and innovation will happen at rapid paces, but the need for choice and private clouds will be here for some time. It's a vision and journey that I'm excited to be part of!

Posted: May 11, 2010

Larry Pizette
Principal Software Systems Engineer
Formerly MITRE

The essential components and capabilities necessary for a private cloud ultimately depend on the system owner's requirements. The ability to control the operational environment is one of the significant factors that Federal IT leaders will likely consider when adopting a private cloud approach. As an example, if an organization requires very high levels of security, they may employ a rigorous architectural approach with comprehensive protections, including highly secure data centers and dedicated networks. Others may not need the same level of security but may require special features for regulatory or statutory compliance. For these two separate cases, private cloud implementations based upon owner requirements may look very different.

A private cloud implementation can offer significant benefit to those Federal IT leaders seeking to realize some of the benefits of cloud computing while maximizing control over their environment. As a starting point, NIST has listed several *essential characteristics* of cloud computing that Federal IT leaders embarking on a private cloud investment would benefit from examining. The essential characteristics that they list are: on demand self-service, broad network access, resource pooling (e.g., multi-tenancy), rapid elasticity (e.g., rapid scaling), and measured service.

In order to securely leverage the capabilities of a private cloud, an organization would need to ensure their data centers have the correct technical underpinnings and implement the appropriate operational processes, governance and management. While new components and legacy components will vary by organization, there are common *essential components* that will likely be needed:

- **Virtualization** allows multiple instances of "guest" operating systems to run concurrently on the same physical infrastructure and enables "multi tenancy," which is the sharing of physical resources. The resulting increase in server utilization can reduce HVAC and electric costs, data center size and other related infrastructure costs. Also, contemporary virtualization offerings can facilitate scalability, self-service provisioning and continuity of operations (COOP).
- **Storage** technology, such as disk arrays, storage area networks (SANs) and storage connection technologies, with supporting software can provide the underlying persistent storage and facilitate COOP and location independent access.
- **Security** capabilities such as identity management, logging and auditing, anti malware software, intrusion detection systems and intrusion prevention systems, and virtual machine isolation should be considered.
- **Provisioning tools, management tools, and metering instrumentation** are key to

providing Federal IT leaders and users with many of the advantages that a private cloud can offer: self- service, burst capability, scheduling, service- level agreement (SLA) monitoring, and if needed, metering for "pay as you go" functionality.

- **Networking** infrastructure for the cloud should be engineered to carry the additional traffic required to connect the service provider to the consumer.

"There is a major trend playing out over the next few years where internal IT providers want to make fundamental changes so that they behave and provide similar benefits (on smaller scale) as cloud computing providers," states Gartner's Thomas Bittman. For those organizations looking to maximize the ROI of their internal IT investment and maintain control, a private cloud may be an attractive option.

Posted: April 29, 2010

Question for March 2010

Cloud computing pilots: In a recent NASA Nebula blog post, it was suggested that the Government perform cloud computing pilots. What are the pros and cons of piloting cloud efforts and what outcomes should the Government look for in cloud pilots?

Ron Knode, Director,
GSS, LEF Research Associate
CSC

Piloting Through Clouds

Pilots! We all love pilots! Not the "wild blue yonder" kind, but the sampling, experimenting, exploring, validating, try it on for size kind. And, it's not just governments that appreciate the value of pilots. Enterprises of all kinds (public or private, supplier or consumer, large or small) have recognized the potential benefits of pilots and generally endorse them as part of larger development or acquisition models. Whenever new products or processing methods or application innovations show up, pilots are among the first techniques chosen to examine the validity and payoff potential of that "new thing".

That's why it was not surprising to see the recent Nebula blog post suggesting that the Government perform cloud computing pilots. What better way to corroborate technology claims and gain experience with cloud deployment and use?

Yet, there is also a seductive danger in pilots. In our zeal to investigate the technologies themselves, and in our eagerness to claim technology's benefits, we often concentrate our pilots on the technologies themselves, while giving little consideration to the operating consequences and governance impacts that are also affected. So, we regularly end up with encouraging results that promise good performance or interoperability or ease of use or maybe even security capabilities. But, when we try to extend the pilot results to an actual operating need, we are surprised that the expected payoffs don't appear, or that we have such difficulties in making the "piloted" technologies work for us as we thought they would.

Research indicates that cloud computing represents an evolution in technology , but a revolution in business. Consequently, pilots that deal specifically with the operating models, IT governance structures, organization, legal, acquisition, accreditation, and human resources alternatives and impacts are as important for government (and industry) cloud services validation as are technology experiments. So, let's encourage the government to initiate cloud pilot efforts that are devoted to the business issues of cloud processing, i.e., the operational application and usage issues. How about a pilot specifically designed to explore what types of certification and accreditation (C&A) doctrine and methodologies are needed and useful in (government) cloud processing With the recent cancellation of the GSA cloud procurement, current presumptions of "whole cloud certification" deserve to be examined. How about other pilots to deal with

acquisition models or government workforce needs for cloud processing There are many different and important non- technology dimensions for cloud processing that beg for pilots. Let's make sure we tend to the most disruptive parts of cloud processing while we're collecting and building clouds with all kinds of technology choices.

In truth, however, technology pilots are far easier to construct and conduct than business issue pilots or consumption pilots or governance pilots. So we must take care not to be distracted by cloud technology. Let's make sure that we don't stare directly into the sun of technology for too long and become blind to the realities and influences of doctrine, policy, legacy investment, organization, training, and concepts of operations.

See the full blog response at www.trustedcloudservices.com.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: March 18, 2010

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP

Yes, Agencies should pilot Cloud Services. You have to determine if you believe that Cloud is here to stay or just a fad. Most would agree that the Cloud Computing train has left the station. Cloud Computing will transform business in a very significant way. Like any new technology, there are early adopters, and there are great skeptics. If you agree with the potential value, then it only makes sense to learn all you can, even if you are not ready to jump in with both feet. One of the main road blocks is the lack of trust in securing this new technology. Pilots are a great way to test any security strategies.

