White paper

context
INFORMATION SECURITY LTD

Assessing Cloud Node Security

Context Information Security
whitepapers@contextis.com

March 2011

Contents

Abstract

Some major Cloud providers currently expose their clients' data to the risk of compromise as a result of serious flaws in the implementation of their technologies. This is the key finding of a major new survey of the security of Cloud nodes completed by Context Information Security.

The growing trend in migrating systems to use Cloud infrastructure to take advantage of the cost savings and flexibility that this form of IT provision can offer has caused concern within the security community, because this virtual and dynamic environment creates a new threat landscape.

This whitepaper is the result of research undertaken by Context into the technical risks associated with Cloud computing infrastructure nodes. Context rented Cloud nodes from four major providers and performed a review of their security, including the limitations imposed by providers on the types of technical security testing allowed to be performed.

The methodology, results, challenges and recommended mitigations are detailed in this whitepaper, which sets out best practices for securing Cloud nodes as a client and will help clients to assess and reduce any associated risk to their systems. Information about the general security issues discovered in actual Cloud nodes has also been fed back to the providers to enable them to resolve these issues.

Executive Summary

Cloud computing can provide major benefits to organisations from a cost, flexibility and scalability perspective, but serious concerns have been raised about the security measures used to protect Cloud environments. This is because the threat landscape associated with this form of IT provision is so different to that associated with traditional dedicated hosting. If organisations no longer have direct control over the hardware or physical locations of their servers, data segregation becomes harder to achieve and regulatory compliance far more difficult to guarantee.

When Context reviewed the current publications covering this subject we found that very little factual information assessing and explaining the reality of these threats was available. We undertook a detailed research project, renting infrastructure nodes from various major Cloud providers and assessing these threats from a practical perspective, with the aim of bridging some of this information gap. The findings of this research reveal that at least some of the unease felt about securing the Cloud is justified.

The major technical issue that must be resolved to secure the Cloud is the separation between nodes. In a traditional dedicated hosted environment any attacker from the Internet must start at the outer firewall and work their way through, onto the web server then an application server and so on. This attack model has existed for decades and most security systems are designed with the repulsion of an attack based on these principles in mind.

But in the Cloud all the systems within the virtualised network reside next to each other, alongside other users' nodes. This means the software that restricts access between nodes becomes pivotal. Instead of facing an infrastructure based on separate physical boxes, an attacker can now purchase a Cloud node from the same provider as used by the organisation they wish to compromise, then start looking for a way to launch an attack on the target organisation from the perspective of their own, fully accessible node, present on the same physical machine and using the same physical resources as a node or nodes used by the target organisation.

Context reviewed the separation of the hard disk, memory, network, hypervisor (the master software that controls the nodes), and remote management. Our aim was to discover how effectively the Cloud providers had secured these aspects of the nodes provided to service users. In order to perform the assessment Context requested permission to perform a security review of our own nodes. This was granted in all cases – but only under certain restrictions. In practice this meant that not all the technical security review activities could be undertaken, so questions remain as to whether or not further security issues exist beyond those that were identified under these conditions.

Our research revealed that certain providers (which will not be named here but which have been informed of our findings) did not securely separate the nodes through the shared hard disk and network. Context was able to view data held on other service users' disks and to extract data including usernames and passwords, client data and database contents. Networks that service users might reasonably assume to be dedicated 'internal' networks are in fact open to attack from other nodes; Context therefore concludes that serious concerns over the use of these Cloud services are

justified and that the current technology used by these providers is lacking from a security perspective.

Alongside the findings of our research, this whitepaper also includes recommended actions for organisations intending to use the Cloud and seeking to counter the security issues that we have identified. This section of the document includes a list of questions that any potential Cloud provider should be asked in order to assess their suitability for the task from a security perspective.

It should be stressed that there is no complete solution to these problems and that any migration to the Cloud will carry an inherent risk. Nevertheless there are practical and effective measures that can be taken to reduce that risk. To significantly reduce the impact of data leakage Context recommends the use of encryption on hard disks and network traffic between nodes. Also, all networks that a node has access to, including internal management and node-to-node networks, should be treated as hostile and should be protected by host based firewalls. Finally, default nodes provisioned by the Cloud providers should not be trusted as being secure; clients should security harden these nodes themselves.

Despite the security concerns associated with Cloud computing, the compelling economic and operational benefits it can offer mean it is inevitable that it will be used in the future by many organisations. Many are likely to become increasingly dependent on this form of IT provision, therefore the security risks will need to be addressed. The aim of this whitepaper is to improve the current understanding of security issues related to the Cloud. We have uncovered serious security flaws within the technology used by certain Cloud providers. While those providers have told us they are working on solving these issues, at the time of publication these issues remained unresolved, therefore no specific details detailing how these attacks can be undertaken or which providers are vulnerable have been included in this document. The issues are discussed as new classes of vulnerabilities and Context recommends that any system in the Cloud applies the recommendations detailed in this whitepaper to reduce the risk of a security breach.

Overview of the Cloud

Introduction to Cloud Computing

Cloud computing services such as Amazon EC2 and Windows Azure are becoming more and more popular but it seems many people are still unclear as to what exactly the buzzword "Cloud computing" actually means. In its simplest form, the principle of Cloud computing is the provision of computing resources via a network.

Cloud computing shifts the responsibility of configuring, deploying and maintaining computing infrastructure from clients to Cloud providers. Providers generally expose an interface for clients to interact with their resources as if they were their own standalone resource; however often a number of resources may be aggregated on the same computer or cluster of computers. The user does not necessarily know the details of the location, equipment or configuration of their resources, rather they are provided with a "virtualised" computer resource hosted in "the Cloud".

Cloud services take care of a lot of the mundane tasks associated with hosting a service (for example, maintenance and backup tasks) and leave developers and IT administrators to concentrate on the specific details of the application they wish to provide.

In addition to public Cloud services, many organisations are implementing internal private Clouds to reduce costs, complexity and consolidate infrastructure. There are a number of different varieties of Cloud services as listed below.

