USAID
FROM THE AMERICAN PEOPLE

CLOUD COMPUTING & FINANCIAL SERVICES FOR THE POOR

PRIMER & PROCUREMENT GUIDANCE

BRYAN BARNETT

OCTOBER 2011

This document was produced for review by the support of the U.S. Agency for International Development (USAID) under the Knowledge-Driven Microenterprise Development (KDMD) project, implemented by the QED Group, LLC. The views expressed are those of the author and do not represent the views of the U.S. Agency for International Development or the United States Government.

Acknowledgments

Special thanks are due to the following people who provided valuable assistance without which this report could not have been completed. Any faults that remain are the sole responsibility of the author.

Antonio Separovic, Mambu.com, Stutgart, Germany
Bernai Velarde, USAID, Quito, Ecuador
Carlos Carafa, Nube Roja, La Paz, Bolivia
Daryl Skoog, Opportunity International, Denver, CO
Denise Fernandez, USAID, La Paz, Bolivia
Lenny Valdivia, ASFI, La Paz, Bolivia
Fabio Alvino, Utiba Americas, Miami, FL
Hugo Ramos, USAID, La Paz, Bolivia
Ivette Espinoza, ASFI, La Paz, Bolivia
Javier Vaca, Red Financiera Rural, Quito, Ecuador
Jesus Luzardo, Utiba Americas, Miami, FL
Jiten Patel, MicroPlanet.com, Princeton, NJ
Karina Duran, USAID, Quito, Ecuador
Maria Stephens, USAID, Washington, DC
Milton Cevallos, Red Financiera Rural, Quito, Ecuador
Santiago Vasquez, Central Bank of Ecuador, Quito, Ecuador

About the Author

Bryan Barnett is an independent business consultant with more than 15 years experience with technology-based businesses, identifying opportunities, planning programs and evaluating results of technology initiatives. Past clients have included the Bill and Melinda Gates Foundation, Consultative Group to Assist the Poor (CGAP), Grameen Foundation and Microfinance Information Exchange (MIX). Previously, he was a Business Manager at Microsoft Corporation, an analyst and project manager with investment firm Vulcan Northwest, and a founder and Vice-President of ApexLearning, a pioneering online learning company. Barnett holds a Ph.D. in Political Science from Rutgers University and a J.D. from the University of Colorado.

CONTENTS

Executive Summary

Just as in the case of the invention of the personal computer or the advent of the Internet itself, what is today known as 'Cloud Computing' involves a major transformation of the way businesses and individuals access and use computing resources. This change is propelled by the dramatic economic advantages of cloud-based services that offer access to computing resources otherwise unaffordable to all but the largest organizations. Cloud computing eliminates the need for end users to invest in expensive hardware and software, allowing them instead to pay for access to sophisticated computing resources on an as-needed basis.

Though the term 'Cloud Computing' is now widely known, behind the familiar experience of accessing software applications over the Internet is a complex technology landscape with combinations of services and providers forming new and different value chains. While the economic advantages of cloud computing are compelling, these come at a price. Users of cloud-based services are entrusting physical custody of their data and control of critical applications to third parties. Without a clear understanding of the details behind the services upon which they are relying, end users of cloud-based services lack the ability to anticipate and address the risks that inevitably accompany this type of outsourcing arrangement.

Cloud computing potentially has a role to play in a wide variety of international economic development contexts, but two prominent areas of development focus provide important illustrative examples of both the promise and perils of cloud-based services. Emerging cloud-based services now offer applications for portfolio management and accounting otherwise unaffordable for the vast majority of microfinance institutions (MFIs). In another case, the powerful software platforms that power mobile money services are now emerging as a cloud-based hosted service. This introduces an important third partner into the system, operating between the mobile network operators and the banks holding trust accounts.

Though there are risks associated with reliance on cloud-based services, these risks are manageable if they are both identified and anticipated. A systematic and thoughtful approach to qualifying and selecting cloud service partners, combined with giving careful attention to key contract issues in service level agreements will normally be sufficient to protect users of cloud services. However, in cases where a disruption or failure of a cloud service might impact the larger economy or financial system, there is an appropriate role for regulatory oversight. It is therefore important to help MFIs, donors, regulators and others involved in the complex mobile financial services supply chain to understand the issues raised by reliance on cloud-based services.

The regulated financial institutions are more likely to have in place the risk mitigation policies, technologies and staff necessary to address most of the issues raised in this report. However, this may not necessarily be the case with unregulated entities. Therefore, the risks identified in this report are more likely to prevail in mobile financial services (MFS) models that do not fall under direct, regulated banking supervision which key off of global Basel bank norms and standards.

ABBREVIATIONS

CPU – Central Processing Unit
GSM – Global System for Mobile Communications, originally *Groupe Spécial Mobile*
ISAE 3402 – International Standards for Assurance Engagements No. 3402
MFI – Microfinance Institution
MFS – Mobile Financial Services
MNO – Mobile Network Operator
SaaS – Software-As-A-Service
SAS 70 – Statement on Auditing Standards No. 70
SSAE 16 – Statements on Standards for Attestation Engagements No. 16

DEFINITIONS

In all cloud computing scenarios there are four distinct roles. Throughout this document these are referenced as indicated below.

> **End User** – the ultimate consumer of a cloud-based service. The end user may be either an individual or a business.
>
> **Service Provider** – the entity that directly owns the customer relationship with the end user.
>
> **Hosting Provider** – the entity that provides the physical facilities, typically a data center, where cloud-based services originate as they are delivered to end users.
>
> **Software Developer** – the entity that originates the software application(s) that end users access via the Internet.

Trust Account – a depository account in a regulated financial instution that holds currency or highly liquid instruments which stand behind the electronic stored value held on mobile phones.

Central Processing Unit – the chip inside a computer that runs programs and processes data. A single physical machine can contain multiple CPUs.

Short Message Service – the text messaging service common to mobile phones. Though there are a couple of technical variants, the term SMS is commonly used to denote any basic text messaging service on mobile phones.

Mobile Network Operator – a telecommnications entity, such as Vodafone, Safaricom, Orange or Verizon, that provides mobile phone service to end users.

INTRODUCTION

From its earliest days, computing technology has been subject to periodic upheavals that disrupt existing patterns of business and communication. The spread of the personal computer and later the Internet, for example, each brought profound changes, creating new business opportunities and altering existing business practices. Today, what is commonly known as 'Cloud Computing' is driving yet another transformation that is having similarly disruptive effects. And while the term is widely used, many people are still uncertain what Cloud Computing actually is. The aim of this document is to provide a concise and accessible explanation of the cloud computing phenomenon and useful guidance for policy officers, project managers, donors and implementation partners. Cloud-based services offer compelling economic advantages for development projects, but they entail certain risks as well. Understanding and anticipating these risks will be key to successful reliance on cloud services and the substantial benefits they provide.