For Pilots we recommend the following principles: 1. Start Small; 2. Learn From Partners; 3. Customize Cloud Service to meet your needs; 4. Build new Private Clouds to start; 5. Expand with Hybrid or Community Clouds. The overall need is to understand which business services and functionality can be abstracted and provided through cloud computing services and which can't. The goal enabling new mission capabilities to develop over time. Cloud Computing has a learning curve to understand how changes will effect operations, organizations, and support of users.

We spoke with several CIOs last week who are early evaluators. Most of their discoveries were as you might expect they were able to evaluate the service, test functionality, responsiveness, types of access, and limitations of the overall service. They also uncovered details that weren't in the original scope, such as how data uploads were throttled by some providers, or existing gaps in training. One side benefit realized was the improved dialogue with their users. Some discovered a mixed benefit that users enjoyed some of the additional capabilities available via the Cloud that they are eager to adopt, but the CIO wasn’t ready to consider.

New strategies should be developed to understanding how the agency's data is protected and controlled, and if the service provider can support the agency's business practices not as the service provider thinks best, but rather how the government needs work to be done. Piloting and

comparing public and private options can help you separate the truly important capabilities from the hype. Remember, at this stage in the evolution we need to understand and accept how things are being done in a "black box" not just taking their word for it. The vendors that are willing to help the customer with this process will be the winners in the Federal marketplace.

There are many lessons to be learned, but one of the first is this: stay open, stay flexible, and don't commit to anything that restricts your freedom to change.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: March 22, 2010

Steven Lebowitz
IT Architect, Federal Cloud Team
IBM

The question is not IF Government agencies should pilot cloud technology, but HOW they should do so. Cloud computing is not the correct answer to every question. Agencies should perform pilots in order to gain experience with the technology and understand WHEN it is appropriate. To that end, agencies and organizations should examine their business processes for "low hanging fruit" opportunities where cloud computing may offer a reasonable and measurable ROI and the opportunity to baseline and repeatedly measure a set of KPIs. This will allow for an objective measure of the value of cloud computing for each pilot project. Pilots should also be tied to real (but not mission critical) programs, and not simply placed in IR&D settings. By placing them in real projects, agencies can encourage their use, as well as gain practical benefits and experience.

A great example of this would be to focus on software development and test. Organizations devote significant resources to acquiring development and test environments. These systems are often under the LEAST amount of management control and are significantly under utilized. But a significant percentage of software defects can be traced to "drift" from baseline standards on these systems. By deploying new "instances" of test resources from managed "templates" as required, this drift can be eliminated. Further, the number of physical systems required to support such test environments can be reduced with a corresponding reduction in capital expenditures. The time to acquire and deploy test environments is often measured in months. By leveraging a centrally managed, virtualized environment, these test environments can be deployed in a matter of hours or minutes. IBM's fundamental strategy for cloud computing offers a "workload" focus as opposed to general purpose IaaS. To that end, we have created offerings for private cloud implementations and services on the IBM Cloud which directly address this development and test problem, thus giving clients a range of options which can immediately offer value across the enterprise and quick return on investment.

In short, YES, Government agencies should pilot cloud technology. However they should do so in a way that offers defined and objective measures of success, has the chance to provide real ROI, and targets real, but not mission critical aspects of their daily work.

For further information, please contact Steven Lebowitz at lebowits@us.ibm.com.

Posted: March 23, 2010

Douglas J. Bourgeois
Director, National Business Center
Department of the Interior

Pilot projects have long been used to manage risk and increase probability of success for full implementation. For this reason, and this reason alone, proceeding with pilot projects across the Federal Government would be a wise course of action. Visibility in the cloud can be hazy with numerous hazards to be avoided along the way. Even the most simple applications (e.g. collaboration, messaging, web 2.0, etc.) will create data, bringing data protection and privacy issues to the forefront. One of the easiest ways to begin a cloud pilot project is to initiate a software development effort via a cloud service. Even this presents eventual data security issues as testing activities progress. In addition, the ambiguity associated with the migration of cloud production applications into a private data center raises important risk based questions prior to production launch. To mitigate these and the numerous other cloud based risks, proceeding with caution is the way to move forward.

Each pilot project will have certain desired outcomes that depend primarily upon the type of project. Yet, there is a set of common criteria that any cloud project should consider. First is ROI. Every project has a cost and the return must be realized for projects to be expanded or duplicated. Second is standards. One of the more substantial risks is the lack of standards. So, pilot projects should aim to identify standards and even establish standard reference architectures that can be reused government wide. This approach would provide a "multiplier effect" for subsequent projects to obtain value from previous efforts. Finally, failure should be expected and mined for important lessons learned to also be shared, thereby reducing risk of subsequent projects as well.

Posted: March 23, 2010

Teresa Carlson
Vice President
Microsoft Federal

Cloud pilots are a great idea, but to maximize their effectiveness agencies have to pick the right projects and define specific success criteria at the outset.

Cloud is here to stay because of the promise it holds in reducing costs and improving efficiency for both government and private sector organizations. But it's a major cultural and technological shift for leaders, and pilots can help increase the speed of adoption by providing familiarity, experience and tangible results.

Cloud isn't an all or nothing proposition, and it's OK to start small. We tell our customers that they can't try to boil the ocean. The key is to identify targeted initial projects that make sense. Agencies probably don't want to jump into the cloud with a pilot involving sensitive or mission-critical data. Security is the concern we hear most often from customers considering cloud adoption, which is why low risk data sets make the most sense for pilots. And it's a trial, so why not start with information that is already publicly available? The Open Government Directive mandates that federal agencies publish high-value data sets registered via Data.gov, and these are great opportunities for cloud pilots.

The closer the pilot experience is to a real-world implementation, the more valuable it becomes in terms of mitigating concerns and showing benefits. A great pilot provides intelligence that can be used for long-term cloud success, and often holds the potential to scale into a full implementation. If an agency is considering creating a dedicated cloud, the pilot should reflect that. If leaders are looking into the pros and cons of a shared environment, create a multi-tenant pilot. The next step is defining success. It can't just be about avoiding breaches. There needs to be quantitative analyses of up time and total cost of ownership. Determine what a project would cost to execute using a traditional IT approach, and then contrast that with pilot results. The GSA recently announced savings of $1.7 million per year by moving USA.gov to the cloud, so it's clear that there are enormous benefits to be realized from cloud adoption.

