reaction
INFORMATION SECURITY

whitepaper

Defending the Cloud

moving to a secure cloud infrastructure

Joseph Sheridan and Chris Cooper
Reaction Information Security

March 2012

Abstract

The following whitepaper aims to assess the security implications of moving resources into a cloud computing environment, and suggests recommendations for IT managers who are considering making this switch. The paper focuses on three major changes that occur in the shift from traditional networks into the cloud, namely the virtualisation layer, multi-tenancy and outsourcing. A number of recommendations are made that should be implemented as a supplement to enterprise security practices, including hardening the hypervisor, securing communication between virtual machines, guarding shared storage and memory, and ensuring the provider can cater for your security and compliance requirements. The paper identifies that a significant amount of risk is directly proportional to the provider's own security posture, highlighting the importance of transparency.

Contents

Introduction

A trend can be observed in both enterprise and personal computing, whereby more and more decision-makers are moving their computing resources and data into the cloud. The following whitepaper will discuss cloud computing as an alternative to traditional enterprise networks, and the effects it can have on an organisation's security posture. It will explore the main differences between traditional computing and cloud computing, and consider the security implications that these changes may have. The paper is targeted towards IT managers and consultants who are considering moving their organisation's resources onto a cloud infrastructure. It will make recommendations that should be employed alongside existing security practices to achieve a secure cloud environment. Additionally, many issues will be highlighted that relate to the cloud service provider, and advise on security measures that should be implemented behind the abstraction of the cloud. This knowledge is intended to assist IT managers in ensuring their provider can meet security requirements.

The definitions and usage cases of cloud computing are very broad, and for the purpose of this paper we will be considering cloud computing as a reasonably uniform and specific service that is most useful to medium and large organisations. Beginning with the NIST definition (Mell & Grance, 2011) it is possible to identify a range of traits that can characterise a cloud computing service. In this case we will be focusing on the Infrastructure as a Service (IaaS) model, which can provide companies and organisations with elastic, on-demand servers and networks in the cloud.

IaaS can provide several advantages over on-site networking and standard data centres. The primary argument in favour of cloud computing is related to the economic and environmental gains (also known as the green argument or cost argument). These gains are a result of elasticity, resource pooling and on-demand provision. Effectively, cloud data-centres do not use any more resources than they need to, rapidly expanding and shrinking with user demand. Townsend (2011) notes that cloud providers are very power-efficient and can shut down unused systems, leading to substantial cost-savings. A magazine article by Wik (2011) further emphasises cost-savings by suggesting that management tasks and provisioning can also be much quicker and more efficient, largely due to the standardisation of virtual machines. The economic benefits of cloud computing are fairly clear (CEBR, 2010).

In addition to being defined in terms of their service models, clouds are also divided into several types of deployment: public, private and hybrid. These models determine whether the cloud is shared between multiple organisations, dedicated to a single organisation, or comprised of a private and public cloud respectively. Essentially, a private cloud would need to be implemented in-house, or outsourced so that the organisation's data was held on dedicated servers. However, this method has received criticism since over-provision would be necessary, negating the main advantage of cloud-computing (Armbrust M. et al., 2010). As such, this paper will be primarily concerned with how organisations can secure their resources in the public cloud, including the public component of hybrid clouds. This also covers providers who offer virtual private clouds, where resources are still shared with other organisations, but further segregation is put into place.

Analysis of the Problem

In order to understand the security implications of moving from a traditional enterprise network to a cloud solution, it is important to consider the differences between these two kinds of computing. The majority of security threats facing cloud computing are the same issues we are already tackling in enterprise networks (Griffin & Jones, 2012). However, some new threats do exist due to the shift in technology and involvement of third parties, and these cannot be overlooked. This analysis identifies three significant changes that occur when transitioning from traditional models to cloud computing.

Firstly, cloud computing is achieved globally by running servers as virtual machines. This enables the most attractive features of cloud computing such as elasticity and resource pooling. Virtual machines are portable across physical servers, and can be easier to administer since they run on top of a manager. In addition to this, the provision of virtual machines is much simpler, and can also provide security advantages. In a TechNet magazine article, Vic Winkler (2011) argues that provisioning physical servers is usually a time-consuming manual process, in which an administrator follows repetitive steps. Virtual machines, on the other hand, can be rapidly provisioned from a template which could be pre-configured with a particular level of security. This enables the accelerated deployment of secure servers.