Cloud computing has grown rapidly in the past several years precisely because it offers access to vast computing resources at significantly reduced cost, with higher levels of security and support than most end users could provide for themselves. At the same time, cloud computing brings expensive software applications within reach of small enterprises that could not otherwise afford this kind of technology. Yet, for all their considerable benefits, end users who rely on cloud-based services are giving physical custody of their data to a third party and relying on mission-critical software applications that they do not manage or control. The risks associated with doing so are genuine, but manageable if identified and anticipated.

Though it has broad implications for many types of development projects, cloud computing is rapidly beginning to assume a central place in the provision of financial services for the poor, both through microfinance and through so-called 'mobile money' (or 'm-money') services. For this reason these kinds of services offer a useful illustration of both the promise and the perils of reliance on cloud services, and they are the subject of special focus in this report. The aim of this document is to outline the evolving landscape of cloud offerings as they are beginning to appear in the microfinance and mobile money domains, and to offer guidance on managing risks. To this end, a systematic approach to procuring cloud-based services is outlined and a check-list of issues to be considered in concluding a contract for cloud-based services is included. These resources are not intended as a substitute for the assistance of qualified legal or technical professionals. Instead, they offer a solid point of departure from which to ask appropriate questions and evaluate options among cloud service providers.

While in the instance of microfinance the risks brought on by outsourcing to third parties in a cloud computing context are primarily borne by individual institutions, in the case of mobile money there is the potential for systemic risk to result from the presence of systemically important payments systems and actors in any cloud-based partnership construct. Accordingly, this document includes a discussion of the regulators' perspective and a brief examination of certain regulatory issues raised specifically by the cloud computing paradigm in mobile financial services.

1. A Cloud Computing Primer

Evolution From Single Computers to Large Data Centers

In the beginning there were isolated computers - some small, others large. Each worked alone taking input from one or more end users and sending back output to a screen or printer. In time, these machines were networked and began communicating with one another, often with a large application (for example, a database or an email server) running on a central server accessed by many individual desktops or laptops. As the number of servers in use grew, they moved from individual offices into central server rooms that in time grew into large data centers often housing hundreds or thousands of machines. While at first connected only to local area networks within individual organizations, the arrival of the Internet ultimately connected all machines and all data centers in one vast worldwide network.

While they were growing in size and being connected to the Internet, data centers were themselves undergoing a transformation. The first data centers were operated by large organizations to house their own computers and data, or by specialized operators who offered to rent space to smaller organizations who wanted to locate their servers in a shared facility. As they grew, data centers became very expensive to operate which created an opportunity for specialized technology outsourcing companies to enter the market offering either to manage data centers owned by others or to provide complete data center operations from facilities owned by the service provider. Thus was born the information technology (IT) outsourcing business, now one of the largest segments of the IT industry.

In the earliest incarnations of outsourcing, the customer (namely the company that was outsourcing some part of its IT operations) continued to own the machines in the data center, if not the entire data center itself; only the management of the data center was outsourced. But in time it became more cost effective in many instances for the outsourcing provider to provide all the machines as well as the physical facility and networking infrastructure and lease these to the customer, who still owned (or licensed) the software they chose to run on the machines they were leasing. Finally, in the last stage of this evolution, software itself began to be provided on a “leased” or “rented” rather than licensed basis and people began talking about computing resources as a kind of utility that would be universally available to be used and paid for only as needed. Some forgotten, nameless individual thought of the vast network of computing resources as a cloud (probably because of an icon used in a PowerPoint presentation) and the name stuck. But there is much more to cloud computing than the simple name implies.

Software-As-A-Service (SaaS)

One aspect of cloud computing is widely familiar (if not well understood) because almost everyone has experienced it. Every time someone searches on Google (or another search engine), sends an email from a Google, Yahoo or Hotmail account, or buys or sells something on eBay, they are using a software application owned by someone else and accessed over the Internet. In one case (eBay) they may be paying a fee, while in

the other cases the service is supported by advertizing. But in all cases there is a complex and powerful application that the end user is accessing as a service delivered over the Internet without the need to own or control any of the machines or software behind the service. Similarly, with online storage services for photos or documents the end user is not using an application as such but is instead accessing computing resources (such as data storage) provided as a service.

Though certainly most visible in the consumer space, this type of arrangement also has a strong and growing place in business-to-business services. The most notable example is SalesForce.com, which has offered customer resource management (CRM) software as a hosted service for several years. Many other such services are now available and more are appearing every day. While it is easy to think of the end user experience of SaaS as equivalent to "cloud computing" this is really a misleading oversimplification. Behind this kind of end user experience is a varied technology landscape that has different implications for businesses that rely on SaaS arrangements.

A data center can be found behind every application accessed as a service over the Internet. Exactly how that data center is operated, and by whom it is operated matters a great deal. In the traditional model of computing, a single organization purchases all hardware and software and hosts these in its own facility. It owns and controls all elements of the IT infrastructure. When the organization leases rack space in a commercial public data center and locates its servers there (referred to as *co-location*) it is purchasing maintenance of the physical environment as a service, while retaining ownership and control of the machines running in that environment as well as the applications running on the machines. A natural evolution of the data center idea is then to let the data center owner provide both the machines and the networking while a client organization just runs its applications there. This stage of evolution is the immediate precursor to a true cloud computing infrastructure.

Pooled Computing Resources: The True Cloud

Once the data center operator is providing all of the hardware and other IT infrastructure, then another critical transformation can take place. When individual customers own the machines on which their applications respectively run, all software applications run on specific machines dedicated to those applications. However, this arrangement is very inefficient because individual machines carry large amounts of unused capacity, and it is very difficult to add or remove capacity when there are sudden changes in demand. In modern large-scale data centers, this arrangement has given way to an alternative in which all machines and all data storage are pooled through a technology called 'virtualization'.

In a process known as 'rapid provisioning' processing power and storage can be allocated in real time to different applications as needed and quickly released when no longer needed. Under this arrangement when an application experiences increased demand, more computing resources are available to meet that demand without any need to purchase, install or configure any new machines themselves. Conversely, when demand slackens, they can release the resources they are using without loss or penalty and those resources immediately

become available to other applications as needed. Finally, adding to the cost efficiencies, the management of all these resources is largely automated through the use of sophisticated management software.

