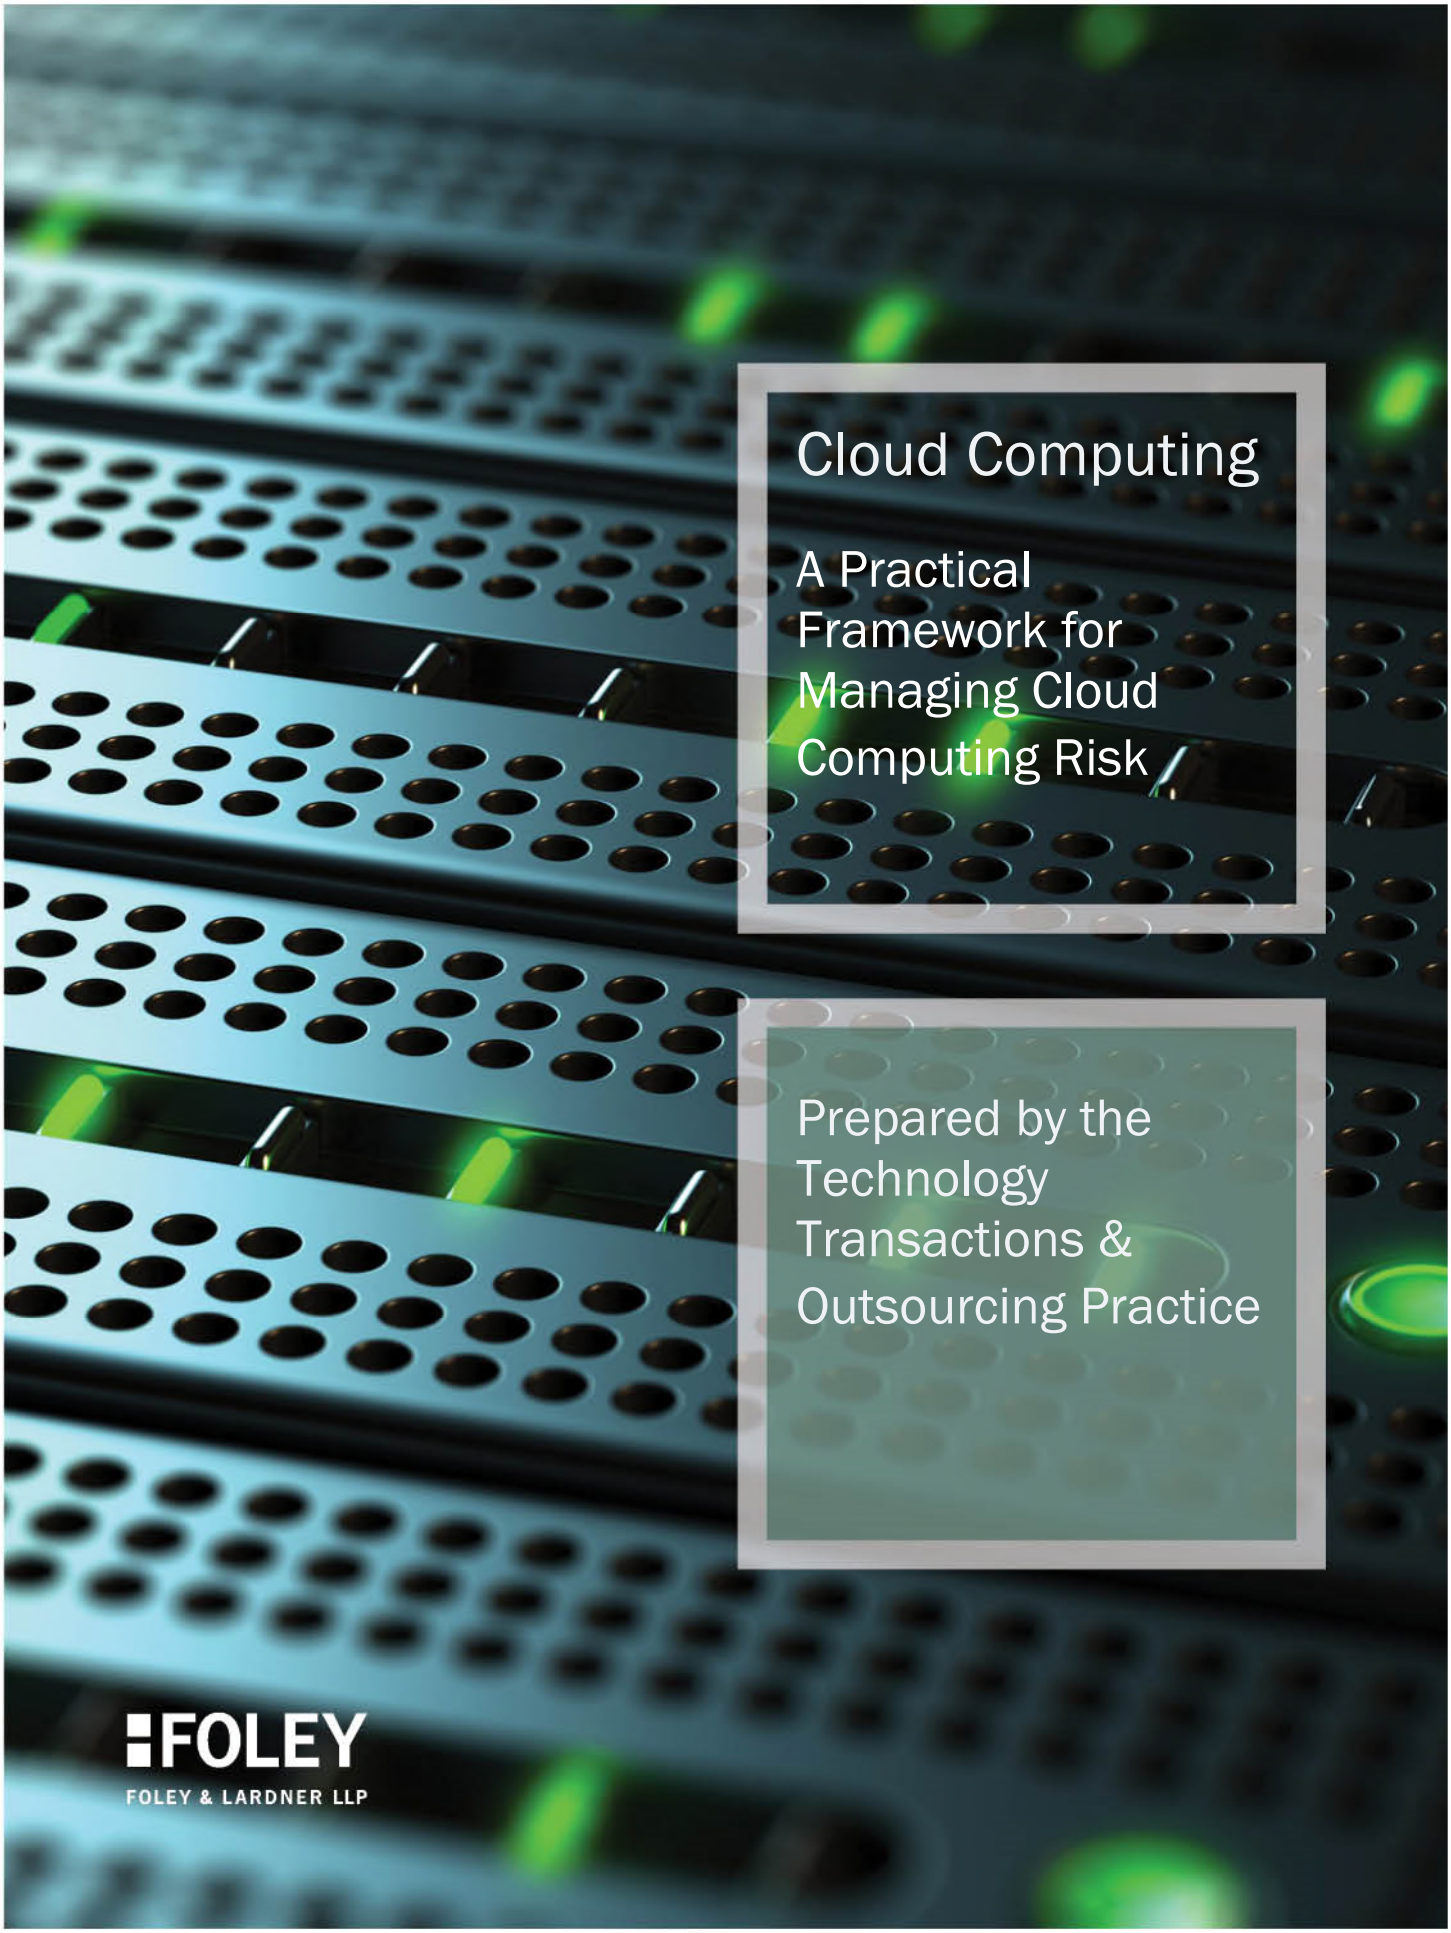

Cloud Computing

A Practical Framework for Managing Cloud Computing Risk

Prepared by the Technology Transactions & Outsourcing Practice

FOLEY

FOLEY & LARDNER LLP

FOLEY

FOLEY & LARDNER LLP

Cloud Computing – A Practical Framework for Managing Cloud Computing Risk

The move to cloud computing is no longer a trend, but a land rush. Gartner, Inc. recently predicted that by 2021, worldwide public cloud services revenue will top an estimated $303 billion.[1]

Already, many entities, including those in highly regulated service areas such as the financial services industry, are utilizing cloud services, in some instances hosting their most sensitive data. By leveraging economies of scale, commoditizing services, geographic distribution, and open source software to lower costs, cloud computing has become an increasingly attractive option for businesses. The question is no longer whether a business will move to the cloud, but when it does will it have the processes in place to effectively balance the risks and rewards of cloud computing.

