On Clouds, Cloud Security and Dependability

Drew Switzer and Raghunath Rajachandrasekar
The Ohio State University
switzer, rajachan@cse.ohio-state.edu

***Abstract*—Cloud computing has been a boon to companies that have heavy processing, storage and infrastructural requirements. Without having to invest in the order of millions of dollars on data-center equipment, companies can utilize the cloud services offered by providers such as Amazon, Cisco, Google, etc., for a small price to satisfy their computing requirements. Inspite of the long list of advantages offered by this paradigm, there are several concerns that prevent decision makers from adopting this model, the majority of which are concerns regarding security. Consumers of cloud computing services will have to trust the cloud providers with their information, their clients' data and other sensitive business logic before they can utilize the cloud platform. But this also allows for potential security breaches by means of authentication compromise, data leaks, etc. As a part of this survey, we aim to: (a) profile a subset of sample providers of different types of cloud services, the services that they offer and their security features; (b) give an insight to the vulnerabilities in the current state-of-the-field of these providers; and (c) provide viable solutions to the vulnerabilities that were identified.**

***Keywords*-cloud computing; cloud security; dependability; fault-tolerance; availability;**

I. Introduction

Providing software, computation and storage services over a network is an old idea. The concept of cloud computing dates back to the 1960's, but only recently have we had the technologies and capabilities to implement such an idea. Today cloud computing is in its infancy with the first cloud computing companies entering the cloud computing market in the mid 2000's. Cloud computing offers endless rewards such as reduced computing costs, better computing performance, and scalability; however, because of the newness of cloud computing, some issues such as privacy and security have not been properly addressed. Many of today's cloud computing service providers, such as Amazon S3 or Azure, have done their best to address some of these issues with cloud computing, but security, privacy, and availability problems still remain for these cloud computing companies.

Cloud computing, as a service, can be broadly classified into three different categories: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Each of these three models cater different classes of customer requirements. In this paper, we would like to give a case study of the security and dependability provisions of these different models by considering specific examples.

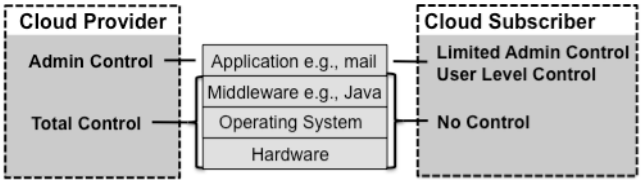

Figure 1. SaaS Architecture

A. Software as a Service

Saas is a low cost cloud computing service that comes with a multitude of benefits. Exploiting the benefits offered by virtualization, SaaS gives IT organizations an alternative to buying, building and maintaining their infrastructures. Not only does SaaS have simpler implementation, but the costs are exponentially smaller due to a pay-per-use model. Painful software upgrades are a thing of the past because the service provider manages all updates.

This particular model is focused on managing access to applications. For example, policy controls may dictate that a sales person can only download particular information from sales CRM applications. For example, they are only permitted to download certain leads, within certain geographies or during local office working hours. In effect, the security policy needs to focus on establishing controls regarding users' access to applications.

B. Platform as a Service

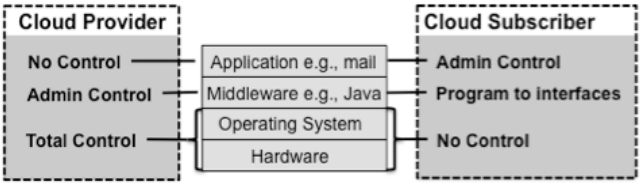

Figure 2. PaaS Architecture

PaaS is the latest in on-demand development services. Implementing a PaaS enables IT organizations to direct their attention towards development and innovation instead of infrastructure maintenance. Benefits of a PaaS solution are lower cost, lower risk, easy collaboration, minimized operational costs, and simple integration with other Web services.

The primary focus of this model is on protecting data. This is especially important in the case of storage as a service. An important element to consider within PaaS is the ability to plan against the possibility of an outage from a Cloud provider. The security operation needs to consider providing for the ability to load balance across providers to ensure fail over of services in the event of an outage. Another key consideration should be the ability to encrypt the data whilst stored on a third-party platform and to be aware of the regulatory issues that may apply to data availability in different geographies.