How resources in the data center are organized is, of course, independent of who owns the data center. For large organizations that operate their own data centers, the use of virtualization technology and the pooling of resources in the data center is economically compelling. These operations are known as 'private' clouds to distinguish them from the 'public' cloud data centers that offer hosting and other services to third parties. For small- and medium-size businesses offering SaaS applications, reliance on public data centers is highly cost-effective and allows these enterprises to offer levels of reliability, scalability and security otherwise unobtainable.

Amazon Web Services

One of the first companies to offer true cloud services to the public was Amazon.com, which began offering '*Amazon Web Services*' to customers in 2006. Today Amazon provides a menu of more than 20 different technology components accessible over the Internet which customers can access and use as needed. These components range from simple storage and virtual machine instances, to payment systems, database and messaging services.

Ownership and Control in the Data Center

When evaluating any SaaS offering, it is important to consider who actually owns or controls various parts of the infrastructure upon which the service rests. Here we explore several dimensions of ownership and control and look at the ways in which cloud computing is disrupting traditional value chains in the technology marketplace.

Taken as a whole, there are three distinct patterns of ownership and organization in the data center that lies behind any SaaS offering and which have important implications for the ultimate customer of the SaaS offering, the end user. These patterns are summarized in the following table.

	SaaS		
	On-premises	Co-location	Cloud
Who owns & controls the data center facility?	End user	Independent data center operator	IT service provider
Who owns & controls individual machines within the facility?	End user	End user	IT service provider
How are IT resources within the data center organized?	Either dedicated or pooled ('private cloud')	All machines dedicated	Pooled ('public cloud')

In the first two scenarios, on-premises and co-location, the organization offering a SaaS application to end users (the 'service provider') retains control of and responsibility for some or the entire IT infrastructure supporting their business. In the third case, the public cloud, the entity that controls the software controls

only this and is dependent on another provider for the entire IT infrastructure to run the software. Note that when relying on these outsourced IT services, the service provider of the SaaS application is able to take advantage of economies of scale otherwise not available, and the service provider can perhaps offer a much more secure and stable infrastructure as well. No one of these three programs for providing IT hosting is necessarily better in a normative sense. How well each of these programs might serve to meet the needs of end users depends upon particular circumstances. Understanding which type of IT infrastructure a particular SaaS offering relies upon is critical to estimating the provider's ability to meet specific end user service requirements. For example, a provider relying on co-location must provide its own hardware and may have limited capacity to service or maintain that hardware if it fails. The provider that relies entirely on its own "on-premises" hosting has immediate access to machines, but may not be able to afford the kind of temperature-controlled secure environment or communication bandwidth of a commercial data center. Generally speaking, relying on outsourced IT hosting, ideally in a cloud-structured facility, is by far the best option for all but the largest companies offering software services over the Internet.

2. The Evolving Value Chain

Microfinance

In the microfinance arena, SaaS takes the form of portfolio management and accounting systems offered to MFIs over the Internet. This is a significant innovation because the very high up-front cost of licensing software for on-premises installation is prohibitive for most MFIs. The emerging SaaS offerings allow MFIs to pay a fee for use of the software that is tied to the number of active accounts, loan portfolio size, or some other metric that allows the MFIs' costs to grow as they grow their loan or savings portfolio.

Any SaaS offering has three distinct components, which may be separate functions within a single organization or components contributed by different organizations. First, there is a software developer who is responsible for creating and maintaining the core application provided by the SaaS offering. Second, there is an IT hosting function which supports the facility that hosts the application and provides connectivity to the Internet. Third, there is a service function. This is the customer-facing component that owns and manages the customer relationship, including sales, support, billing, etc. SaaS offerings with varying arrangements are now available to MFIs from both established companies and new entrants.

New Opportunities for MFIs

While the advent of cloud services has created opportunities for new companies, it has also created new opportunities for established organizations in the microfinance industry. Red Financiera Rural (RFR) is a network of more than 40 microfinance institutions in Ecuador that has formed a subsidiary, Sí RED, to provide core banking software to RFR member institutions. Sí RED has licensed the FitBank system from BanTec and contracted with Global Crossing for data center services. Charges for use of the FitBank application will be based on the number of active user accounts at each institution. Initial deployments at four institutions are scheduled from November 2011 to January 2012.

Utilizing this type of model a traditional microfinance association, acting on behalf of its members, is able to offer sophisticated back office software that would otherwise be unaffordable to those institutions on an individual basis.

The most traditional arrangement is one in which a single company develops software and offers it to customers from its own data centers. In this instance a single company provides all three basic components of the SaaS offering. Google is a classic example, writing all its own software and delivering its applications from its own data centers.

In another arrangement, the company that develops software offers it online and maintains the customer relationship while relying upon a third party for the IT hosting component. Mambu is an example of such an arrangement. The company develops its own software but purchases hosting services from Amazon, which allows Mambu to concentrate resources on software development while taking advantage of the capacity and flexibility of the Amazon cloud infrastructure. The Mifos open-source portfolio management software is similarly available as a SaaS offering from several organizations that host it on the Amazon cloud infrastructure.

In yet another arrangement, the company behind a SaaS offering concentrates only on the customer relationship, licensing the software it offers and contracting for data center services. By way of example, MicroPlanet offers software developed by InfrasoftTech and uses data center services provided by Opportunity International. MicroPlanet, as a non-profit organization, aspires to be a leading provider of low-cost portfolio management software to small and medium MFIs. By partnering for both software and IT hosting it is able to apply limited resources most efficiently to address the needs of small institutions that are very cost conscious.

These examples illustrate the ways in which cloud computing is changing the value chain in information technology, as different providers of individual components are brought together to create a complete service offering. The flexibility of these arrangements lowers the cost of entry for new providers and allows for innovative new services to emerge quickly yet deliver high quality offerings based on best-of-breed components. At the same time, however, it is important for the end user of these new services to be able to determine how well the SaaS provider will be able to meet its performance requirements. Knowing what other partners the service provider is relying on is critical to evaluating how reliable the providers offering will be. In many instances a provider who relies on outsourced IT hosting or licenses software from a specialized software developer will in fact have a stronger offering because the outsourced components are superior to anything the service provider could create on their own.

Mobile Financial Services

Introducing an outsourced cloud component in the mobile financial services context is more complex owing to the ways in which mobile money schemes are structured as partnerships between Mobile Network Operators (MNOs), banks, and payments platform providers.

The core of any mobile money system is a piece of software that maintains customer accounts for stored value on mobile phones, settles transactions between mobile account holders, and communicates the resulting balances to agents and customers as well as banks holding trust funds that generally stand behind the

electronic value stored on the phone[1]. This software is developed by technology companies which specialize in mobile payments software. In principle, the software might be licensed by an MNO or a bank and operated by either institution in their own data centers, but in many instances this software is actually provided as a cloud-hosted service by independent third party providers. Figure 1 below illustrates a typical arrangement.

