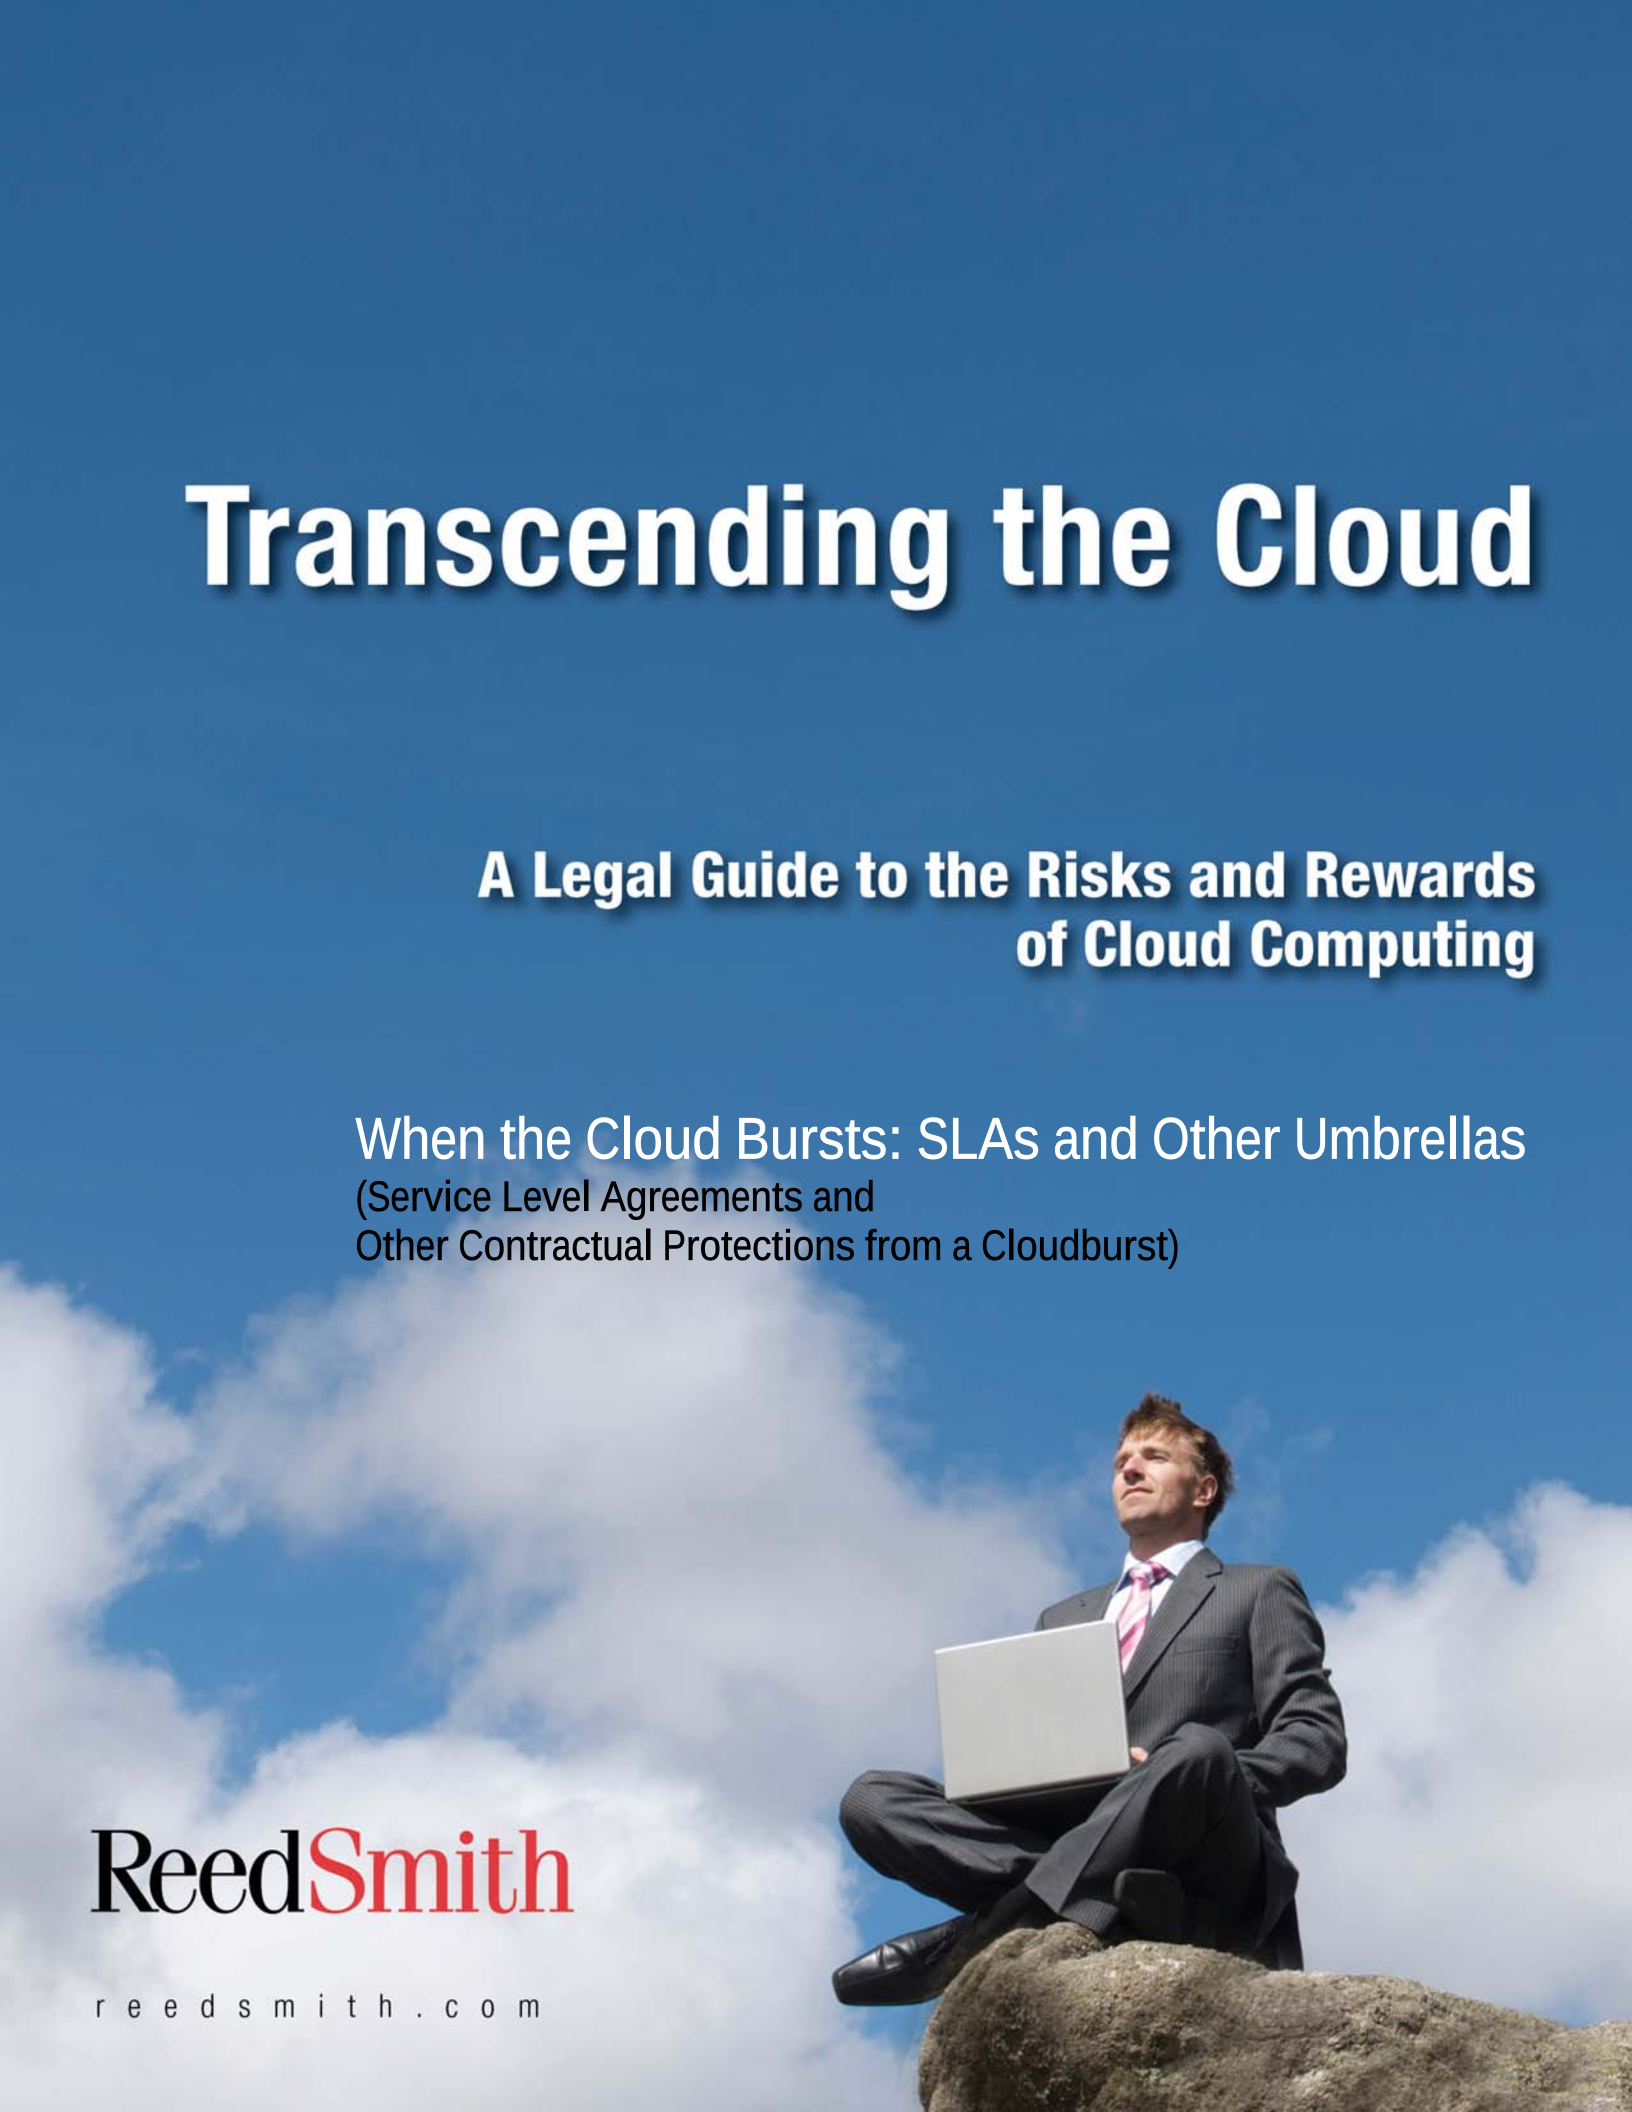

Transcending the Cloud

A Legal Guide to the Risks and Rewards of Cloud Computing

When the Cloud Bursts: SLAs and Other Umbrellas

(Service Level Agreements and
Other Contractual Protections from a Cloudburst)