However, this introduces a whole new layer to the security model. The virtualisation layer sits between the hardware and guest operating systems, and usually consists of a hypervisor – also known as a virtual machine manager (VMM). The virtualisation layer can be seen in Figure 1 below, with two types of hypervisor. While virtualisation has been used by some organisations within their enterprise networks, for many it will be a completely new addition. If any organisation is moving into a cloud environment, then they will probably be transferring a great deal of their computing resources onto this technology – if not all of their resources. It is imperative that we consider how to best secure this layer.

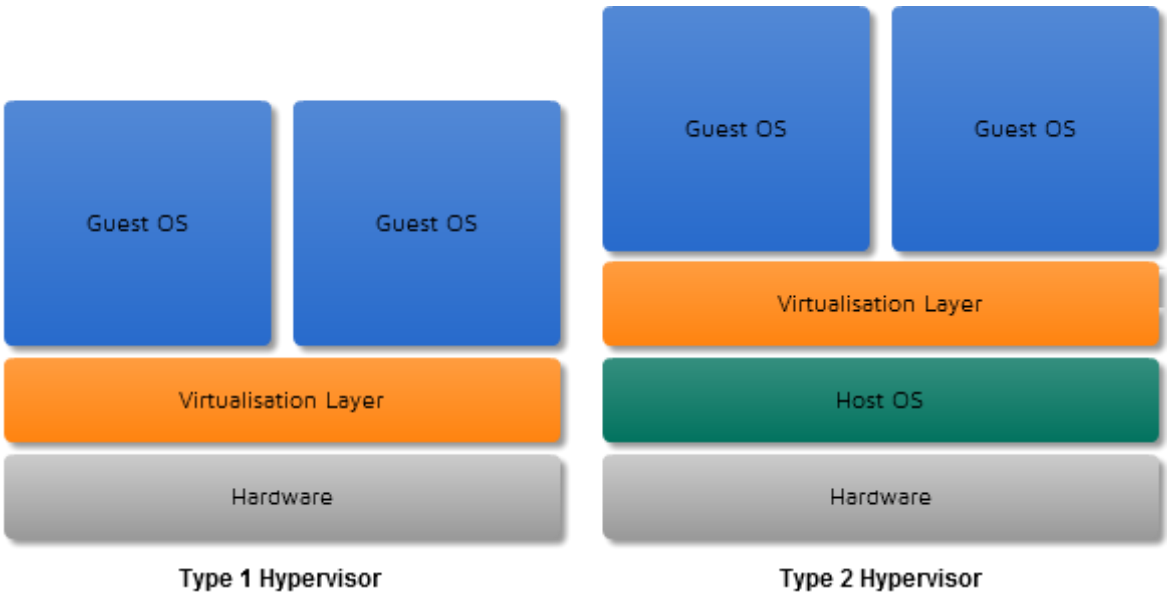

Figure 1. Types of hypervisor.

The primary issue in the virtualisation layer is the hypervisor, since it holds a highly privileged position in the stack. Ultimately, this component is in control of every guest operating system hosted on the physical machine. If the hypervisor is compromised, the guest operating systems are effectively compromised in succession. Proofs of concept have been published which show how a rogue hypervisor can execute code underneath a guest operating system in order to access sensitive data (King & Chen, 2006; Rutkowska, 2006). This can be achieved using advanced rootkits that are very difficult for virtual machines to detect, putting hypervisor security almost solely in the hands of the provider.

The second difference between traditional networking and cloud computing is that you're sharing the virtualized environment with various unknown third parties. Multi-tenancy is a given when we consider the nature of cloud computing (Wang, 2009) and not knowing who you are sharing with is often the biggest concern for organisations considering the move. It calls into question how we can ensure that untrusted hosts, sharing the same physical machine, cannot gain unauthorised access to our systems.

One of the primary issues with multi-tenancy is networking between virtual machines. Cloud data centres are secured with firewalls and intrusion detection systems (IDSs) which restrict access between different physical servers. However, if a virtual machine attempts to communicate with

another guest on the same physical host, the connection may be established over a virtual network, effectively bypassing any physical firewall or IDS (Cloud Security Alliance, 2011). Hardening the virtual network managed by the hypervisor is another critical aspect of cloud security.