In this arrangement, transactions originate with customers or agents who send SMS messages[2] to a mobile operator which then routes these messages via the Internet to a mobile payments service provider. The transactions are completed (referred to in the finance industry as 'clearing') using mobile payments technology provided by the service company. The service provider then sends transaction data to the bank which holds the funds that back the electronic stored value on the phones.

Figure 1: Cloud Service for Mobile Money

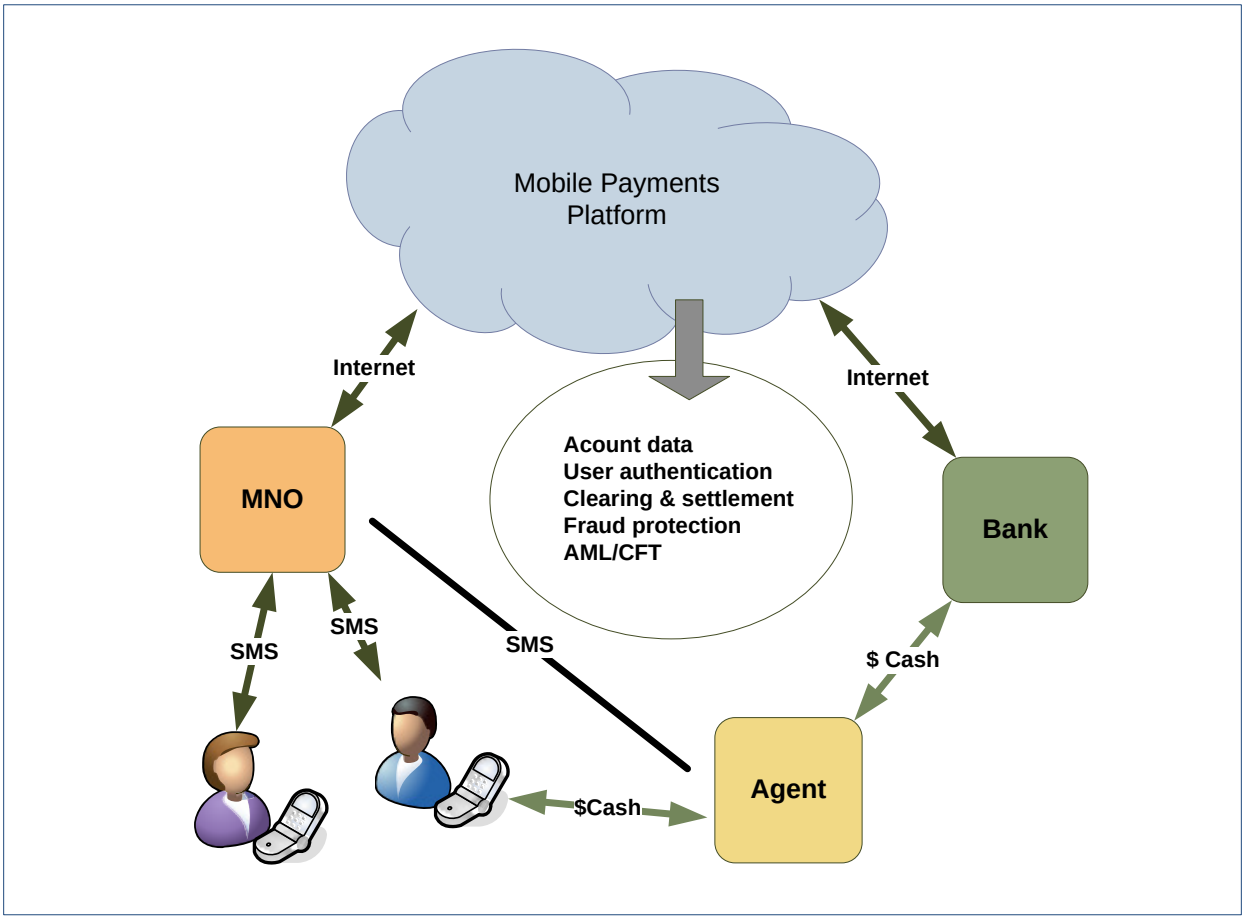

[1] The trust account maintains 1:1 parity between the electronic stored value on the phone (e-money) and highly liquid financial instruments or cash.

[2] In practice there are five distinct mechanisms by which transaction data can be sent from mobile phones to a mobile money platform. These are summarized in Appendix B.

As with other cloud-based offerings, the service provider may rely on software that it has developed or it may rely on software licensed from others. Similarly, the provider may rely on its own data centers or may procure this service from a third party. For example, data center services for the M-PESA service in Kenya have been provided by Vodafone in the U.K. In Bolivia, a new company, Nube Roja (a local affiliate of Red Cloud Technologies), proposes to offer a cloud-based mobile money transfer service using software licensed from Iceni Mobile. At the time they were interviewed for this report Nube Roja had made no final decision on hosting, but in early 2011 a successful pilot test of the Red Cloud platform was conducted from a data center in the U.K. Regardless of the ultimate choice, Nube Roja does not plan to create its own data center, but will rely on contracted hosting services.

Ecuador's National Mobile Money System

While mobile money transfer or payment systems are most commonly sponsored and managed by either mobile network operators or banks, in Ecuador the Central Bank has undertaken to create a nationwide mobile money system by establishing a fully interoperable national mobile money system backed directly by the Central Bank. The planned system will be fully interoperable among mobile networks in Ecuador and will operate through multiple agent networks licensed by the Central Bank.

Because they were not in a position to create the technology for such a system themselves, the Central Bank published a set of requirements and invited proposals from potential technology partners. Untiba Americas was selected to provide the software platform for the Ecuadorian system as a hosted service. At the time Utiba was interviewed the service was hosted in Miami and linked to Ecuador via the Internet. Utiba Americas is a joint venture between Alternet Systems of Miami, Florida and Utiba Technology of Singapore. Alternet provides the hosting and customer facing elements of the offering, while Utiba provides the software.

The first pilot rollout of the Ecuadorian mobile money service is scheduled for November 2011.

3. Risks and Regulatory Concerns Associated with Cloud Computing

Local v. Systemic Risk

Reliance upon SaaS offerings can pose risks for the individual microfinance institutions that rely on them. But the consequences of a failure of a cloud service in such a case is limited to the customers of that service. By contrast, failure in a cloud service that underpins a large scale mobile payments service can pose a systemic risk to the larger financial system. In their analysis of risks attendant upon any mobile money service, the authors of the USAID/Booz Allen Hamilton Mobile Financial Services Risk Matrix identify several risks that have special significance in the cloud computing context.

