CHAPTER 4

Data Security and Service Reliability

Consider this quote by a cloud industry executive: ''Any business leader worried about the security and reliability of their data in the cloud should remember that they've been trusting, saving, and storing their personal financial assets in an external, virtual banking cloud for years.''

Will Your Cloud Service Provider Be Here Next Year?

This is a good question, and one that every executive should be asking. On one hand, there are highly reputable cloud service providers who are so well established—either from their previous lives, or because they are new but have grown to widespread use—that they are quite likely to be in this business for the foreseeable future. Amazon Web Services, IBM, HP, Salesforce.com, among many others, are strong players with solid futures.

On the other hand, for every one of the large players in cloud computing, there are hundreds of smaller players offering some type of cloud solution. Whether the provider you consider is large or small, every potential customer should exercise proper due diligence in the selection process that would include analysis of the company's financial stability, future prospects, and viability as a long-term and going concern.

Jeff Kaplan, Managing Director of THINKstrategies, a SaaS and cloud computing advisor, feels that key cloud providers are taking security and reliability quite seriously. In his white paper ''The CIO's

Guide to Software as a Service,''[1] he says that unlike the traditional, on-premise software model that puts the burden of success on the customer, the SaaS subscription model places the burden on SaaS vendors to deliver reliable and secure services that meet the needs of their customers. According to Kaplan, the vendors' business depends on delivering quality services and safeguarding their customers' valuable data, so the leading SaaS vendors invest in state-of-the-industry service delivery and security technologies and certifications programs that include SAS 70 (the Statement on Auditing Standards 70), ISO standards, and Payment Card Identification (PCI). This requires the SaaS providers to implement extensive and well-documented security practices that govern their data center operations and personnel—including processes that regularly test facilities and staff.

All of this said, the most critical step in investigating the viability, reliability, and security of a cloud solution provider for the long term is the creation of a well-prepared request for proposal (RFP), which has several objectives. At the highest level, the RFP allows your organization to concisely compile all of the requirements of a particular technology initiative. Once those requirements are all agreed upon and determined, it then gives your organization an opportunity to formally request proposals from various bidders, and for the bidders to submit their proposals all using the same requirements.

A standard portion of any RFP is the request for information about the financial standing and viability of the vendor company. Gathering this information from bidding vendors is important with any RFP, but is particularly important with newer technologies where the players being considered are a mix of traditional vendors with long-standing brand recognition and reputation, and newer start-ups with smaller customer bases, less financial resources, credibility, and reputation.

Relevant information requested in the RFP can and should range from how long the company has been in business, the growth of the customer base and revenue, the company's available credit, and references from other customers. Ultimately, this information helps the potential customer assess and balance the value of the vendor's potential solution—which could appear to be significant—against the backdrop of the company's financial strength and future prospects, which could be shaky.

All of this might help a company make a decision like the following: Two vendors among a field of six are the strongest contenders based on the strength of their solution. One vendor is more diversified and has a longer-standing presence and track record of success. Its future viability, as demonstrated by information reported in the RFP as well as customer references, is relatively strong. The other vendor, with a more focused portfolio, less customers, and less experience in the field, has an even stronger potential solution, however, its prospects of weathering the ups and downs of challenges in the IT market, aren't nearly as great as the other vendor.

All of this makes selecting a vendor difficult. With new cloud computing vendors and solutions emerging (either from within larger, more established vendors, or as small, venture-funded startups), this makes the financial viability of all considered vendors particularly important.

What to Look for in a Good Service Provider

Once a vendor is selected to support a cloud computing initiative, a smart way for the organization to proceed is with a pilot project—a limited initiative that helps the organization move some computing to the cloud and minimizes risk because the project is smaller and manageable as a first-time effort. A pilot project also allows the organization to learn in the process. Simply put, if and when issues arise during a pilot, they will be much easier to resolve, and far less risky, if the footprint of the project is relatively small.

Preparing for a pilot project requires research and planning—and ultimately helps the organization understand the specifics to identify in a suitable provider. Naturally, this research will form the basis for preparing the RFP described earlier. As important as selecting a vendor, the analysis will force the organization to assess the ultimate business value of moving to the cloud in the first place.

Any move to the cloud should be based on a thorough situation/business analysis. The standard questions should be asked. Can a return be calculated on the investment? Can a longer-term return be anticipated? Does the technology perform as well as, if not better than, the existing internal platform? Is the cloud solution efficient, if not more efficient, than the existing environment?