C. Infrastructure as a Service

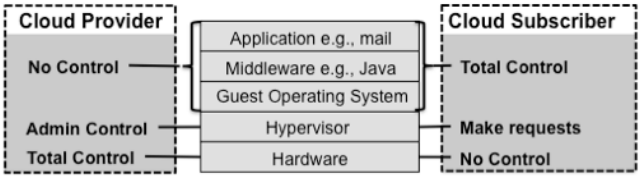

Figure 3. IaaS Architecture

IaaS is a fully outsourced infrastructure service that can provide many cost saving benefits if done the right way. IaaS allows your IT infrastructure to be fully scalable and works with high and low peak times. As with SaaS and PaaS, IaaS is also a cost reducer by means of lowering energy expenditures, driving efficiency with IT resources, and allowing for automation of everyday tasks.

Within this model the focus is on managing virtual machines. The security policy's priority is to overlay a governance framework to enable the organization to put controls in place regarding how virtual machines are created and spun down thus avoiding uncontrolled access and potential costly wastage.

II. Case Studies

A. Amazon

Amazon Simple Storage Service is a simple, scalable, and reliable Cloud Infrastructure as a Service system. Amazon's main goal with the Amazon Simple Storage Service was to provide an Infrastructure Service that is fast, reliable, and cheap, but in doing this, they may have created some data security flaws. S3 provides a very reliable web service that allows it users to access data anytime and anywhere on the Internet. Amazon S3 allows users to write, read, and delete files as large as 5 terabytes, and each user is can have an unlimited amount of storage space. Each file is stored in a bucket, which is assigned a unique key. These unique keys are chosen by the users of the Amazon S3 system. Users of the Amazon S3 system are also give the option to choose where their data is stored. There are several regions that data can be stored in, and regions can be selected to optimize for latency, minimize costs, and address regulatory requirements. Amazon S3 users are also given the option of auto scalability, so that if if an Amazon S3 user runs out of storage space, it will automatically increase their storage capacity. Amazon S3 also provides authentication mechanisms to make sure only authorized users can access data from a bucket. Users are given the ability to make files public or private, and they can choose specific user access of a file. Amazon S3 does not provide any encryption of data, but Amazon does provide a client that can be used to encrypt and decrypt data.

User data on Amazon S3 can be secure, but it all depends on whether or not the user wants to go the extra mile to make their data secure. Each bucket used for storing files is assigned a unique key that only the owner knows, but the owner needs to assign a unique key that is difficult to guess. If the owner assigns a key that is easy to guess, then other users may able to gain access to the owner's bucket and their files in that bucket. Assigning a unique key that is difficult to guess will help ensure that a users data is secure and hard for malicious users to access. Amazon S3 also allows users to create Access Control Lists and Identity and Access Management policies. Identity and Access Management policies controls user access for organizations and groups of multiple users, and it also can control user access policies for specific files within a bucket. Access Control Lists allow owners to specify what users can access there data. Owners of Amazon S3 storage buckets should use the Access Control Lists and Identity and Access Management policies to their advantage because it will help keep their data more secure if they do not hand out their unique bucket key to users who need access to the data inside their buckets.

Owners can also take advantage of query string authentication. This form of identification allows owners to share data through URLs. Each query string authentication is valid for a predefined expiration time. One of the largest security problems of the Amazon S3 system is that it does not encrypt data before it is stored in the cloud. Users of Amazon S3 are responsible for encrypting and decrypting data on their own, which is the problem with Amazon S3. If users do not encrypt their data, then their private data could be easily acquired if a malicious user can gain access to their bucket. However Amazon does provide a client, called the Amazon S3 Encryption Client, which can be used to to encrypt and decrypt data that is written or read from the Amazon S3 cloud. The Amazon S3 Encryption Client provides default encryption algorithms, but users are also able to create their own encryption algorithms. It is important that users choose an encryption algorithm that is difficult for malicious users to decrypt. If users choose a weak encryption method, then their data may still be insecure if a malicious user gains access to their bucket.