- A customer may be unable to cash out electronic stored value because of a network outage preventing communication with the mobile money service. (§1.1 p.14; §3.4 p.26)
- An employee of the mobile money service "manipulates agent credit allowances, agent e-money balances, or customer e-money balances for financial gain" (§4.1 p.30; §7.10 p.51; §7.11 p.52)
- "System availability is not maintained by the account provider" (§ 4.6 p.34; § 7.9 p.54; §7.10 p.51)
- "Inadequate transaction records impair investigation of fraud or criminal activity" (§7.4 p.45; §4.5, p.33)

Six Risk Categories

From a cloud perspective, these risk scenarios can be conveniently grouped into six categories:

1) **Internet service interruption**. Because interactions with any SaaS application are conducted via the Internet, any disruption in Internet traffic, including both increased latency or outright interruption of service, will adversely affect the ability of customers to use the software on which they rely. Service interruptions are common enough when individual customers are unable to access the Internet because of problems with their local Internet Service Provider (ISP)[3], but these problems are very localized, affecting only a small number of users at a time. A more significant risk occurs when Internet access is interrupted for a whole nation, as in the case of political conflict or civil unrest. Egypt recently experienced interruptions related to the social unrest in the country.[4] Moreover, as local Internet traffic grows, particularly at times of peak demand, local Internet services can become overwhelmed and slow dramatically, significantly degrading the quality of service for financial applications that require rapid processing of financial transactions in real time.

2) **Data center security and stability.** The data center is the physical location of the hardware and software upon which a SaaS application relies. The data center is subject to four types of risks.
   a) **Physical Security.** A data center is maintained by professional personnel who have access to all machines and software in the center. Improper or unauthorized access is a security risk. (See below).
   b) **Power and Cooling.** Any failure of power or cooling (data centers generate a lot of heat that will damage processers if not controlled) can result in compromised performance or data loss.
   c) **Disaster Recovery.** Data centers are subject to the same risks of natural disaster as any other physical facility. These should be anticipated in the design of a data center location.
   d) **Automation.** Cloud computing is delivered from data centers, but not every data center is host to a cloud infrastructure. Perhaps the most notable distinguishing feature of cloud data centers is the degree of automation upon which they rely. Because vast numbers of physical devices (software instances, CPUs, storage drives, routers, etc.) are allocated as necessary in response to actual processing needs, all cloud infrastructure relies on complex automation software to manage the resources in the data center. A recent incident at an Amazon data center outside of Washington, DC dramatically illustrates what can happen when this automation technology malfunctions. Such a malfunction, triggered by a software update initiated by data center staff, produced a cascading failure

[3] The ISP provides Internet access to end users. The most common challenge at the local level is traffic congestion where a large number of local users attempting to use a limited amount of network bandwidth cause slower response times (known as increased latency). These delays will cause some applications to time out before responding to user input.
[4] See: http://www.oafrica.com/broadband/libyan-Internet-service-interrupted-for-hours/, retrieved July 18, 2011; and Richtel, Matt. "Egypt Cuts Off Most Internet and Cell Service," New York Times, January 28, 2011. Retrieved from http://www.nytimes.com/2011/01/29/technology/Internet/29cutoff.html on March 13, 2011.

that brought down the entire data center for an extended period of time, interrupting the business of many Amazon cloud customers.[5]

3) **Data Privacy and Security.** Unauthorized access to applications or data is a major security concern for all cloud services. Cloud systems are vulnerable to compromise because data from multiple enterprises is housed jointly and employees of the cloud provider have access to systems they administer. Even when there is no malicious intent, SaaS providers who have immediate access to the data of a large pool of customers will have an interest in mining that data for marketing purposes or otherwise using data to advance their own business interests.

4) **Political Instability.** Data centers are subject to political conditions in the jurisdictions where they are located. Instability can affect data center physical security as well as the privacy of data.

5) **Bankruptcy or Transfer of Control.** Any company providing SaaS services is subject to normal business risks which can result in the provider ceasing operations (going out of business) or being acquired by another company. In the case where an end user has license software directly from a vendor for their own use, they can continue to use the software and retain control of their data even if the original developer goes out of business or is acquired. But if a company or individual is contracting with a SaaS provider, the end user does not have custody of their data nor do they have the ability to continue to use the software should the provider cease operations. This is a particular risk where the SaaS provider is a small startup operation, as many in the microfinance or mobile money fields tend to be. The prospect of a SaaS provider being acquired by another company is certainly less consequential than going out of business, but it does nonetheless present risks for the SaaS customer, who may find terms of service altered and previous contractual terms voided. Moreover, the SaaS customer who wishes to migrate to another application or provider may face significant difficulty in obtaining copies of their data in an accessible form.

Mifos: A Cautionary Tale

While business failures among cloud-based SaaS offerings are not common, the case of the Mifos project sponsored by the Grameen Foundation offers a cautionary tale. Mifos is an open-source portfolio management system for MFIs. Some users of Mifos paid the Foundation for access to the software as a SaaS offering.

When the Foundation decided it was no longer able to continue support for Mifos, customers of the SaaS offering were left to find another provider and migrate all their data in a relatively short period of time. When the Foundation ceased paying for the hosting of the application it would simply cease to be available.

Because Mifos is an open source project users have the option to install and run it as an on-premises solution, though few are equipped to do so. And, of course, this option is available only because Mifos was an open source project to begin with, which most SaaS offerings are not.

[5] See: Bright, Peter. "Amazon's lengthy cloud outage shows the danger of complexity," ARS Technica, April 2011. Retrieved from http://arstechnica.com/business/news/2011/04/amazons-lengthy-cloud-outage-shows-the-danger-of-complexity.ars on April, 29 2011.

6) **Contract Enforcement.** Outsourcing of technology services relies on contractual agreements between parties to allocate responsibility for risks and provide remedies. In the case of cloud services, and especially in the developing country context, the ability to enforce contractual agreements cannot be assured. First, service providers and their customers are often located in different countries. Second, customers in developing countries who contract for services may have a very difficult time pursuing judicial remedies either in their local court systems (because those are weak) or in the courts of the host country of the service provider (because it is too costly). Finally, while contracts are a valuable means of setting mutual expectations, they may not be easily enforceable in a cloud-outsourced context involving multiple jurisdictions, especially when developing countries are concerned.