In order to answer these questions, draw a circle around the users, applications, and business processes that will potentially use the cloud solution. From there, it shouldn't be difficult to analyze the costs of the existing solution compared to the cloud solution. Naturally, one should be careful to consider all costs and savings related to hardware, software, personnel, and any ancillary and ongoing expenditures.

Then comes the harder part—assessing performance. While it's relatively easy to understand performance of the existing systems—their availability, redundancy, backup, recovery, latency, and the like—identifying the same for a cloud solution is a bit more difficult. The cloud services providers being considered need to allow a thorough audit of their application and systems performance. In a cloud solution, fundamentals like hardware and software environments, ISPs, and server locations can change over time, and so will the corresponding system response time and throughput. In order to determine if the cloud solution will perform at least as well as the existing solution, the cloud provider will need to enable a cooperative performance audit. Ultimately, the audit will prove to be pivotal in determining if a particular vendor's solution is at least as efficient as the existing environment.

Elements of Good Data Security Policy

Security is arguably the top concern of companies when considering a move of data or computing resources to the cloud. Companies have grown accustomed to safeguarding data sitting in their own data centers, so getting used to the concept of proprietary data and applications sitting outside of traditional company jurisdictions presents worries, concerns, and challenges to not only data policies, but a company's well-entrenched culture and values.

With so many companies moving data and applications to the cloud, however, much thought is naturally being invested in addressing these concerns. One organization in particular has addressed the matter in order to ensure that data in the cloud is widely safeguarded. The Cloud Security Alliance was formed in 2008 by a group of industry leaders in order to promote best practices in assuring security with cloud computing initiatives. The alliance intends to:

- Promote a common level of understanding between the consumers and providers of cloud computing regarding the necessary security requirements and attestation of assurance.
- Promote independent research into best practices for cloud computing security.
- Launch awareness campaigns and educational programs on the appropriate uses of cloud computing and cloud security solutions.
- Create consensus lists of issues and guidance for cloud security assurance.[2]

A major product of the Cloud Security Alliance's efforts is their publication of and updates to a document available on their web site entitled ''Security Guidance for Critical Areas of Focus in Cloud.''[3] The document provides users and vendors alike with a key source of specific recommendations in managing security policies and is a must read for technical teams investigating and implementing cloud solutions. The report's detailed guidance suggests the following:

- Determine exactly what data or function is being considered for the cloud.
- Assess how important the data or function is to the organization.
- Determine which of the following cloud options are acceptable: public; private (internal); private (external); community; hybrid.
- Evaluate the degree of control available to implement risk mitigations.
- Map out the flow of data in and out of the cloud to identify points of exposure to risk.[4]

In June 2008, analyst firm Gartner published a report called ''Assessing the Security Risks of Cloud Computing'' which identified seven security issues prospective buyers of cloud services should raise with potential vendors. An article in *Network World* entitled ''Gartner: Seven Cloud-Computing Security Rrisks'' recaps the advice:

1. *Privileged user access.* Sensitive data processed outside the enterprise brings with it an inherent level of risk, because

outsourced services bypass the ''physical, logical, and personnel controls'' IT shops exert over in-house programs. Get as much information as you can about the people who manage your data. ''Ask providers to supply specific information on the hiring and oversight of privileged administrators, and the controls over their access,'' Gartner says.

2. *Regulatory compliance.* Customers are ultimately responsible for the security and integrity of their own data, even when it is held by a service provider. Traditional service providers are subjected to external audits and security certifications. Cloud computing providers who refuse to undergo this scrutiny are ''signaling that customers can only use them for the most trivial functions,'' according to Gartner.
3. *Data location.* ''When you use the cloud, you probably won't know exactly where your data is hosted. In fact, you might not even know what country it will be stored in. Ask providers if they will commit to storing and processing data in specific jurisdictions, and whether they will make a contractual commitment to obey local privacy requirements on behalf of their customers,'' Gartner advises.
4. *Data segregation.* Data in the cloud is typically in a shared environment alongside data from other customers. Encryption is effective but isn't a cure-all. ''Find out what is done to segregate data at rest,'' Gartner advises. The cloud provider should provide evidence that encryption schemes were designed and tested by experienced specialists. ''Encryption accidents can make data totally unusable, and even normal encryption can complicate availability,'' Gartner says.
5. *Recovery.* Even if you don't know where your data is, a cloud provider should tell you what will happen to your data and service in case of a disaster. ''Any offering that does not replicate the data and application infrastructure across multiple sites is vulnerable to a total failure,'' Gartner says. Ask your provider if it has ''the ability to do a complete restoration, and how long it will take.''
6. *Investigative support.* Investigating inappropriate or illegal activity may be impossible in cloud computing, Gartner warns. ''Cloud services are especially difficult to investigate, because logging and data for multiple customers may be co-located and may also be spread across an ever-changing set of hosts