Attacks and Outages: Amazon S3 has had its share of availability problems. One such example occured on April 21, 2011. On this day, Amazon planned a configuration

change to upgrade the capacity of the primary network. To perform an update such as this without causing an outage to its users, Amazon usually routes all incoming traffic through another router on the primary network, but on this occasion the rerouting was performed incorrectly, and it shifted all the incoming traffic to a network that could handle the volume of the incoming network. This caused all the affected nodes using Amazon S3 to become disconnected from the Amazon network. When Amazon fixed the problem with re-routing the incoming traffic and restored the network, all the disconnected nodes began searching for available server space on Amazon's system. Because of the large number of nodes affected, server space quickly became exhausted, and it left many nodes in a loop searching for server space.

Amazon S3 experienced another large outage on July 20, 2008. On this day Amazon engineers observed that many of the datacenters were experiencing high error rates. Amazon S3's servers use a "gossip protocol" to relay server information to other servers throughout the Amazon network. This protocol allows servers to inform other servers how to route around failed or unreachable servers. On this day the engineers found that many of the servers were spending most of their time "gossiping", which caused many of the servers to fail. With a large number of the servers failing, Amazon S3 was not able to process user requests. To fix this problem, Amazon had to shut down all server communication and restore the system.

B. Google Apps

Google Apps is a Cloud Software as a Service company that offers a group of messaging and collaboration applications for Businesses, Education, and Government. Google Apps offers a reliable and powerful web applications that is available 99.9% of the time. Google Apps includes software applications such as Gmail, Google Calender, Google Docs, Google Reader, and more. In creating these powerful and reliable software application, Google may have ignored important security issues with Google Apps.

Google is dedicated to keeping its users' data stored on Google Apps safe, secure, and private. Google achieves this goal by ensuring that data centers are physically safe and secured, ensuring that only authorized users are able to gain access to Google App accounts, and by using the best software and server architecture to reduce security and privacy risks.

Google is very thorough when it comes to keeping data centers physically safe and secured. All of Google's data centers are constantly under video surveillance and protected by security guards. All employees at a Google data center undergo security inspections before they are allowed into Google's data centers each day. Google stresses that it can ensure 99.9 percent availability to its users. Google does this by storing user data in multiple locations so that if one location becomes inaccessible, then the user can still access their data from another one of Google's data centers. Google also has mechanisms in place that protect data stored in their databases from physical harm. For example if there was a fire in one of Google's data centers, their databases have fire detection mechanisms that are designed to transfer data to another data center before any data is destroyed. Google ensures that only authorized users gain access to their accounts by their two step identification process. In this process, users can only gain access to Google Apps by first providing their user name and password, and then by providing a verification code that is sent to the user's mobile phone upon request. The two step identification process is an optional feature. If users do not use this feature, they could be at risk of having their accounts compromised if users choose a weak password that is easy to guess. Using the two step verification process is highly recommended, but it also has some problems. The two step verification process gives users the option to remember their verification code for up to days,so users that allow this feature will have the same potential problems as users who do not use the two step verification process. The two step process also gives users an option to use a back up verification code to log in to their account so that users are still able to gain access to their account if they lose their mobile phone. If a user's is back up verification code becomes compromised, then their account once again has the same potential problems as a user who does not use the two step verification process.

Google does its best to reduce security risks by hand-picking the software that it uses for its databases and servers. For example, Google's machines that run the Google Apps server use the Linux operating system, but only the necessary functions of the Linux operating system are installed to help reduce possible security risks. Google keeps Google Apps data private by encrypting data that is stored in their data centers so that if someone gains direct access to their databases, none of the information they see is in human readable form.

Google is thorough when it comes to deleting user information from its data center. If a user chooses to remove their account and information from Google Apps, then Google will completely overwrite the hard drive containing their information. After completely overwriting the user's old hard drive, they will check the hard drive to make sure all the old information has been completely overwritten before putting the hard drive back into their system. When hard drives reach the end of their life cycle, Google destroys its hard drives so that no one can access information from their old hard drives.