Regulatory Issues

As noted above, cloud computing as applied to microfinance represents localized risks, while mobile money services represent potential systemic risks in light of their unique partnership structure and, as a consequence, are an appropriate concern of regulators. Issues specific to situations where mobile money services rely on cloud-based services include:

- **Scrutiny of Outsourcing Agreements**. When the core component of a mobile money system is outsourced to a third party provider who delivers the service over the Internet from a remote location, it is unclear if any regulatory authority extends to the provider. The normal response is to treat the bank or MNO who contracts with the cloud service as the responsible party, but regulators may not be aware of the extent to which these institutions are reliant upon outsourced providers for critical components of the system. While regulators do not normally inquire into the specifics of outsourcing agreements, greater scrutiny may be warranted in the case where the outsourcing agreement pertains to a software service that manages a very substantial percentage of financial transactions in a given jurisdiction.

Nube Roja: Illustrating the Regulatory Challenges

Among many emerging providers of mobile money platform services, the model proposed by Nube Roja in Bolivia illustrates the challenges confronting financial regulators posed by mobile money platforms offered as a service. Nube Roja is a newly formed company that wants to offer a complete mobile money transfer service as a stand-alone business operating in partnership with banks and MNOs. As proposed, Nube Roja would operate a cloud-based mobile money platform interoperating with MNOs (for text messaging) and with banks (for management of accounts holding funds to back the electronic stored value on mobile phones). Nube Roja would also recruit, train and manage the agent network for cash-in/cash-out.

Nube Roja's contention has been that their proposed operation is allowed under existing regulation. However, when interviewed in June 2011 the position of the Bolivian financial regulator (ASFI) was that Nube Roja (or any similar business) represented a new type of financial institution requiring new regulations.

At the same time there appeared to be confusion concerning Nube Roja's offering reflected in the comment of one member of the regulator's staff that Nube Roja was "just a software provider" to the banks. That Nube Roja represented a complete outsourcing of the mobile money platform rather than a software licensing arrangement was apparently not yet appreciated by the regulator at the time they were interviewed.

- **Access to Customer Data by Foreign Jurisdictions**. A cloud-based mobile money platform will potentially store and manage customer data in a location outside the country where the mobile money service is used by customers. There are unresolved questions concerning the ability of law enforcement personnel or regulators to access customer data held in their jurisdictions when the data belong to customers located in another jurisdiction.

4. Guidelines for Procurement of Cloud-Based Services

Daunting as they might seem, the risks associated with cloud computing are in reality similar to other outsourcing risks faced by most businesses. The key, as in all mitigation situations, is anticipating risks and accounting for them in business planning. At the same time, regulatory authorities have a role to play in maintaining an operating framework designed to ensure the overall integrity of financial systems and to avoid systemic risk.

Specification of Requirements

In order to have a meaningful standard against which to evaluate different SaaS or other cloud-based offerings, it is essential to have clearly defined requirements consistently applied to all potential providers. This is often the most difficult, but nonetheless the most important phase of any procurement process because there will be many different views of what is needed within any organization. It is critical that a final statement of requirements be in written form in order to provide consistent communication to all potential providers. Requirements must be prioritized because it is unlikely that any provider will be able to meet every requirement fully and a final choice should, at a minimum, fulfill all the most important requirements. If requirements are determined in light of clearly established business needs and goals then prioritizing the features and capabilities that a SaaS offering must meet will be notably easier. The assistance of an experienced technical expert will often be helpful in drafting requirements.

Solicitation of Proposals

Based on the requirements specified, it is normal to solicit proposals from multiple vendors. In addition to addressing those requirements, some additional questions should be addressed by any vendor offering cloud-based services.

- Do they develop their own software or license it from a third-party developer? If software is licensed, are there any significant conditions attached to the license (limiting the markets where the service provider can operate, for example)?
- Is data encrypted? Are data transmissions sent over the Internet encrypted? What type or level of encryption is employed?

- What, if any, strategies are employed to overcome bandwidth constraints on the Internet (for international data traffic)?
- Where and how is data hosted? What terms of service apply to their data hosting (if provided by a third party)?
- Have data center facilities received any sort of security audit by a reputable third party? Has any sort of certification been obtained?
- How do customers obtain copies of their data? Are there any conditions or limitations on access?
- What types of guarantees or warranties are provided? Do they cover liability for security breaches in addition to service availability?
- Does the business carry any form of insurance covering losses to customers?

Due Diligence

When requirements are published, proposals are received and an initial screening is complete, the next important step is to undertake due diligence analysis of the prospective providers. Some of the required information will come from the provider themselves, but it is generally valuable to get information from informed third parties as well. Asking appropriate questions and obtaining satisfactory answers is essential to identifying the best partner (who may not be the company that most completely meets all the stated requirements).

Typical questions one might ask of a potential cloud services provider include:

- How long have they been in business?
- Is the business properly licensed in the jurisdiction where it maintains its official place of business?
- Does the business have outstanding debts? Is it financially sound?
- How many current customers does it have? How many did it have a year ago? Is the customer base growing, static, or in decline?
- How large or adequate is the staff? How is the staff deployed? What percentage of the current staff is available on a regular basis for customer support relative to other functions?

Note that there is no definite 'right answer' to these questions. The fact that a vendor does not use encryption, for example, may not be significant depending on the circumstances. It is nevertheless important to have this information before making final decisions.

As part of a normal due diligence process it is advisable to ask for the names of current customers who can be contacted for references. Though the names provided will invariably be those of satisfied customers, it is still worth talking to them to get details.

Data Center Audit Standards

Operations in the data center are critical to the reliability of any cloud-based service. Because customers of these services have little or no opportunity to assess the quality of data center facilities or operations employed by their outsourcing partners, there is a growing trend toward independent audits. There are two aspects to evaluating a data center and there are distinct (but overlapping) standards applicable to each. One type of standard applies to the physical construction of the data center. The other type of standard applies to operations in the data center. Data center auditing practices and standards are still not fully mature and there are a number of competing standards of both types. Most important for the cloud services consumer, however, is a basic distinction between two types of audit reports.

The first type of report merely describes the physical features or operations in the data center, or attests that operations in the data center conform to whatever description of them is provided by the management of the data center. This type of report does not apply any independent standard in evaluating data center operations. The second type of audit report applies an independent standard of performance and assesses the degree of compliance with that standard.

The most widely known type of data center audit, the Statement on Auditing Standards No. 70, Service Organizations (known as 'SAS 70')[6] report, is of the first type. Developed by the accounting profession in the United States, it applies to all service organizations, not just to data centers. SAS 70 looks at internal controls and merely seeks to verify that any internal controls established by management are in fact being followed. The Statement on Attestation Standards for Assurance Engagements No. 16 (SSAE 16) and the International Standards for Assurance Engagements No 3402 (ISAE 3402) are similar standards.