Cloud computing is not without substantial risk, particularly at a time when businesses of almost every kind are finding themselves subject to an ever expanding range of state and federal data security and privacy laws, document retention requirements, and other standards of accountability. This is reflected in the primary area of cloud computing that has failed to resonate with businesses: using cloud-based security services (e.g., user access and provisioning, two factor authentication services, etc.). A recent semi-annual survey of information security professionals at large and midsize firms in North America asked respondents about whether they would consider using cloud-based security services, less than 15% responded they would be likely to do so.

Cloud providers often tout their services as having superior availability and security. Recently however, numerous high profile failures have caused many in the business community to give additional scrutiny to these claims. These incidents reinforce the need for pre-agreement due-diligence to verify whether a provider's systems are capable of meeting those promises and whether those promises are reflected in the provider's service level agreement.

A portion of Amazon's service serving a portion of the eastern United States was unavailable over a period of several days. Although Amazon did manage to recover almost all data, media reports indicated that some customer data had been permanently lost as a result of the failure.

Although reports indicate that Amazon provided its customers credits beyond its contractual obligation, the incident placed a spotlight on the fine print often included in service level commitments. Amazon's most frequently cited EC2 uptime commitment applied only to unavailability of services that were hosted across multiple "Availability Zones" that Amazon divided its service areas into. A failure confined to a single Availability Zone may not constitute unavailability under this calculation, and would not warrant a service credit for the outage. Although Amazon chose not to do so, such an approach would permit a provider to deny credits to customers that had not paid the additional fees to have their services duplicated in other Availability Zones.

As the criticality of cloud-based application increases, even minor failures can become extremely costly for enterprise customers. While the adoption of cloud based productivity software is in its early stages, such applications are being heavily marketed as a

replacement for locally installed applications. An enterprise client that has widely adopted such a cloud solution without a locally installed alternative could experience significant financial and other productivity costs from even minor outages.

Threats to the security of cloud solutions have become more numerous and brazen recently as well. The high-profile security breach of Sony's PlayStation Network prompted scrutiny of cloud providers' ability to secure and protect sensitive personal information.

Underscoring these concerns is a recent survey of large companies using cloud services which found that nearly half of the respondents experienced a data security lapse or issue in the last twelve months.

While cloud computing risks primarily revolve around data confidentiality, integrity, and availability, they extend to other key relationship issues of provider accountability, price control, international data transfers, and due diligence. There have been industry efforts to address cloud computing issues, such as the CloudAudit Working Group which is working to provide a common interface for cloud computing providers to automate the audit, assertion, assessment, and assurance of their cloud computing environments and allow authorized customers access to that information. Another organization, the Cloud Standards Customer Council, focuses on developing best practices and providing a forum for members to discuss common issues. In addition, because in many cases there are limits on the protections a customer can negotiate with the provider, industry insiders, privacy advocates, and others are calling for legislation and regulation in the cloud computing area. Those protections, if they ever come, are well in the future and unlikely to address the wide-range of risks presented by the cloud today. In fact, a draft report prepared by the National Institute of Standards and Technology found that for the typical customer most areas of the cloud contract are "nonnegotiable."

For the foregoing reasons, our firm, working with our clients, has developed the framework provided in this White Paper for risk management and mitigation in cloud engagements. What Is Cloud Computing?

There have been many different definitions for the term "cloud computing" proposed by technology experts and a wide range of organizations, including service providers, IT research firms, government agencies, and educational institutions. Several popular definitions are provided in Appendix A (Definitions for "Cloud Computing").

However, the different definitions and lack of agreement around a definition have created confusion as to what cloud computing really means. Many of the proposed cloud computing definitions have been criticized as being too broad and vague, which has allowed the term to be applied to almost any technology developed today and leading many to view cloud computing as simply a new marketing term to describe existing service delivery models. For purposes of this White Paper, we attribute the following high level characteristics to a cloud computing service delivery model: (1) delivery over the Internet (cloud), (2) software, platform, or infrastructure resources provided as services, (3) scalability on-demand, and (4) utility and/or subscription billing (i.e., payment based on the customer's actual use and/or a period of time).

Appendix B (Cloud Computing Features and Comparison) provides additional descriptions of common features of cloud computing, and compares cloud computing to some existing delivery models, specifically Application Service Provider (ASP) and Software-as-a-Service (SaaS).

Executive Overview of a Practical Framework for Managing Cloud Computing Risk

A cloud computing approach to IT services can offer many benefits, including cost reduction and service flexibility. By moving software and infrastructure to the provider's remote data center, customers can lower some of the up front risks and complexity associated with realizing the benefits of new technology. Customers may achieve a reduction in capital costs, including the up front investment in new infrastructure, new software licenses, implementation services, and

personnel hiring and/or training. In addition, less equipment means that less physical space at a customer site is needed to store such devices. Further, there may be lower costs as a result of reductions in planning, purchasing, installing, maintaining, managing, and supporting the software and infrastructure, as well as hiring, training, and managing an IT staff.

In addition, customers are attracted to the flexibility of (a) being able to quickly set up and implement an IT solution, (b) being able to access services from anywhere at any time via the Internet, and (c) being able to quickly add and remove IT resources on-demand, so that customers can effectively respond to internal business requirements and changing market conditions. Despite these benefits, businesses must protect against a non-disciplined rush to cloud computing, as the risks and costs to a customer's business of selecting cloud computing without appropriate protections to support a critical business process or to store sensitive data can greatly outweigh these benefits.

The objective of this White Paper is to provide an understanding of the risks in cloud computing, approaches to mitigate those risks, and a framework for evaluating risk within your organization.

We recommend that risk associated with the decision to implement a cloud computing solution should be evaluated primarily by assessing two variables: (1) the criticality of the business process being supported by the cloud computing solution, and (2) the sensitivity of the data that will be stored in the cloud. The graph below illustrates this approach.

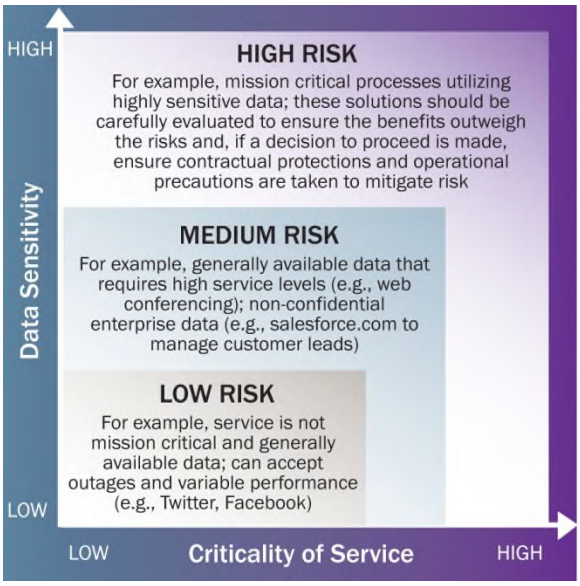

When presented with a potential cloud computing solution, simply plot the corresponding points for data sensitivity and criticality of business process on the graph above to quickly get a read on the overall risk profile for the solution. The risk level assessment should be equated with the highest risk plotted for either variable.

Specific Recommendations for Managing Cloud Computing Risk

Cloud computing involves accessing a provider's software and infrastructure remotely and oftentimes includes storing the customer's data with that provider.

To that end, cloud computing agreements have *some* similarity to traditional software licensing agreements but often have more in common with hosting or application service provider agreements. As such, the most critical issues and concerns identified with respect to hosting and application service provider agreements are equally applicable to cloud computing agreements.

In a traditional software licensing or hardware purchase engagement, the vendor installs the software or equipment in the customer's environment. The customer has the ability to have the software or hardware configured to meet its particular business needs and retains control over its data. In a cloud computing environment, the software, hardware, and

the customer's data are hosted by the provider, typically in a shared environment (i.e., many customers per server) and the software and hardware configuration is much more homogeneous across all customers. Accordingly, the customer's top priorities shift from configuration, implementation, and acceptance to service availability, performance (i.e., service levels), and data security and control. However, like a traditional software licensing agreement or hardware purchase agreement, provisions such as insurance, indemnity, intellectual property, limitations of liability, and warranties remain important as well.

Understanding and Managing Risk Factors

The following identifies areas of risk that arise in Cloud Computing transactions and which we discuss in detail in this paper. Click any area of interest and you will be linked to a summary of the risk and the detailed discussion of how to manage the risk.

