Research Paper

Computer Science

Cloud Computing: Security Concerns, Risk Issues & Legal Aspects

Dr. A.M.Khan JJT University, Jhunjhunu Rajasthan (India)

Kazi Hazim Ali JJT University, Jhunjhunu Rajasthan (India)

ABSTRACT

The present paper deal with cloud computing and security concerns, Assessing Risk Tolerance in Cloud Computing and Legal and Regulatory Issues also discussed, we covered many of the qualities and promises of cloud computing. In addition, we examined the three models for cloud services (SPI) and the four models for cloud deployment (public, private, community and hybrid).While developing a background in cloud computing, we also discussed many security aspects of clouds.

Since the provider's investment in achieving better security costs less per consumer. For the same reasons, a private cloud can obtain significant advantages for security. But there are wrinkles. We won't get the benefit without investment and not every model is appropriate for all consumers. But, regardless of which services delivery model or deployment model we select, we will transfer some degree of control to the cloud provider—which would be completely reasonable if control is managed in a manner and at a cost that meets our needs.

KEYWORDS: Cloud Computing, Security Concern, Risk Issues, Legal issues.

1. SECURITY CONCERNS IN CLOUD COMPUTING

In this section we first introduce some major security concern-

- **Network Availability** The value of cloud computing can only be realized when our network connectivity and bandwidth meet our minimum needs: The cloud must be available whenever we need it. If it is not, then the consequences are no different than a denial-of-service situation.
- **Cloud Provider** Viability Since cloud providers are relatively new to the business, there are questions about provider viability and commitment. This concern deepens when a provider requires tenants to use proprietary interfaces, thus leading to tenant lock-in.
- **Disaster Recovery and Business Continuity** Tenants and users require confidence that their operations and services will continue if the cloud provider's production environment is subject to a disaster.
- **Security Incidents** Tenants and users need to be appropriately informed by the provider when an incident occurs. Tenants or users may require provider support to respond to audit or assessment findings. Also, a provider may not offer sufficient support to tenants or users for resolving investigations.
- **Transparency** When a cloud provider does not expose details of their internal policy or technology implementation, tenants or users must trust the cloud provider's security claims. Even so, tenants and users require some transparency by providers as to provider cloud security, privacy, and how incidents are managed.
- **Loss of Physical Control** Since tenants and users lose physical control over their data and applications, these results in a range of concerns:

(a) Privacy and Data With public or community clouds, data may not remain in the same system, raising multiple legal concerns.

(b) Control over Data User or organization data may be comingled in various ways with data belonging to others.

(c) A tenant administrator has limited control scope and accountability within a Public infrastructure-as-a-service (IaaS) implementation, and even less with a platform-as-a-service (Paas) one. Tenants need confidence that the provider will offer appropriate control, while recognizing that tenants will simply need to adapt their expectations for how much control is reasonable within these models.

(d) New Risks, New Vulnerabilities There is some concern that cloud computing brings new classes of risks and vulnerabilities. Although we can postulate various hypothetical new risks, actual exploits will largely be a function of a provider's implementation. Although all software, hardware, and networking equipment are subject to unearthing of new vulnerabilities, by applying layered security and well-conceived operational processes, a cloud may be protected from common types of attack even if some of its components are inherently vulnerable.

A Closer Examination: Virtualization

Before we consider some of the security concerns around the use of virtualization in cloud computing, we need to understand how virtualization is implemented.

Starting at the level of our objective, a virtual machine (VM) is typically a standard operating system (OS) instance captured in a fully configured and operationally ready system image. This image essentially amounts to a snapshot of a running system including space in the image for virtualized disk storage. Supporting the operation of this VM, we need some form of enabling function, typically called a hypervisor that represents itself to the VM as the underlying hardware.

Vendor implementations of virtualization will vary, but in general terms, there are several types of virtualization:

- **Type 1** also native or bare metal virtualization is implemented by a hypervisor that runs directly on bare hardware. Guest OSs run on top of the hypervisor. Examples include Microsoft Hyper-V, Oracle VM, LynxSecure, VMware ESX, and IBM z/VM.
- **Type 2** or hosted virtualization has a hypervisor running as an application within a host OS. VMs also run above the hypervisor. Examples include Oracle VirtualBox, Parallels, Virtual PC, VMware Fusion, VMware Server, Xen, and XenServer.
- **OS implemented virtualization** is implemented by the OS itself taking the place of the hypervisor. Examples of this include Solaris Containers, BSDjails, OpenVZ, Linux-VServer, and Parallels Virtuozzo Containers.

There are many interesting security concerns around the use of virtualization even before we consider using it for clouds.

First, by adding each new VM, we are adding an additional OS—which itself entails security risk. Every OS should be appropriately patched, maintained, and monitored as appropriate per its intended use.