Types of Cloud Providers

Cloud services are usually divided in the three main types, Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS).

Software as a Service (SaaS)

SaaS clients rent usage of applications running within the Cloud's provider infrastructure, for example SalesForce. The applications are typically offered to the clients via the Internet and are managed completely by the Cloud provider. That means that the administration of these services such as updating and patching are the provider's responsibility. One big benefit of SaaS is that all clients are running the same software version and new functionality can be easily integrated by the provider and is therefore available to all clients.

Platform as a Service (PaaS)

PaaS Cloud providers offer an application platform as a service, for example Google App Engine. This enables clients to deploy custom software using the tools and programming languages offered by the provider. Clients have control over the deployed applications and environment-related settings. As with SaaS, the management of the underlying infrastructure lies within the responsibility of the provider.

Infrastructure as a Service (IaaS)

IaaS delivers hardware resources such as CPU, disk space or network components as a service. These resources are usually delivered as a virtualization platform by the Cloud provider and can be accessed across the Internet by the client. The client has full

control of the virtualised platform and is not responsible for managing the underlying infrastructure.

IaaS Cloud Technology

As this study focuses on Infrastructure as a service, we'll briefly discuss the technologies used in this Cloud service.

IaaS commonly makes use of virtualisation technology to provide computing resources as virtual private servers (VPS). These VPS's are functionally equivalent to separate dedicated physical servers however they share computing resources with other VPS nodes. Virtualisation allows multiple VPS nodes to be hosted on a single physical host.

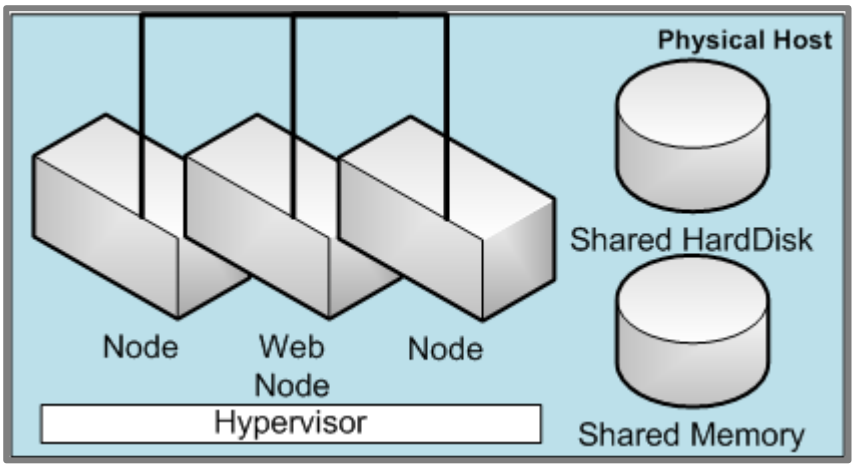

Figure 1
High level overview of IaaS Cloud components

A specially designed Operating System (OS) called a Hypervisor provides an abstraction interface between the physical hardware and the virtual nodes. The most common hypervisor seen during this review is an open source product called Xen.

Each hypervisor has a single supervisor virtual machine (called a Node0 for Xen) which has privileged access and controls access to volatile memory(RAM), hard disks, network interfaces and all other physical hardware for other virtual machine nodes. These supervisor nodes are also responsible for nodes being isolated from each other.

Using VPS's allows the service provider to quickly and efficiently deploy and manage independent computing nodes using common hardware. It can also be scaled to use multiple hardware devices to increase capacity. Products like OpenStack can be used to manage large clusters of these server hosts.

The Business Case for Cloud Computing

In an economic environment where organisations must achieve more with less, Cloud computing has become one of the buzzwords of the moment. The flexibility and potential cost savings offered by the Cloud are just two of the reasons that can make it an attractive business proposal to many organisations. In this section of our whitepaper, Context shall provide a high-level summary of the key business drivers for Cloud computing as conveyed to us by our client base which comprises multinational organisations across a variety of verticals.

Benefits

Faster Implementation

A frequently mentioned perceived benefit of Cloud computing is the ease and speed with which organisations can get a solution up and running on the Cloud. Compared with the amount of time required to set up a new solution internally within many organisations, a Cloud-based implementation can be achieved relatively quickly, accelerating the time required to bring new services to market or to make them available to internal users. Some of our clients pointed out that moving services to the Cloud reduced their overall administrative overhead, thereby further speeding up the process.

Lower Cost

For many organisations, especially new businesses, smaller companies and those embarking on new projects, Cloud computing can be less expensive than hosting systems and services internally. The on-demand nature of the Cloud is such that organisations only pay for services when they are required, so where an organisation has not already invested in the IT infrastructure needed to host their systems, the Cloud can represent a lower total cost of ownership. Several clients pointed out that they will benefit from the lack of direct maintenance costs and increased innovation by more intimate collaboration with their clients, partners and suppliers.

Security

Much has been said about Cloud security or the lack of it and this topic shall be dealt with in detail in the course of this document; however for many small and medium-sized organisations that do not necessarily possess dedicated security resources or know-how; a move to the Cloud can result in improved security. As part of the contract between the Cloud provider and the client, the provider should have an obligation to provide the client with a degree of security. The level of security provided may trump that which the client organisation itself has the means to provide. Context's recommendation is to read the SLA extremely carefully; more will be said about contractual issues later.

Greater Flexibility

As mentioned previously, Cloud computing provides clients with considerable flexibility. Clients have the option to scale-up or down their usage of Cloud services in accordance with demand, giving them more elasticity in terms of how they manage their operations. The flip side of the coin is that organisations that have invested heavily

in IT in order to cope with heavy demand during key periods can now sell their resources to others at periods of lower demand; thereby deriving more benefit from their investment.

Benchmarking

As a final point in this high-level summary, one motivator for turning to Cloud computing is to benchmark internal providers against Cloud providers. By comparing internal levels of service against those of the Cloud providers an organisation can identify areas where internal service provision can be improved and made more efficient. In some cases organisations may conclude that the service provided internally is actually of a higher level than the organisation is currently able to obtain from the Cloud.