1. Identifying and Fixing All Contract Documents

While it may seem fundamental, the task of identifying all relevant contractual documents may not be straightforward. In many instances, all or some portion of the cloud contract may, itself, be hosted in the cloud. That is, the entire agreement or key portions (e.g., the service specifications, support services, and service levels) may be provided online through designated Web pages. This means the contract under which the services will be provided may not be "fixed," but rather may change at any time by the provider changing the relevant Web pages. In most instances, the provider may even decline to provide notice of the changes, placing the burden on the customer to continuously monitor the relevant pages for modifications.

As a first step in assessing the relevant contract documents, the customer should make every effort to "fix" the entire contract, including support obligations, service specifications, etc. in a single document to be signed by both parties. Only then can the customer have confidence as to the service it is actually purchasing.

For example, if the body of the provider's form agreement references a Web page with the support obligations, the customer should ask that the Web page be printed and attached as an exhibit to the contract.

In many cases, however, the provider will refuse to "fix" all elements of the contract, arguing that it is constantly evolving its services, including support and service levels. In those cases, a technique to mitigate risk is to add language to the contract making clear that any future changes in those elements must not (i) materially decrease the level of protection, service, performance, etc. existing as of the effective date of the contract and (ii) impose any materially new or different obligations on the customer. The agreement should also make clear that the provider is obligated to provide affirmative notice to the customer of any changes.

In some cases, the provider will not even make the commitment described above. In those instances, a final risk mitigation measure is to include a termination right for the customer if a later change materially decreases the level of protection, service, performance, etc. existing as of the effective date of the contract.

1. SERVICE AVAILABILITY AND CONTINUITY

AT A GLANCE

- In the event that a provider stops delivering services to a customer, the customer will have no access to the services (which may be supporting a critical business function) and, perhaps more importantly, no access to the customer's data stored on the provider's systems.
- A customer needs to be able to continue to operate its business and have access to its data at all times.
- To mitigate risk in this area, the customer should ensure it obtains: (a) an appropriate uptime service level, (b) data ownership rights and provider's commitment to regular data backups, (c) disaster recovery and business continuity assurances, (d) provider's agreement not to withhold services, and (e) protections against provider financial instability.

In the event that a provider stops delivering services to a customer, whether due to (i) a server being down, (ii) the failure of a telecommunications link, (iii) a natural disaster causing damage to the provider's data center, (iv) the provider withholding services because of a fee dispute, or (v) the provider ceasing to do business because of financial difficulties, the customer will have no access to the services (which may be supporting a critical business function) and, perhaps more importantly, no access to the customer's data stored on the provider's systems and no means for readily retrieving that data.

A customer needs to be able to continue to operate its business and have access to its data at all times. As such, the customer must ensure that it has the proper contractual protections to address the various risks relating to service availability.

(A) SERVICE LEVELS

Appropriate service levels are needed to ensure that service availability is aligned with the customer's expectations. Also, appropriate remedies should be available to ensure the provider is incentivized to perform in accordance with the agreed upon service levels. Item 2 below, titled "Service Levels," discusses the uptime service level and the corresponding remedies in more detail.

(B) CUSTOMER DATA

Data protection provisions commensurate with the sensitivity of the information being placed at risk should be included in the agreement, including a provision that explicitly specifies customer's ownership of any information stored by the provider for the customer, and a provision that requires provider to (i) perform regular data backups to an off-site storage facility and (ii) either deliver periodic copies of all data to customer or provide customer ongoing access to such data. Item 3 below titled "Data – Security, Redundancy, Ownership and Use Rights, and Conversion" discusses data ownership and redundancy in more detail.

(C) DISASTER RECOVERY AND BUSINESS CONTINUITY

Disaster recovery and business continuity provisions should be included requiring the provider to demonstrate and promise that it can continue to make the services available even in the event of a disaster, power outage, or similarly significant event. In the event of a prolonged outage, continuity of services should be provided through a secondary server, data center, or provider, as appropriate. Too often the customer does not request these provisions or, even if it does, it does not read the actual provider policies and procedures. The customer should review any related provider policies and procedures, and obtain contractual assurance regarding disasters and continuity.

By way of illustration, the following is a sample disaster recovery and business continuity provision:

> *Provider shall maintain and implement disaster recovery and avoidance procedures to ensure that the Services are not interrupted during any disaster. Provider shall provide Customer with a copy of its current disaster recovery plan and all updates thereto during the Term. All requirements of this Agreement, including those relating to security, personnel due diligence, and training, shall apply to the Provider disaster recovery site.**

(D) WITHHOLDING OF SERVICES

Also, a general provision prohibiting the provider's withholding of services should be included in any cloud computing agreement. The provider should not be withholding services because of a fee dispute. If the provider believes the customer is in breach of the agreement, the remedy is notice of the breach and, if not cured, termination. It should not have the ability to unilaterally cease performance.

An example of a provision prohibiting the withholding of services is provided below:

> *Provided Customer continues to timely make all undisputed payments, Provider warrants that during the Term of this Agreement it will not withhold Services provided hereunder, for any reason, including but not limited to a dispute between the parties arising under this Agreement, except as may be specifically authorized herein.**

(E) BANKRUPTCY; FINANCIAL WHEREWITHAL

Typically, an agreement may include a provision providing the customer the right to terminate the Agreement in the event of a provider bankruptcy, and may also include a separate provision requiring the provider to assist in transition of the services to a third party provider or to the customer, in the event of expiration or termination of the agreement. However, once the provider has declared bankruptcy, the provider's ability to assist the customer will be limited.

Therefore, if a customer is not confident of a provider's financial stability, the customer should consider adding a provision that enables the customer to identify issues in advance. For example, a provision requiring the provider to deliver periodic reports on its financial condition would enable the customer to assess ahead of time whether the provider will be able to continue to provide services. If the customer identifies any issues, the customer has an opportunity to take the appropriate action to minimize any negative impact.

Provided below is a sample "financial wherewithal" provision:

> *Quarterly, during the Term, Provider shall provide Customer with all information reasonably requested by Customer to assess the overall financial strength and viability of Provider and Provider's ability to fully perform its obligations under this Agreement. In the event Customer concludes that Provider does not have the financial wherewithal to fully perform as required hereunder, Customer may terminate this Agreement without further obligation or liability by providing written notice to Provider. **

(F) IN-HOUSE SOFTWARE SOLUTION

In the event that a provider stops providing "infrastructure" services, the customer may be able to switch to another third party provider with comparable services or purchase the required equipment to replace the infrastructure services. However, the provider's "software" services may be unique and more difficult to replace. Therefore, for critical applications provided as a service, the customer should consider requiring the provider to make available or develop an in-house solution. The inclusion of this provision is very much dependent on the nature of the software provided as a service. The more critical the application, the more important it becomes that the provider be required to develop a long term in-house solution.

2. SERVICE LEVELS

AT A GLANCE

- In addition to obtaining an appropriate uptime service level, it is critical that the customer establish an acceptable service level for performance and responsiveness of the services, as services that fail to provide timely responses to its users are effectively unavailable.
- Other common service level issues that the customer should address are: (a) simultaneous visitors, (b) problem response time and resolution time, (c) data return, and (d) remedies.

One of the most critical aspects in contracting for cloud computing services is establishing appropriate service levels in relation to the availability and responsiveness of the services. Because the software and infrastructure are hosted by the provider, outside the control of the customer, service levels serve two main purposes. First, service levels assure the customer that it can rely on the services in its business and provide appropriate remedies if the provider fails to meet the

agreed service levels. Second, service levels provide agreed upon benchmarks that facilitate the provider's continuous quality improvement process and provide incentives that encourage the provider to be diligent in addressing issues. The most common service level issues that the customer should address are: (i) uptime, (ii) service response time, (iii) simultaneous visitors, (iv) problem response time and resolution time, (v) data return, and (vi) remedies.

(A) UPTIME SERVICE LEVEL

The provider needs to provide a stable environment where the services are available to the customer during the customer's "normal" business hours, and at all other times as needed to support the customer's business. The uptime service level addresses this issue by having the provider agree that the services will have an uptime (i.e., availability) of a certain percentage, during certain hours, measured over an agreed upon period.

By way of illustration, here is an example of an uptime service level provision:

> *Provider will make the Services Available continuously, as measured over the course of each calendar month period, an average of 99.99% of the time, excluding unavailability as a result of Exceptions, as defined below (the "Availability Percentage"). "Available" means the Services shall be available for access and use by Customer. For purposes of calculating the Availability Percentage, the following are "Exceptions" to the service level requirement, and the Services shall not be considered Un-Available, if any inaccessibility is due to: (i) Customer's acts or omissions; (ii) Customer's Internet connectivity; and (iii) Provider's regularly scheduled downtime (which shall occur weekly, Sundays, from 2 am – 4 am central time).**

The specific service level targets depend on the facts and circumstances in each case, including the relative leverage during negotiation. Customers should not simply accept the default provider positions on uptime percentages, measurement periods, and exceptions, but should instead negotiate terms that address the customer's business needs.