Google has multiple layers of defense to help protect their network from external attacks. Google hand-picks services and protocols for their network that meet their security requirements. Google uses industry standard firewall and ACL technologies to protect their network, and only authorized personnel have access to all their network devices. To help

detect and stop malicious attacks over their network, Google routes all incoming traffic through a custom front-end server that is designed to detect malicious attacks. Google also monitors network logs to check for exploitation of possible programming errors.

Attacks and Outages: Google Apps usernames and passwords are in constant danger of being attacked by malicious users. One example of such an attack was the "spear phishing" attack. This year, some users of Google Apps have had their accounts attacked by the "spear phishing" attack. In this attack, users are presented with a mock Google Apps web page. When they enter their username and password on this mock web page, the malicious users set up the users account to forward its email to another address to allow them to monitor the attacked user's email. The Chinese government was blamed for this "spear phishing" attack, and soon after this attack on Google Apps, the United States made legislation to make Cyber attacks an act of war. Google added the two-step authentication mechanism to help battle this form of attack.

Availability can always be a problem because unexpected events, such as natural disasters or hard drive destruction, may cause outages of service. Google Apps has had problems in the past with availability. On February 27, 2011, some Google Apps users experienced an outage after a software update eliminated all online copies of some users' data. The users who were affected by this lost access to all of their saved emails. For many users key parts of the Gmail service went missing, including email, chat, contacts, etc. After this error from the software update, Google did not allow affected users to sign into their accounts until they were able to repair all of the affected accounts.

C. Windows Azure

Windows Azure is a cloud operating system that aids in application development, hosting and service management. In essence, Azure provides platform as a service (PaaS). Azure aims at providing three fundamental dimensions: confidentiality, integrity and availability. It abstracts away much of the underlying architecture, such as database server, operating system, scheduler, web service, etc. so that developers can focus on building applications. Azure provides compute and storage functionality using a subscription mechanism. Each subscription is linked to a billing account, the access for which is controlled using the Windows LiveID credentials. A subscription can include zero or more hosted services and storage accounts, each of which contains one or more deployments. Customers can manage their application development either using the Windows Azure Portal or their Service Management API. The authentication to this Service Management API is based on a client-generated public/private key combination. The clients also have to register a self-signed certificate at the Windows Azure portal. This certificate is used for subsequent accesses through the API.

Azure provides confidentiality by means of access management, minimizing interaction with data by isolating them logically in containers, and optional encryption for rigorous data protection capabilities. Azure also guarantees integrity by providing client-controlled ACLs and configuration files which dictate components/data accesses. Integrity of the storage accesses is controlled by applications using storage keys to associated data. Accounts management encompasses of several components such as monitoring, logging and reporting to provide safe accesses to customers and to prevent misuse of trusted accounts.

In summary, the following policies and measures are in place to provide security across the stack:

1) Physical security of the data centers (locks, cameras, biometric devices, card readers, alarms)
2) Firewalls, application gateways and IDS to protect the network
3) Access Control Lists (ACLs) applied to virtual local area networks (VLANs) and applications
4) Authentication and authorization of persons or processes that request access to data
5) Hardening of the servers and operating system instances
6) Redundant internal and external DNS infrastructure with restricted write access
7) Securing of virtual machine objects
8) Securing of static and dynamic storage containers

D. SalesForce

The force.com cloud service focuses on three main aspects of security. The first aspect is users and security, looking at how users are authenticated, network-based security that determines the IP ranges from which a user may access the network, sessions and auditing. The second aspect is programmatic security. Any software client that needs to log in to the platform does so through a Web services SOAP interface. The third aspect is the Force.com platform security framework, which one can use to offer different access permissions to authenticated users within an organization. This security framework lets clients grant security permissions to users or profiles, determine access control over a wide range of components, and configure data sharing, which limits access to individual records.