Client Concerns

Although the Cloud can provide a number of benefits, many of the organisations Context has spoken with have concerns about Cloud computing. This white paper will look into those concerns from a technical security perspective, however at this point we believe it is worthwhile touching upon another key concern; that of the contractual aspects of a obtaining a service from a Cloud provider.

Many of the organisations Context has dealings with have commented upon what they perceive to be the often vague nature of the Cloud provider's contractual documentation. In fact, Gartner has released a report[1] on this subject, highlighting the chief concerns that CIOs should be aware of when considering the pros and cons of moving to the Cloud; these are summarised at a high level below:

Contracts are not Mature for all Markets

In its report, Gartner points out that Cloud computing contracts "lack descriptions of Cloud service providers' responsibilities and do not meet the general legal, regulatory and commercial contracting requirements of most enterprise organisations."

Contract Terms Generally Favour the Vendor

Gartner stated in its report that current contracts favoured the Cloud provider and that potential clients need to be "clear about what they can accept and what is negotiable."

Contracts are Opaque and Easily Changed

Gartner's research indicated that many Cloud contracts lack detail. Clauses within the Cloud contract are not necessarily detailed, and the documents themselves contain links to web pages that contain further details as well as additional terms and conditions. Some of the content lacking was found to relate to critical elements such as terms for service and support and QoS. A key concern is that the clauses on the web may be changed at short or no notice. Gartner points out that clients need to ensure that "terms cannot change for the period of the contract and, ideally, for at least the first renewal term without forewarning."

[1] The title of the report by Frank Ridder and Alexa Bona, which was released on the 9th February 2011, is "Four Risky Issues When Contracting for Cloud Services ". The report can be obtained/purchased here: http://www.gartner.com/DisplayDocument?doc_cd=210385

Contracts do not have Clear Service Commitments

According to Gartner, many Cloud providers do not clearly define their service commitments, in most cases; limiting their responsibility to their own network.

It is therefore critical that any organisation considering a move to Cloud computing should scrutinize the provider's contract and ensure that it meets requirements, where it does not; potential Cloud clients are strongly encouraged to negotiate better terms, consider other suppliers, or think again about their move to the Cloud.

The Project

Overview

Context reviewed a number of Cloud computing service providers offering infrastructure as a service (IaaS) from a security perspective. Context conducted an assessment against a number of Linux and Windows nodes for each provider. This assessment included analysing the security of the node, the virtualised environment and the Cloud provider's management interface. Context also evaluated the security of the network connections between nodes in the Cloud and between nodes and the greater external network (Internet).

Objectives

The primary objective of this project was to assess and identify common security problems that clients could potentially encounter when using an external Cloud provider. We examined the following broad categories of infrastructure as a service (IaaS):

- Node Assessment
- Virtualisation
- Internal Infrastructure
- External Infrastructure

Scope

Cloud providers place strict restrictions on what is allowed to be reviewed from a security perspective. Restrictions were applied to the security assessment due to the various legal agreements imposed by different providers. In some cases Context was restricted to the extent where we were obliged to follow the Cloud provider's own penetration testing guidelines.

In most cases, Context was limited to testing the security of the host operating systems relating to our nodes. Any tests deemed to be potentially destructive or intrusive were not permitted. Actively attacking the hypervisor or the underlying infrastructure was prohibited by most providers. Restricted or disallowed cases are clearly marked in the findings contained in this whitepaper. In these cases more security issues may exist.

It is worth noting that a normal client can request a penetration test and normally has restrictions similar to those mentioned above placed upon them. It has been known however that Cloud providers have lifted these restrictions for client organisations with greater purchasing power.

Specific vulnerabilities discovered during this research have been fed back to the providers for them to rectify. Due to the sensitivity of these issues the specific providers have not been named within this whitepaper. Context is committed to responsible disclosure and will release details when the Cloud providers have fixed the specific issues.

Attack Vectors

The new threat landscape created by Cloud computing adds a number of potential attack vectors; this is due to the shared nature of the technical resources such as memory and disk space.

The following attack vectors were assessed within this whitepaper.

- Public (Internet)
- Internal (Malicious Node-to-Trusted Node)
- Hypervisor subversion
- Shared physical resource (Memory and Hard disk)
- The Cloud Provider (Hosting Company Breach)

The following diagram provides a graphical representation of the attack vectors. The red lines show the threats to a node within the Cloud.