A customer should carefully consider the outage measurement window (e.g., daily, monthly, quarterly). Providers tend to want longer measurement periods because they dilute the effects of a downtime and thus make remedies less available to the customer. Customers should also avoid minimum downtime requirements before an outage will be counted toward overall unavailability for a given measurement period. For example, some providers now require that an outage last at least two hours before it is counted toward their availability requirement. This means the provider could be down for repeated instances of an hour and fifty-nine minutes without ever being liable for a service level failure.

Customers should receive written documentation of a provider's scheduled downtime and ensure the window creates no issues for the customer's business. Moreover, the contract should include a limit on the amount of scheduled downtime in any, for example, thirty day period (e.g., four hours in any thirty day period). Most contracts have no limit on scheduled downtime. In fact, many cloud agreements are written such that as long as the provider gives prior notice, the service may be unavailable for an unlimited period of time without a service level failure.

Customers may also request the provider be pro-active in detecting downtime by explicitly requiring the provider to constantly monitor the "heartbeat" of all its servers through automated "pinging." Requiring the provider to do this should result in the provider knowing very quickly that a server is down without having to wait for a notice from the customer. Finally, the concept of "unavailability" should also include severe performance degradation and inoperability of any service feature – see the next section on Service Response Time Service Level.

(B) SERVICE RESPONSE TIME SERVICE LEVEL

Closely related to and, in fact, often intertwined with the uptime service level is the response time service level. This service level sets forth maximum latencies and response times for a customer's use of the services. Services that fail to provide timely responses to its users are effectively unavailable. As with the uptime service level, the specific service level target depends on the facts and circumstances in each case, including the complexity of the transaction at issue, the

processing required, and how critical speed is to achieving the customer's business objectives.

For example, if a customer is accessing Services over an Internet connection, then it is recommended that the service level is set in terms of the Keynote Business 40 Internet Performance Index, which measures the average download time for 40 important business web sites. However, if the services are accessed over a leased line, then the Keynote Business 40 Internet Performance Index may be supplemented or replaced by imposing a response time requirement measured at the provider's external router.

An example provision for a response time service level is provided below:

> *The average download time for each page of the Services, including all content contained therein, shall be within the lesser of (a) 0.5 seconds of the weekly Keynote Business 40 Internet Performance Index ("KB40") or (b) two (2) seconds. In the event the KB40 is discontinued, a successor index (such as average download times for all other customers of Provider) may be mutually agreed upon by the parties.**

If the provider, does not commit to some form of service response time service level, then the customer should ask that the provider at least share its history of response time measurements and should establish some ongoing management of risk in this area. For example, the parties may agree to conduct an end user satisfaction survey, and agree to take action to improve any user dissatisfaction with respect to service response.

(C) SIMULTANEOUS VISITORS SERVICE LEVEL

If the customer expects the services to support multiple simultaneous users, then a service level should be included to explicitly specify such requirement. That is, the contract should make clear that the provider will achieve all required service levels (e.g., uptime, response time, etc.) while supporting up to a designated number of simultaneous users.

(D) PROBLEM RESPONSE TIME AND RESOLUTION TIME SERVICE LEVELS

The provider's obligation to resolve issues in a timely manner needs to be included in any cloud computing agreement. Providers often include only a response time measurement, meaning the time period from when the problem is reported to when the provider notifies the customer and begins working to address the issue. These obligations typically fall short of what is necessary. The agreement should also include a resolution time measurement, meaning the time period from when the problem is reported to when the provider implements a fix or acceptable workaround.

(E) DATA RETURN SERVICE LEVEL

For services involving a critical business function or sensitive customer information, the customer should also consider adding a service level that measures the time between the customer's request for data and the provider's return of such data. This will incentivize the provider to provide the customer data in accordance with the timeframe requirements of the agreement, and provide additional assurance to the customer that it will be able to operate in the event that the provider stops providing services. This service level will, of course, tie directly to the timing of required data backups. That is, the value of the returned data will turn on its currency. If the last backup was a week prior to the incident, the data will likely be useless.

(F) REMEDIES

Typically, remedies for failure to hit a service level start out as credits towards the next period's service. For example, a remedy might provide: for every X increment of downtime below the agreed upon level in the measurement period, or for every Severity Level 1 support issue provider does not resolve within the stipulated time, customer receives a credit of 5% of the next month's bill, up to a maximum credit of 75%. The remedies should scale such that if repeated failure occurs, the customer should have the right to terminate the agreement without penalty and without having to wait for the current term to expire.

Equally important to specifying realistic credits is ensuring that those credits are not worded in terms of the customer's "sole and exclusive remedy." Credits, even if heavily negotiated, seldom will compensate the

customer for its true damages resulting from a service failure. Rather, as discussed above, they are intended only to give the provider an incentive to quickly correct the problem. As such, if the credits are provided as the customer's sole remedy, what appears to be a protection for the customer will, in fact, be a protection for the provider, limiting its liability to a very small amount. The customer's remedies should include, but not be limited to, credits. In some instances, the provider may argue that allowing the customer to also sue for damages will result in a double recovery. While this this argument is without merit, it is frequently possible to reach a middle ground where the customer, on the occurrence of a service level failure, has the option to elect either to receive a credit as its sole remedy or to pursue damages for breach of contract, but not both.

Here is a portion of a sample remedy provision for a service level failure:

> *In the event the Services are not Available 99.99% of the time but are Available at least 95% of the time, then in addition to any other remedies available under this Agreement or applicable law, Customer shall be entitled to a credit in the amount of $_____ each month this service level is not satisfied. In the event the Services are not Available at least 95% of the time, then in addition to any other remedies available under this Agreement or applicable law, Customer shall be entitled to a credit in the amount of $_____ each month this service level is not satisfied. Additionally, in the event the Services are not Available 99.99% for (a) three (3) months consecutively or (b) any three (3) months during a consecutive six (6) month period, then, in addition to all other remedies available to Customer, Customer shall be entitled to terminate this Agreement upon written notice to Provider with no further liability, expense, or obligation to Provider.**

3. DATA — SECURITY, REDUNDANCY, OWNERSHIP AND USE RIGHTS, AND CONVERSION

AT A GLANCE

- The security of a customer's data in a cloud computing environment has been recognized as one of the largest areas of concern for a customer, as the customer is ultimately accountable for complying with privacy and security regulations, and data security breaches have proven to be costly events for organizations.
- In addition, to ensure it is able to continue using its data, the customer should confirm ownership of all data stored by the provider for the customer, and require regular data backups to provide for data redundancy and require the appropriate data conversion to address any data format issues.
- Also, because the provider will have access to, and will be storing, the customer's sensitive information, the agreement should contain specific language (a) regarding the provider's obligations to maintain the confidentiality of such information and (b) placing appropriate limitations on the provider's use of such customer information.

Ensuring customer ownership of its data, and addressing the provider's use of such customer data and the security and confidentiality of customer data are very important in a cloud computing agreement. The provider should provide detail regarding, and agree to reasonable provisions addressing, its competency and its policies and procedures related to: (i) protection against security vulnerabilities, (ii) data backups, (iii) the use of customer data, and (iv) data conversion.

(A) DATA SECURITY

The need for data security is obvious. A cloud computing provider may possess a customer's most sensitive data, including data that may be subject to state and federal regulations (e.g., personally identifiable financial and healthcare information). Loss of data or unauthorized disclosure of such data is a significant concern, because the customer is ultimately accountable for complying with security and privacy laws, regardless of where the data is stored, and data breaches have proven to be costly events for an organization.

Customers should be aware that unique data security issues arise in a cloud computing environment. For example, in an ASP environment, a single physical server may be dedicated to the customer for hosting the application and storing the customer's data. However, in a cloud computing environment, technologies and approaches used to facilitate scalability, such as virtualization and multi-tenancy, may result in customer data being stored on a physical server that also stores data of the provider's other customers, which may increase the risk of unauthorized disclosure. Companies are recognizing the unique security and privacy risks related to a cloud computing service delivery model, and calling on the government for legislation to enhance and strengthen security and privacy protections.

To address data security issues, customers should conduct due diligence regarding the security practices of a provider and include specific contractual protections relating to information security. Part of a customer's due diligence should include identifying the location of the data center where the data will be physically stored and who may have access to the data. If the data center is located in a foreign country, then the customer should be concerned as it may not have an opportunity to inspect the foreign location to ensure it complies with customer's information security requirements. Even if the data center is located in the United States, help desk personnel accessing the data could be located in a foreign country with limited or different security and privacy laws. In addition, the location of the data and the ability of data to be widely distributed across different jurisdictions present complex issues of which law is applicable in a given transaction. At this time, there is very little guidance from courts on these conflict of law issues. For example, if personally identifiable information is located in Europe, then European law may govern that information regardless of what is provided for in the contract. Also, a Vendor may have multiple data centers, each located in a different state in the United States, with each state having its own law regarding data privacy and security. Therefore, to minimize potential issues, the customer should consider adding a restriction against offshore work and data flow to foreign countries, including a requirement that the data center (including the hosted software, infrastructure, and data) be located and the services be performed in the United States, and that no data be made available to those located outside the United States.