Not every workload is right for the cloud, and pilots are a great way to gain the skills and experience needed to make informed decisions. We encourage agencies to examine a variety of models and vendor solutions in order to increase efficiency and maximize taxpayer dollars.

Posted: March 31, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

Pilot projects are a consistently effective means of exploring the cloud computing model. Cloud initiatives are not hampered by the up- front capital requirements that rob on premise experiments of critical startup momentum. Rapid construction of cloud applications, especially with the drag- and- drop immediacy of metadata based customization, provides a tight feedback loop that promotes energetic engagement by the owners of the problem being solved: this improves the initial quality of the solution, and increases the likelihood of adoption and subsequent feedback.

When process stakeholders see the speed, capability, and cost- effectiveness of the cloud, many common objections make a crucial shift from "obstacle" to mere "issue." When the benefits of the cloud are hypothetical, reasonable concerns

about security and integration can all too easily become an excuse to do nothing. When cloud benefits become real and tangible, present for all to see, the beneficiaries go into problem solving mode—and will almost invariably find affordable and satisfactory solutions.

Cloud pilot projects should be evaluated on their speed in solving problems and their flexibility in accommodating changing requirements: two criteria on which traditional IT models have often proved disappointing. Even if the initial and operational costs of the cloud were merely competitive with those of earlier models, faster and better solutions are a compelling advantage.

Cloud pilot projects should also be fully compared, though, against earlier IT models on the basis of total life- cycle cost. It is essential that these comparisons recognize, fairly and fully, the ongoing costs of maintaining and updating the in place legacy technology—rather than treating the status quo as being "bought and paid for." Space, power, cooling, hardware refresh, software patching and regression testing are huge fractions of the cost of on premise IT: these costs must not be overlooked when the "all- in" costs of cloud service subscriptions are being weighed.

A successful cloud pilot project will generate strong demand from other process owners for equal access to the outstanding outcomes and compelling economies that the cloud will typically provide. A successful cloud pilot project will also identify opportunities for effective integration among processes and resources that have previously been separated by boundaries of technology, and will pave the way for outcomes- based evaluation of activities that have previously been able to measure only the inputs that they consume. Finally, a successful cloud pilot project will result in people asking, "Why aren't we doing everything this way ?"—instead of asking, "Why is it so much easier to get stuff done at home than it is at work?"

People who use Gmail and Facebook at home; who buy merchandise on Amazon.com and sell collectibles on eBay; who see a startup business using salesforce.com, and never needing to buy a server: these are people who want a chance to work this way. Proposed pilot projects in the cloud will find people more than ready to make it happen.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: March 31, 2010

Geoff Raines
Principal Software Systems Engineer
MITRE

Thanks to all the submitters for their insightful contributions this month. One can see a theme in the responses above pilots are considered a useful tool for risk mitigation and requirements clarification in Federal information technology (IT) programs. David Wyld suggests, "IT managers should pick one area—even one specific project—to "cloud pilot" and assess their ability to manage and bring such a project to fruition. […] These are important efforts, and they should be supported—and reported within and outside the organization—so that others in IT and the wider community can learn of the successes and the downsides of operating in the clouds."

A good pilot is framed by the resolution of key issues or concerns for a Federal program. The issues may focus on the technical underpinnings of a new solution, answering particular questions, such as the performance, scalability, or security of a cloud service. Or the pilot might focus on cultural issues, such as new training requirements, or the day-to-day use of a remote Software As A Service (SAAS), instead of a locally installed desktop application. In either case, piloting benefits from the clear articulation of the program questions and risk areas to be explored. The National Archives offers, "Pilots reduce the risk of investment by identifying technical risk (e.g., compatibility problems with existing systems and infrastructure), areas for policy and procedure changes (workflow issues), and information for production planning (e.g., providing information for developing a realistic cost estimate and implementation and training schedule)." A good pilot also follows a structured process, with a plan defining pilot objectives, an engineering life-cycle, and a documented evaluation of the pilot results.

A key consideration for a Federal leadership team is the fidelity of the pilot to the final implementation architecture. For example, the results of the pilots will provide greater risk mitigation if they more closely mirror the architecture of the target enterprise solution. While some degree of compromise can be expected in the pilot for characteristics such as scale, or resources, care must be taken making the pilot as true to the target architecture as is reasonable. Cloud computing pilots can have a unique potential benefit in that the pilots can often run on the very same provisioned infrastructure as the final target architecture,

maintaining high architectural fidelity from the pilot through to the operational implementation.

We can expect to see several cloud computing-focused pilots in the next year. The Analyst Perspectives document with the White House's 2010 budget request describes several upcoming pilot efforts for cloud computing, including the use of Software As A Service on an enterprise scale, secure virtualized data centers, and improved end-user communication for a mobile workforce. 2010 may be the year of the Federal cloud pilot.

Please join us in April when we will focus on the essential capabilities to consider in constructing a private cloud.

Posted: April 1, 2010

Question for February 2010

What can Government do to facilitate the adoption of cloud computing to more effectively provide IT services? Please list specific actions that you'd recommend Government should take.

Ron Knode, Director,
GSS, LEF Research Associate
CSC

The question for February 2010 is clearly just a short step from January's question. So, let's deal with both of them:

- First (Jan 2010): "What's most significant cloud computing concern for federal orgs?"

The most authoritative (and accurate) answers would indeed come from "federal orgs" themselves. But, the three primary "lacks" in cloud computing that are encountered by "orgs" of all kinds, i.e., lack of standards, lack of portability, and (most importantly) the lack of transparency, are only intensified in government needs for cloud computing. (See www.csc.com/security/insights/32270- digital_trust_in_the_cloud for more discussion.) When we consider: (1) that security approval doctrine (certification and accreditation) is mandatory in the government (not an item to be traded off as part of a risk/reward equation); (2) that government data can be nationally classified, and therefore directly subject to laws and consequential impacts of non- compliance (not just a policy violation); and, (3) that the government uses IT as an element of national policy projection, including combat (and therefore must include stakeholder impacts far beyond those traditionally considered by commercial enterprises), then we can see how the impact of the three "lacks" becomes intensified.