and data centers. If you cannot get a contractual commitment to support specific forms of investigation, along with evidence that the vendor has already successfully supported such activities, then your only safe assumption is that investigation and discovery requests will be impossible.''

7. *Long-term viability.* Ideally, your cloud computing provider will never go broke or get acquired and swallowed up by a larger company. But you must be sure your data will remain available even after such an event. ''Ask potential providers how you would get your data back and if it would be in a format that you could import into a replacement application,'' Gartner says.[5]

Cyber Threats and Perimeter Security in Cloud Computing

> ''. . . it is critical that cloud customers select the right cloud formations for their needs, to ensure they remain secure, able to collaborate safely with their selected parties as their evolving business needs require, and compliant to applicable regulatory requirements–including on the use and location of their data. The joy of the cloud model is that it can deliver great advantages, but only if you know where in the different formations of cloud you need to be in order to achieve the right flexibility for your business needs . . .''[6]
>
> —*The Jericho Forum's Cloud Cube Model*

Prior to cloud computing, companies created perimeter security by installing hardened firewalls to block unwanted traffic trying to access the corporate network. They also established restricted access through passwords and education to block malicious access to their traditional data centers. With far fewer mobile users and with virtually all data resident in-house, this strategy made sense. But times have changed. Since today's cloud computing model moves company-owned data outside the traditional corporate security boundaries, and since professional hackers have proven that they will continue to explore and exploit weaknesses, companies need to take a fresh look at their security strategy, objectives, and defenses.

While perimeter security has been a foundation of corporate information security planning and implementation, some corporate security officers have advocated the need for a more contemporary philosophy that embraces the changes resulting from widespread Internet access and the resulting corporate security vulnerabilities. As early as 2001, some security experts began discussing the need for what they call ''deperimeterization.'' Today, a group called the Jericho Forum has advanced that philosophy for cloud computing models. They have put forth a deperimeterized approach to security. As shown in Figure 4.1, this approach, called the Cloud Cube Model, identifies and defines four criteria to differentiate cloud formations from one another.

The four criteria include:

1. *Internal versus external.* If it is within your own physical boundary, then it is internal; if it is not within your own physical boundary, then it is external.
2. *Open versus proprietary.* Proprietary means that the organization providing the service is keeping the means of provision