In addition, the customer should identify who will be operating the data center. If the provider is not operating the data center itself (e.g., the provider is the owner of the software and will be providing support, but is using a third-party data center to host the software), then the provider should be required to (i) ensure that the third-party host complies with the terms of the agreement (including the data security requirements), (ii) accept responsibility for all acts of the third-party host, and (iii) be jointly and severally liable with the third-party host for any breach by the third-party host of the agreement. Also, the customer should consider entering into a separate confidentiality and non-disclosure agreement with the third-party host for the protection of the customer's data. If the provider ever desires to change the host, the provider should be required to provide the customer with advance notice, and the customer should be given time to conduct due diligence with regard to the security of the proposed host and the right to reject any proposed host.

Providers should be required to provide specific details in the agreement regarding baseline security measures, security incident management, and hardware, software, and security policies. These details need to be reviewed by someone competent in data security – either someone within the customer's organization, a data security attorney, or a third-party consultant. The provider's policies should address security risks particular to cloud computing, and services being delivered over the Internet and accessible through a Web browser (e.g., security risk relating to Adobe Flash which allows hackers to upload malicious Flash objects and launch attacks on users). Some providers will not distribute copies of their security policies but will allow customers to come to the provider's site and inspect them. Such policy inspection should be done if the customer information at issue is very sensitive or mission critical. A customer should compare the provider's policies to its own, and in fact, many customers demand the provider match the customer's policies. The customer should also consider verifying the provider's capabilities via a physical visit or SAS 70 audit (IT internal controls audit) conducted by a third party, or both. It is becoming far more expected that providers regularly demonstrate to

their customers that their security controls remain intact and robust.

Consider the following sample of a typical data security provision:

a. In General. Provider will maintain and enforce safety and physical security procedures with respect to its access and maintenance of Customer Information (1) that are at least equal to industry standards for such types of locations, (2) that are in accordance with reasonable Customer security requirements, and (3) which provide reasonably appropriate technical and organizational safeguards against accidental or unlawful destruction, loss, alteration, or unauthorized disclosure or access of Customer Information and all other data owned by Customer and accessible by Provider under this Agreement.

b. Storage of Customer Information. All Customer Information must be stored in a physically and logically secure environment that protects it from unauthorized access, modification, theft, misuse, and destruction. In addition to the general standards set forth above, Provider will maintain an adequate level of physical security controls over its facility. Further, Provider will maintain an adequate level of data security controls. See Exhibit A for detailed information on Provider's security policies protections

*c. Security Audits. During the Term, Customer or its third party designee may, but is not obligated to, perform audits of the Provider environment, including unannounced penetration and security tests, as it relates to the receipt, maintenance, use, or retention of Customer Information. Any of Customer's regulators shall have the same right upon request. Provider agrees to comply with all reasonable recommendations that result from such inspections, tests, and audits within reasonable timeframes.**

Also, it is critical to require the provider to notify the customer in the event the provider is required by law, lawful order of a court (e.g., request for production of documents), or governmental authority to disclose the customer's data (unless the notification is specifically precluded by such law, lawful order, or government authority). The provider should be required to provide the customer with written notice of the request sufficiently in advance of the date specified for production of the records so that the customer can act to protect its data (e.g., by seeking a protective order from the court). In addition, the provider should be obligated to use reasonable efforts not to release the data pending the outcome of any measures taken by the customer to contest, otherwise oppose, or seek to limit disclosure by the provider.

Lastly, the cloud computing agreement should require that if a breach of security or confidentiality occurs, and it requires notification to customer's customers or employees under any privacy law, then customer should have sole control over the timing, content, and method of such notification. The agreement should also provide that if the provider is culpable for the breach, then the provider must reimburse customer for its reasonable out-of-pocket costs in providing the notification.

(B) DATA REDUNDANCY

Because the customer relies on the provider as the custodian of its data, the customer should demand the cloud computing agreement contain explicit provisions regarding (i) the provider's duty to back up customer data and the frequency of that back up, and (ii) the customer's ongoing access to such data or the delivery of such data to the customer on a regular basis. A good place to start is for the customer to compare the provider's backup policies to its own and make sure they are at least as stringent.

Below is a sample data redundancy provision:

Provider will: (i) execute (A) nightly database backups to a backup server, (B) incremental database transaction log file backups every 30 minutes to a backup server, (C) weekly backups of all hosted Customer Information and the default path to a backup server, and (D) nightly incremental backups of the default path to a backup server; (ii) replicate Customer's database and default path to an off-site location (i.e., other than the primary data center); and (iii) save the last 14 nightly database backups on a secure transfer server (i.e., at any given time, the last 14 nightly database backups will be on the

*secure transfer server) from which Customer may retrieve the database backups at any time.**

(C) DATA OWNERSHIP AND USE RIGHTS

Detailed provisions should be added to clarify that customer owns all data stored by the provider for the customer. In the event that the provider stops providing services and a customer requests the return of its data, there should be no separate dispute as to ownership of the data that resides on the provider's servers.

Because the provider will have access to, and will be storing, the customer's sensitive information, the agreement should contain specific language (i) regarding the provider's obligations to maintain the confidentiality of such information and (ii) placing appropriate limitations on the provider's use of such customer information (i.e., confirming that the provider has no right to use such information except in connection with its performance under the cloud computing agreement).

Many cloud computing providers want to analyze and use the customer data that resides on their servers for their own commercial benefit, in particular, the data customers create as they use the services. For example, the provider may wish to use a customer's data, aggregated along with other customers' data, to provide data analysis to industry groups or marketers. The provider may limit its use to de-identified customer data. These uses are very similar to what businesses and individuals have been dealing with while surfing the Internet while "cookies" follow where a user goes and what a user does.

Here however, the customer data in the cloud is proprietary and confidential to the customer and its business. As such, the customer should consider such use of any of its data very carefully and, if the agreement does not mention these sort of uses, then the customer should ask the provider about its uses and add a provider representation about which uses, if any, are permitted. Most customers should conclude that the provider should not have any right to use the customer's data, whether in raw form, aggregated, or de-identified, beyond what is strictly necessary to provide the services. An example where commercial use might be acceptable is where the provider provides a service that directly depends on the ancillary use of such data, such as aggregating customer data to provide data trending and analysis to customer and similarly situated customers within an industry.

Another area of concern is the practice by some cloud computing providers called de-duplication which removes redundant data from customer files to save storage space in the provider's network. If a customer uploads a file to the provider's network and then later retrieves that file, while it may not appear the content of the file has been altered, the de-duplication process may have removed "meta data" from the file (i.e., data about the file, such as who created it, when it was created or last modified, etc.). The removal of this "hidden" information can result in many issues in the event of litigation. For example, if the customer has agreed to produce meta data in response to an electronic discovery (or e-discovery) request and later finds the data is missing or has been altered, the customer may find itself subject to sanctions in the litigation. In addition, meta data (such as dates and comments) may be useful as evidence at trial, and the customer may not be able to rely on such evidence if it is removed or altered by the cloud computing provider. Further, the file itself may not be admissible as evidence, as the removal of the meta data may bring into question the authenticity of the electronic document in its entirety. The customer should discuss these issues with the provider, and ensure its cloud computing agreement does not contain terms and conditions allowing removal of meta data from files stored in the provider's network.

(D) DATA CONVERSION

Data conversion, both at the onset and termination of the cloud computing agreement must be addressed to avoid hidden costs and being "locked in" to the provider's solution. Going into the relationship, the customer should confirm that its data can be directly imported into the provider's services or that any data conversion needed will be done at provider's cost or at customer's cost (with customer's agreement). A customer should consider conducting a test run of provider's mapping scheme to see how easy or complicated it will be (likewise when checking provider's references, a customer should ask about data migration experiences). Lastly, the customer does not want to be trapped into staying with provider because of data format issues. To that point, the

agreement should include explicit obligations on the part of the provider to return the customer's data, both in provider's data format and in a platform-agnostic format, and thereafter destroy all of the customer's information on provider's servers, all upon expiration or termination of the agreement.

A sample data conversion provision is provided below:

> *At Customer's request, Provider will provide a copy of Customer Information to Customer in an ASCII comma-delimited format on a CD-ROM or DVD-ROM. Upon expiration of this Agreement or termination of this Agreement for any reason, Provider shall (a) deliver to Customer, at no cost to Customer, a current copy of all of the Customer Information in the form in use as of the date of such expiration or termination and (b) completely destroy or erase all other copies of the Customer Information in Provider's or its agents' or subcontractors' possession in any form, including but not limited to electronic, hard copy, or other memory device. At Customer's request, Provider shall have its officers certify in writing that it has so destroyed or erased all copies of the Customer Information and that it shall not make any use of the Customer Information.**

4. INSURANCE

AT A GLANCE

- The customer should self-insure against IT risks by obtaining a cyber-liability policy.
- In addition, the provider should be required to carry the following forms of liability insurance: (a) Technology Errors and Omissions Liability Insurance, and (b) Commercial Blanket Bond, including Electronic & Computer Crime or Unauthorized Computer Access Insurance.