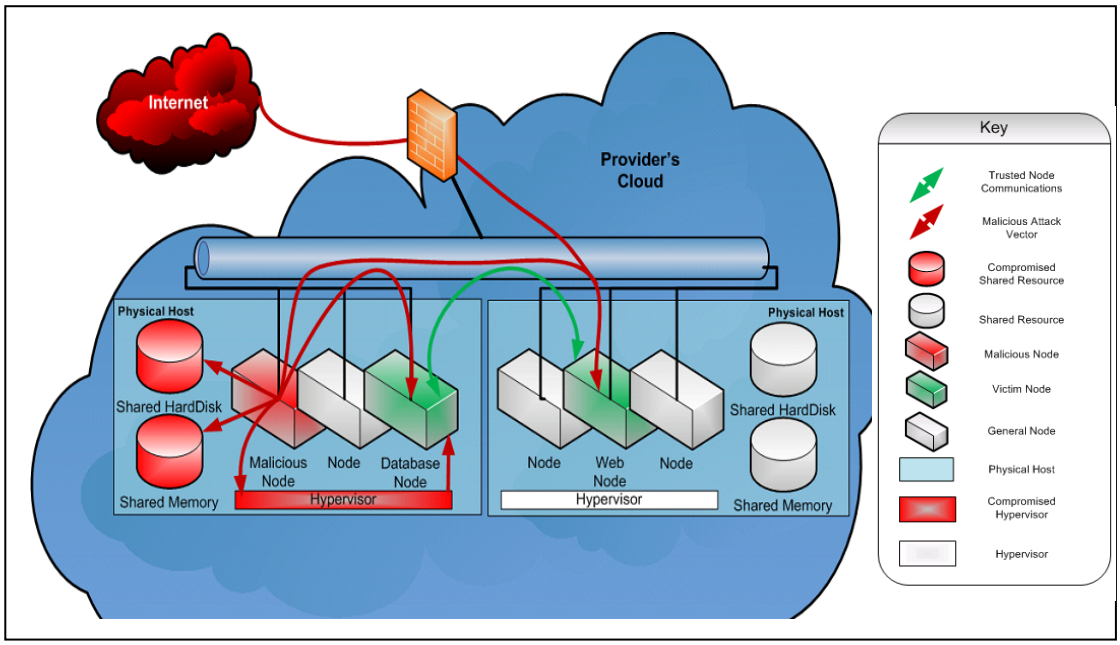

Figure 2
Diagram of the new threat landscape

As can be seen in the diagram there are more threats than just the attacks from the Internet. With a shared Cloud environment an attacker has the ability to place a malicious node onto a shared VM hosting platform or network segment, effectively allowing an attacker to connect to the internal network from the Internet.

In addition, the shared resources are potential attack vectors for an attacker. Memory and hard drives can be imaged from a malicious node potentially disclosing memory segments or file content.

The hypervisor is software that controls all the nodes running on that physical machine. The communications between the hypervisor and the nodes exposes an attack surface which if compromised would enable an attacker to subvert the hypervisor and take control of all the nodes on that server.

It is also worth noting that by using a 3rd party to provide Cloud services, there is an implied trust of that 3rd party and their security practices. It is entirely possible that providers could utilise improperly configured or flawed software, improper security controls or even have legitimate backdoor access for maintenance reasons. There is the potential therefore, that a malicious employee, another node client or outside attacker could compromise a node's security due to 3rd party failings.

Results of Security Review

In this section the details of the security review are explained along with the general results seen across the providers. Overall it was found that the providers failed in 41% of the tests undertaken indicating that Cloud providers currently have serious security issues. Where the tests could not be undertaken due to the restrictions imposed the result is marked as unknown. This makes up 34% of the tests cases.

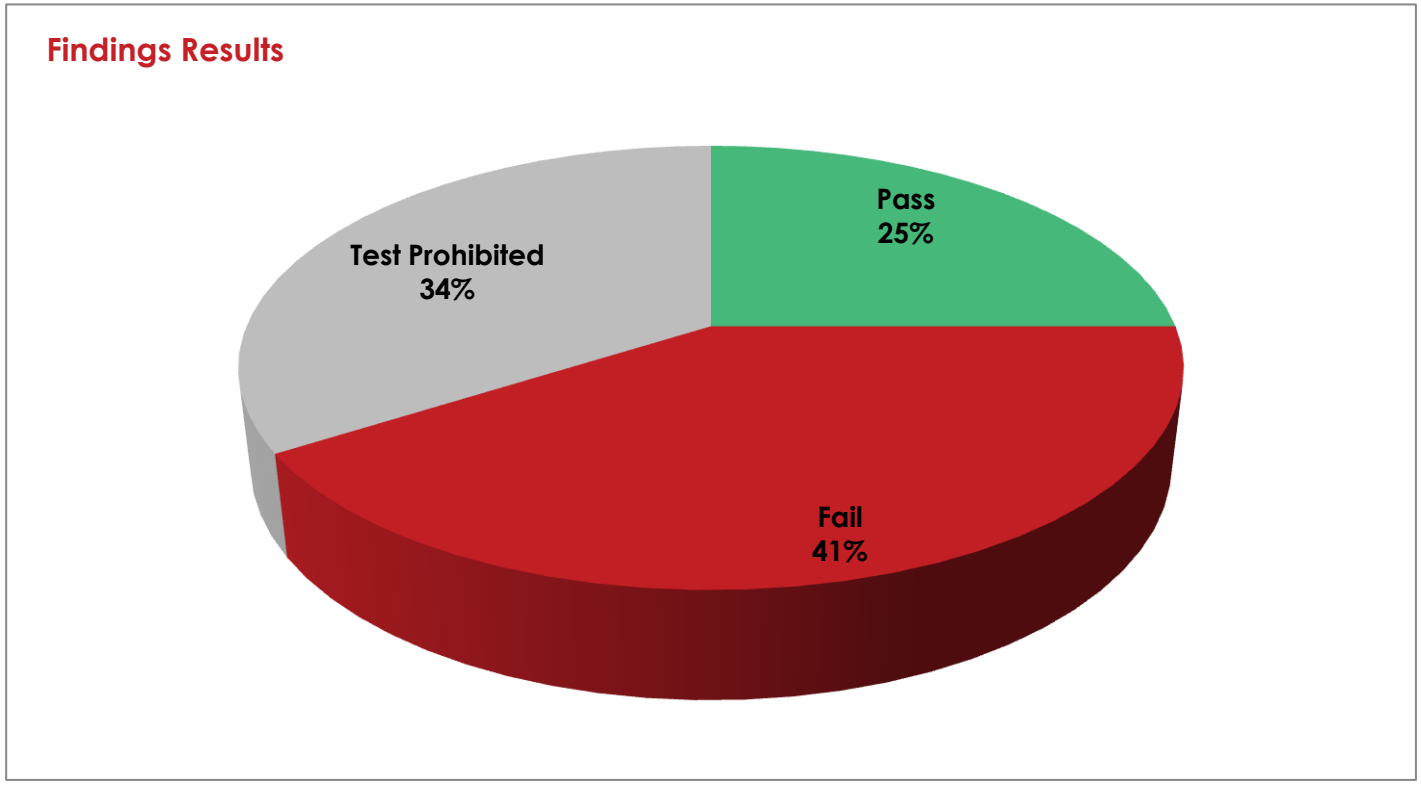

Figure 3
Results of the security review showing the percentage of providers which passed, failed and where the test was prohibited

Node Assessment

During this project Context performed reviews of the default nodes issued, both Windows and Linux based systems were assessed. The reviews were performed in-line with our standard host based build review methodology.