Second, typical network-based intrusion detection does not work well with virtual servers that are collocated on the same host, consequently advanced techniques are needed to monitor traffic between VMs. When data and applications are moved between multiple physical servers for load balancing or failover, network monitoring systems cannot yet assess and reflect these operations for what they are. This is even more the case when clustering is used in conjunction with virtualization.

Third, the use of virtualization demands the adoption of different management approaches for many functions, including configuration management to VM placement and capacity management. Likewise, resource allocation problems can quickly become performance issues; thus, performance management is critical to run an effective virtualized environment.

A Closer Examination: Cloud Operation, Security, and Networking

The cloud model brings benefits for the IT operations and support teams. Every step required to build and operate a traditional IT solution is overhead for the underlying goal. It entails expensive skills and often times inefficient repeated effort. Furthermore, traditional IT infrastructure can be dwarfed by the scale of cloud computing. Infrastructure at massive cloud scale demands automation. But even with a small cloud, automation is critical if IT processes are to be performed in a cost- and time-effective manner.

Cloud infrastructure demands efficient structure and organization. By defining and following patterns, at every step from racking individual computers to cabling them, from operations to security, savings recur and processes can be tuned and refined. An intelligently planned and organized cloud infrastructure can be more effectively and more efficiently built and operated by a smaller staff then if we take the same computers and disperse them to many server rooms.

The aggregation of components into patterns is not limited to computers, storage, and network. Power and network cabling also benefit from regular patterns, this includes their labeling or nomenclature and it is empowering to the configuration management and change management processes. These patterns have value when they are optimized to eek even small margins in the build stage of a cloud, but they have recurring benefit at every stage afterward: from provisioning VMs to managing and operating cloud infrastructure. Objectives such as lights out management, remote operations, and fail in place contribute to the further refinement of patterns.

Has security come up yet in this discussion on scale, structure, and organization?

The combination of automation and structure also means that immensely large clouds can be managed and operated by smaller staff. This, along with the technologies used in cloud computing, will drive expansion of the skill set of cloud engineers. Simply put, we gain the advantage of graduating from a series of systems administrators associated with typical infrastructure or server closets to a dedicated team of cloud administrators and a dedicated security team. Even with a private cloud implementation, the aggregated scale of a private cloud implementation accrues benefits. The benefits of intelligently conceived patterns and automation can include fault tolerance and reliability, along with greater resiliency. There is little question that a well conceived and correctly implemented cloud network can offer a tenant or other customer better networking security than many could otherwise achieve if they instead attempted to build, configure, and operate a traditional network infrastructure.

First, the implementation patterns make for a more predictable and disciplined network than the typical infrastructure network or data center network.

Second, most enterprises cannot afford the level of networking expertise that a cloud provider can deliver indirectly when they hire their staff. There is no question that the cloud customer benefits from this.

Third, maintaining the security of a network involves constant learning and intelligent response to new and emerging threats. It is simply more cost effective to benefit indirectly from the work that the cloud provider performs on behalf of countless customers beside our self.

Among the many advantages of a cloud provider delivering network security is the tendency for a provider to employ carrier grade network gear that has more sophisticated capabilities than typical enterprise networking gear. Sure we can buy the same gear, but its cost will likely exceed the cost of all our other data center costs! Such carrier grade gear requires expertise to install, configure, and operate. But the benefits are truly substantial since the security functionality will afford greater resilience to dedicated attacks, better automated traffic inspection among many other capabilities. Besides strong perimeter security, benefits include protection against a distributed denial of service along with sophisticated VLAN capabilities.

2. ASSESSING OUR RISK TOLERANCE IN CLOUD COMPUTING

A frequent question about cloud computing goes like this: Is it safe to use a public cloud? This is a fair question that is begging for information. But answering it depends on a clear understanding of our acceptance of risk. And understanding how much risk we can tolerate depends on assessing our security requirements and how we value our information assets (data, applications, and processes).

Only when we understand these issues can we make an informed decision as to which deployment models and which service delivery models are appropriate for our needs and risk tolerance.

Identifying information assets is important before we adopt a public or hybrid model because these will involve at least some degree of ceding control over how that information will be protected and where it might reside (location/jurisdiction).

Assessing the Risk

In this section we will briefly examining risk analysis.

- Threat Categorization : What can happen to our information assets?
- Threat Impact :How severe could that be?
- Threat Frequency: How often might that happen?
- Uncertainty Factor :How certain are we in answering these three questions?

The central issue with risk is uncertainty that is expressed in terms of probability. But what we really want to know is what to do about it (countermeasures or risk mitigation). So, once we analyze and address risks, we can ask several further questions:

- Mitigation what can we do to reduce the risk?
- Mitigation Cost what does risk mitigation incur?
- Mitigation Cost/Benefit Is mitigation cost effective?