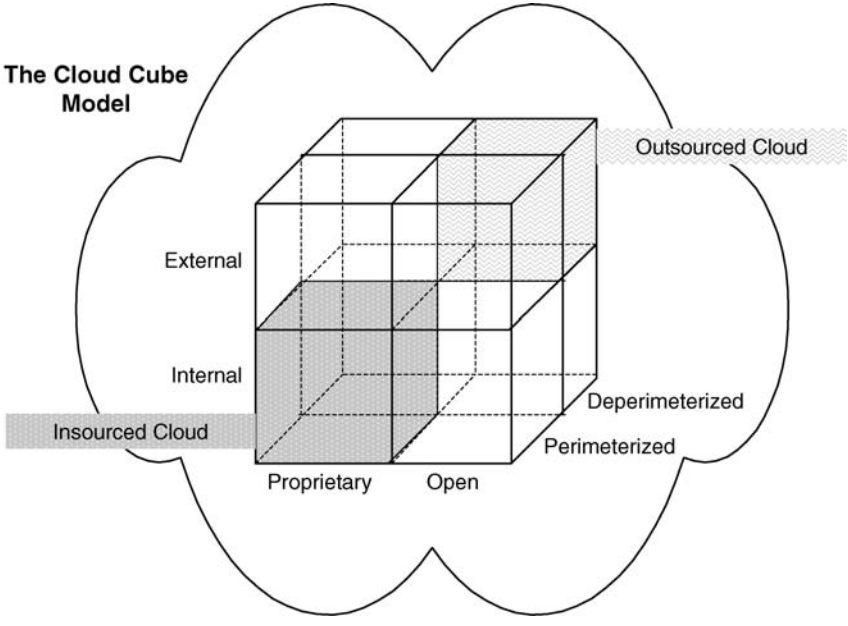

Figure 4.1 The Cloud Cube Model

under their ownership. As a result, when operating in clouds that are proprietary, the Jericho Forum suggests you may not be able to move to another cloud supplier without significant effort or investment. Often the more innovative technology advances occur in the proprietary domain. As such the proprietor may choose to enforce restrictions through patents and by keeping the technology involved a trade secret. Clouds that are open are using technology that is not proprietary, meaning that there are likely to be more suppliers, and the Jericho Forum suggests you are not as constrained in being able to share your data and collaborate with selected parties using the same open technology. Open services tend to be those that are widespread and consumerized, and most likely a published open standard, for example, email (SMTP).

3. *Outsourced versus insourced.* Outsourced means the service is provided by a third party; insourced means the service is provided by your own staff under your control.
4. *Perimeterized versus deperimeterized.* Perimeterized implies continuing to operate within the traditional IT perimeter, often signaled by ''network firewalls.'' As has been discussed in previous published Jericho Forum papers, this approach inhibits collaboration. In effect, when operating in the perimeterized areas, the Jericho Forum suggests you may simply extend your own organization's perimeter into the external cloud computing domain using a VPN and operating the virtual server in your own IP domain, making use of your own directory services to control access. Then, when the computing task is completed you can withdraw your perimeter back to its original traditional position. The Jericho Forum considers this type of system perimeter to be a traditional, though virtual, perimeter. Deperimeterized, assumes that the system perimeter is architected following the principles outlined in the Jericho Forum's Commandments and Collaboration Oriented Architectures (COA) Framework. The terms *micro-perimeterization* and *macro-perimeterization* will likely be in active use here—for example, in a deperimeterized frame the data would be encapsulated with metadata and mechanisms that would protect the data from inappropriate usage. COA-enabled systems allow secure collaboration. In a deperimeterized environment an organization can collaborate securely

with selected parties (business partner, customer, supplier, and outworker) globally over any COA capable network.[7]

Encryption: The Next Frontier of Data Security

Many organizations have been reluctant to move to cloud computing due to regulatory restrictions that prohibit them from using the cloud for storing sensitive data, or due to concerns about the privacy and security of data in the cloud. A recent article in *CIO* magazine put it like this, ''There's no doubt that cloud computing is dominating today's IT conversation among C-level security executives. Whether they're lured by its compelling cost savings or its perceived advantages, security leaders are probing the capabilities and restrictions of the cloud. At the same time, security and compliance concerns remain issues holding large enterprises back from capitalizing on the cloud's benefits.''[8]

To address these concerns a new class of data security products is appearing. These new data security products employ data encryption to keep the data secure. Even if unauthorized parties get through the other security measures and get access to the data, it will be to no avail if they can't read that data. Even if data is copied or stolen, if it is well encrypted, it will be useless unless those stealing the data can also secure a copy of the encryption key that will unlock the data and enable them to read it. There is a new class of encryption-based data security products that puts the enterprise administrator in control of powerful measures to protect an organization's data.

These products specifically address the limitations and concerns of using SaaS and other cloud applications to handle sensitive corporate data. Implemented as an appliance on the enterprise LAN/WAN or as a cloud service, these encryption products encrypt data used by SaaS application before it is transmitted from the enterprise to the SaaS provider. Authorized people using these SaaS applications are not affected by the encryption and remain largely unaware of this process. Their data is encrypted as it leaves their internal company systems and, when stored on the SaaS provider's servers, it is unreadable by anyone without the encryption key. Database theft and regulatory compliance issues are then addressed as all sensitive data remains undecipherable when in transmission and at rest outside the enterprise firewall.

