



Can Enterprises Trust the Cloud?

Evaluating the security, reliability and recoverability of cloud environments for business-critical applications.

A Business White Paper

Mike Carpenter and Alex Hamerstone

TOA Technologies

February 2013

Copyright TOA Technologies, 2013. This document is the sole property of TOA Technologies and is strictly confidential. It may not be reproduced, either in part or whole, it may not be transmitted or manipulated, in any form or way, may be it electronic, mechanical, photocopied or recorded, without TOA Technologies' expressed written permission. It may not be lent, rented or in any way transferred without the previous written permission of TOA Technologies, the holder of the copyright. Any breach of these conditions committed by any individual or organization who has access to the documentation will be prosecuted to the full extent of the law.

The Popularity of the Cloud

Over the past decade, more and more enterprises have embraced cloud computing in various forms for various purposes. In a recent survey conducted by the global financial advisory firm KPMG, about two out of every three companies from a broad cross-section of industries reported using, transitioning to or planning cloud-based solutions for key business functions. Asia-Pacific leads all regions in businesses adopting cloud services (68 percent), followed by EMA (Europe-Mediterranean-Asia) at 64 percent and the Americas at 55 percent.

When asked “why turn to the cloud?” decision-makers offer a variety of reasons:

- **Easy, speedy entry** – Cloud-based solutions require no capital outlay for equipment and leave no physical footprint on premises, yet they give businesses access to massive data storage and computing power. They also deploy with exceptional speed. Easy entry and speed of deployment partially explain cloud services’ high adoption rate in the Asia-Pacific region. Young companies in the emerging economies there haven’t invested heavily in information technology (IT) hardware or infrastructure. But they’re turning to the cloud as a way to level the playing field with older, more established and better-equipped competitors elsewhere.
- **Simple budgeting** – Like a utility, cloud services are metered, or priced by usage. The pay-as-you-go model turns data storage, processing and other computing functions from a capital expense to an operating expense, which simplifies budgeting and streamlines purchasing decisions for an enterprise. The scalability and elasticity of cloud-based solutions allow organizations to adapt to changing business conditions.
- **Low maintenance** – Maintaining servers and other hardware in a corporate data center, licensing software and updating it represent significant ongoing expenses. By contrast, cloud-based software (SaaS, for Software as a Service) updates automatically, usually at no charge to the users. Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) employ off-premises hardware, minimizing an enterprise’s need for IT personnel.
- **Greener operations** – Reliance on the cloud helps enterprises achieve their sustainability and energy efficiency goals. A recent study by Colorado-based Pike Research concludes that business, government and consumer use of cloud services will reduce energy consumption and greenhouse gas production by as much as 27 percent over the next decade. “Simply put, clouds are better utilized and less expensive to operate than traditional data centers,” report Pike researchers. “...Only the very largest organizations – both commercial and governmental – will have the capital and expertise to achieve a similar level of efficiency at a comparable cost.”

The Shadow of a Doubt

Despite these widely acknowledged benefits, clouds still cast a vaguely ominous shadow for a sizable minority of companies. Some senior IT executives think of cloud services as maturing technology, not quite ready for enterprise applications. The KPMG study reveals the extent of this doubt. Globally, roughly one out of four businesses surveyed (27 percent) continues to evaluate its options or simply refuses to consider cloud-based solutions.

A 2012 survey of business IT professionals conducted by Information Week offers some reasons for enterprises' wariness. Generally, security issues topped the list of cloud concerns. Specifically, 55 percent of the respondents worried about the integrity of proprietary data, 54 percent expressed anxiety about the integrity of customer data and 45 percent suspected security defects within cloud technology itself. Concerns about the reliability of a cloud environment and the recoverability of data also figured heavily among holdouts.

Given this degree of concern, then, it comes as a surprise that 40 percent of *Information Week's* survey respondents using, planning to use or considering cloud services have absolutely no process for assessing a cloud provider's services. An additional 9 percent of these IT professionals admit that they don't know whether their organization has such an assessment process in place. The information and analysis below serve as a starting point for enterprises to evaluate different types of cloud environments with respect not only to one another, but also to on-premises corporate data centers.