To be clear, these three questions are more rhetorical for a public cloud than for a private or hybrid one. In a public cloud we get what we pay for, and the cloud provider is the party that is responsible for answering these three questions above. Similarly, these questions are also less relevant for SaaS than they are for PaaS, but they are more relevant yet for IaaS.

Information Assets and Risk We stated above that the central issue with risk is uncertainty, and applying that to our question, we must examine our information assets a bit more. Identifying information assets can be elusive, especially so with the create-once, copy-often aspect of digital systems. The typical organization rarely has sufficient control over its information in terms of assurance that if we control a given copy we can rest assured there are no other copies. From the standpoint of protecting digital data (a leaky sieve in the ocean?), that may be the worst of it. But organizations have many other problems managing their information assets.

In the real world, this is organizationally controlled along the lines of information classification and additional handling caveats (such as Project X Only). In the world of computers, the appropriate controls are usually insufficient to prevent digital duplication and intended or unintended information hemorrhaging.

Remembering the security triad (confidentiality, integrity, and availability), we can ask a series of targeted questions around information assets along the lines of what would the consequence be if:

- The information asset was exposed?
- The information asset was modified by an external entity?
- The information asset was manipulated?
- The information asset became unavailable?

If these questions raise concern about unacceptable risk, we might approach the overall problem by limiting risk-sensitive processing to a private cloud and by adopting use of a public cloud for non risk-sensitive data. But adopting a private cloud does not obviate the need for appropriate controls.

In that regard, let's consider what we might get:

- By mixing outsourcing in a public cloud for non sensitive data and reserving internal systems for sensitive data we might gain some cost advantages without assuming new risk.
- Where use of a private cloud would pose no new risks to our information assets, use of a hybrid or public cloud model may.
- Switching from a traditional IT model for internal processing to a private cloud model may reduce risk.

These are reasonable statements that constitute a start toward aligning the importance of our information assets toward both deployment models and service models.

Privacy and Confidentiality Concerns

Beyond the information asset risks we discussed above, we may be processing, storing, or transmitting data that is subject to regulatory and compliance requirements. When data falls under regulatory or compliance restrictions, our choice of cloud deployment (be it private, hybrid, or public) hinges on an understanding that the provider is fully compliant. Otherwise one will risk violating privacy, regulatory, or other legal requirements. This obligation usually falls on the tenant or user. It should go without saying that the implications for maintaining the security of information are significant when it comes to privacy, business, and national security information.

Privacy violations occur often enough outside cloud computing for us to be concerned about any system—cloud-based or traditional—storing, processing, or transmitting such sensitive information. In 2010, several cloud privacy information exposures occurred with a number of cloud-based services, including Facebook , Twitter, and Google.

3. LEGAL AND REGULATORY ISSUES

Cloud computing which employs a hybrid, community, or public cloud model "creates new dynamics in the relationship between an organization and its information, involving the presence of a third party: the cloud provider. This creates new challenges in understanding how laws apply to a wide variety of information management scenarios." The impact of this is that it creates practical challenges in understanding how laws apply to the different parties under various scenarios.

Regardless of which computing model we use, cloud or otherwise, we need to consider the legal issues, specifically those around any data we may collect, store, and process. There will likely be state, national, or international laws that we (or preferably, our lawyers) will need to consider ensuring that we are in legal compliance.

If the tenant or cloud customer operates in the United States, Canada, or the EU, then they are subject to numerous regulatory requirements. These include Control Objectives for Information and related Technology and Safe Harbor. These laws may relate to where the data is stored or transferred to, as well as how well this data is protected from a confidential aspect. Some of these laws will apply to specific markets only, such as the Health Insurance Portability and Accountability Act (HIPAA) for the health care industry. However, often companies may store health-related information about individual employees, which means that the company may have to comply with HIPPA even if they are not operating in that market.

The failure to adequately protect our data can have a number of consequences, including the potential for fines by one or more government or industry regulatory bodies. Such fines can be substantial and potentially crippling for a small- or medium-sized business. For example, the Payment Card Industry (PCI) can impose fines up to $100,000 per month for violations to their compliance.

Laws or regulations will typically specify who within an enterprise should be responsible and held accountable for the accuracy and security of the data involved. If we are collecting and holding HIPAA data, then we must have a security position designated to ensure compliance. The Sarbanes–Oxley Act designates the Chief Financial Officer (CFO) and Chief Executive Officer (CEO) to have joint responsibility for the financial data. The Gramm–Leach–Bliley Act (GLBA) is broader, specifying the responsibility for security with the entire board of directors. Less specific is the Federal Trade Commission (FTC), who just require a specific individual to be accountable for the information security program within a company.