Test ID	Test Title	Pass	Fail	Unknown	Findings
1.1	**Node Hardening** System hardening is a broad subject and is specific to the application of the node. In general Context assessed the following aspects of the default nodes issued by the providers from a security perspective: Unnecessary Package/software removal File System Hardening Boot/Start-up Configurations Service/Demon Hardening Password Policy Hardening Network Exposure Hardening User Privileges Auditing Configuration Available Encryption Mechanisms Patch Management Anti-Virus Host Based Firewall Configuration	0%	100%	0%	**Default Nodes are Insecure** The default build issued by the providers should not be considered fully secure. The build should undergo stringent security hardening as recommended for any production system that is publically accessible. One provider was found to use a default system account configured with a password and login rights. This could possibly be used as a backdoor to other nodes. All the builds were outdated and were missing security patches. Most of the operating systems had third party software installed that is not required from an operational system and increases the risk to the security of the system. One provider installed a custom service on Windows which could be abused to escalate privileges. None of the providers offers disk encryption by default. This puts the node's data at risk of being read by another node if such an attack vector is present. Of the observed Windows and Linux systems none had an antivirus system installed and the password policies were system default. Central auditing was not implemented for any providers tested.

Virtualisation

Security considerations directly related to the virtualised environment.

Test ID	Test Title	Pass	Fail	Unknown	Findings
2.1	**Host Platform Configuration** Access from the virtualised node to the management system must be prohibited. The introduction of a virtualised environment introduces new systems to a network which can control the nodes they are virtualising. If an attacker can gain control of the host then all nodes would be compromised.	0%	0%	100%	**Not Permitted to Test** Testing of the host platform was forbidden by the providers because it could interrupt the production system.
2.2	**User Management Methods** Remote node management (web interface or other APIs) (Not testable). Considerations which were taken into account: Secure Channel Encryption Input Validation Access Controls Context was unable to perform a security assessment of the remote management facilities associated with each provider due to legal restrictions.	0%	0%	100%	**Not Permitted to Test** A secure remote method is required in order to manage Cloud nodes via a web application or API (Application Program Interface). The majority of providers implement a custom web application which can be remotely accessed allowing for various actions to be performed, for example provisioning of nodes, console access and power-cycling. A full security assessment for any management method should be undertaken. It is paramount that a secure authentication method is used for the remote Cloud management method. As these methods tend to be public facing, weak authentication methods could lead to a full compromise of the Cloud infrastructure.

Test ID	Test Title	Pass	Fail	Unknown	Findings
2.3	**Provider Node Management** Some providers use backdoors and insecure practices in order to manage the nodes. This is potentially a security issue if not implemented in a secure way. A number of providers offer the option to aid in support issues by logging directly into nodes hosted within the Cloud. This is done usually through an active service (Backdoor).	50%	50%	0%	**Certain Providers use Backdoors** Context identified a web service used for remote configuration on one provider and a system account on another. These remote administration interfaces leave the nodes vulnerable to exploitation. For example if the connection between the node and the web service is intercepted, or if access to the system account is granted. Some providers perform the password reset function by powering off a node and accessing the file system to change the password from the hypervisor or the storage system; this mechanism is more secure than providing a backdoor into the node.

Test ID	Test Title	Pass	Fail	Unknown	Findings
2.4	**Hypervisor Exploitation** The hypervisor is the software that provides the virtualisation of the nodes and is responsible for ensuring that the nodes are securely segregated while in operation. A security flaw within this software would undermine the whole environment. This is the case of any system that must trust the core components that they are built upon, whether this is the operating system or web server. All hypervisors have suffered from security vulnerabilities in the past, and will continue to have issues in the future. Therefore it is important to check that the software is up-to-date and locked-down as per best practice. Due to the fact that nodes can migrate from one host to another, it was not possible to check all instances; only the ones present during the review.	0%	0%	100%	**Hypervisor Versions Out-of-Date** All of the hypervisors examined during the review utilised Xen for virtualisation. All the vendors reviewed were found to have outdated Xen versions installed. It should be noted that this information was gathered based on string analysis performed on the node. Ideally Context would prove whether or not these versions are actually vulnerable to the known exploits; however this was not possible due to the restraints imposed on the testing by the Cloud providers reviewed.
2.5	**Node Memory Separation** Can one node access the physical memory of another node or the physical memory of the underlying host? The memory of a virtualised system is shared between nodes. The memory can be inspected in various ways. Most of the Linux systems tested were found to prevent direct access to the memory; however, ways exists to bypass this protection and it must be ensured that memory does not leak information to and from other nodes.	100%	0%	0%	**Memory Securely Separated** None of the nodes tested were found to leak memory from one system to another. Context performed this test on Linux systems using a proprietary kernel module that we designed to access the memory without the host memory protection in place.

Test ID	Test Title	Pass	Fail	Unknown	Findings
2.6	**Virtual Disk Separation** Can the node access the virtual disk of another node or the host? The physical disk can leak data from other systems. Typically this occurs if the allocated hard disk space has not been wiped securely. This could result in allowing access to other data, or having the node's data accessed by unknown parties.	50%	50%	0%	**Flawed Disk Management** Two of the providers reviewed were found to have a fundamental flaw in their implementation of hard disk separation. Context was able to access other node's data from their virtual disks. Full details have not been included in this report because the providers are currently resolving this issue. The remaining providers did not leak other client's data during the test.
2.7	**Resource Exhaustion** When sharing a host the nodes on this host should not be able to consume all of the available resources. Especially on IO activities this could lead to a Denial of Service for other hosted systems. The Cloud provider must ensure that exhausted systems are separated during runtime so as to prevent the performance of other nodes being affected.	0%	0%	100%	**Not Permitted to Test** Due to restrictions imposed by all of the Cloud providers reviewed, Context was unable to test resource exhaustion. Context believes the most likely exhaustion could occur at the IO level. The underlying storage system could not be inspected in order to complete this test.