The customer should always address insurance issues in cloud computing situations, both as to the customer's own insurance policies and the provider's insurance. Most data privacy and security laws will hold the customer liable for a security breach whether it was the customer's fault or the provider's fault. Thus, the customer should help self-insure against IT risks, including data and privacy issues, by obtaining a cyber-liability policy.

Cyber liability insurance can protect the customer against a wide range of losses. Most cyber insurance policies will cover damages arising from unauthorized access to a computer system, theft or destruction of data, hacker attacks, denial of service attacks, and malicious code. Some policies also cover privacy risks like security breaches of personal information, may apply to violations of state and federal privacy regulations, and may provide reimbursement for expenses related to the resulting legal and public relations expenses.

Requiring the provider to carry certain types of insurance enhances the likelihood that the provider can meet its obligations and provides direct protection for the customer. The primary forms of liability insurance that a provider should be required to carry are: (a) Technology Errors and Omissions Liability Insurance and (b) Commercial Blanket Bond, including Electronic & Computer Crime or Unauthorized Computer Access Insurance. These types of insurance will cover damages the customer or others may suffer as a result of the provider's professional negligence and intentional acts by others (provider's employees, hackers, etc.). It is critical that the customer require the provider have these sort of policies and not just a general liability policy. Many commercial general liability policies contain a professional services exclusion that precludes coverage for liability arising from IT services as well as other exclusions and limitations that make them largely inapplicable to IT-related risks. The customer should also consider requiring the provider to list customer as an additional insured on its polices; doing so allows the customer to go directly against the provider's insurance company in the event of a claim.

5. INDEMNIFICATION

AT A GLANCE

» The provider should indemnify the customer from third party claims relating to provider's breach of its confidentiality and security obligations and claims relating to infringement of a third party's intellectual property rights.

The provider should agree to defend, indemnify, and hold harmless the customer and its affiliates and agents from any claim where the provider breaches its obligations in regards to the confidentiality and security of the customer's data. Any intentional breach should be fully indemnified, meaning that the customer will have no "out of pocket" costs or expenses related to recovery of the data and compliance with any applicable notice provisions or other obligations required by data privacy laws. In the event the data breach is not intentional, the provider may require a cap on its potential liability exposure, which may be reasonable depending on the type of customer data in question.

The provider should also agree to defend, indemnify, and hold harmless the customer and its affiliates and agents from any claim that the services infringe the intellectual property rights of any third party. This means that the customer will have no "out of pocket" costs or expenses if some third party claims infringement. Providers often try to limit the intellectual property indemnification only to infringement of copyrights. That is not acceptable, as many infringement actions arise out of patent or trade secret rights. The indemnity should extend to infringement claims of any "patent, copyright, trade secret, or other proprietary rights of a third party." In addition, customers should avoid any restriction to patents "issued as of the Effective Date" of the agreement. Providers usually also limit the indemnification to "United States" intellectual property rights, and that is generally acceptable, but the customer should consider whether its use of the services will occur overseas.

6. INTELLECTUAL PROPERTY

AT A GLANCE

» Whether the provider will be performing significant implementation services or simply making modifications to configurable screens based on customer's direction, the customer should be aware of its intellectual property rights and the impact of those rights on its business.

The customer needs to understand the impact of intellectual property rights on its business. In the event the provider will be performing significant implementation services in connection with the cloud computing services, the intellectual property ownership structure proposed by a provider may not effectively address the customer's business needs. If the provider's intellectual property is incorporated into work product delivered to the customer, then such provider intellectual property may be embedded in the customer's business processes as a result. This could encumber the customer's business by creating uncertainty about the customer's rights to such processes on which the business depends. Therefore, the customer should obtain ownership of any "work product" and a very broad license to use any provider intellectual property incorporated into any work product, so that it is able to remain in sole control of the direction of its business and each of its underlying processes.

Even in the case where significant implementation services are not being provided, and the customer is merely providing direction as to configurable screens that will be used by the customer, the customer should realize the potential impact on its business. As a provider may benefit from such ideas provided by the customer, the customer should consider adding a restriction against the provider using those same ideas in services being delivered from provider to any of customer's competitors.

7. LIMITATION OF LIABILITY

AT A GLANCE

- Limitation of liability clauses must be carefully scrutinized.
- While a customer will not likely be able to eliminate the limitation of liability in its entirety, the customer should seek the following concessions: (a) mutual protection (i.e., application of the limitation of liability to both provider and customer, not just provider), (b) carve-outs for particularly important areas (e.g., confidentiality, security, indemnity), and (c) a reasonable liability cap for direct damages.

The provider's limitation of liability is very important in a cloud computing engagement because virtually all aspects of data security are controlled by the provider. Thus, the provider should not be allowed to use a limitation of liability clause to unduly limit its exposure. Instead, a fair limitation of liability clause must balance the provider's concern about unlimited damages with the customer's right to have reasonable recourse in the event of a data breach or other incident.

A provider's limitation of liability clause usually (a) limits any liability of provider to the customer to the amount of fees paid under the agreement or a portion of the agreement (e.g., fees paid for the portion of the services at issue), and (b) excludes incidental, consequential (e.g., lost revenues), exemplary, punitive, and other indirect damages. While a customer may not be able to eliminate the limitation of liability in its entirety, the customer should ask for the following concessions:

- The limitation of liability should apply to both parties. The customer should be entitled to the same protections from damages that the provider is seeking;
- The following should be excluded from all limitations of liability and damages: (a) breach of the confidentiality and security provision by either party; (b) the parties' respective third party indemnity obligations; (c) either party's infringement of the other party's intellectual property rights; and (d) breach of the advertising/publicity provision (see item 12 below titled "Publicity"); and
- The overall liability cap (usually limited to fees paid) should be increased to some multiple of all fees paid (e.g., two to four times the total fees paid or the fees paid in the twelve months prior to the claim arising). Customer should keep in mind that the overall liability cap should not apply to the exclusions in the bullet point above.*

8. IMPLEMENTATION

AT A GLANCE

- When there will be significant implementation services, the client should consider establishing a broad definition of "Services" in the cloud computing agreement.
- This is useful in limiting provider claims of "out of scope" activity and requests for additional money.

In the event significant implementation services are being provided (e.g., extensive software or hardware installation, configuration, or customization services), the definition of "Services" in a cloud computing agreement should be broadly worded to capture all of the services being provided. For example, "'Services' shall mean Provider's provision of software and infrastructure services described in Exhibit __ (Software and Infrastructure Services) and implementation services described in Exhibit __ (Implementation Services), and any other products, deliverables, and services to be provided by Provider to Customer (a) described in a Statement of Work, (b) identified in this Agreement, or (c) otherwise necessary to comply with this Agreement, whether or not specifically set forth in (a) or (b)." A broad definition of "Services" such as the one above is recommended, as it is useful in limiting provider claims of "out of scope" activity and requests for additional money.

In addition, the customer must fully understand its requirements and the capabilities of the services being provided to determine if any additional features or

functionality is needed. Any additional work required to support such features or functionality should be discussed and identified up front, as typically a cloud computing offering may have more limited configuration and customization options (e.g., multi-tenant application) in order for the provider to more efficiently manage the services and provide a more scalable solution. Any additional work agreed upon to support such features or functionality should be included in the description of services.

9. FEES

AT A GLANCE

- The cloud computing agreement should provide customer the ability to both add and remove resources, with a corresponding upward and downward adjustment of the service fees.
- In addition, the customer should identify all potential revenue streams and make sure that the identified fees are inclusive of all such revenue streams.

Typically, a cloud computing service will be offered on a "pay-as-you-go" or "pay-per-use" cost structure (e.g., per virtual machine each hour, per gigabyte of storage each month, per active user each month). Accordingly, the agreement should provide for the ability to both add and remove resources, with a corresponding upward and downward adjustment of the service fees. The best time for the customer to negotiate rates for incremental and decremental use is before signing the agreement. Customers should attempt to lock in any recurring fees for a period of time (one to three years) and thereafter an escalator based on CPI or other third-party index should apply.

In addition, the customer should identify all potential revenue streams and make sure that the identified fees are inclusive of all such revenue streams. For example, the provider may attempt to charge additional fees for additional storage after a certain amount of data, or additional fees for software updates. The customer should ensure that these are included as part of the negotiated fees.

10. TERM

AT A GLANCE

- The customer should be able to terminate the agreement at any time without penalty upon reasonable notice (14 to 30 days).

Because the software and infrastructure are being provided as a service, like any service, the customer should be able to terminate the agreement at any time without penalty upon reasonable notice (14 to 30 days). The provider may request a minimum commitment period from the customer to recoup the provider's "investment" in securing the customer as a customer (e.g., sales expenses and related costs). If the customer agrees to this, then the committed term should be no more than one year and the provider should provide evidence of its up-front costs to justify such a requirement.

11. WARRANTIES

AT A GLANCE

- The customer should seek to obtain warranties related to conformance to specifications, provider performance, third party intellectual property infringement, no pending litigation, etc.

There are several warranties that are typically included in a cloud computing agreement.