Third Parties

If we use a cloud infrastructure that is sourced from a cloud service provider, all legal or regulatory requirements that apply to our enterprise must be imposed on this supplier as well—this is our responsibility, not the providers. Taking the HIPAA regulations as an example, any subcontractors that we employ (for example, a cloud service provider) must have a clause in the contract that they will use reasonable security controls and also comply with any data privacy provisions. In the United States, both federal and state government agencies such as the FTC and various Attorney Generals have made enterprises accountable for the actions

- Does the physical security of their data centers meet our legal, regulatory? and business needs?
- Are their business continuity and disaster recovery plans consistent with our business needs?
- What is their level of technical expertise within their operations team?
- How long have they been offering the service and do they have a track record with verifiable customers?
- Does the provider offer any indemnification?

Once our enterprise has performed such due diligence we can begin serious evaluation of providers. This will reduce the time we will spend overall in the negotiations and ensure that the correct level of security is in place for our particular needs. The cloud supplier cannot be expected to know our business requirements in detail and may well be unaware of the regulations that need to be adhered to. If there is a breach in regulations, it will be our enterprise that is penalized and not the cloud supplier we have selected.

Contract Negotiation

Once we have narrowed our selection of cloud service providers, the actual contract needs to be agreed to. Depending upon the service we are contracting for, this may not be negotiable at all, and our contract may be limited to an online click-through agreement which we can either accept or not. The results of the due diligence will obviously play a part in deciding what we need in the form of a contract. If we need to have a tailored contract, we can immediately eliminate a number of suppliers. But to be clear, the bulk of cloud services are less likely to involve tailored contracts than traditional hosting or outsourcing contracts—the economics of the model (for both provider and client) make that the case.

Where we can and want to negotiate the contract, ensure that our requirements are defined in a way that the provider can understand and agree to. Specifying that data is to be held according to HIPAA regulations.

Litigation

Litigation may affect either the cloud service provider or client, where our data needs to be accessed or given to a government agency or a lawyer. We will need to be satisfied that if we are asked to deliver specific data, our cloud provider can access and deliver the necessary data to the depth required. As the data owner, we will be held responsible if we cannot deliver it. If we, as the cloud service client, are in litigation with a third party, we must know how our cloud provider will react to requests for data, and in what timeframe. There are a number of compliance regulations related to e-discovery that will need to be met and will apply to both the provider and client.

4. SUMMARY

As systems, clouds are massively complex in terms of scale and orchestration of resources. But as we stated in section Cloud Scale, Patterns, and Operational Efficiency, massive scale, a disciplined appearance, and repeated patterns are three qualities of successful cloud implementations. The complexity of clouds is in part an illusion, as much of a cloud amounts to repeated patterns at massive scale, or in other words, multiplied simplification: The security benefits of this are significant. Likewise, security achieves additional operational advantages as all management is done using common functional units.

The resulting homogeneity contributes to simplified security testing and security assessment. It also makes for simplified auditing and monitoring, except that these functions now need to incorporate additional information sources if the monitoring in a highly dynamic cloud is to both correctly reflect the relationship between entities and if automated analysis is to be accurate and complete in its indications and warnings.

In contrast to traditional IT implementations, with cloud we have multi tenancy combined with elasticity and abstraction away from physical infrastructure. The most significant consequence is that when we use a public cloud we can no longer have a sense of comfort that we know where our data and applications are located. Although this may raise concern, the fact is that with the cloud model, and even with our use of public clouds for non sensitive data, we can actually achieve greater security and better IT management of our information resources at a lower overall cost. The cloud model also enables redundancy and disaster recovery.

REFERENCES

1. Antonopoulos A. A risk analysis of large-scaled and dynamic virtual server environments, Nemertes Research. http://www.nemertes.com/issue_papers/virtualization_risk_analysis. | 2. Brunette G, Mogull R. The Cloud Security Alliance, Security Guidance for Critical Areas of Focus in Cloud Computing V2.1, p. 35; 2009. | 3. EU Directive 95/46/EC – The Data Protection Directive; 1995. | 4. http://newslite.tv/2010/04/06/7500-shoppers-unknowingly-sold.html; 2010 [accessed 21.03.11]. | 5. http://www.infolawgroup.com/2010/10/articles/cloud-computing-1/cloud- computingcustomers-bill-of-rights/ [Abstracted]; 2010 [accessed 21.03.11]. | 6. Mann A. Five Steps on the Journey from Virtualization to Private Cloud, CA Community site.http://community.ca.com/blogs/automation/archive/2010/08/02/five-steps-on-thejourney- from-virtualization-to-private-cloud.aspx; 2010 [Adapted] [accessed 21.03.11]. | 7. Plummer D. Rights and Responsibilities for Consumers of Cloud Computing Services, Gartner Global IT Council for Cloud Services; 2010. | 8. William "Bill" Meine, in private communication; 2010.