Security

Security involves keeping data safe from theft and unauthorized access. Whatever the cause, lost or compromised data can spell catastrophe for an enterprise. The Ponemon Institute and Symantec Research calculate the cost of a typical data breach at \$194 per compromised customer, or an average of \$5.5 million per organization.

Internal Threats: According to a 2012 report by the global IT analysts at Forrester Research, Inc., employees account for 75 percent of all corporate security breaches. Most employee incidents (63 percent) involve mundane carelessness, such as misplacing a laptop, thumb drive or other device. Certainly, security breaches such as these can occur in any hosting environment. No deployment model completely protects against human error. But breaches due to carelessness become less likely and less significant with cloud-based applications accessible through the Web, because no data resides in the device.

Carelessness also has a flip side: purposeful, malicious theft. So-called "inside jobs" – security breaches perpetrated by executives, employees, interns and other personnel who occupy positions of trust and privilege in an organization – constitute a significant percentage of data losses, and they depend on access to data. Cloud-based solutions present virtually no additional risk.

A cloud-hosting facility may be located hundreds or even thousands of miles away from corporate offices, sometimes in two or more centers widely separated from one another. The personnel operating, maintaining and monitoring in those facilities have far less motive for theft than do an organization's insiders. Statistics bear out that claim. According to a recent annual report compiled by Verizon and the U.S. Secret Service on a broad spectrum of data breaches, business partners and other third parties accounted for less than 1 percent of all problems, and that percentage is falling year after year.

External Threats: Few corporate data centers can match the physical security of the best cloud-hosting centers. For example, facilities meeting the Telecommunication Industry Association's TIA-942/Tier IV standards as a primary hosting site must comply with strict design, construction and procedural specifications. They include entrances with a single-person interlock accessible only with security cards, constant monitoring

from hardened rooms and barriers to prevent vehicles from approaching within 60 feet (18.3 meters) of the building. Only the largest global enterprises can make the investment required to design, build and maintain a data center with such formidable safeguards.

Cyber security is another matter, however, and here the difference between public clouds and private clouds becomes important. (See “Parting the Clouds,” at right.) By their very nature, private clouds remain private by constructing a security perimeter (firewalls and other measures) and controlling access to the resources within that perimeter. The means of controlling access include virtual privacy networks (VPNs), firewall rule sets and encryption, among others. Essentially, these boundary controllers serve as the guarded gates to a private-cloud compound, just as human guards and laser security systems restrict physical access to hosting facilities.

The National Institute of Standards and Technology (NIST) notes that an appropriately strong security perimeter can “protect private-cloud resources against external threats to the same level of security as can be achieved for non-cloud resources.” In other words, private clouds can offer secure exclusivity equal to that of a private data center – so long as the private cloud has strong access control and authentication mechanisms, such as strong passwords.

By contrast, public clouds have no such security perimeter. They are public by definition, open to a multitude of diverse users via the Internet, with a subscriber’s data and applications running right alongside those of other subscribers. As NIST notes, “This introduces both reliability and security risk, and a failure or attack could be perpetrated by any subscriber. Scaling to larger sets of subscribers and resources is one of the important strategies for public clouds to achieve low costs and elasticity; if this scaling is achieved, however, it also implies a large collection of potential attackers.”

Compliance and Auditing: In evaluating the security of any cloud environment, public or private, an enterprise should ask prospective cloud providers whether they comply with the following benchmarks, among the strictest for IT security and integrity.

Parting the Clouds

“The cloud” has become the popular term for a network of shared computing resources (such as servers, hardware, applications and processing services) available on demand in a scalable, measured way to users. But the term fails to distinguish between different types of clouds. In evaluating the security, reliability and availability of cloud environments, enterprises must differentiate between public and private clouds.

- **Public cloud** – The National Institute of Standards and Technology (NIST) defines a “public cloud” as one wholly owned by a business, academic or government institution and provisioned for open use by the general public. It exists on the premises of the cloud provider, but anyone who pays a fee can use it. Amazon Web Services, Google Drive and Dropbox represent well-known examples of public clouds.
- **Private cloud** – In contrast to a public cloud, a private cloud is built and provisioned for the exclusive use of a single organization, even if that organization comprises multiple users or business units. It might be owned, operated and managed by that organization, by a third party or by some combination of both, and it might exist on or off the organization’s premises.