Virtual machines also share computing resources such as storage and memory, and these can often contain sensitive information. These resources do, therefore, need to be handled even more carefully than they are in physical servers, or data could be exposed. For example, if a machine were to be moved off a physical server, or if it were to crash unexpectedly, then its memory allocation may be reassigned to a new virtual machine. If the data held in memory was not cleared, it could be briefly exposed to an untrusted guest (Winkler, 2011). Similarly, data in storage is at a higher risk in a shared environment than it would be if it was on-site, and this must also be considered.

The final difference between traditional enterprise networks and cloud computing is outsourcing. As discussed earlier, to reap the benefits provided by cloud computing, off-site public clouds are desirable, and this means outsourcing the whole infrastructure to a third party. In many ways, the use of an external provider can deliver increased security. Cloud computing providers are often large companies with extensive experience, and their security standards tend to follow suit. Cloud providers are usually held to higher scrutiny than enterprise security teams, and security patching can be much faster (Mills, 2009).

However, while many organisations already outsource some or all of their IT requirements, using a cloud provider does have some further security implications. Unlike outsourcing to a standard data centre, going to a cloud provider “decouples data from infrastructure” and “obscures low-level operational details” creating a layer of abstraction between you and your systems (Wang, 2009). This means that it is difficult to pinpoint where your data is being stored, for example, and how it is replicated (Heiser & Nicolett, 2008). With such a lack of transparency inherent in cloud computing, a greater amount of responsibility for the security and integrity of your data falls to the cloud provider (Armbrust et al., 2010). We need to consider how we can ensure our data is secure in the hands of the provider. Furthermore, the abstraction of the cloud may make it harder

to comply with regulatory requirements, since many details about data and processes are unknown. How can organisations ensure their data is handled in a manner that is compliant?

This analysis has considered the differences presented by a move from enterprise networking to cloud computing. Three significant differences were found to have security implications:

- The addition of a virtualisation layer increases the attack surface available to an attacker, and the impact of a compromise to the hypervisor would be very high.
- Multi-tenancy increases the risk of virtual machines being compromised. Virtual networks, memory and storage need to be secured from untrusted guests sharing the physical server.
- Outsourcing an organisation's infrastructure creates a layer of abstraction that makes it difficult to know how data is being stored and handled. Many security responsibilities fall to the provider.

These three issues have implications in terms of compliance as well as data security. They have presented several concerns which may be negligible in many traditional networks, but are vital to maintaining a strong security posture during a transition into the cloud.

The Virtualisation Layer

The consequences of a compromised hypervisor can be huge, with the possibility of an attacker gaining complete access to the guest operating systems. As such, in addition to securing the physical servers and guest virtual machines, it is vital that we consider the hypervisor itself. A Gartner research paper (MacDonald, 2010) considered the security risks presented by virtualisation and the hypervisor. It emphasises virtualisation as the most important platform from a security and management perspective, and suggests that organisations should require cloud providers to adequately address the risks involved.

Treating this platform as critical will assist in creating a locked-down environment with a minimal attack surface. Employing established security practices to the hypervisor in their most secure forms is a primary action point. This should include:

- regular and rapid patch management processes
- monitoring configurations, and blocking changes where necessary
- scanning regularly for vulnerabilities and/or missing patches
- hardening administrative access to the hypervisor, including the adoption of two-factor authentication

(MacDonald, 2010; Cloud Security Alliance, 2011)

While these practises are widespread in enterprise computing, they are not always transferred successfully to the virtualisation layer. This is often because hypervisors are marketed as being hardened operating systems. While hypervisors are generally more secure than regular operating systems, they are still vulnerable when they are misconfigured or out-of-date. IT managers should be aware of their providers' policies on the security practises above, ensuring they are to the highest standard.

Hypervisors are often considered to be secure due to their thin profile. They tend to only run processes that are absolutely necessary and open a minimal number of networking ports. This is important since it greatly reduces the attack surface available to an attacker. However, some hypervisors may not operate in such a thin environment, increasing the likelihood of a

compromise. In particular, type 2 hypervisors run on top of a traditional operating system (as shown back in Figure 1) which reveals a significantly greater number of potential attack vectors. When running a type 2 hypervisor, a breach in the host operating system could allow an attacker to use the virtualisation layer as a 'stepping-stone' onto the guest machines. This is why MacDonald (2010) favours the use of thinner architectures, as opposed to general-purpose operating systems. The use of a bare-metal, type 1 hypervisor can greatly increase the security of all the guests running on top of it. Most modern hypervisors are now bare-metal and providers who use older type 2 architectures should be avoided.