Attacks and Outages: Nevertheless, SalesForce has had its share of outages and security breaches too. In November 2007, SalesForce customers were hit by a phising attack. A SalesForce employee had been tricked into disclosing a password and allowed a customer contact list to be copied. Information in the contact list included first and last names, company names, email addresses, telephone numbers of Salesforce.com customers, and related administrative data belonging to Salesforce.com. A small number

of Salesforce.com customers' end users also revealed their passwords as a result of the attack. A few days later, a new wave of phishing attempts that included attached malware that secretly installs viruses or key loggers appeared and seemed to be targeted at a broader group of customers.

III. Security Issues and Threats

In addition to the above mentioned service-specific threats, all the cloud service providers are vulnerable to the following common threats.

Data Protection
Cloud computing poses several data protection risks for cloud customers and providers. In some cases, it may be difficult for the cloud customer to effectively check the data handling practices of the cloud provider and thus to be sure that the data is handled in a lawful way. This problem is exacerbated in cases of multiple transfers of data, e.g., between federated clouds. On the other hand, some cloud providers do provide information on their data handling practices. Some also offer certification summaries on their data processing and data security activities and the data controls they have in place.

Direct breach, Side channel attack
Malicious users may pierce hardware, software, or network isolation boundaries to compromise the confidentiality or integrity of another user's data, code, or communications. These malicious users can also exploit potential shared-cache accesses to gain control of sensitive data from applications that use these shared caches.

Denial of Service
Such malicious users might also compromise availability by consuming too many resources or exploiting vulnerabilities exposed through tenant-accessible APIs. For example, one such vulnerability in the Xen hypervisor has been studied in the literature.

Denial of Shared Resources
A denial-of-service attack on one tenant, or on the cloud provider itself, may impact other tenants who rely on shared resources targeted in the attack.

Isolation Failure
Multi-tenancy and shared resources are defining characteristics of cloud computing. This risk category covers the failure of mechanisms separating storage, memory, routing and even reputation between different tenants (e.g., so-called guest-hopping attacks). However it should be considered that attacks on resource isolation mechanisms (e.g.,. against hypervisors) are still less numerous and much more difficult for an attacker to put in practice compared to attacks on traditional OSs.

Lock-In
There is currently little on offer in the way of tools, procedures or standard data formats or services interfaces that could guarantee data, application and service portability. This can make it difficult for the customer to migrate from one provider to another or migrate data and services back to an in-house IT environment. This introduces a dependency on a particular Cloud Provider for service provision, especially if data portability, as the most fundamental aspect, is not enabled.

Compliance Risks
Investment in achieving certification (e.g., industry standard or regulatory requirements) may be put at risk by migration to the cloud: if the Cloud Provider cannot provide evidence of their own compliance with the relevant requirements if the Cloud Provider does not permit audit by the cloud customer. In certain cases, it also means that using a public cloud infrastructure implies that certain kinds of compliance cannot be achieved.

Management Interface Compromise
Customer management interfaces of a public cloud provider are accessible through the Internet and mediate access to larger sets of resources (than traditional hosting providers) and therefore pose an increased risk, especially when combined with remote access and web browser vulnerabilities.

Insecure Data Deletion
When a request to delete a cloud resource is made, as with most operating systems, this may not result in true wiping of the data. Adequate or timely data deletion may also be impossible (or undesirable from a customer perspective), either because extra copies of data are stored but are not available, or because the disk to be destroyed also stores data from other clients. In the case of multiple tenancies and the reuse of hardware resources, this represents a higher risk to the customer than with dedicated hardware.

Malicious Insider
While usually less likely, the damage which may be caused by malicious insiders is often far greater. Cloud architectures necessitate certain roles which are extremely high-risk. Examples include Cloud Provider system administrators and managed security service providers.

Under-provisioning
When infrastructure is shared, tenants must rely on the cloud provider to properly estimate their collective peak resource needs and provision appropriately. Under-provisioning may cause performance degradation or outright failure at times

of peak demand. Cloud providers that fail to account for correlated usage bursts (economies of scale in provisioning require assuming some level of independence in resource needs) might unintentionally under-provision. So might a provider who fails to anticipate the exponential growth of a tenant; Imagine a cloud provider that discovers too late that it is hosting the Internet's next Google or Facebook. Competitive pressures might lead cloud providers to intentionally provision fewer spare resources than their tenants expect, as every unused resource cuts into profit margins. Customers and monitoring firms are already measuring performance degradation in Amazon's EC2 service and interpreting it as evidence of under-provisioning.