Obviously, the key distinction between a public cloud and private cloud is access. Think of a public cloud like public transportation, a public golf course or a public swimming pool; none of those businesses denies access to anyone who can pay the user fee. By contrast, a private cloud operates more like a corporate car service, a country club’s golf course or health club’s pool, with access limited to “members only” – even if those members must pay a user fee. The distinction becomes critical in analyzing not only the security, but also the reliability and availability of cloud services.

- **PCI-DSS** – Developed jointly by major credit card companies, Payment Card Industry Data Security Standards (PCI-DSS) comprise policies and procedures to ensure the secure transmission, storage and processing of data. They include standards for firewalls, access controls, encryption, processing architecture, monitoring and testing. PCI-DSS compliance assures enterprises that a cloud provider employs the same security measures used by the international banking industry for financial transactions.
- **ISO27001** – Promulgated by the International Organization for Standardization (ISO) and International Electrotechnical Commission, ISO27001 addresses the overall management of security controls. It takes a holistic approach to assessing and addressing risks, ensuring the thoroughness and co-ordination of various protective measures.
- **DIACAP** – The U.S. Department of Defense Information Assurance Certification and Accreditation Process (DIACAP) governs transmitting, processing, storing and managing data used in the defense industry. Military branches and their contractors must adhere to these standards when dealing with sensitive or classified information.
- **Regulatory requirements** – Depending on the nature and location of its business, an enterprise should ensure that a cloud provider complies with all applicable regulations – for example the U.S. Health Insurance Portability and Accountability Act (HIPAA) governing medical information, the Sarbanes-Oxley Act covering certification of financial information or the Data Protection Directorate regulating the processing and handling of personal data within the European Union.
- **SSAE 16** – None of the above standards means much unless a cloud provider can demonstrate independent confirmation of the fairness and accuracy of its claims. The Statement on Standards for Attestation Engagements (SSAE 16) represents the most current guidelines published by the American Institute of Certified Public Accountants for verifying the policies and procedures of service organizations, such as cloud providers. Audits performed under SSAE 16 SOC 2 report on the provider's security, availability, processing integrity, confidentiality and privacy controls. SOC 2 Type 1 audits cover the design and suitability of the organization's controls; SOC 2 Type 2 audits go further, reporting on the operating effectiveness of such controls. As such, they require actual testing.

Cloud providers that cannot demonstrate compliance with security standards or external audits deserve an enterprise's suspicion. Gartner Inc. – a global IT research and analysis firm – notes that providers who do not submit to external audits are “signaling that customers can use them only for trivial functions.”

Reliability

Typically, public-and private-cloud providers advertise uptimes ranging from 99.5 to 100 percent, and enterprises usually demand availability to “four nines” (99.99 percent) for mission-critical applications. However, advertised claims can differ from guarantees offered in service-level agreements (SLAs) – the legal documents that hold providers to their promises and specify remedies for breaking or falling short of those promises.

In evaluating advertised claims or SLAs, enterprises need to examine the precise definition of “uptime” as used by the cloud provider. Often, providers calculate uptime using a specified time interval for loss of service during a year or a billing cycle; service must be unavailable for the entire interval to count against its uptime percentage. So, for example, if a cloud provider specifies a 1-hour interval and suffers a service loss lasting as long as 59 consecutive minutes, it can still claim 100 percent uptime.

Definitions of “up” can become similarly confusing. In some cases, a single cloud subsystem can fail without affecting uptime; multiple subsystems specified by the cloud provider must fail simultaneously before service is considered “down” or unresponsive.

For example, in April 2011, Amazon Web Services – widely regarded among the best public clouds – lost services in its Elastic Compute Cloud (EC2) in a region for four days; the failure temporarily brought down Quora, Foursquare and other websites.

But, because the outage affected only two subsystems of EC2 (Amazon’s Elastic Block Store and Relational Database Service) in only one region, the company didn’t technically breach its SLA. And, by its definition, Amazon could still honestly boast a 99.95 percent uptime for EC2.

None of this suggests that cloud environments cannot or should not support mission-critical applications for business. Rather, it shows that enterprises need to exercise due diligence in reading and interpreting a cloud provider’s SLA. Many service providers calculate uptime by a simple, straightforward method, without including intervals. Furthermore, enterprises should do more than trust a prospective provider’s promises regarding service levels, even legally binding promises. They should investigate and verify the provider’s experience in recent history – the number of actual outages, their duration and their effects on customers.