ReedSmith

reedsmith.com

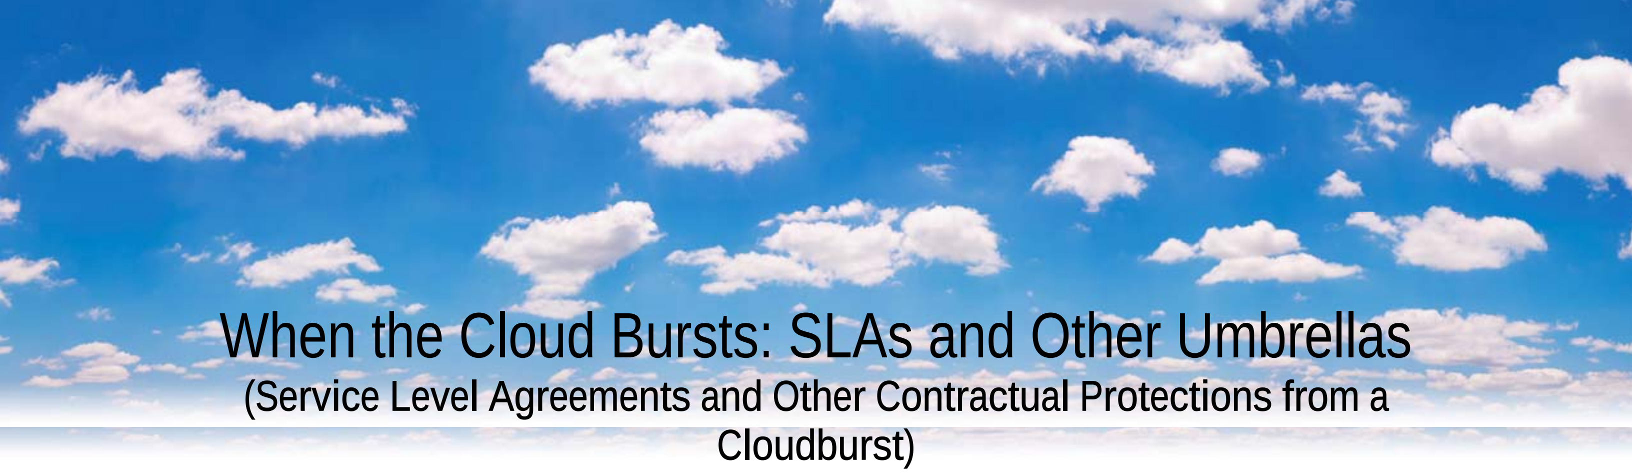

When the Cloud Bursts: SLAs and Other Umbrellas

(Service Level Agreements and Other Contractual Protections from a Cloudburst)

Author

Rauer L. Meyer, Partner – rlmeyer@reedsmith.com

New Benefits, New Risks

Cloud computing is increasingly becoming an appealing method of obtaining computing services, as it offers both dramatically lower costs and scalability, which in turn are the result of features that are inherently double-edged. Among the realities that customers/users of cloud computing must reconcile are:

- Their data, applications and infrastructure are stored and managed by others in remote locations
- Their proprietary data can be stored with the data of other tenants (some of whom may even be competitors) on shared infrastructure (at least in the public cloud)
- Access and use is through the Internet, and hence, depends on its bandwidth and availability
- Hosting facilities are often sited in low-cost locations with cheap power
- Cloud computing providers often subcontract and outsource the provisioning of their services to unknown third parties in unknown locations

New Risks, New Concerns

As customers and providers alike now begin to realize the benefits offered through cloud computing, they must also face a series of new risks and fears. Granted, while some of these concerns existed prior to the onset of cloud computing in the context of third-party services, many are most definitely new. The following is just a sampling of these risks:

- **Loss of service as a result of provider outages.** There have been several well publicized cases recently in which customer data was lost. In fall 2009, a server-failure affected some of T-Mobile's Sidekick customers, resulting in the loss of considerable contact and calendar data. Google Apps has been down on several occasions over the past couple years for several hours at a time, obviously impacting business customers. Amazon S3 was down for almost an entire day in 2008. Back in September 2009, Workday, a provider of human resource, financial, and payroll applications, suffered a 15-hour outage and had to resort to a long backup data center transition.
- **Slow performance and response times** because of connectivity and bandwidth problems and insufficiencies
- **Loss of data privacy and security breaches**. Many surveys of information technology and data processing professionals have put this concern atop the list, even ahead of performance, provider financial liability and business continuity.
- **Ineffective/inadequate disaster recovery.** With many small and mid-size cloud computing providers opting to establish facilities and infrastructure in countries that offer less expensive power and utility resources, more favorable tax laws, and often less stringent business and labor laws and regulations, onsite expertise and oversight may be minimal. Hence, when the cloud goes down, those customers with critical data at risk may not get the fixes, attention and information they need to effectively manage the situation.
- **Uncertain regulatory compliance**. Although customers in regulated industries (*i.e.*, financial

services, health care, broker/dealer, etc.) have the same desire to migrate their networks and systems to a cloud environment for all the benefits available to them via cloud computing, they must be acutely aware of the unique set of risks that other customers in non-regulated industries may not necessarily face. By its distributed nature, cloud computing often blurs the location of and security measures associated with data. These customers or their advisors must be familiar enough with the regulations that govern their business in order to assess the viability and risk levels of putting their data, network services and processing into the cloud.