Multi-Tenancy

As we have described, there are two primary issues that occur in multi-tenancy environments that should be tackled when moving into the cloud. These relate to virtual networking which may manifest between guest operating systems on the same physical host, and shared resources such as storage and memory. Protecting these key areas ensures that untrusted virtual machines are not able to access your data.

Virtual Networks

The main problem with virtual networking is the fact that physical defences could be bypassed. Rules and policies created to manage network traffic could potentially be negated in these circumstances, opening a server to attack from other hosts (Winkler, 2011). As a client, it may be possible to employ your own defences to prevent access from unauthorised machines. Host-based firewalls and IDSs running as software could be used to limit network access to trusted machines. MacDonald (2010) warns, however, that this solution would be at the expense of a high management overhead. Administering individual firewalls and IDSs on a number of systems would be laborious, increasing the likelihood of insecurity due to mismanagement. This certainly needs to be considered, and segregation implemented universally by the provider could potentially be more effective.

The use of Virtual Local Area Networks (VLANs) is another method that may assist in virtual machine segregation. Establishing VLANs across physical and virtual networks is an effective strategy for isolating an organisation's network from third-parties, and is supported by most modern virtualisation products (Winkler, 2011). However, it is important to note that a VLAN is not, in and of itself, a security solution. They provide segregation and make the management of separate organisations' networks much easier, aiding security greatly. However, they were not designed from the ground-up as a security feature, and it is imperative to combine them with firewalls to prevent unauthorised communication between different VLANs (Cloud Security Alliance, 2011).

As a premise for network security in a multi-tenancy environment, the rule of thumb must be to ensure all network traffic between virtual and physical servers is filtered and processed against a

single set of agreed policies. This way, every connection between every device can be controlled by the provider, and different organisations' resources can be kept separate.

With this in mind, the next thing to consider is how we can implement security devices and policies that are consistent throughout the virtual and physical networks. It may be possible to use virtual devices such as firewalls and IDSs, running directly under the hypervisor (MacDonald, 2010). These could potentially be managed by the client organisation within their own VLAN, although they are more likely to be used by the provider so they can span both the physical and virtual. It is possible to implement these in such a way that they enforce rules governed by a central, physical firewall. Such products are often developed specifically for use in multi-tenant environments like the cloud.

Alternatives exist (suggested by MacDonald, 2010) that involve directing traffic from the virtual switch to a physical device for inspection ('tapping' for example), although these generally have consequences with regards to network load. Given that firewalls tailored specifically for cloud environments are readily available and represent a less complicated deployment, they may be the most sensible option for providers. Furthermore, the use of a known firewall vendor may deliver a more transparent service to the client organisation.

Shared Resources

In sharing the same physical machine, virtual guests may also have access to the same resources such as memory and storage as each other. In general, hypervisors should divide these resources up between virtual machines, thereby maintaining separation. Nevertheless, access to the same physical resources leaves less protection between the virtual machine's data and other guests. As such, providing additional layers of security to protect against data leaks is an important step, especially for more sensitive processes.

As discussed earlier, one issue with virtualisation is that memory can be reallocated regularly to new virtual machines, and these will probably belong to untrusted organisations. This means that if data is not cleared from memory before being reallocated, the new machine may be able to access sensitive information (Winkler, 2011). To some extent, ensuring applications are written to

manage their memory efficiently can assist here, although ultimately it comes down the hypervisor itself. Hypervisors should be zeroing-out all memory before it is allocated to a virtual-machine, eliminating the risk of data getting through (Cloud Security Alliance, 2011). This feature can be seen in many hypervisors that are available, and providers should be using vendors who support it (VMware, 2010; Szefer & Lee, 2012).

In cloud environments, sensitive data is often stored on the same physical hard drive used by multiple parties. While the virtualised environment should keep this data separate, we need to consider what might happen if an unexpected data leak was to occur (due to a bug in the virtualisation software, for example). An additional layer of security is required to minimise the risk of this shared space. This may also assist in protecting sensitive data from the provider themselves, although they will be discussed further in the next section.