IV. Countermeasures

Hypervisor protections on the confidentiality and integrity of tenant software and data have been the subject of a great deal of research, both offensive and defensive. Current directions include making hypervisors smaller and easier to verify using formal tools. Microprocessor manufacturers are adding support to simplify and improve the design of hypervisors.

When tenants do communicate, they are more likely to be using existing network protocols than to communicate via the hypervisor. Similarly, shared database or storage services are likely to be accessed via network protocols as well. However, shared networks, storage, and other resources will also be potential sources of isolation failure. IPsec and virtual private networks are already heavily-relied upon in order to protect data in transit on potentially-unfriendly networks. Since these tools are insufficient to prevent availability attacks on the infrastructure itself, control networks may also be necessary. Building authentication, encryption, and integrity protection into intra-tenant and intra-cloud communications could actually simplify the work of application developers, who today must implement myriad authentication solutions to support different services. Finally, for those systems that do not use hypervisors or want to add defense-in-depth, application-level sand-boxing and API-limitations may further aid isolation, at the cost of interfering with existing tools that may assume access to resources now restricted.

The threat of other tenant's hogging or stealing resources will also challenge technologists to develop appropriate fairness algorithms and accounting mechanisms. This poses many challenges. For example, if tenants control the TCP stack they will also control the network's fairness mechanisms and the code that determines whether a transmission was in fact completed. Accounting imprecision due to coarse measurement units or outliers such as process swap time may be the target of arbitrage by malicious tenants.

Collateral damage from denial-of-service attacks may be limited through the use of resource quotas. For example, a tenant might have as a default quota that is equal to its average resource consumption plus the resources available when the load on the cloud is at its 99.9th percentile.

Alternatively, when resources are scarce, the cloud provider could calculate ratio of current resource consumption to average resource consumption for each application or tenant. Scarce resources could be granted to those with the lowest utilization ratios. Resources could even be moved away from those with the highest ratios.

Attestation-based audit mechanisms could also be used to verify that spare capacity has been provisioned to the levels promised in a service level agreement contract. Resource accounting modules on each system could report total usage over time without revealing anything about the tenants who used those resources. Again, attestation could be used to safeguard against cloud providers' attempts to tamper with resource usage reports. Availability reporting poses a greater challenge than usage reporting. Attestation-based mechanisms to prove that resources were available if needed, at times when they were powered down, are an open research problem. Attestation-based capacity auditing may not be necessary if spare capacity levels are not made verifiable, but instead guaranteed via penalties paid by the cloud provider to the tenant in the event that resources are unavailable. However, large outages could result in penalties so large as to require the cloud provider to obtain insurance. Insurers would then likely find it necessary to audit spare capacity.

V. Hybrid Cloud Framework

Based on the above discussion on security threats and dependability issues, it is clear that a new cloud service model needs to be developed. In this context we propose a Hybrid cloud computing framework which provides a cost-effective cloud management solution without compromising on the security or dependability of tenant data and services.

The motivation for this framework comes from the predominant security concern - trusting the Cloud service provider. When moving from self-owned infrastructure to cloud-hosted ones, the client has to rely on the provider to secure the actual facility physically, and also secure the hardware, software and the workforce. But with the proposed approach, the client still possesses control and governance over the fashion in which data is presented and exposed to the cloud. The hybrid approach also addresses concerns raised in the case of shared-tenancy in clouds.

We envision the Hybrid cloud framework to be applicable to all the three service models - Saas, Paas and Iaas. It will be a combination of private infrastructure that can house sensitive data and control, and a public cloud service which can provide computational capabilities at a price.