- These circumstances lead naturally into a response for February's question: "What can Government do to facilitate the adoption of cloud computing to more effectively provide IT services?"

If the lack of standards, portability, and (especially) transparency are, indeed, the largest obstacles to the effective provision of cloud- based IT services for government use, then the government can certainly move powerfully to reduce the impact of those "lacks".

1. Publish the government's interpretation of certification and accreditation (C&A) in cloud computing. We know that NIST is working hard on a publication that delivers the U.S. government definition of a cloud (http://csrc.nist.gov/groups/SNS/cloud computing/index.html), and which is expected to provide recommendations on how cloud computing might be safely used by the government. By expanding this publication to include C&A doctrine and process for government cloud computing, much of the speculative ambiguity about what is and isn't acceptable could be eliminated.

2. Actively join in the standards bodies that are attempting to define protocols and techniques that can reclaim visibility/transparency into and through cloud processing. Such participation could also come via the issuance of government criteria, but interactive dialogue with industry around such efforts would be even better. For example, the A6 effort (www.rationalsurvivability.com/blog/? p=1276) and industry offerings like the CloudTrust Protocol(www.csc.com/security/insights/32270- digital_trust_in_the_cloud) offer ready-made places to start.

3. Identify a government cloud research, development and approval "center of excellence". While the pre- eminence of NIST as the cloud standards leader for the government is unquestioned, the initiation of a parallel development, test, and deployment lead would centralize and speed knowledge collection, including actual (trial) implementations and case studies. New applications for government use could emerge more quickly, even including a special emphasis on their "C&A ability". Such agencies as NASA or DHS could well organize and lead this effort on behalf of the entire government.

For the complete blog response, visit www.trustedcloudservices.com/3- things- government-can- do for- cloud- adoption.

For further information, please contact Ron Knode at rknode@csc.com.

Posted: February 5, 2010

Navin Sharma, Prashant Shenoy, David Irwin,
and Michael Zink
Laboratory for Advanced Software Systems
University of Massachusetts

Over the last five years the idea of cloud computing- using remote on demand computation and storage- has emerged as a dominant paradigm for the next- generation of Internet- enabled network services. Yet, despite their growth, the development of cloud infrastructures, especially commercial clouds, remain in a nascent state, providing an opportunity to significantly impact their evolution moving forward. As with the initial development of the Internet 40 years ago the Government has a role to play in ensuring that clouds advance key societal goals, in addition to commercial ones. To accomplish broader societal goals, we believe the Government should focus on at least three areas: encouraging cloud standardization and interoperability, incorporating networks into networked clouds, and more closely linking the clouds that process data to the sensors that produce it.

First, the Government must support the standardization of both internal and external cloud interfaces to enable consumers to more efficiently shift their computation and storage between independent clouds. The Internet never would have grown into what we know today if Internet Service Providers had built and commercialized closed networks at an early stage of its development. The Government run Internet forerunners, e.g. ARPANET and NSFnet, were crucial in developing the standards and protocols that made the commercialization of an open

Internet possible. Likewise, the Government will now play an important role in determining whether or not independent clouds, like the Internet, are interoperable and open.

Second, the Government should build on efforts to include network resources as part of cloud infrastructures, as embodied by NSF's GENI initiative (http://www.geni.net). The ability to reserve isolated network links, in addition to computing and storage, will promote the development of more secure distributed services and enhance researcher capabilities for wide-area Internet experimentation. Ironically, the "network" has been a key element missing in the growth of networked cloud computing. The Government should increase support to cloud-enable Government sponsored national networks, and allow researchers to study how to include the "network" in networked cloud computing.

Finally, the Government should encourage expanding clouds to incorporate the programmatic sensors and actuators that already serve key societal functions. The Government already operates an array of data- producing sensors, such as the NEXRAD radar network, that monitor the environment to serve near-term, e.g., predicting hurricanes and tornadoes, and long term goals, e.g., longitudinal studies of potential climate disruption. Meanwhile a key motivator in the development of cloud infrastructures has been the capability to quickly and efficient harness vast numbers of computers for tasks requiring massive, parallelized data processing. Closely linking these sensors, as well as the data they produce, to cloud infrastructures will enhance both investments by providing a scalable platform to drive future sensing based on processed data.

ViSE (http://geni.cs.umass.edu/vise) is an open testbed we are building as part of GENI to study these concepts by closely linking sensors, such as radars and cameras, to the GENI prototype, a national testbed based on open and interoperable edge cloud testbeds at multiple Universities linked by Government sponsored national networks including the National Lambda Rail and Internet2.

Posted: February 15, 2010

Nicklous Combs
Chief Technology Officer
EMC, Federal

Cloud Computing is the most overused term in IT today. The cost benefits of moving to a cloud type environment are just too beneficial to avoid. The important thing for federal organizations is to understand how they can get the characteristics of a cloud environment yet still meet the security requirements to protect the information. The private cloud is the only way federal organizations can address this issue today. Although security is at the top of the list, standards is something that has not yet been adopted for cloud computing. If you believe that virtualization is the foundation of a cloud like I do, then we need to adopt a cloud operating system that follows a standard that all vendor's can support. This will prevent vendor lock-in and provide a baseline for clouds to become federated enabling private clouds to match the public cloud cost models. As we move to this new environment we must move from perimeter security to an information centric approach to security.

Perimeters and bolt on security are still going to be important but will not solely address the needs of data protection in a federated cloud environment. When it comes to DoD, technology tends to be pretty reactive and behind the technology curve, this is due in large part because of the acquisition process. DoD 5000 was written for the development and acquisition of tangible products like trucks and planes not networks and technology. Programs today must follow a rigorous process that do not allow them to keep up with the changes in technology. We must take action to modify these acquisition rules.

Here is our web site: http://www.emc.com/ ?fromGlobalSiteSelect.