The following is a list of warranties that the customer should seek to obtain:

- The services will materially conform to the specifications and, to the extent not inconsistent with the specifications, provider's documentation;
- All services will be provided in a professional, competent, and timely manner by appropriately qualified provider personnel in accordance with the agreement and consistent with provider's best practices;

- The provider will provide adequate training, as needed, to customer on the use of the services;
- The services will comply with all federal, state, and local laws, rules, and regulations;
- The customer's data and information will not be shared with or disclosed in any manner to any third party by provider without first obtaining the express written consent of customer;
- The services will not infringe the intellectual property rights of any third party;
- The services will be free from viruses and other destructive programs;
- There is no pending or threatened litigation involving provider that may impair or interfere with the customer's right to use the services; and
- The provider has sufficient authority to enter into the agreement and grant the rights provided in the agreement to the customer.*

12. PUBLICITY

AT A GLANCE

- The customer should consider including restrictions on publicity relating to the agreement and use of the customer's name.

The customer's reputation and good will are substantial and important assets. This reputation and good will are often symbolized and recognized through the customer's name and other trademarks. Accordingly, every agreement should contain a provision relating to any announcements and publicity in connection with the transaction. The provider should be prohibited from making any media releases or other public announcements relating to the agreement, or otherwise using the customer's name and trademarks without the customer's prior written consent.

13. ASSIGNMENT

AT A GLANCE

- The cloud computing agreement should allow the customer to assign its rights under the agreement to its affiliates and successor entities.

The customer should be able to assign its rights under the agreement to its affiliates and other entities which may become a successor or affiliate due to a reorganization, consolidation, divestiture, or the like. Any concerns the provider may have about an assignment can be addressed by the requirement that the assignee will accept all of the customer's obligations under the agreement. Similarly, the customer should also obtain assurance that any provider assignee will agree to be bound by all of the terms and conditions of the agreement, including without limitation, service level obligations.

14. EXCLUSIVITY

AT A GLANCE

- In order for customers to obtain the best pricing, providers are asking customers to contractually commit to an exclusive engagement.
- Before considering entering into such an arrangement, the customer should ensure that it has the proper protections in the agreement, including excellent service levels, appropriate exceptions to exclusivity, and the right to transition in anticipation of a termination.

More and more providers are seeking exclusivity in their cloud computing contracts. That is, to obtain the best pricing, providers are asking customers to contractually commit to an exclusive engagement in which the customers may not seek similar services from another provider. The challenge of these types of arrangements is that if the contract does not provide excellent service levels and other protections, the customer could find itself bound to an agreement with a poorly performing provider which it cannot terminate and which prohibits customer from seeking supplemental services from an alternate provider.

There are three primary areas to consider in entering into an exclusive arrangement:

- **Is the provider offering strong service levels?** To commit to an exclusive agreement, the customer must have confidence the cloud services will be available when needed and achieve all other performance requirements. Those service levels must be very clearly defined and not be qualified with dozens of vague exceptions, and there must be realistic credits to ensure the provider has sufficient incentive to achieve required performance levels and a customer termination right for continuing or substantial service level failures.
- **Are there appropriate exceptions to exclusivity?** There are situations that may arise in which the cloud provider cannot perform as required under the agreement, but would not be in breach. For example, the provider may be subject to a force majeure event or other circumstance that temporarily relieves the provider of its performance obligations (e.g., a period in which the provider is operating under its business continuity and disaster recovery procedures). The problem is that the customer may still need to conduct its business during the pendency of the event. In such cases, the customer should be relieved of its exclusivity obligations to the extent necessary to obtain temporary services from an alternate provider. Depending on the type of services at issue, if the event continues for more than a few days, the customer should have the right to terminate and permanently transition to an alternate provider.
- **Does the agreement permit transition in anticipation of a termination?** Every cloud agreement will have a defined duration or term (e.g., an initial term of two years, with certain renewal terms). As that term comes to an end, the customer may want to explore a relationship with an alternate provider. To ensure a smooth transition, the customer will likely need the right to enter into an agreement with the alternate provider well before the existing agreement expires. The exclusivity provision must be drafted to include the right for the customer to enter into an agreement with an alternate provider in anticipation of expiration.*

Exclusive engagements can provide the customer with potentially substantial pricing advantages. Nevertheless, any time a customer enters into an exclusive relationship, it is increasing the difficulty of making a change based on performance or pricing or other changes in circumstance, and the advantages of such agreements must be carefully weighed against the overall risk of the contract.

15. PRE-AGREEMENT PROVIDER DUE DILIGENCE

AT A GLANCE

- Pre-agreement due diligence on the provider can provide the customer valuable insight as to whether the provider can meet the customer's expectations.

The customer should consider performing pre-agreement due diligence on the provider. In many instances, the due diligence may be the customer's strongest protection in entering into a cloud engagement. This is particularly so when the cloud contract is presented as largely non-negotiable. In those cases, the customer's only protection is to thoroughly vet the provider prior to entering into the contract.

Diligence can take many forms: site visits, product demonstrations, discussions with vendor personnel, reference site visits, discussions at user groups, industry groups, etc. In addition, diligence can be more formal. This generally takes the form of a customer developing a diligence questionnaire for the provider to complete.

By crafting and using a provider questionnaire, the customer can, at the outset, get a good idea of the extent to which the provider can meet the customer's expectations and, where gaps exist, eliminate them or negotiate through them. Examples of the items to cover in such a due diligence questionnaire include provider's financial condition, insurance, existing service levels, capacity, physical and logical security, disaster recovery, business continuity, redundancy, and ability to comply with applicable regulations. The questionnaire should include language that makes clear the customer will be relying on the responses in

making its decision to enter into a contract with the provider.

Where possible, the completed questionnaire should be attached to the cloud contract as an exhibit. In addition, the contract should require the parties to meet on a periodic basis to discuss updates to the questionnaire responses.

16. POST-EXECUTION ONGOING PROVIDER ASSESSMENT

AT A GLANCE

» Establishing a regular program of evaluating the provider's performance would allow the customer to perform ongoing risk assessments during the term of the agreement.

Lastly, it is recommended that the customer and provider agree to implementation of a regular program of evaluating the provider's performance, under which the provider would be required to supply the requisite information to assess the services, notify the customer of any changes with regard to the provider, and provide any recommendations to improve the services. This information could then be used by the customer to perform ongoing risk assessments, and determine whether to continue the provider relationship.

If possible, post-execution assessments should be coupled with express audit rights under the cloud contract.

Negotiations

If the customer has substantial leverage when negotiating a cloud computing agreement, then the customer should seek to obtain the protections described above. However, in circumstances where the customer does not have such leverage, providers may be resistant to such protections and any modification of its form contract provisions. Therefore, it may not be realistic to expect that the customer can obtain all of the protections listed above.

The customer must then evaluate the business risks, including whether the services support a critical business function, involve sensitive customer information, or are customer facing. If the customer is not able to obtain the level of protection needed in the most significant areas of risk, then the customer should consider walking away from the transaction. If walking away is not an acceptable option, then the customer needs to focus on risk mitigation. For example, if the provider refuses to modify its uptime service level, arguing that it can not separately administer such a service level for different customers, then the customer should negotiate improved remedies and exit rights for a failure of such service level. In this type of situation, where a customer is unable to obtain the appropriate contractual protections and chooses to proceed, the post-execution ongoing assessment of the provider relationship described above becomes even more important.

Conclusion

In conclusion, as businesses are rushing to the cloud to lower costs and achieve service flexibility, there has been a growing recognition of the substantial risks that come with a cloud computing solution. Unlike traditional software licenses and hardware purchase agreements, but similar to hosting and application service provider agreements, the customer needs to focus less on configuration, implementation, and acceptance and more on service availability, performance, and the security and control of the customer's data. By keeping these areas in mind, along with the other risk factors and recommendations identified in this White Paper, customers can more effectively manage and substantially reduce the risks presented by cloud computing relationships.

FOLEY

FOLEY & LARDNER LLP

For More Information

Our Technology Transactions & Outsourcing attorneys possess a unique breadth and depth of experience structuring and negotiating transactions and analyzing issues to increase the value to your enterprise associated with the acquisition, implementation, use, and management of technology.

- Core areas of practice include:
- IT and business process outsourcing
- Software Licensing, Cloud Computing, IT Infrastructure, and Professional Services agreements
- Privacy, Security, and document retention
- Data and voice services
- Online business issues
- Technology procurement and vendor management

For more information about our Technology Transactions & Outsourcing, please contact:

Matt Karlyn
Co-Chair, Technology Transactions Industry Team
617.502.3239
mkarlyn@foley.com

FOLEY

FOLEY & LARDNER LLP

Appendix A — Definitions for "Cloud Computing"

Forrester Research

"a standardized IT capability (services, software, or infrastructure) delivered via Internet technologies in a pay-per-use, self-service way"

Gartner

"a style of computing in which scalable and elastic IT-enabled capabilities are delivered as a service to external customers using Internet technologies"

IDC

"an emerging IT development, deployment and delivery model, enabling real-time delivery of products, services and solutions over the Internet (i.e., enabling cloud services)"

National Institute of Standards and Technology (NIST)

"a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction"

U.C. Berkeley Reliable Adaptive Distributed Systems Laboratory (RAD Lab)

"refers to both the applications delivered as services over the Internet and the hardware and systems software in the datacenters that provide those services. The services themselves have long been referred to as Software as a Service (SaaS) ... The datacenter hardware and software is what we will call a Cloud. When a Cloud is made available in a pay-as-you-go manner to the public, we call it a Public Cloud; the service being sold is Utility Computing ... We use the term Private Cloud to refer to internal datacenters of a business or other organization that are not made available to the public. Thus, Cloud Computing is the sum of SaaS and Utility Computing, but does not normally include Private Clouds."