In effect, under the SAS 70 regime the data center operator defines the criteria for an audit. For this reason these audits have limited value and, notwithstanding certain vendors' claims, do not constitute any sort of certification. For these reasons, newer standards are attempting to establish a more concrete external standard of evaluation. The American Institute of Certified Public Accountants has established the Service Organization Controls Nos. 2 & 3 (SOC 2, SOC 3) reporting standards. Reports under these standards are intended to provide assurance about controls related to

- security,
- availability,
- processing integrity,
- confidentiality, and
- privacy of a system and its information.

[6] There are technically two different SAS reports: Type I and Type II. A Type I report relies upon management descriptions of data center operations while a Type II report involves an actual on-site evaluation of data center operations.

Moreover, an audit under these standards is governed by pre-defined controls criteria designed specifically for evaluating the design and operating effectiveness of controls at a data center or other service organization.

While these standards have been developed in the United States for use by U.S. auditors, they reflect a larger trend toward more sophisticated assessment of data center operations that will likely gain momentum around the world as more and more businesses come to rely on public cloud data centers.

For the business considering entering into a relationship with a cloud-based service, it is important to know that it should be possible to obtain some form of independent assurance regarding the quality of data center operations maintained or contracted by their SaaS vendor.

5. Contract Issues

The economic and other advantages of cloud-based IT services are compelling in many instances. However, when outsourcing IT services to a cloud provider, one is entrusting data and critical business functions to a third party under a contractual agreement. It is important, therefore, to consider carefully the terms of that agreement and the ability to enforce those terms, particularly in a cross-jurisdictional context.

Most providers of cloud services offer standard terms of service, often posted on their website. Needless to say, these terms are frequently favorable to the interests of the cloud provider and should not be assumed to be acceptable as offered. Listed in Appendix A are a number of issues that should be addressed in a services agreement with a cloud provider. How specifically each of these issues is handled will vary according to the specific circumstances, but at minimum these issues should be considered and assessed before entering into any agreement. Many of these issues are common to any IT outsourcing contract and are not necessarily specific to contracts for cloud services.

In the international economic development context, the potential enforceability of agreements is a special consideration. In most instances the provider of a SaaS application is located in a different country from the customer and the data are hosted with the provider. The provider has no formal presence in the legal jurisdiction where the customer resides, but merely provides access to services over the Internet. Therefore, the provider is not subject to legal recourse in the country in which the customer is located, forcing the customer to pursue legal claims in a foreign court. This may not be too much of a burden for large entities like MNOs or banks which contract with m-money platform providers, but is an insuperable barrier to almost any individual MFI.

Conclusion

Because of the compelling cost and convenience advantages that they offer, cloud-based services will continue to be ever more pervasive. Business-to-business services will proliferate across all industries and all markets. The services that today are helping to propel microfinance and m-money are advancing and proliferating

rapidly with more and wider choice of providers all the time. Particularly because this is a relatively new and rapidly changing marketplace it is especially important for all key stakeholders, including donors, regulators and program managers, to understand what cloud computing really involves in order to become effective advisors, investors, administrators and consumers of these services. Like any outsourcing arrangement, reliance on cloud-based services entails certain risks. But if these are identified and anticipated, and properly weighed against the advantages of cloud-based services, the risks can be successfully mitigated. Where m-money is concerned, risks will be limited when services are mediated through otherwise regulated institutions, but regulators will ultimately need to take notice of the structure of cloud-based arrangements and adapt regulatory regimes accordingly.

Appendix A. Cloud Services Contract Checklist

When preparing to negotiate any contract with a cloud service provider there are a number of issues that ought to be addressed. It is not essential that all issues are explicitly addressed in the contract, not all potential issues are relevant to all situations. This checklist is suggestive only. It is not designed to be exhaustive, nor is it a substitute for the advice of a qualified expert with knowledge of specific business circumstances. It is intended as a guideline, a starting point from which to consider the terms of any particular agreement under consideration.

Issues/Questions	Notes/Comments
Data Management	
Where will data be physically housed? Will there be any restrictions on physical location?	Some countries will not allow certain types of data to be stored outside the country. Nonetheless, the ability to replicate data to multiple locations is important to data security.
How is each customer's data isolated from other customers' data?	Well constructed applications (called 'multi-tenant applications) can reliably store multiple users' data in the same database.
How long will the provider be required to retain customer data? What will become of customer data if and when the contract is terminated?	If a contract is terminated and all data returned to the customer, the provider should be required to purge all customer data from its files. Regulations may require that certain data be retained for a period of time.
Data Security	
What security arrangements are in place at data centers where customer data is kept?	The service provider should be explicitly made responsible for physical security and integrity of the data center and any loss resulting from unauthorized access or physical damage.
Is data encrypted?	Encrypted data is more secure but may cause unacceptable degradation in application performance so should be required only where genuinely necessary.
What policies and procedures govern backup of data? Is there any provision for escrow of data or software with a third party?	Data should be backed up or replicated to a different geographic location on a regular (at least daily) basis. If there is any concern about the future viability of the SaaS provider, providing for the escrow of critical software code with third parties can provide a limited guarantee of the ability to continue operations if the SaaS provider terminates its business.
Who, other than customers' employees, will have access to the data center or to customer data? Under what conditions will this be allowed?	Access to customer applications and data by data center personnel may be necessary for maintenance purposes.
Will the service provider be allowed, for example, to perform research on customer data?	Specialized providers of cloud services will have an interest in mining customer data that is under their physical control to discover important trends in usage of their services. This will cause no harm to service users so long as the service provider is prohibited from sharing any confidential information.
What customer information is the service provider allowed to share?	Service agreements should have strong non-disclosure provisions prohibiting the service provider from sharing privileged or proprietary customer information with any third party without the customer's consent.
Service Levels	
How often will services be unavailable? How quickly will hosted applications respond to user input?	The Internet itself is often the cause of slow responses to user input (known as 'latency') and this is beyond the control of the service provider.
What types of technical issues will the service provider be responsible for and how quickly will they respond to customer issues?	If the service provider offers use of software created by others they may have little influence over the resolution of customer issues traceable to software defects.