General Risk Mitigation

As described above, cloud computing can pose potentially serious risks to customers. Thus, how can they reap the benefits of the cloud while minimizing the risks? Cloud computing needs effective and credible risk management, and remedies for failures. Information technology and data processing professionals recommend several approaches to avoid bad outcomes, among them: Recognize that some things may not belong in the cloud (or at least a public cloud) in the first place, such as critical business data, legacy enterprise applications, ERP, personal data, and highly transactional systems or latency-sensitive data. Customers should think twice before moving critical data into the cloud without an effective backup plan.

- Plan a good mix of public, private, and hybrid clouds[1], depending on a customer's risk analysis.
- Conduct a reasonably thorough due diligence of the cloud computing providers being considered. Get references and talk to existing customers. Seek to conduct pilot tests of the provider's system.
- Establish one's own disaster recovery and backup capabilities for anything sent to the cloud, thereby not relying exclusively on the cloud provider.
- Reserve the right and establish a mechanism for the customer to terminate its cloud computing agreement, and confirm (i) one's ability to retrieve its data from the cloud (don't take this for granted), and (ii) one's right to transition from the provider's cloud to another service or to its own data center.

But for all these measures and precautions, bad outcomes may still happen. Accordingly, the customer owes it to itself to be proactive and seek out the best remedy available to it in the service contract—if the cloud should burst.

The SLA Solution

The service level agreement (SLA) part of contracts between providers and customers is a familiar part of almost every computing or information processing service arrangement. In cloud computing, while the SLA serves similar purposes, it requires some adaptation to the new risks of the cloud, and its benefits should get a fresh evaluation in the overall risk management analysis.

Service providers typically offer SLAs as a limited remedy for their customers for failures in the provider's own systems. An SLA specifies service level metrics (*e.g.*, system uptime of 99.99 percent each month, average help desk service response time of 15 minutes). The provider's actual performance is monitored, measured against the standards, and reported to the customer. Substandard performance triggers credits against fees or services, in the nature of liquidated damages, within limits that the provider can live with, especially if (as is usual) many customers will be affected by the same failure. Notably, only failures within the provider's control, will trigger the credits. Providers understandably disclaim responsibility for things out of their control such as Internet connectivity. Finally, often (but not always) the provider requires the customer to agree that these credits are the customer's sole and exclusive remedy for the failure. In other words, even if a customer suffers considerably greater losses as a result of some information technology or data processing failure, it's essentially stuck with the credits and the credits alone.

The SLA is supposed to provide a customer with two kinds of protections:

- An incentive for the provider to perform as promised, giving it skin in the game
- Some compensation for the customer's losses from a failure

However, SLAs are increasingly viewed by customers as unsatisfactory forms of protection that weigh heavily in the provider's favor. First and foremost, disputes often arise over the monitoring of performance and fault, especially when the governing records are those of the provider. Also, if the provider's skin in the game is modest and less than its cost to provide better service, it is not much of an incentive. Moreover, the compensation for customer loss is inherently unpredictable, and in those rare instances in which a customer will be compensated for its actual damage through the SLA, it will generally be coincidental. As a result, customer information technology and data processing departments often view SLAs as more trouble than they're worth.

Without an SLA or an equivalent liquidated damage provision, a customer is left to its general contract remedies, which have their own shortcomings. A customer is in theory entitled to recover its entire loss if it can prove that the provider was at fault and in breach. Information technology and data processing provider contracts invariably disclaim consequential damages (*e.g.*, lost profits) and put a cap on direct damages (*e.g.*, fees paid to the provider). Add to these uncertainties the certain cost and delay of litigation, and it's not a pretty remedy for the customer.