The solution that will most reliably prevent data from being read by any third-party is encryption (Heiser & Nicolett, 2008). This method may need to be applied more stringently than in enterprise computing, given the increased risk from the provider and other tenants. It is also important that the encryption applied protects the data both at rest and while it is in use (Cloud Security Alliance, 2011), sometimes referred to as On-The-Fly Encryption. Potentially, encryption is a service that could be offered by the provider, although an organisation would also be able to implement data encryption from within the infrastructure provided.

However, data encryption only provides protection if the keys are kept secret. Credant (2011) suggest three different ways of managing encryption keys in the cloud environment.

1. Store the keys in the enterprise, and allow the provider to temporarily hold the keys while data is being accessed.

2. Trust the provider to manage the keys for you.

3. Allow a trusted third-party to manage the keys for you

If sufficient resources are available within the organisation, then the first option is always the best since it leaves you in complete control, and only you have access to the data. However, if it is

necessary to store the keys with someone else, then employ strategies outlined in the 'Outsourcing' section of this document

Even so, this encryption will not necessarily protect the virtual image itself. Often, providers might encrypt data in cloud storage, but not the virtual machine images used to access them. The Cloud Security Alliance (2011) recommends encrypting these images additionally since the operating system is likely to contain some sensitive information. A blog post (Hughes, 2009) suggests how booting an encrypted image may be possible, and some providers may be able to deliver a feature like this.

Outsourcing

The most significant difference between security in enterprise computing and the cloud, is the considerable amount of responsibility which falls to the provider. Many of the solutions suggested so far involve understanding the resources and processes implemented by the provider. Engaging with them in this way makes their service more transparent, allowing us to see past the abstraction of the cloud. This is a vital mechanism that can guarantee an organisation's data security.

One way to increase this transparency is to understand the software and hardware being used by the provider. Virtualisation in particular is a technology dominated by certain vendors, and the software in use can have a profound effect on security. As well as delivering the kinds of security features we have suggested so far, organisations can ensure the products in use have relevant certifications, and adhere to high standards (Winkler, 2011). Security certifications such as the Common Criteria Evaluation and Validation Scheme (CCEVS) can be useful in determining a particular product's level of security.

However, the use of a particular product over another does not account for misconfiguration. The vendor's standards would be negligible if the provider did not have strict security principles and processes. Furthermore, this information concerning the provider is a necessity for organisations requiring a highly compliant system. A primary way that providers should be delivering this kind of transparency, is submitting themselves to external audits and security certifications, providing details of the processes and systems that were evaluated. As Heiser & Nicolett (2008) suggest, a provider unwilling or unable to do so should not be trusted with any kind of confidential or private information.

Furthermore, a provider that is known to deliver a secure environment must be able to guarantee that they will continue to do so, and specifically in relation to your organisation. This is especially important in terms of compliance since there needs to be an assurance that any data moved onto the cloud infrastructure will be completely secure. This assurance can be provided in the form of a Service Level Agreement (SLA), which should define the provider's responsibility in terms of

security measures. Such an agreement can help an organisation assess the security provided in detail, allowing increased transparency that will also be maintained in the future.

Recommendations

The following table summarises the solutions suggested in the previous sections of this document, categorised as action points for the client organisation, or recommendations for the level of service that should be sought from a provider.

Action Points for the client organisation	**Service Provided** requirements for the cloud provider
Establish a detailed service-level agreement that provides transparency, and an assurance of security.	
	Submit to external security reviews and audits, detailing the aspects of the service which were evaluated.
	Use products (such as virtualisation technology) from security-certified vendors.
As an additional layer of security, or as you feel it is necessary (given the provider's security posture), consider implementing host-based or virtual firewalls and IDSs.	Implement virtual appliances that can operate consistently with physical firewalls/IDSs and provide a single, strict rulebase that separates organisations' systems.
	Implement VLANs in addition to firewalls and IDSs to provide segregation.
Encrypt data stored in the cloud, both at rest and in use.	Encrypt data stored in the cloud, both at rest and in use.
If possible, manage all encryption keys onsite.	
	Encrypt virtual machine images.
	Apply established security principles to the virtualisation layer (patching, monitoring, scanning and hardening administrative access).
	Use a bare-metal, type 1 hypervisor, as opposed to one that runs on top of a traditional operating system.
	Use a hypervisor that clears (zeros-out) memory before allocating it to a new machine.