Developers of cloud applications don't need to implement any special code in their applications to provide them with this level of security and regulatory compliance. This encryption software can be added to an application as a service (SOA or SaaS) and it will encrypt and secure data as it flows through the application. Some people have coined the acronym VPS or ''virtual private SaaS'' to describe this new category of encryption-based data security software.[9]

Contracts, Service-Level Agreements, and Guarantees

As with any newly contracted service—whether it's for information technology or not—organizations committed to moving to a cloud-based solution should have a knowledgeable lawyer review their cloud service provider contract both during the negotiations and before final signature. There are many law firms that specialize in legal considerations related to information technology implementations, and some are familiar with the nuances and new ground encountered by deployments in the cloud.

Service-level agreements (SLAs) have become more widely used in many industries over the last few decades and, of course, they can and should be used well in managing service-level expectations and requirements from cloud service providers. Simply put, SLAs are comprised of the language—in the context of an overall master services agreement—that clearly specifies for the customer and the service provider what's expected of the provider and customer. These details are valuable to both parties because they provide a legally binding reference document to help manage the ongoing service level, including specific metrics and measures of performance along with pricing tables. Like any legally binding agreement, the objective is to protect both parties and to prevent challenges and disputes in managing service levels. A well-written SLA will actually prevent problems before they can significantly impact ongoing business operations.

Customers should be aware that providers will have engaged legal advice in constructing their proposed master agreement and SLA. With their own best interests in mind, customers should also involve a corporate legal representative to review the document and provide advice prior to signature. In this process, it is important for

the agreement to be written plainly and clearly, and for the customer parties who will be managing the business enabled by the services to have provided input, especially in the form of asking ''what if'' questions that may impact business performance.

SLAs can be constructed in several ways, and there is much guidance available on the best way to approach them, but generally, there is consensus that an SLA document should cover at least the following basic sections:

- *Overview.* This section should briefly identify the parties entering the agreement and concisely describe the general nature of the agreed-upon services purchased.
- *Scope of work.* A more elaborate overview, this section is also commonly called the SOW and forms an important part of the agreement that clearly details the services provided to the customer.
- *Performance measures.* Good preparation and input from both parties is required to create this critical portion of the agreement. Measures appearing in this section should fairly and reasonably identify metrics that will be continuously monitored throughout the term of the agreement including items as varied as uptime, throughput, and the number of end-user customers that can be served simultaneously. While the provider will have ideas on what's best to measure, the business buying the service should incorporate and negotiate to include critical measurements that help it meet its business objectives—and that serve customers reliably and securely. These items can and should be very specific to the business enabled by the cloud provider. In cases where the service is extremely critical to second-to-second business success, it may be worth investing in a third-party service that can help monitor and report on service levels.
- *Managing problem resolution.* This section should detail the agreed-upon process whereby the customer can alert the provider to problems along with the timeliness of response and the procedures for how the problem will be resolved. By putting this process in writing, both parties are then clear on what will take place when a problem arises. The provider will likely want protections here as well that ensure they can address any problems that might be inadvertently caused by the customer.

- *Fee structure.* This section should very clearly and simply state the fees being charged by the provider to the customer along with payment terms.
- *Customer obligations.* So that the provider is given the best opportunity to meet all of their obligations of the agreement, the provider will require that the customer remain obligated to providing needed information on a timely basis. This area of the agreement will be highly specialized and should detail the specific information and related processes to exchange information between the customer and the provider.
- *Warranties.* This is the area of the agreement where the customer wants the provider to make guarantees and to specify how they will make good on those guarantees if, for any reason, the provider can't meet the obligations of the service they guarantee. Essentially, it's where the provider is held accountable for nonperformance during the agreement and where the customer can seek relief for that non-performance. Naturally, this can be a highly negotiated point before the agreement is finalized where the provider wants to minimize the guarantees and, conversely, the customer wants some ironclad remedies where they may have specific concerns if or when the provider doesn't fulfill service obligations. This is also an area where the customer can ask the provider to warrant simple but important facts. An example would be warranting that the provider has the legal ability to provide the needed business services in all of the applicable geographic locations.
- *Security.* Since security is of major concern to customers when implementing cloud solutions, it's a good idea for the SLA to clearly describe the security capabilities of the solution along with the procedures that will take place in the event of a security breach.
- *Compliance.* Some customers are governed by industry-specific regulatory requirements restricting how information is shared and any processes to deal with these issues, and they should be detailed in this section to ensure that the provider will be conforming–and how it will be dealt with if a problem arises.
- *Confidential information and intellectual property.* This section provides the opportunity for both parties to clearly define the respective intellectual property they own that is not the right

of the other party. Importantly, it is also where they require the other party to treat specified information that may be exchanged as confidential. If information is to enter the hands of third parties on behalf of the provider, the customer will want the provider to secure confidentiality guarantees from the third party.