Posted: February 17, 2010

Barry X. Lynn
Chairman and CEO
3Tera

I am going to take a counter- position here? Why?

Well, yes. Government will have to make changes to adopt Cloud Computing, as will any large organizations and enterprises. But, I am certain that mitigation of the most important challenges facing government IT in general is inherent in Cloud Computing done right.

So, what are these challenges, and how are they mitigated by Cloud Computing done right.

Security and Privacy the common belief that Cloud Computing creates problems in this area is a myth. In fact, it improves security and privacy. When Cloud is done right, and applications can be abstracted from the physical resources they require, ergo can run anywhere, any time, they

can be set up as moving targets, rather than sitting ducks always running the same way in the same place like they do today.

Deployment Many government applications, especially those that support very tactical operations, require very fast deployment in multiple geographies. Cloud done right, where applications are fully encapsulated and abstracted from their data centers, enables instantaneous deployment in any geography.

Standards Cloud done right creates the ability to manage services and applications, regardless of what infrastructure they run on. When services can run independent of infrastructure, the need for standardized infrastructure for those services to be used anywhere, goes away. In fact, my advice to the government here, when adopting Cloud Computing is don't wasted tons of time standardizing infrastructure. If the Cloud is done right, you shouldn't have to.

And, all of the above aside, the most important point to make here is this.

There is an accepted belief that government IT is slow and not very innovative. But this is another myth. Government IT folks are thought leaders. Their technologists are ahead of the game, capable and nimble. The misconception arises because the approval and procurement processes are soooooooo slow.

So, sure, the most obvious recommended action item is to streamline these processes. Wishful thinking, right Well, where I come from, hope is NOT a strategy. It would be wise, instead of hoping that these processes become streamline some day, to mitigate the negative effect that the current processes have.

That's exactly what Cloud Computing can do.

You see, Cloud Computing done right enables the most granular scalability, and, more importantly, the ability to scale in an instant. So, when implementing things on a large scale, if they can be implemented in many small pieces without huge up front initial investments, approval of gigantic projects and the ensuing procurements becomes several procurements for several much smaller projects. And we all know that the overall hassle with regard to procurement has little to do with the number of procurements done, and everything to do with the size of the procurements.

So, the single most important action item I can recommend to the government is Move to Cloud Computing done right in a hurry. You ARE ready NOW!

Posted: February 22, 2010

Gregg (Skip) Bailey, Ph.D.
Director
Doloitte Consulting LLP
with Contribution from
Paul Krein
Deloitte Consulting LLP

Cloud Computing is touted as the holy grail of computing technology for the 21st century. It may prove so, but technology usually isn't so much a revolution as an evolution. The enterprise and the mission may take bigger leaps forward as the business side is what we expect to be disrupted, even reinvented. Cloud Computing, alongside operational efficiency mandates, may be just the catalyst we need for this change.

The good news for the CIO is that the technology change is truly evolutionary combining virtualization, better management tools, tremendous bandwidth and innovations around aggregating capacity. Pressures facing the next federal CIO include being bombarded with competing technologies and users who are more enabled and demand greater, faster, simpler access to their favorite technologies. The mission still needs agility and ever increasing quality, while the expectation is for a steady decrease in the cost of services each year, all coupled with a sea of changing demands. Cloud Computing brings promises of commodity pricing, high resiliency, and immediate sign- up for anyone willing to take the leap. However, Cloud Computing is more of a business opportunity than a technology change. The CIO now, more than ever, needs to have a clear understanding of where the organization is going from a business perspective, the challenges confronting the organization and the critical success factors of the business.

The CIO who embraces the opportunity and focuses on re orienting his organization with a clear vision and supported expectations will be out in front. But, to stay out in front, the CIOs role will quickly change from How good is your operation, to What have you done for me today In order to help the business deliver new value, CIOs will have to step up with a well defined plan and a personal extreme makeover role wise.

To sell the plan, CIOs will have to get out of the traditional IT box and create a vision and a roadmap to accommodate and leverage future choices. The future CIO must figure out a go forward strategy which embraces the rapid trajectory of the technologies, while enabling greater success for the organization. Their organizations core IT disciplines must be solid, but the nature of those disciplines will change. The opportunity is about preparing for the various toolsets of the future, and setting the right vision to intersect with the future demands of the business, even if the capabilities are not fully defined today.

The makeover is challenging; transitioning from the role of Efficient Operator to purveyor of capabilities and services is a major shift. At the same time the organization is demanding more innovation. In a nutshell either the CIO can act like the services broker to the business, and offer an optimized portfolio of services to the client, or the users will take on the broker role for themselves.

All things considered, it is no surprise that major enterprises are finding that deploying Cloud Computing models to be non- trivial and wrought with peril.

For further information, please contact Gregg (Skip) Bailey at gbailey@deloitte.com.

Posted: February 24, 2010

Gretchen E. Curtis
Director of Communictions
NASA Nebula Cloud Computing Platform

There are several key actions that the Government should take to accelerate the successful adoption of Cloud Computing. First, it should invest in Cloud Computing pilots to gain a better understanding of the technology and how the Cloud operating model impacts costs. Pilots hasten the adoption of technology standards and best practices and allow the Government to test, with a limited level of risk, the impact that the Cloud model has on budget and infrastructure procurement. The experience gained from Pilots will help the Government be a smarter, more informed buyer of Cloud technology.

Next, it should push the adoption of Cloud standards through open collaboration with the private sector. Public- Private collaboration maximizes the use of each sector's strengths, reduces risk, lowers capital investment, and improves efficiency. The Private sector sometimes has some advantages over Government, such as greater management efficiency, access to newer technologies, increased mobility as well as a broad perspective of the actions needed to meet public demands. Partnering with Industry allows Government Agencies to tap into this knowledge and leverage their expertise to better serve American citizens.

Finally, Agency CIOs should actively collaborate and participate in Federal Cloud Governance bodies, such as the Cloud Computing Advisory Committee and Cloud Computing Working Groups. Open communication and inter- Agency collaboration allows Government Agencies to share valuable experiences and insights and build upon a common body of knowledge, preventing a duplication of effort and leading to greater efficiency.