Conclusion

This paper has analysed a range of new security threats that can be encountered when moving from an enterprise computing environment into the cloud. Three major shifts were identified which could significantly impact security. These changes in technology are uncharted territory for many enterprise security teams, and they each bring new risks to an organisation's doorstep.

However, one of these changes has the potential to affect the whole security environment. Most of the issues identified, relating to virtualisation and multi-tenancy, are most effectively resolved by the provider, not the client organisation. A significant amount of risk associated with cloud computing is directly proportional to the cloud provider's own security posture, potentially making this one of the largest contributing factors to how safe your organisation's data is. This paper has made a number of recommendations that can be used to comprehend how effectively a provider can defend your systems, as well as how you might be able to further protect it through specific action points. Client organisations should initially pursue the highest level of transparency that can be made available, and this can be used to assess the provider's security against strict criteria. As a general point, no single factor, such as an SLA, can provide a complete guarantee, although combing the methods suggested for increasing transparency can help ensure your provider will deliver a hearty defence to any potential attacker.

Bibliography

Armbrust, M. et al. (2010, April). A view of cloud computing. *Communications of the ACM, 53*(4).

CEBR. (2010). *The Cloud Dividend: Part One - The economic benefits of cloud computing to business and the wider EMEA economy.* EMC. London: Centre for Economics and Business Research (CEBR).

Cloud Security Alliance. (2011). *Security Guidance for Critical Areas of Focus in Cloud Computing.* Retrieved March 02, 2012, from Cloud Security Alliance: https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

Credant. (2011). *Cloud Key Management.* Retrieved March 23, 2012, from Credant: http://www.credant.com/downloads/CREDANT_Key_Cloud%20Management_WP_1221w.pdf

Griffin, D., & Jones, T. (2012, January). *Cloud Computing: The First Trip to the Cloud.* Retrieved March 02, 2012, from http://technet.microsoft.com/en-us/magazine/hh771030.aspx

Heiser, J., & Nicolett, M. (2008). *Assessing the Security Risks of Cloud Computing.* Stamford: Gartner.

Hughes, J. (2009, July 23). *Encrypted Storage and Key Management for the Cloud.* Retrieved March 23, 2012, from Crypto Clarity: http://www.cryptoclarity.com/CryptoClarityLLC/Welcome/Entries/2009/7/23_Encrypted_Storage_and_Key_Management_for_the_cloud.html

King, S. T., & Chen, P. M. (2006). SubVirt: implementing malware with virtual machines. *2006 IEEE Symposium on Security and Privacy* (p. 327). Berkeley/Oakland: IEEE.

MacDonald, N. (2010). *Addressing the Most Common Security Risks in Data.* Stamford: Gartner.

Mell, P., & Grance, T. (2011, September). *The NIST Definition of Cloud Computing.* Retrieved March 01, 2012, from http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Mills, E. (2009, January 27). *Cloud computing security forecast: Clear skies.* Retrieved March 02, 2012, from CNET News: http://news.cnet.com/8301-1009_3-10150569-83.html

Rutkowska, J. (2006, June 22). *Introducing Blue Pill.* Retrieved March 02, 2012, from The Invisible Things Lab's blog: http://theinvisiblethings.blogspot.com/2006/06/introducing-blue-pill.html

Szefer, J., & Lee, R. B. (2012). Architectural Support for Hypervisor-Secure Virtualization. *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS).* London: Association for Computing Machinery.

Townsend, K. (2011). *Spotlight on Cloud Computing: The Great Data Center Debate*. Retrieved March 01, 2012, from http://www.infosecurity-magazine.com/view/16692/spotlight-on-cloud-computing-the-great-data-center-debate/

VMware. (2010). *Understanding Memory Resource Management in VMware ESX 4.1.* Retrieved March 23, 2012, from VMware: http://www.vmware.com/files/pdf/techpaper/vsp_41_perf_memory_mgmt.pdf

Wang, C. (2009, November 18). *Cloud Security Front And Center.* Retrieved March 02, 2012, from Forrester Blogs: http://blogs.forrester.com/security_and_risk/2009/11/cloud-security-front-and-center.html

Wik, P. (2011, October). Thunder Clouds: Managing SOA-Cloud Risk - Part I. *Service Technology Magazine*(LV).

Winkler, V. (2011, December). *Cloud Computing: Virtual Cloud Security Concerns.* Retrieved March 02, 2012, from http://technet.microsoft.com/en-us/magazine/hh641415.aspx