- *Liability protection.* Since the cloud involves customer data that can sit in multiple locations and be processed by multiple companies, the technology industry finds itself on new legal ground with liability protections. Of significant concern to the customer are the implications of a breach to a cloud provider's security, and what court-ordered restrictions might potentially be placed on processing that data in the event of an unforeseen legal challenge. With that in mind, the agreement should specify all of the locations where customer data might potentially reside—and provide guarantees that all of those handlers of the data (which may include other providers) are in compliance with regulatory requirements. Customers should require language that helps avoid any legal interruption to ongoing business operations or access to their own data should a breach or legal challenge materialize. Given the changing cloud landscape and emerging precedents, it's critically important for customers to engage relevant legal advice at the earliest stages of the RFP process to ensure the most contemporary protections are applied to the agreement—and that providers are aware of the customer's legal requirements.
- *Regular review.* To make sure that unforeseen issues have an opportunity to be aired throughout the contract period, a scheduled set of review meetings provides an opportunity to not only maintain a good working relationship, but to adjust the agreement as necessary should changes in business circumstances warrant.
- *Termination.* As is standard in legally binding agreements, the SLA should provide language describing how the agreement can be terminated by either party, and the procedures to follow, including how data is transitioned to a new environment along with the related schedule.
- *Implementation.* This portion of the agreement describes the schedule the parties have agreed upon to start the transition

to the new service, the date intended to launch the service to end-user customers, and key milestones and deliverables required in either direction along the way.

Negotiating Service and Pricing

Many cloud computing observers argue that cost is often the primary driver for considering cloud initiatives. They are correct that the cloud can offer some cost-effective alternatives to, say, owning and operating 100 servers to achieve a particular business objective. They also point out that the amount of time and energy consumed by worrying about hardware, maintenance, uptime, and reliance on internal data center resources can represent a sizable opportunity cost for the business. That said, there is every reason to take a very close look at cloud computing's delivery and pricing models despite what may initially appear to be a good value. All of this starts with a careful negotiation that can prove to be a strategic factor in not only developing a good working relationship with a cloud provider, but in establishing fair pricing models for both parties.

While contracts can be laden with seemingly clinical language and endless legal terms, written agreements should be based on the outcome of very practical business discussions between the two parties. The best place to start is making sure that the potential provider understands that the customer is looking for a partnership, rather than what can often become a standoffish procurement/supplier relationship. This approach can be critical to an ongoing, fair, and friendly working relationship.

Once a fair and friendly tone of negotiations is established, it becomes much easier for cloud service customers to state their expectations of the negotiation process and what they feel the steps in the process should be. In their most basic form, they can include:

- Identifying requirements in plain business language.
- Folding those requirements into a more formal requirements document.
- Reserving time for the customer to seek some third-party expertise.

For the provider, knowing these steps up front becomes valuable as it minimizes the amount of guesswork and related time the

cloud service provider might waste in trying to understand the customer's preferred negotiation process. It also gives the provider an opportunity to weigh in on that process and add value based on their own experience. The key here is to avoid dwelling on specific wording of contract language (which will come later) and to keep the conversations friendly, practical, and based on business goals.

Once negotiation expectations are established, it's time to develop the relatively short list of customer requirements. Frank discussion with the potential provider should occur, focusing on items essential for business success. This is the critical time to sort them out—and to understand if there are nuances from the provider regarding what can and cannot be delivered. It is also the juncture where the customer has to critically assess if some requirements are unrealistic. The customer has to decide on whether—by accommodating the provider's particular capabilities on a specific negotiating point—it will (or will not) seriously impact the customer's ongoing business. Again, the key here is for the customer to be pragmatic about what must eventually appear in the contract versus what they would like to see in the contract. This means compromising on some hoped-for service requirements that are not actually critical requirements. A consistent look at ''what really matters,'' and what other, possibly noncritical business factors are actually motivating a particular requirement, is important here.

Once the general business requirements are established and fundamentally agreed upon by the customer and the provider, the provider will want to drop the requirements into their standard agreement framework. The customer should then audit the requirements in that legal context to make sure that the intent and meaning aren't altered by legal terminology.