For further information, please contact Gretchen E. Curtis at gretchen.e.curtis@nasa.gov.

Posted: March 2, 2010

Larry Pizette
Principal Software Systems Engineer
Formerly MITRE

Thank you to the February submitters who provided insightful responses on the challenges facing government IT leaders with the adoption of cloud computing. The knowledge of all the submitters from academia, government, and industry – and their variety of perspectives sheds light on the steps that government leaders can take.

Consistent with past IT innovations, government leaders need to determine whether cloud computing concepts meet their IT needs and how they can best be leveraged to maximize the benefit and minimize risk. In government, the range of IT needs is broad. Needs range from highly secure systems that *always* need to be available for national security, to systems that contain information destined for public dissemination that do not always need to be available. These systems vary in their requirements based upon operational needs, statutory requirements and levels of security. As a result, government IT leaders' trade-offs for cost savings, scalability, location independence, security, application portability and tolerance for risk will vary.

Similar to the breadth in government requirements, cloud computing capabilities are also quite broad. For example, NIST defines three service models for cloud computing which are infrastructure-as a-service (IaaS), platform-as a service (PaaS), and software-as a-service (SaaS) and four deployment models which are private, community, public and hybrid. The service models have implications for Government IT leadership in many program focus areas such as development timelines and portability. The deployment models have different characteristics for cost reduction, type of costs (e.g., capital costs vs. operating expense), acquisitions, security, risks, and scalability.

The cloud computing choice is not binary there are many options. The challenge for Government IT leaders will be to match their requirements and system and data characteristics to the cloud computing capabilities that can best provide value to them within risk tolerances for the type of data, applications, and users they have. For this, we suggest a **structured decision process** that incorporates the following general steps:

- Determine which cloud services will provide benefit
- Establish a business case
- Define detailed requirements for a cloud solution
- Determine when to use internal private clouds or external public clouds
- Assess when to use cloud offerings provided by other Government entities

My colleague Geoff Raines and I will talk more about this decision process in an upcoming white paper that we are currently preparing.

In stepping through this decision process, Government IT leaders can consider their needs against the benefits and risks of different cloud options. They can also look for process

"accelerators." IT leaders can look for cloud offerings that are available via already negotiated buying schedules or that have already been certified and accredited. Similarly, they can look to place select capabilities and data that are intended for public consumption in cloud environments.

In order to mitigate risks, government IT leaders can employ pilots or move capabilities to cloud offerings incrementally to learn as they go. How to employ pilots is an open topic for discussion. In fact, this will be the topic of our March 2010 blog question. Please check back in March for in-depth thoughts from our submitters!

Posted: February 24, 2010

Question for January 2010

What do you perceive as the most significant concern for federal organizations who want to use cloud computing?

a) Acquisitions, b) Availability, c) Performance, d) Scalability, e) Security, f) Solution maturity (bugs/defects), g) Vendor lock- in, h) Other

Please address how commercial offerings are addressing this concern.

Steve Oberlin
Distinguished Engineer
CA

Certainly these are all valid concerns at one level or another, depending on the organization (their constituency, mission, IT requirements, etc.), the applications contemplated for cloud deployment, and the current state of IT operations and efficiency. While security is often raised as the most significant concern, cloud computing isn't just shorthand for outsourcing IT; internal or private clouds can provide the same or better efficiency and agility benefits as external clouds without incurring new exposure. Scalability, performance, and availability gains are actually all benefits one should anticipate reaping from cloud computing (internal or external). If these are concerns, you may be starting with the wrong applications, suppliers, or expectations. Similarly, vendor lock-in, solution maturity, and acquisitions are not problems unique to cloud computing and are amenable to the same sorts of mitigation strategies one uses to evaluate, procure, and deploy conventional IT technology and processes.

Instead, I think the biggest organizational cloud computing challenge is a larger one than those on this list, one that supersedes all of these and more. That issue is change readiness, the ability of the organization to culturally and operationally adapt to new paradigms for IT management, provisioning, accounting, and trust.

Cloud computing fundamentally is about enabling agility. Cost savings are nice, but of secondary value to an organization (the limit of savings is the IT budget, while the potential top-line gain from increased organizational agility is boundless). Agility results from trusting and enabling users IT's constituents to self-serve their own changing needs, to dynamically dial up and down and transmute the nature of their consumption, and to enjoy real time feedback on the cost and value of IT. Agility comes from disintermediation of traditional IT management and stripping stultifying process from between users and the IT services they consume (and increasingly want to directly control).

To deliver agility, IT management (the CIO and his organization) needs to embrace a different role. Disintermediation does not mean they can simply abdicate their ultimate responsibility for the safety and security, capacity and performance, and compliance and accountability of all

IT services required by the constituent organization. Instead, they need to provide an IT environment that enables users to manage their own services while ensuring all the above responsibilities are met behind the curtains. This means IT must employ invisible actors to enact policy constrained processes in real time, using technologies like dynamic optimizing automation, self-service portals, and catalogs of customizable template services.

IT's new role is a move to higher-level management. The compound cloud in our future will be managed by policy, not by rote process, and new skills and organization structures will be required before we're through. Change can be frightening, even when it is a promotion. Are you ready? Your users are.

Many commercial cloud technology offerings only deal with one narrow slice of the combinatory complexity of applications, platforms, and infrastructure. Though the cloud is still in its infancy, robust heterogeneous cloud management solutions are emerging. Check out www.ca.com/cloud for more.

For further information, please contact Steve Oberlin at steven.oberlin@ca.com.

Posted: January 12, 2010

Bruce W. Hart
Chief Operating Officer
Terremark Federal

If we focus our discussion around Infrastructure as a Service (IAAS, the foundational level of the Cloud), I believe that three concerns are of equal urgency Availability, Security, and Performance.