In the current cloud computing market, providers typically promote "reliable service," since this is a common customer concern, and offer SLAs of one variety or another. As an example of current offerings, the SLAs of most providers "guarantee" some uptime metrics ranging from 99.95 percent to even 100 percent availability each month. Amazon EC2 offers 99.95 percent, AT&T Synaptic Hosting offers 99.7 percent, and 3Tera commits to 99.999 percent for a virtual private data center. Many providers offer options at different percentage rates for different prices. But these numbers by themselves translate into small comfort for the customer in the typical case as they measure cumulative downtime (*i.e.*, not per-incident) and their true value turns on the nature and size of the credits. These solutions to the remedy problem will no doubt evolve as customers demand more assurances from cloud providers.

Can cloud computing SLAs even be negotiated? Many public cloud services are available only through non-negotiable click-wrap contracts that cannot be negotiated and strictly limit the provider's liability, since the model is based on a low-cost, one-size-fits-all offering that avoids customization. In this case, the SLA remedy is not worth much. SLAs play a more important role in the private cloud model, where customers can do several things to improve their remedies. Private cloud SLAs are usually negotiable, since the provider is only negotiating with a single user for a single hosting environment, rather than having to guarantee different service levels to different users of the same cloud. The more a customer brings to the provider, such as large upfront fees (*e.g.*, for migration and implementation) or a large volume of services, the more power it will typically have to negotiate. The customer should always try, keeping in mind that better protection will come with higher fees.

Here are some tips a customer should consider:

- Adapt your SLA remedies to your use case. As mentioned above, if you are merely developing a new system that is not overly time- or data-sensitive, you might not need the tightest SLA possible. The provider's standard SLA could very likely suitable. But if a service failure will harm your business significantly, the standard offering will not be enough.

- The basic model of the common SLA is inadequate and should be rethought for cloud service risks. In a given metric (*e.g.*, availability), a single percentage of uptime is specified on a cumulative basis over a month and a single credit is provided if the standard is missed. If it is missed, however, typically a singe credit ($X) or discount is given to the customer against its hosting costs, which constitute the customer's sole remedy. But what if a single outage continues for many multiples of the metric? The customer still gets only its $X, nothing more.

The incentives and compensation in this structure haven't seemed to evolve as quickly as the technological offerings. Customers instead should ask for graduated credits that increase over time with each incident. For example:

Downtime per Incident	Credit
First Hour	$X
Next 2 hours	2$X
Next 2 hours	4$X

By tying the credits to single incidents, the provider is motivated to fix each one and, by increasing the credits over the time of the failure, to fix it quickly. It also better measures and compensates actual loss to the customer. This way, the interests of both provider and customer are better aligned. In return for this more favorable SLA, the customer can more easily accept that these credits will constitute its sole and exclusive remedy for the failure in question.

- Who should be monitoring the provider's performance? The customer should ask that a pre-agreed, third-party, performance-management provider (such as Cloudkick, Gomez, or Apparent Networks) monitor and report provider performance against the SLA's metrics. Many providers will not accept a third-party's measurements when credits are claimed, but even if they do not, a customer is advised to conduct its own monitoring. This, at least, enables the customer to verify the provider's reporting data and detect problems early on, often before the provider takes action.

- The typical information technology and data processing SLA measures availability and customer service response time. The customer should develop additional metrics in a cloud SLA for its own use. If security is critical, the customer should measure security failures. If scalability is critical, the customer should build a metric to measure this. If a provider uses geographically distributed servers in the cloud to serve a global, broad market, the customer should measure the metrics on a region-by-region basis. And, as always, the provider should provide a periodic report of performance against these metrics.
- Customers are strongly encouraged to facilitate proof of the failures that trigger the credits, and evaluate their own internal risks and likelihood of failure. To the extent practicable, the customer must seek to measure the traffic, bandwidth levels, and connectivity in its own network before expanding to the cloud. If a customer understands the points of failure in its own environment, these can be separately mitigated and also facilitate a cause analysis vis-à-vis the cloud provider in the event of failure. This applies especially to the experience of remote workers who are connecting from home networks.