Test ID	Test Title	Pass	Fail	Unknown	Findings
2.8	**Remote Alternative Booting** The providers often rely on booting an alternative operating system via the network, either for disaster recovery, or for the initial operating system install. This should happen on a network separated from the production system and the provider needs to ensure the environment is especially hardened for this purpose.	75%	25%	0%	**Generally No Alternative Booting Supported** None of the providers were found to offer remote booting capabilities; alternatively these were not visible to the Cloud node. Context discovered that one provider was using self-assigned IP addresses to retrieve the node configuration files via the boot process. This could put the node at risk of an internal attack, or possibly privilege escalation attacks. However, this would involve other systems and the underlying infrastructure, which Context was not allowed to test during the review. Before acquiring a Cloud node Context recommends the process of the configuration mechanism should be reviewed. In general no remote booting should be offered prior to the booting mechanism.

Infrastructure (Internal)

Security considerations between nodes hosted on the same virtualisation platform.

Test ID	Test Title	Pass	Fail	Unknown	Findings
3.1	**IP Network Separation** Communications between nodes should be restricted as per connections from the Internet. Some providers offer multiple interfaces; those intended for public access and those to facilitate private configuration. This can lead to node-to-node attacks form the external and internal interfaces.	25%	75%	0%	**Internal Networks Insecure** Context found that certain providers were making two network interfaces available; one externally facing with an Internet address, the second on an internal management LAN. The internal LAN provided no separation between nodes and therefore could be used to launch a node-to-node attack. This demonstrates the importance of host-based firewalls and use of encrypted protocols (see mitigation section). One provider was found to offer a centrally-managed firewalling solution by default that filtering both internal and external traffic in the same way. This is a good example of how to provide node-to-node IP separation.
3.2	**Routing Protocols** This comprises a review of the layer 2 and routing protocols in use in order to assess the configuration-related security implications. Layer 2 communication must be protected against internal threats. Typically ARP is in use and an ARP proxy should be used by the Cloud provider to prohibit layer 2 communications between nodes on the same network segment. Other protocols, for example, spanning tree or Cisco Discovery Protocol should not be visible to the node at any time.	0%	50%	50%	**Limited Testing Permitted** Due to the threat of attack from potentially malicious nodes on a shared LAN, Context planned to perform layer 2 network penetration testing within the Cloud environment. However due to the restrictions imposed by the providers Context was prohibited from performing this type of testing. Context cannot therefore comment on the effectiveness of layer 2 attacks. It was noted that two providers implemented an ARP proxy which prevents nodes from accessing other nodes on that layer. One provider was found not to filter any layer 2 traffic in the network segment, whereas another provider was found to leak switching protocols.

Test ID	Test Title	Pass	Fail	Unknown	Findings
3.3	**TCP/UDP/ICMP filtering** The node should block TCP/UDP/ICMP network traffic by default and it should allow only predefined sources and/or ports to communicate with the Cloud node. This should be provided by a firewall that can be configured through the node web management interface.	25%	75%	0%	**Most Providers Do Not Provide an Internal Virtual Firewall** Context identified that only one provider provides a virtual firewall which allows for full filtering of TCP/UDP and ICMP protocols. Others do not, and nodes are responsible from ensuring that a host-based firewall and service lockdown have been configured. This is a weaker security posture than would be expected within a traditional dedicated host-based environment.
3.4	**IPv6 Support** In the case of IP version 6 (IPv6) being enabled and configured on the node, for example, via auto-configuration, this service should be secured; especially if the Cloud provider supports inter-node communications using IPv6; and IPv4 is blocked by a central firewall.	50%	50%	0%	**Certain Providers Allow IPv6 Unrestricted** All nodes auto-configure an IPv6 address. However, none of the vendors route IPv6 traffic by default and therefore no connectivity via IPv6 can be achieved from the external network. Two vendors allow internal IPv6 communications. As there was neither a central firewall, nor a host-based firewall controlling internal IPv6 traffic, communications could occur without any filtering. Several services were identified on the default build allowing inbound IPv6 communications. This puts the services at risk of being attacked by another internal node via IPv6.
3.5	**Other Protocol Support** Protocols other than TCP/UDP/ICMP should be investigated to determine the level of support available. A typical example is the need for IPsec protocols, or SCTP. The Cloud provider should secure these services by default (deny), but, where required, must allow access for secure VPN communications to the Cloud.	0%	100%	0%	**Insecure Support for Other Protocols** One provider was found to prevent access to the node for protocols other than ICMP, TCP and UDP. This means it was not possible to use secure communications like IPSec tunnels or SCTP. The three other providers allowed communications via all protocols from the Internet; however they did not provide any means to control these protocols via a virtual firewall.

Infrastructure (External)

Security of the guest nodes from external entities on the Internet.

Test ID	Test Title	Pass	Fail	Unknown	Findings
4.1	**Supporting Boundary Infrastructure** The Cloud infrastructure network contains several devices used for load balancing, switching and firewalling. These devices must be reviewed for security vulnerabilities that could affect the Cloud node. Changes to the devices affecting a node in use should be communicated to the node owner. Ultimately, due to the transient nature of the nodes within the Cloud, a review of the node in one particular configuration could change to another location with a different boundary and therefore potentially raise new security issues.	0%	0%	100%	**Not Permitted to Test** Context was not permitted to perform an assessment of the external devices during the review. At the time of the review none of the providers were willing to provide Context with information regarding maintenance performed on the network, or on the hypervisor.
4.2	**Infiltration/Exfiltration Configuration** The nodes should be protected with an effective firewall with a strict rule sets to ensure communications from only trusted sources are permitted. The rule set should be under the control of the client's of the node. In addition, the Cloud providers' boundary devices (e.g. firewalls) should be independently audited to ensure an adequate level of physical boundary protection. Providers should notify clients of changes to boundary device architectures including rule-base changes.	25%	75%	0%	**Most Providers do not Provide an External Virtual Firewall** Three of the four providers did not provide a firewall to restrict network traffic from the Internet to the nodes. For these three providers Context did not identify any firewalling in place at any stage; however, not all of these aspects could be tested due to restrictions imposed by the providers.