In this model, the client has complete control over his data. Consider the case of a SaaS cloud. In a naive cloud-only approach, the data required by the application, such as user records, key-value pairs, etc., will be stored on the cloud

itself. This makes the data vulnerable to direct breaches, side-channel attacks, leaks to malicious users and other such attacks discussed in Section III. With the proposed framework, clients can maintain shadow-copies of the data that has to go into the cloud. Data encryption can be carried out locally in the private infrastructure without involving the cloud APIs or services. This dilutes the tight-coupling between the client data and the service provider. This framework also allows for the possibility of obfuscating the data, key-value pairs for instance, before sending uploading the data to the cloud. The clients can maintain a mapping table, which maps the right key to the obfuscated key, on the local system. Whenever the software application running on the cloud needs a particular key, it can send a light-weight reverse-mapping request to the local cloud. The channel established between the private infrastructure and the public cloud can be secured using well-known communication security protocols. This hybrid approach can help clients realize a right balance between the amount of funds that needs to be available for in-house infrastructure and the amount that can be invested in cloud offerings.

However, in spite of having such benefits, there are a lot of issues that need to be looked into before this model can be adopted. One if it being performance degradation. Having round-trips between the private cloud and the public one will increase the latency of the application functions. Given the fact that latency-sensitive applications such as stock-market predictions and weather forecasting simulations choose the High-Performance Computing (HPC) clusters as their platform, this might not yet be a threat to cloud computing applications. Another issue that needs to be worked out is that of consistency. The shadow copies of data that are being maintained by clients on their local systems need to be periodically synced with that on the cloud. This includes taking in the dirty data that has been modified by the application running on the cloud. Furthermore, one other issue that has to be considered is scalability. Shadowing the entire data-set might not certainly be beneficial to the client. This redundancy might prove to be more expensive that just having the entire data on the cloud and trusting the service provider with the security. However, this can be overcome by implementing policy managers that selectively shadows data based either on priority or on Quality-of-Service (QoS) levels determined by the customers of the client's software. The client need not explicitly protect the data using the hybrid framework unless the customer explicitly asks for it. This will also lessen the space constraints on the client's in-house infrastructure.

VI. Conclusions

As a part of this survey, we have studied the security infrastructure and policies adopted by four unique cloud service providers - Google Apps, Amazon S3, Windows Azure and SalesForce.com. We have also identified the potential threats to security and dependability in these services, and have also suggested possible countermeasures. In particular, we have proposed a Hybrid Cloud framework which addresses most of the threats in a holistic manner.

Self-hosting gives complete control over computing and storage infrastructure in terms of security, availability, personnel, etc. But moving to a cloud-hosted infrastructure for software, development platforms or hardware might not lead to an entire loss of security either. The issue lies in identifying the right cost / security trade-off threshold for each infrastructure. This depends on a lot of factors like the business model adopted by a company, the applications that need to run on the hosted infrastructure and the pricing model which they offer their customers. Once such a threshold is identified, a framework that is similar to the one proposed in Section V can be adopted.

References

[1] Cloud Hypermarket, "Google Email Attacks," http://www.cloudhypermarket.com/blog/google-cloud-email-security-attacks.

[2] Amazon Web Services, "Overview of Security Processes," http://s3.amazonaws.com/aws-blog/.

[3] Reese,George, "Key Security Issues for the Amazon Cloud," http://broadcast.oreilly.com/2008/11/key-security-issues-for-the-am.html.

[4] Google, "Software-as-a-service Has Built-in Security Advantages: Google Apps," http://www.google.com/apps/intl/en/business/infrastructure-security.html.

[5] Amazon Web Services, "Summary of the Amazon EC2 and Amazon RDS Service Disruption," http://aws.amazon.com/message/65648/.

[6] Amazon Service Health Dashboard, "Amazon S3 Availability Event: July 20, 2008," http://status.aws.amazon.com/s3-20080720.html.

[7] Jon, R. Greenwald, "An Overview of Force.com Security," SalesForce.com Whitepaper.

[8] Deb Shinder, "Microsoft Azure: Security in the Cloud," http://www.windowsecurity.com/articles/Microsoft-Azure-Security-Cloud.html.