Cloud-based IAAS comes in various flavors. The most popular conception the shared, commodity based Cloud, where you simply order up some compute capacity, swipe your credit card, build your virtual machine and you're off and running has a certain appeal, but it does not scale to the enterprise-level, mission critical demands of Federal agencies. Indeed, we have seen instance after instance recently where commodity based Cloud services simply "crash." Federal IT leaders cannot afford to ignore the underlying physical architecture from which Cloud offerings are launched and just hope for the best. They must assure at least the level of availability, security and performance that they realize from traditional hardware-based IT architectures ideally, they should be able to interconnect those traditional systems to the new Cloud services that they acquire. This creates leverage from all of the benefits of Cloud infrastructure on demand capacity and massively scalable elastic architecture, which can bring a new level of flexibility and agility to IT leaders, and with it a compelling economic model that eliminates lumpy capital expenditure and precisely aligns IT infrastructure spend and capacity with the real time needs of the organization but it does not sacrifice the power and reliability of controlled, standards-based systems.

Further, the IAAS services they acquire should be integrated with a suite of security features that preserves the integrity of the Government's data. This must go beyond simple firewalls and intrusion detection/prevention utilities to more comprehensive capabilities, beginning with the

physical security of the site where the Cloud services originate and moving through multi- factor authentication to sophisticated forensics capabilities, including memory forensics, network analysis, end user analytics, and Certification and Accreditation support.

Finally, the performance of the virtual systems created on Cloud resources for Federal missions must be excellent Federal leaders should ensure that they have access to dedicated resources, rather than sharing resources in an over- subscription service model, while also retaining the ability to surge as needed on a pay- as you- go basis. Enterprise class, federal dedicated clouds hold the promise of economy, agility, and most important of all, elasticity, at not just the system level, but right down at the individual server's ability to expand and contract according to real time need. Commercial service providers should be held accountable by informed, inquiring Federal leaders for delivering on that promise.

For further information, please contact Bruce W. Hart at bhart@terremark.com.

Posted: January 15, 2010

Peter Coffee
Director of Platform Research
salesforce.com inc.

As federal organizations accelerate their adoption of cloud computing, there's nearly universal consensus on the cloud's compelling advantages:

- lower capital requirements
- rapid, scalable deployment of high- function solutions
- radical reduction of cost, schedule, and technical risk

What remain are two sets of concerns:

- issues of perception that must be addressed to satisfy stakeholders
- issues of technology and practice that must be addressed to maximize value

During acquisition phase, organizations should think of the cloud as an extension, not a replacement, of current IT assets. It's a common misperception that the cloud must be adopted in whole, or not at all; it's a vital component of cloud success to recognize opportunities for integration among services of multiple providers, and between cloud and on premise resources.

Many opportunities in the cloud come from liberating latent value of current IT systems to deliver accurate, actionable information in a secure and reliable manner to points where that information can best be used.

All responsible parties demand assurance of the availability of cloud-based systems. Cloud availability is often superior to that of on premise systems: the scheduled maintenance alone of many on premise systems exceeds the total non-availability, from *all* causes, of an enterprise- grade cloud service. Providers such as salesforce.com and Amazon Web Services operate public Web dashboards reporting all departures from normal operation, however slight, with performance monitoring beyond what's available to most in house operations.

Cloud security and governability are routinely assumed to be less stringent than that of local systems, but this can not be generalized. There are consumer cloud services designed for easy sharing, and there are enterprise cloud services designed for precise and granular privilege assignment with robust and auditable management.

Security is not a technology, but a combination of culture and process. Actual data loss or security breach in federal systems, as reported each year by the GAO (ref. 2009 report at gao.gov/new.items/d09546.pdf), is most often the result of accidental or deliberate misuse of privileges intentionally assigned to systems' users. Enterprise grade cloud services offer rigorous separation of duties; world-class security teams, tools and practices; and superior ability to monitor and report the actual time and manner of users' and administrators' actions.

With these issues candidly addressed, IT leaders in federal agencies should proceed with cloud adoption bearing three strong guidelines in mind:

- What works well now should be measured against cloud alternatives, and should be complemented rather than replaced unless the cloud is measurably better.
- What doesn't work well now should never be merely relocated to the cloud, but should rather be reenvisioned in a way that takes maximum advantage of cloud connection capabilities and proven cloud services.
- Detailed analysis, not facile generalization, should be applied to all questions of security, availability and capability. In many cases, cloud offerings are already superior in these respects, and rapidly improving as well but in all cases, the solution should be chosen based on the specific need.

For further information, please contact Peter Coffee at pcoffee@salesforce.com or see his blog at http://cloudblog.salesforce.com/.

Posted: January 15, 2010

Steven Lebowitz
IT Architect, Federal Cloud Team
IBM

While many of the items listed are of concern to Federal agencies to a greater or lesser extent, there are perhaps, larger issues which need to be addressed. Certainly, there is a great deal of interest in cloud. Many organizations have yet to be convinced that they will receive improved service and reduced costs by moving from dedicated infrastructure into a shared cloud. There are also issues regarding who can participate in a cloud due to a number of security and privacy concerns. How will the Government address these security and accreditation policies and practices in order to adopt a highly virtualized, automated, and shared infrastructure? Are the potential cost savings (and other benefits) significant enough to overcome the organizational "stove piping" and fear of losing control of both infrastructure and data? Finally, organizations need to determine what applications are appropriate for being deployed in a cloud. There is a process of application portfolio rationalization, and an analysis of "cloud readiness" that should be done in advance of making technology choices and deployments.

We at IBM have had significant interest from our clients in deploying fully integrated, easy to deploy, self- service, test and development environments into their organizations. Software test and development is an application area which is typically decentralized. With this, they are attempting to gain first- hand insight into the benefits of cloud computing, and a better understanding of its impact to their organization without making large monetary investment, or a giant leap of faith with a public cloud provider.

For further information, please contact Steven Lebowitz at lebowitz@us.ibm.com.

Posted: January 19, 2010

Teresa Carlson
Vice President
Microsoft Federal

Federal agencies have to consider many factors when it comes to if, when, and how to move to the cloud. Most of the agencies we have been meeting with are smartly planning to walk before they run, and they know there are serious concerns around availability, performance, privacy and security. From our perspective, these are all important, but security is probably the most significant cloud concern for federal leaders.