Recommendations

In general it is advisable to perform a review of all components that affect the Cloud node. Some aspects may not be testable due to legal reasons, in which case the risk needs to be accepted or alternative providers found. The key recommendations are to lockdown the default nodes that are provided, ensure that encryption is used for both virtual hard disks and network traffic, and treat all network interfaces as un-trusted. This section details these recommendations.

Node Hardening

System hardening is a broad subject and is specific to the application of the node; however, as a base node is provided by the majority of Cloud providers, some clients may assume they have been hardened to a reasonable standard. This is an incorrect assumption as Cloud nodes are often not hardened to the required level, and rely on the user to make each node secure.

Clients should always perform stringent hardening in the same way as they would do on any of their own publically-accessible production servers. Due to the number of potential attack vectors, all aspects of the node should be considered public-facing, and therefore hardening should be performed accordingly.

The following aspects regarding node hardening within a shared Cloud environment should be taken into account:

- Patch management
- User privileges
- Enforce least privilege
- Service/demon hardening, all unnecessary services/demons should be disabled
- Unnecessary packages/software should be removed
- Anti-virus scanning should be performed
- Host-based firewall configuration
- File system permissions hardening
- Password policy hardening
- Network exposure hardening
- Available encryption mechanisms
- Auditing/monitoring configuration
- Boot/start-up configurations ensure only required start-ups are active

It is outside of the scope of this whitepaper to fully detail server hardening guidelines. Several guidelines exist for device hardening including operating systems. Examples include the US National Security Agency (NSA) and the UK Centre for the Protection of the National Infrastructure (CPNI); both of which provide some comprehensive guides:

http://www.nsa.gov/ia/guidance/security_configuration_guides/current_guides.shtml

http://www.cpni.gov.uk/advice/infosec/business-systems/

Virtualisation

Due to the nature of virtualisation, shared resources of the physical device will be used. An attacker may be able to leverage this relationship and gain access to resources allocated to other victim nodes. Ensure that the provider has an adequate patch management policy in place. The virtualisation framework should be updated regularly.

The use of software-level encryption, supported by most modern operating systems, should be used to encrypt content written to the disk, including SWAP space. Encryption ensures that no sensitive data is retrievable from physical disk space, should it be reallocated and subsequently exposed to malicious nodes.

Management of nodes should be done via a secure mechanism. The majority of providers use a web-based application or an API to manage nodes. These management methods should have undergone independent security assessments which should be available to clients on request. In addition, to regular security reviews, the use of two-factor authentication to the management interfaces is recommended.

Custom software used on nodes should zero-out encryption key values stored in memory once used. During boot time the virtualised BIOS configuration should be configured to clear memory during the POST (Power on self test) process during the boot process.

Resources reported to be available on the node system may not reflect what is actually available on the Cloud provider's hardware. If possible, use provider APIs and management interfaces to monitor physical resource allocation, ensuring the correct resource allocation is assigned to your node. This allows for monitoring of resource exhaustion which can affect availability.

Where available, ensure that the virtual BIOS is locked down and that booting from remote sources is not possible.

Infrastructure

Internal connectivity, and node-to-external sources such as the Internet, should be treated the same. All connections to the node are potentially hostile.

Limiting the connectivity between nodes restricts the attack surface. Host-based firewalls should be enabled and correctly configured to limit connections to only those that are permitted. A white-list approach is recommended to deny all traffic (ICMP/UDP and TCP) then implicitly define rules for authorised communications channels to trusted hosts.

In addition to host-based firewalls, various providers allow clients to modify boundary firewall rule-sets for their specific nodes. It is recommend this be utilised and be locked-down with a strict rule-base giving a more granular control to routing.

Inter-node communication channels should also be considered when securing a node. Sensitive data between nodes should be encrypted using a transport-level encryption protocol such as SSL. SSL certificates should be used to detect man-in-the-middle attacks.

Only use the protocols necessary for the operation of the node. Reduce the network protocol footprint to only those which can be filtered and are supported by the provider. For example, only allow TCP communications on specific ports and remove support for the ICMP protocol wherever possible.

Network services should be limited to the services required for business operations. TCPWrapping and host-based firewalls should be used to limit these services to trusted hosts only. Adequate auditing for all services should be configured and reviewed regularly.

Ensure the provider uses adequate layer 2 network separation, static routing and limits the size of network segments. Adequate boundary device protection between segments should also be in use. Consider the implementation of ARP static routing on nodes or ARP proxies to aid against ARP spoofing and IP spoofing. Providers using Xen networking and routing were found to provide a good level of network segregation between nodes however this is not always correctly configured.

Other Considerations

This paper takes a technical approach to securing a node within a Cloud service infrastructure, however, there are a number of other considerations, primarily legal and procedural that should be taken into account when selecting a Cloud provider.

Standardisation of Cloud computing is in its infancy, hence identifying providers certified in established industry security standards such as ISO 27001, PCI and SAS70 is paramount to ensure the correct internal policies are being adhered to within the provider's remit.

Standards are not the only consideration, as only certain parts of an organisation may be accredited. Other considerations when selecting a Cloud provider are:

- Service Level Agreement (SLA)
- Incident response procedures
- Disaster recovery procedures
- Backup management policy
- Node management option
- Physical device access controls
- User validation

There are a number of non-technical security issues to consider when using a Cloud computing provider. You are at the mercy of their security practices; without rigorous security processes, your nodes are vulnerable to social engineering attacks or compromise as a consequence of the provider being infiltrated.

Questions to Ask Your Provider

This section details certain pertinent technical questions that any potential user of a Cloud provider should request of them. Defining a definitive list of questions for Providers is difficult; the nature of the Cloud is extremely diverse and questions should be specific to the application of the nodes being implemented. The following are general questions which can be used as part of a Provider assessment.

General

What standards does the Provider adhere to?