Recoverability

Earthquakes, hurricanes and other natural disasters can endanger data as much as human carelessness and malice can. Even relatively common events, such as severe thunderstorms, can knock out power and lead to crashed servers and lost data.

As explained earlier, many cloud-hosting facilities have better intrusion-prevention and -detection features than do all but the most expensive corporate data centers. Similarly, their structural designs resist natural disasters better than all but elite corporate data centers. For instance, hosting centers constructed to TIA-942/Tier IV standards must have the ability to sustain at least one worst-case event (such as a Category 5 hurricane on the U.S. East Coast, with its accompanying storm surges) with no critical load impact. Furthermore, Tier IV facilities include doubly redundant critical components; for example, they have a second uninterruptible power supply if the first fails, and a third if the second fails.


For mission-critical enterprise applications, double redundancy won’t suffice. According to NIST’s recommendations, “Subscribers should also investigate whether a candidate provider offers redundancy for the sites they operate and opt for providers not tied to a specific geographic location in case of natural disasters or other disruption.” The best cloud providers – even those hosted in TIA 942/Tier IV-compliant centers – network with at least one other hosting center in another part of the country. Such geographical failover prevents data loss or downtime even in the unlikely event of a catastrophe at the primary hosting center.

Like every good corporate data center, good cloud providers have a disaster recovery plan – one that they will share with their customers. The best providers, in fact, not only provide visibility into the disaster recovery plans, but also practice them to ensure the stated recovery time objective (RTO, the time required to restore service) and the recovery point objective (RPO, the acceptable window of data loss). For example, some cloud providers consider losing 24 hours' worth of data acceptable – although their enterprise clients may not.

Conclusion

Compared with installed software or large corporate data centers, cloud environments offer many compelling advantages for businesses, large and small. They require no initial investment in hardware, deploy quickly, minimize or eliminate maintenance, scale up or down as needed and help enterprises achieve green goals.

However, cautious enterprises should not dismiss concerns about the security, reliability and recoverability before entrusting sensitive data or mission-critical tasks to cloud-based systems.

- **Security** – With the proper hosting facilities, perimeters, access controls, system architecture and procedures, private clouds can offer security equal to or better than all but the most capital-intensive corporate data centers; by contrast, public clouds present a broader spectrum of risks that prove more difficult to manage.
- **Reliability** – In comparing the reliability of either a public or a private cloud to that of an on-premises data center, enterprises should investigate a provider's claims carefully, demand visibility into their prospective partner's policies and procedures and verify the provider's actual performance record.
- **Recoverability** – With multiple, geographically dispersed hosting sites, as well as transparent and practiced disaster-recovery plans, the best cloud providers offer less risk of total failure than most on-premises solutions. 

About TOA Technologies and ETAdirect

Recognized as a leader and the foremost visionary for field service management by Gartner, TOA Technologies has pioneered the convergence of mobile, social and cloud-based technology. ETAdirect, TOA's patented platform, delivers immediate and lasting return on investment through a holistic application that supports the entire field service life-cycle, from the booking of an appointment all the way to completion of the service event, including capacity management, routing optimization, advanced mobility and field management.

As the industry's only complete on-demand solution, ETAdirect...

...learns. Rather than using averages or guesswork, ETAdirect learns the time that each individual employee takes to perform tasks and to travel. And, it keeps learning every second of every day.

...predicts. Using advanced and proprietary pattern recognition and predictive analytics technology, ETAdirect's patented algorithms predict travel times, arrival times and job durations with ever-greater accuracy and precision.

...optimizes. Using time-based data, ETAdirect takes a holistic approach, automatically assigning the right employees to the right jobs, setting the best schedules and mapping the best routes to make sure your business is increasingly optimized every day, enterprise-wide.

...empowers collaboration. Going beyond simple instant messaging and social networking, ETAdirect provides the first platform for truly mobile, context-aware collaboration. It intelligently connects the office to the field and the field staff to each other in ways never before possible.

...communicates. From order to delivery, ETAdirect brings the customer into the loop. It provides customers with end-to-end service visibility, giving them dynamic, valuable information and a real voice for the first time.