Will the service provider be required to provide evidence that their data center facilities have been reviewed by an independent auditor? How often will the provider be required to produce evidence of a successful audit?	This is still not a widespread practice but worth considering if the provider is not hosting with a known reputable hosting provider.
What regulatory regimes is the customer subject to? How will the service provider comply with these requirements?	Regulatory requirements can include specification of the geographic location of data, audit requirements and certifications.
How will the service provider respond to demands from local law enforcement agencies or courts (i.e. to search warrants or discovery requests)?	With cloud services, data most often resides in a legal jurisdiction different from the point of service (where end users are located). It is vital that there be a clear understanding between a cloud provider and their customers how various demands by government officials for access to data will be handled, at minimum requiring the provider to notify the customer when any such demand is received.
How will usage of cloud services be charged? Will the customer have a right to an audit of usage charges if those are not otherwise transparent?	Cloud services are sometimes offered on an annual subscription basis, but increasingly are offered on a "metered" basis, according to which charges are tied to actual usage. A variety of measures of usage is possible and should be specified clearly.
Warranties, Penalties and Damages	
What liability does the service provider have for damage to an end users' business as a result of failure to meet service levels specified in the contract?	The service agreement should clearly state what the service provider is responsible for and what liability they will incur for failing to meet their responsibilities. Because SaaS providers rely on the Internet, their services are subject to degradation or interruption caused by the Internet itself and cannot assume any responsibility for these. It is also important to note that service providers relying on third parties for software development or IT hosting may have limited ability to address problems caused by those partners.
How will damages be calculated?	In circumstances where assessing the real costs of a service failure would be difficult or costly, the parties may agree in advance on a specific formula for calculating damages or a specific amount. This is often sufficient to avoid litigation to assess damages.
Does the contract contain a non-disclosure provision?	In the process of supporting a SaaS customer, the service provider may come into possession of proprietary information concerning the customer's business. Any agreement should provide protection against disclosure of the customer's proprietary information and provide penalties for breach of this responsibility.
What law will govern the contract? Where will disputes be heard?	Very often a cloud provider and customer are in different countries. Moreover, the cloud provider may not have any physical presence in the country where the customer is located and may therefore not be subject to the jurisdiction of courts accessible to the customer. Contracting parties can specify a choice of governing law, however local law or regulation in some countries may supersede any such contractual provision.
Termination of Relationship	
Under what circumstances may either party terminate the cloud service relationship?	Termination may be for cause or without cause. Often there is a required notice period if termination is without cause.

What arrangements exist for the return of customer data in the event that the relationship is terminated?	Customers should be entitled to the immediate return of a complete set of most current data delivered in an accessible format and the provider should be required to delete any customer data upon return of customer data.
Will the customer be able to access the SaaS application if the provider goes out of business?	In some circumstances it may be possible for the customer to obtain access to the provider's software at least long enough to continue business operations while arranging for a replacement. Some agreements require the escrow of copies of software code with a trustworthy third party to as insurance against loss of access in the event of a business failure.

APPENDIX B. TECHNOLOGY ALTERNATIVES FOR MOBILE FINANCIAL SERVICES

The table below briefly summarizes the five available mechanisms for communicating transaction data from a mobile phone to a mobile money platform. Of these, only the first three actually involve an MNO. The last two bypass the mobile network to connect directly to the Internet. These latter technologies are only available on so-called 'smart phones' which are not yet widespread in developing countries and, due to higher cost, are largely absent among the poor. Most existing mobile financial services programs therefore rely on the mobile network using one of the first three mechanisms listed below. This situation may change, however, as cheaper smart phones begin to appear in parts of the developing world, expanding the penetration of these more capable devices.

SIM – based application	A small application is loaded into the SIM card on the phone which presents a menu to the user. GSM only. Uses encrypted SMS.
Structured SMS	Requires no special application on the phone. Not secure. Uses direct connection from mobile money platform to SMS gateway.
USSD (Unstructured Supplementary Service Data)	No special application required. GSM only. Uses continuous open channel for two-way communication so is more responsive than SMS. Requires no application on SIM.
Java J2ME	Java applications can generate either internet data communications or SMS messages, but Java not supported on all phones (especially low-cost feature phones).
Mobile Internet (Wi-Fi)	Independent of MNO. Communication direct with mobile money platform, via 802.11 wireless networks.

APPENDIX C. CLOUD COMPUTING USER RESOURCES

There are a wide variety of industry organizations devoted to promoting cloud computing and furthering best practices in the development and adoption of cloud computing services and infrastructures. While several of these organizations have developed putative standards for cloud computing, in practice none of these is yet widely adopted. Cloud computing is yet too young and fast moving for established standards to have taken hold on a widespread basis.[7]

The list of organizations below is by no means exclusive but is representative of resources available to developers and consumers of cloud services.

- Cloud Customer Standards Council (www.cloud-council.org): Notwithstanding the name, CCSC is not a standards setting organization but a consortium of stakeholders interested in furthering adoption of cloud computing and best practices.
- Distributed Management Taskforce (dmtf.org): Though devoted to distributed computing generally, DMTF has a working group that promotes a management interface standard for managing cloud infrastructures (with emphasis on interoperability). The Open Virtualization Format is a standard allowing portability of virtualized appliances, giving customers of cloud services a means to move applications among different cloud vendors.
- Open Grid Forum (www.ogf.org): Another organization devoted to promoting cloud computing with its own (competing) standard for cloud interoperability.
- OMG.org (Originally the Object Management Group): OMG is a large established industry organization devoted to establishing standards across a wide variety of enterprise integration technologies. They sponsor industry conferences and workshops around cloud computing interoperability but currently do not appear to have any formal standard under development.
- Cloud Security Alliance (https://cloudsecurityalliance.org): The Alliance promotes "the use of best practices for providing security assurance within Cloud Computing" through conferences and educational activities.
- Cloud Computing Interoperability Forum (www.cloudforum.org): The Forum is devoted to communication and collaboration among vendors and customers, but does not maintain any standards setting activity.
- InterNational Committee for Information Technology Standards (https://standards.incits.org/kwspub/home/): Committee 38 on Distributed Applications and Services (DAPS38) is responsible for standards in the areas of web services, service oriented architectures and cloud computing. INCITS is the primary U.S. focus of standardization in the field of Information and Communications Technologies (ICT), and serves as the Technical Advisory Group to the American

[7] See: Fogarty, Kevin. "Cloud Computing Standards: Too Many, Doing Too Little," CIO Magazine, April 6, 2011. Retrieved from http://www.cio.com/article/679067/Cloud_Computing_Standards_Too_Many_Doing_Too_Little on April 6, 2011. Fogarty offers a review and critique of the multiple standards efforts underway currently.

National Standards Institute (ANSI) for the International Organization for Standardization/ International Electrotechnical Commission (ISO/IEC) Joint Technical Committee 1. The JTC 1 is responsible for international standardization in the field of Information Technology.

- Storage Industry Networking Association (www.SINA.org): This Association maintains the Cloud Storage Initiative and promotes the Cloud Data Management Interface Standard.
- Institute of Electrical and Electronics Engineers Intercloud Working Group (http://standards.ieee.org/develop/wg/ICWG-2302_WG.html): The ICWG has its own proposed standard for Intercloud Interoperability and Federation.