Industry recognised security standards such as ISO 27001, PCI or SAS70 Type I/II should be adhered to by the provider. This helps ensure that correct internal procedures around access controls and security are being adhered to.

If required, perform an independent security audit against the standards that should be adhered to.

Are regular security assessments carried out against the provider's infrastructure and applications?

The provider has a responsibility to ensure components owned and run by them meet a reasonable level of security. Regular security assessments should be carried out and the results should be available to clients on request.

What data recovery/backup procedures are in place?

Ensure that adequate disaster recovery procedures are in place and that a backup facility for nodes is provided.

Is the SLA adequate for the organisation's requirements?

Ensure that the provider's support structure for technical issues is acceptable for your organisations requirements.

What incident response procedures are in place?

A good provider should have an Incident response procedure in place and clients should be kept informed of any security related issues which may affect their nodes.

In the event that a node is suspected or found to be malicious, a policy to isolate and remove the node should be in place. Providers should have adequate network monitoring and logging to help identify malicious activity from an internal and external perspective.

A good provider should keep clients informed of the latest security threats relating to the Cloud and should be able to aid in security incidents.

Node Hardening

Are default nodes hardened using recognised guidelines?

Default nodes should be hardened to a reasonable standard, including the latest, stable patches and general hardening by default. The majority of providers state that it is the client's responsibility to ensure the security of the node.

Has a security audit been previously been performed against the default builds?

A security assessment of a default node build should be carried out regularly by the provider. The results should be available to the client on request.

Virtualisation

What version of the virtualisation platform is being used? Is it up to date?

The latest, stable version of the hypervisor should be installed. Custom implementations of a virtualisation framework should be security assessed and risk balanced.

Any known security issues associated with the version deployed?

Known security vulnerabilities may lead to the compromise of the hypervisor or the physical hosting platform itself. A risk balance case should be defined for any vulnerability known to affect the platform installed should no patch be available at present.

What is the patching procedure/policy?

The virtualisation platform on all nodes and the physical hosting platform should be regularly updated. An adequate patching procedure should be in place to ensure critical and important patches are installed quickly following a security advisory. Patching must be applied to all base hosts.

Can multi-factor authentication be used to access management?

The majority of providers allow users to manage their node infrastructure via a web interface. Context found that access is typically provided via username and password authentication. A good provider may offer two-factor authentication as an additional security measure.

In addition, direct node access should be done by means of multi-factor authentication using the concept of lowest-level user privileges. Direct Root/Administrator access should not be authorised.

What is the node overwriting policy? How are disks managed after failure?

Following hardware failure or device reallocation, disks should be securely wiped ensuring no sensitive data can be retrieved.

Infrastructure

What network segmentation is being used?

The virtual architecture should take node separation and routing into account. Separation of network segments should be done via secure methods such as static routing and subnetting.

In addition, boundary devices should be configured between network segments, providing IPS and strict rule-base configurations.

Xen networking proved to be a good method of node network separation within Xen based virtualised platforms.

Does the boundary firewall configuration provide enough security?

Some providers allow client access to rule-bases associated with their own nodes. This offers a far better level of granular control for client nodes, providing a suitable rule-base has been applied.

The rule-base should be independently verified if access to the boundary device is not possible. This can be done through a firewall configuration review.

Externally-facing firewalls and those which provide protection between network segments should be configured with adequate Intrusion prevention and Denial of Service capabilities. This should also be verified through an independent firewall configuration review.

Conclusions

In conclusion, Context found that the security concerns related to this emerging technology are real. Although Cloud computing can offer significant benefits, this new set of technologies presents challenges to those wishing to secure it. Cloud computing is a relatively immature technology and experience in securing it is limited given the short time it has existed. Context found serious security flaws, which allowed, in some cases, full compromise of Client's nodes. In total, around half of the tests conducted identified security issues and a quarter of the tests could not be conducted due to contractual restrictions placed upon us by the providers. As a result, further vulnerabilities could exist that could not be tested for. It should also be noted that certain issues can never truly be tested: as the node can be moved around the Cloud by the provider, the surrounding infrastructure, and the security posture of that infrastructure, can change.

The aim of this whitepaper is to raise awareness of the technical security risks associated with migrating or using the Cloud, to provide practical advice on how best to defend against these risks, and to suggest questions that a prospective Cloud client may wish to pose their would-be provider. Context suggests that clients should adopt a specialised set of security policies relating to Cloud environments to ensure that they remain secure when migrating to and using the Cloud. Among other issues, clients should not assume that Cloud providers automatically provide security; in fact, our research indicates that there are more security risks associated with the Cloud than with a traditional dedicated hosting solution.

Contributions

The project was a collaboration between the following members of the Context team:

- Michael Jordon
- Sven Schlüter
- Karl Madden
- James Cormack
- David Oruba

About Context

Context Information Security is an independent security consultancy specialising in both technical security and information assurance services.

The company was founded in 1998. Its client base has grown steadily over the years, thanks in large part to personal recommendations from existing clients who value us as business partners. We believe our success is based on the value our clients place on our product-agnostic, holistic approach; the way we work closely with them to develop a tailored service; and to the independence, integrity and technical skills of our consultants.

The company's client base now includes some of the most prestigious blue chip companies in the world, as well as government organisations.

The best security experts need to bring a broad portfolio of skills to the job, so Context has always sought to recruit staff with extensive business experience as well as technical expertise. Our aim is to provide effective and practical solutions, advice and support: when we report back to clients we always communicate our findings and recommendations in plain terms at a business level as well as in the form of an in-depth technical report.

context
INFORMATION SECURITY LTD

Context Information Security Ltd

London (HQ)
4th Floor
30 Marsh Wall
London E14 9TP
United Kingdom

Cheltenham
Corinth House
117 Bath Road
Cheltenham GL53 7LS
United Kingdom

Düsseldorf
Adersstrasse 28
1. Obergeschoss
D-40215 Düsseldorf
Germany

Melbourne
Level 9, 440 Collins St
Melbourne
Victoria 3000
Australia