A critical step takes place next: getting outside perspective. A lawyer familiar with these types of agreements should review it to make sure that critical points are not missed. Just as important, however, is for the customer to get some peer review (depending upon desired confidentiality) of the business requirements, the pricing, or the entire agreement, from some trusted colleagues who have purchased similar services. This process can reveal not only pitfalls to avoid, but insight on pricing models successfully negotiated by other companies.

Performance Penalties and Restitution Clauses

The sheer volume and popularity of different contractual arrangements whereby companies buy technology services (as opposed to buying just software and hardware that they then operate on their own) has created an environment full of providers adept at negotiating, if not limiting, penalties they might incur for missed performance. On the other hand, all of this provides valuable lessons in how to apply simple penalties and incentives that deliver the best value. The following are some critical tips customers should consider as they navigate this part of their contract negotiation with a cloud provider:

- *Simplicity is better for all.* In this classic application of ''less is more,'' customers should focus on incorporating only a few pragmatic penalties and incentives and describe them clearly and succinctly. In addition, by defining performance penalties in business terms and relating service levels directly to business processes (rather than to the related technology performance), the customer is able to hold the provider accountable on what really matters to business success. While there may be temptation to focus on uptime and throughput because these are easy performance levels to measure, it is also important to keep the cloud service provider clearly aware of what performance levels or potential problems can truly have a negative impact on the customer's ongoing business operations and become serious threats to joint success.
- *Review potential penalties before signature.* It's one thing to establish the penalties in the contract, but by reviewing the established penalties face-to-face with the provider prior to contract signature, the provider is given a clear understanding that the customer takes performance seriously. In addition, it makes sense to discuss performance and penalties in a regular review cycle that can be specified in the contract.
- *Use care in balancing incentives and penalties.* No one understands the customer's business better than the customer. With that in mind, the customer should not assume that well-meaning but poorly defined incentives and penalties can prevent a well-intentioned contract from running into trouble over time. For example, after a contract is signed, it

can be very easy for the customer to ignore contract details and focus on their business, and for the provider to focus on simply providing service (and spending time on other customers). During this time, seemingly minor, but regular, performance shortfalls and penalties can begin to occur. All of this can be easy to ignore, but it's a sign of a problem that needs to be fixed.

- *Focus on preventing problems.* Well-designed penalties should focus on identifying and alerting all parties so as to react quickly and avoid problems before they become serious. If a problem occurs, it should be a signal for the provider to solve the problem for the longer term, rather than allowing it to continue by incurring a soft but ongoing penalty. One way to accomplish this is to accelerate penalties for problems that persist over time or that the provider is lax in addressing. While this may seem harsh, it causes the provider to clearly understand that the agreement has no tolerance for unresolved problems.
- *Understand that you get what you pay for.* The customer should adopt a stance of being a business realist and avoid temptations to hold the provider to unrealistic service levels that aren't truly required to maintain business operations. And the customer should be willing to pay a fair price for the service levels sought. Demands for flawed and excessive requirements with stiff penalties will not only reduce the field of willing providers, but can result in the customer's overpaying for a service that could be had at a lower cost if requirements were simply more realistic.

Notes

1. Jeff Kaplan, ''The CIO's Guide to Software as a Service: A Primer for Understanding and Maximizing the Value of SaaS Solutions,'' white paper from THINKstrategies http://thinkstrategies.com/researchpublications/whitepapers.html.
2. Cloud Security Alliance, www.cloudsecurityalliance.org/About.html.
3. Cloud Security Alliance, ''Security Guidance for Critical Areas of Focus in Cloud Computing V2.1,'' December 2009, www.cloudsecurityalliance.org/csaguide.pdf.
4. Ibid.
5. Jon Brodkin, ''Gartner: Seven Cloud-Computing Security Risks,'' *Network World,* July 2, 2008.

6. ''Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration, Version 1.0,'' April 2009, www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf.
7. Ibid.
8. Jim Hietala, ''Compliance under a Cloud,'' CIO.com February 24, 2010.
9. This concept was recently recognized as one of 10 finalists for the RSA Innovation Sandbox contest at the RSA Conference 2010 (https://365.rsaconference.com/docs/DOC-2392). Navajo Systems at www.NavajoSystems.com is one of the companies now offering VPS security software.