The move to cloud is a huge cultural shift you're allowing someone else to host your data and trusting that they'll protect it. That's a big deal for complex government organizations that work with highly sensitive information, often with national security implications. Cloud providers must move forward with solutions that meet the best industry security standards that exist today,

and earn the trust of the organizations we serve. Earning that trust starts with transparency, and government organizations are rightly demanding full view into the processes we're implementing to protect their data. Over time, the best solutions and processes will inform quality regulation, so that governments can confidently take advantage of the tremendous cost and efficiency benefits that cloud computing offers.

Fortunately, moving to the cloud isn't really starting from scratch. Hosted services, Service Oriented Architectures (SOA) and Web applications have been around for a while, offering us a good foundation for best practices that we can carry into the discussion on cloud security. The same holds true from a standards perspective. We believe that over time standards will emerge based on industry and customer demand which has often been the case throughout the history of IT. Establishing standards will be a joint- effort between industry and government. We live in a mixed/hybrid IT world and government customers need the freedom to choose the best solutions and locations for their data. Interoperability avoids vendor lock- in and ensures choice and competition.

We should note, too, that datacenters are a key foundation of any organization's approach to cloud computing, and should be built in compliance with the best security and privacy standards that exist today. These standards include International Organization for Standardization (ISO) 27001, FISMA, ITAR, Health Insurance Portability and Accountability Act (HIPAA), Sarbanes Oxley Act of 2002 and SAS 70 Type 1 and Type II. All are examples of widely accepted US and International standards today.

The cloud will evolve and mature over time, but we can't stop the innovation while we wait. A good analogy is one from the car industry. There are a lot of important federal safety compliance standards for manufacturers, but there are a lot of "extra" safety measures (ie. additional airbags) that appeal to some consumers at a price. Cloud options will be similar. There will be baselines and there will be extra assurances. Forcing everyone to buy the most feature rich IT system would be cost- restrictive. We need mandatory standards and self- regulation. Brad Smith, Microsoft's General Counsel recently spoke at the Brookings Institution on this very topic, and urged congress to move forward with a "Cloud Computing Advancement Act" to foster innovation, protect consumers and establish rules and regulations around data privacy and security.

For further information, please contact Teresa Carlson via FutureFed.

Posted: January 28, 2010

Bert Armijo
VP Marketing and Product Management
3Tera, Inc.

All of these topics come up in any conversation about cloud computing, but they're really all part of one larger concern control.

IT is tasked with delivering reliable and secure services to users, and safeguarding their organization's data assets. Therefore, it's understandable IT managers are concerned whether

cloud computing gives them sufficient control to ensure the viability of those services. When performance lags, will they be able to respond quickly? Is data secure and under their control If the provider can't meet SLAs, can they move applications and data quickly to another If security requires it, can they bring applications back in house? Will their processes and monitoring work with cloud deployments? In short, IT managers want to know they have the control to perform their jobs.

Furthermore, there's no need for compromising control. Advances in cloud computing platform technology afford IT managers similar control in the cloud to what they have in their own data centers. As IT managers evaluate potential providers and projects they should look specifically provisions for their control not just of provisioning a virtual machine, but of storage, networking, security and critical infrastructure. Cloud computing needs to be more than just a way to save money, it needs to be a better way to run your service.

Posted: February 4, 2010

Geoff Raines
Principal Software Systems Engineer
MITRE

First, I want to thank the group of respondents above for participating in our initial cloud computing question and providing in depth and thoughtful answers. One of the primary goals of this forum is to bring together thought leaders in the cloud computing marketplace to solve the IT Government's challenges, and it is clear this month's contributors gave this topic a lot of consideration.

This month's cloud computing question focused on the perception of common risk areas for Federal cloud computing efforts. Understandably, one can see from the responses above that few people identify one single cloud computing risk area as their sole concern. Similarly, a survey on cloud computing transition by Kelton Research for Avanade, between December 2008 and January 2009, focused on C- level executives (e.g., CEO, CFO, CIO, CTO) and suggested that even though "nearly two in three IT execs worldwide and four of five in the United States believe cloud computing reduces up- front costs," there is still "strong reluctance to change driven by fears of security threats and loss of control. The survey goes on to state that, "In this economic environment, costs are not a top barrier to change." A 2008 survey by CIO Research suggested that while "58 percent say cloud computing will cause a radical shift in IT," 45% of respondents cited security concerns as their greatest concern for cloud adoption, followed by concern over integration with existing systems, loss of control over data, and availability. (Please see the inset figure for the broader list of their concerns.) Further, surveys by F5 Networks and Unisys in 2009 each suggested that not only security but data privacy are key concerns regarding cloud computing. All these surveys suggest that, like any traditional large- scale infrastructure effort, cloud computing efforts bring with them a series of program risks to be addressed, whether as a service provider, or a service consumer.

Federal leadership teams are familiar with actively managing risks to ensure program success. Risk management techniques suggest that for each risk, mitigations and courses of action can be developed and put in place to improve the outcomes for a program. As noted in the contributors' examples above, the marketplace of commercial and Government cloud service offerings is evolving to address commonly perceived risks to make cloud computing capabilities a viable alternative for many Government IT needs. Consequently, the Government decision maker has a group of options that range from wholly commercial cloud services, and Government run community clouds, to the creation of internal private clouds. Characteristics such as the expected costs, agility, scalability, and an organization's data and system requirements will suggest which path is most appropriate for individual Government programs. MITRE is currently developing a whitepaper to address this decision process in greater depth.

The MITRE cloud blog is a new forum, started in January 2010, to provide a mechanism to effectively connect industry and Government. Each month we ask thought- leaders and market leaders to offer their ideas on cloud topics important to Government decision makers. As a reader of this blog, if you have comments on future questions, or on the answers above, please feel free to email us through the link provided below.

Next month we focus on what Government can do to facilitate the adoption of cloud computing to more effectively provide IT services. Please bookmark us and check back with this site, or subscribe to our MITRE RSS feed above to stay informed on new cloud computing postings.

Posted: January 29, 2010