Wikipedia

"Internet- ('cloud-') based development and use of computer technology ('computing') ... a new supplement, consumption and delivery model for IT services based on the Internet, and it typically involves the provision of dynamically scalable and often virtualized resources as a service over the Internet"

FOLEY

FOLEY & LARDNER LLP

Appendix B — Cloud Computing Features and Comparison

Cloud Computing Features

Although there may be no single widely accepted definition of cloud computing, based on various definitions, the following features are generally common to a cloud computing approach to delivering IT services.

CLOUD

- IT resources are delivered over the Internet "cloud"
  - Note: For purposes of our analysis, we focus on the "public" cloud (i.e., the Internet, which includes the hardware and software systems in the provider's remote data center), as opposed to a "private cloud" (i.e., hardware and software systems in a customer enterprise)
  - IT resources are managed in the provider's remote data center, rather than on the customer's local computers and servers
  - The provider has responsibility for the IT resources, including design, development, procurement, installation, testing, deployment, provisioning, and management

DEPLOYMENT MODELS

- *Private cloud*. The cloud infrastructure is operated solely for any organization. It may be managed by the organization or a third party and may exist on premise or off premise
- *Community cloud*. The cloud infrastructure is shared by several organizations and support a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise
- *Public cloud*. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- *Hybrid cloud*. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds)

SERVICE

- IT resources are delivered to the customer via the Internet and consumed by the customer as a "service"
  - IT resources include software, software development platforms, and infrastructure (including virtual servers, memory, processors, storage, network bandwidth)
  - The customer accesses the services using a Web browser or other interface

SCALABLE ON-DEMAND

- IT resources are scaled up and down at the customer's demand
  - scalability – ability to scale up to "unlimited" resources
  - elasticity – ability to quickly add and remove resources (within seconds or minutes, as opposed to days or weeks)

- Scalability on-demand is often facilitated through use of the following technologies/approaches, which improve resource utilization and allow for a more scalable approach:
  - virtualization – relates to creating a layer of abstraction that converts physical computing resources into a virtual pool of resources, which can be shared by different users (e.g., through server virtualization, a single physical server is partitioned into multiple virtual machines each running a separate operating system, such that the computing resources of the underlying physical server are used as a pool of resources by each of the virtual machines)
  - multi-tenant software architecture – a single instance of software serves multiple customers at the same time, with each customer sharing hardware resources

UTILITY/SUBSCRIPTION BILLING

- utility billing
  - payment is based on the amount of resources used, similar to how one is charged for water, electricity, or gas (e.g., per virtual machine each hour, per gigabyte of storage each month, per active user each month)
- subscription billing
  - payment is based on a period of time, similar to how one is charged for a newspaper or magazine subscription (e.g., per month)

Is Cloud Computing Just Another Name for Existing Service Delivery Models?

There has been much discussion as to whether cloud computing is just another name for existing service delivery models, including Application Service Provider (ASP) and Software-as-a-Service (SaaS). The confusion over whether cloud computing is any different than ASP or SaaS naturally arises from the fact that cloud computing, as with ASP and SaaS, involves the remote hosting of software and delivery of software services over the Internet. Although there is this significant overlap among the terms, there are some subtle distinctions. While SaaS is often considered a type of service under cloud computing, ASP can either be considered a type of cloud computing service as well or be distinguished as something very different depending on how the term "ASP" is defined.

However, regardless of whether a solution is ultimately identified as ASP, SaaS, or cloud computing, it is important to note that these solutions share similar critical risk issues and, as a result, a very similar risk analysis applies to each of these service delivery models.

SOFTWARE-AS-A-SERVICE IS A TYPE OF SERVICE UNDER CLOUD COMPUTING

Cloud computing has generally been broken down into three types of services: (1) Software-as-a-Service (SaaS), (2) Platform-as-a-Service (PaaS), and (3) Infrastructure-as-a-Service (IaaS).

- Software-as-a-Service (SaaS) – refers to provider's software being delivered over the cloud to the customer as a service (e.g., Salesforce CRM).
- Platform-as-a-Service (PaaS) – refers to provider's software development platforms being delivered over the cloud to the customer as a service (e.g., Google App Engine). For example, a customer may use the service to develop, test, and deploy applications that are then hosted on the provider's infrastructure.
- Infrastructure-as-a-Service (IaaS) – refers to virtual servers, memory, processors, storage, network bandwidth, and other types of infrastructure resources, being delivered over the cloud to the customer as a service (e.g., Amazon Elastic Compute Cloud (EC2)). For example, a customer may use the service to obtain multiple virtual servers to enable a scalable deployment of the customer's own applications.

As a result, SaaS is considered a part of the cloud computing service delivery model.

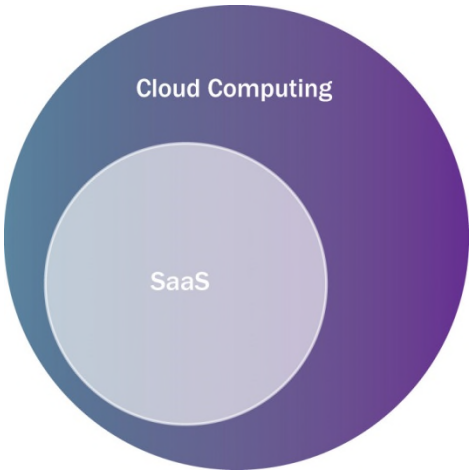

Many take the position that there is no difference between SaaS and ASP. Therefore, if the term "ASP" is defined as a business model that provides software as a service over the Internet (similar to SaaS), then ASP would also be included as a part of the cloud computing service delivery model.

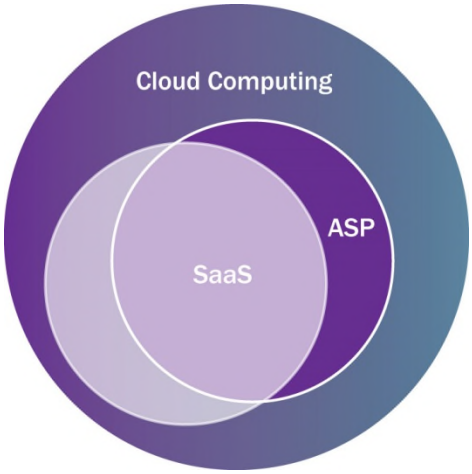

However, cloud computing may be distinguished from both SaaS and ASP in that it also includes providing software development platforms and infrastructure as a service. Through PaaS, a customer is provided access to a provider's software development platform as a service so that the customer can develop, test, and deploy applications. Through IaaS, a customer is provided the option of purchasing IT infrastructure resources (e.g., servers, storage, etc.) as a service delivered across the Internet, instead of the customer having to procure such resources for implementation in its local enterprise network.

ASP MAY BE DISTINGUISHED AS A NON-SCALABLE SOLUTION

Some take the position that the ASP and SaaS service delivery models are different. Those taking such a position often define ASP as a business model that provides only a "single-tenant" approach to delivery of software services (i.e., a single instance of software serves a single customer, with the ability to customize the application features and functionality as required by the customer), and defined as such, the ASP service delivery model is quite different from both SaaS and cloud computing. Cloud computing, as with SaaS, is often described as including a "multi-tenant" approach to delivery of software services (i.e., a single instance of software serves multiple customers at the same time with each customer sharing hardware resources, but resulting in more limited customization options). A "multi-tenant" approach is used in cloud computing to improve resource utilization and scalability, with scalability being one of the most important distinguishing features of cloud computing and essential to making it possible for large-scale data centers to be able to cost-effectively provide computing resources to millions of users as a utility.

UTILITY BILLING FOR CLOUD COMPUTING

Cloud computing may also be distinguished from ASP with respect to billing. While there are cloud computing services that are billed on a subscription basis, cloud computing is also billed on a utility basis, such that the customer only pays for those resources used. Moreover, even if cloud computing services are billed on a subscription basis, usage above a set level usually triggers a utility billing model for such excess usage. Cloud computing solutions can provide customers the ability to request, and enable providers to deliver, only those resources needed, resulting in more efficient use of provider's IT resources (by avoiding over-utilization and under-utilization) and thereby enabling the provider to deliver services at a lower cost. If the provider passes such benefits through to the customer, then utility billing may be a more attractive option for customers.

[1] *Gartner Forecasts Worldwide Public Cloud Revenue to Grow 21.4 Percent in 2018*, Gartner, Inc. (April 12, 2018), at https://www.gartner.com/newsroom/id/3871416.