Conclusion

Like most things in life, cloud computing can very much be a double-edged sword. Further compounding some customers' reluctance to entertain and/or migrate into a cloud environment, most cloud computing contracts to date leave customers much to desire. It is essential, therefore, for a customer to have its cloud computing contract reviewed by competent counsel who is knowledgeable and familiar with his/her client's issues and concerns, the technology and services involved, and industry standards. Again, the goal of any contract (and cloud computing contracts no less) should be to capture a fair, balanced and realistic set of terms that depict the transaction, deter complacency, protect that which is most vulnerable, and incentivize the parties to do their best work at all times. This may not be easy to accomplish in the early days of cloud computing, but whoever said the business of technology should be easy?

— Biography of Author —

Rauer L. Meyer, Partner +1 213 457 8124 rlmeyer@reedsmith.com

Rauer's practice focuses on deals for the development, protection, licensing, and other commercialization of information technology (IT), web-based services, cleantech, and other technologies, and the manufacture, procurement and distribution of technology products and services. This includes negotiating and drafting complex outsourcing transactions and other IT procurement and licensing deals. His internet transactions experience includes web site development, hosting, and maintenance arrangements, co-marketing, content acquisition, and customer sale transactions. Rauer also advises clients in franchising goods and services through networks of retail, operations. This includes designing franchise systems, compliance with state and federal laws regulating sale, termination and changes of franchises, and agreement and disclosure documentation, as well as the legitimate avoidance of the franchise laws where appropriate.

— CLOUD COMPUTING TASK FORCE LEADER —

Joseph I. Rosenbaum
Partner and Chair, Advertising Technology & Media Law Group
jrosenbaum@reedsmith.com
+1 212 702 1303

Joe is a member of Reed Smith's global Advertising Technology & Media Law practice, and has more than 30 years of international experience across a wide range of sophisticated and complex commercial transactions, in industries including advertising, entertainment and media, financial services, travel-related services, technology and many more. Joe specializes in the law and policy arising at the intersection of technology and online and behavioral advertising, social media, entertainment, finance, e-commerce, information security and digital rights, online gaming, promotions, privacy and data protection, among others. Joe's experience includes virtual worlds, mobile marketing, digital payments and PCI compliance, digital broadcasting, co-branded credit and gift cards, loyalty rewards programs, branded entertainment, online product placement and endorsements, user-generated content, buzz, word-of-mouth and viral marketing, licensing, software development and outsourcing. Joe lectures and writes extensively and, among others, has authored a book on outsourcing (*Outsourcing Agreements Line by Line*; Aspatore Publishing, 2004) and a seminal law journal article on privacy ("Privacy on the Internet: Whose Information Is It Anyway?"; *Jurimetrics Law Journal*, 1998). Joe's work has been cited by appellate courts, law reviews and journals, industry and trade periodicals. Joe is regularly quoted in widely respected publications such as the *National Law Journal*, Advertising Age, the American Banker, Euromoney and has been interviewed and appeared as a commentator on CNBC's *Squawkbox* and CNN Financial's *Business Unusual*. Joe is General Counsel & Secretary to the Interactive Advertising Bureau and a member of the Advisory Board of the Center for Law, Science and Technology at the Sandra Day O'Connor College of Law at ASU.

— Endnotes —

[1] A public cloud, where data of multiple customers is hosted in a shared environment offering significant economies of scale, is appropriate for non-business critical applications that do not involve core processes, such as the archiving of non-critical data, disaster recovery, and HR. A private cloud, involving dedicated computing environments, is preferred where the quality of service and reliability are critical. Hybrid models combine public and private clouds for a given customer. A development project in which you are merely building and testing a new app with no time sensitivity could be rescheduled and doesn't suffer mightily from an outage; it is appropriate for the public cloud. If on the other hand your data is sensitive to privacy concerns, don't send it to a public cloud, but instead to a private cloud with dedicated servers, or keep it in your data center.