Negotiating the Highway to the Cloud

Dr David Ross

Security Practice Manager, CISO
Bridge Point Communications, Australia

David_Ross@bridgepoint.com.au

August 19, 2011

Abstract

Data centre investment dollars are increasingly moving to virtualisation technologies and the cloud – whether that is an in-house private cloud, external provisions, or a hybrid. This stampede to take advantage of the benefits of virtualisation also comes with its own set of headaches for management and system administrators alike. With many refurbishments and new disaster recovery plans utilising virtualisation and cloud computing, it is now essential to have a comprehensive toolset to assist in the secure design, implementation and provisioning of new virtual environments and the assurance of SPI (SaaS, PaaS, IaaS) offerings. The author presents the considerations and foremost advocacy in virtualisation and cloud computing security, including details of the Cloud Security Alliance's (CSA) Certificate of Cloud Security Knowledge (CCSK), along with a full update on the Cloud Security Alliance's Australian activities.

Introduction

The Cloud has received a large amount of press in recent times, not all of it favourable.

Many incidents with public cloud service providers produced notable outages, such as those of Amazon's Elastic Compute Cloud (EC2), taking out FourSquare, HootSuite, Indaba Music, Quora, Reddit, and others [1], just after midnight (PDT) on 21st April this year [2] and another massive outage, in Ireland, with some permanent data loss due to power and UPS failures on 7th August [3]. Microsoft's North American data centre outage on 16th August just past, downed Microsoft Office 365, SkyDrive and CRM Online services [4]. Google's numerous Gmail outages include the more notable ones of 24th February 2009 [5] and 1st September 2009 [6]; and the Gmail "lost emails" incident occurred after a faulty software update between 6:00 PM PST on 27th February and 2:00 PM PST on 28th February this year, that (permanently) rejected mail deliveries during the outage and denied access to all emails, for "some 40,000" of its 193 million user accounts for up to three days [7].

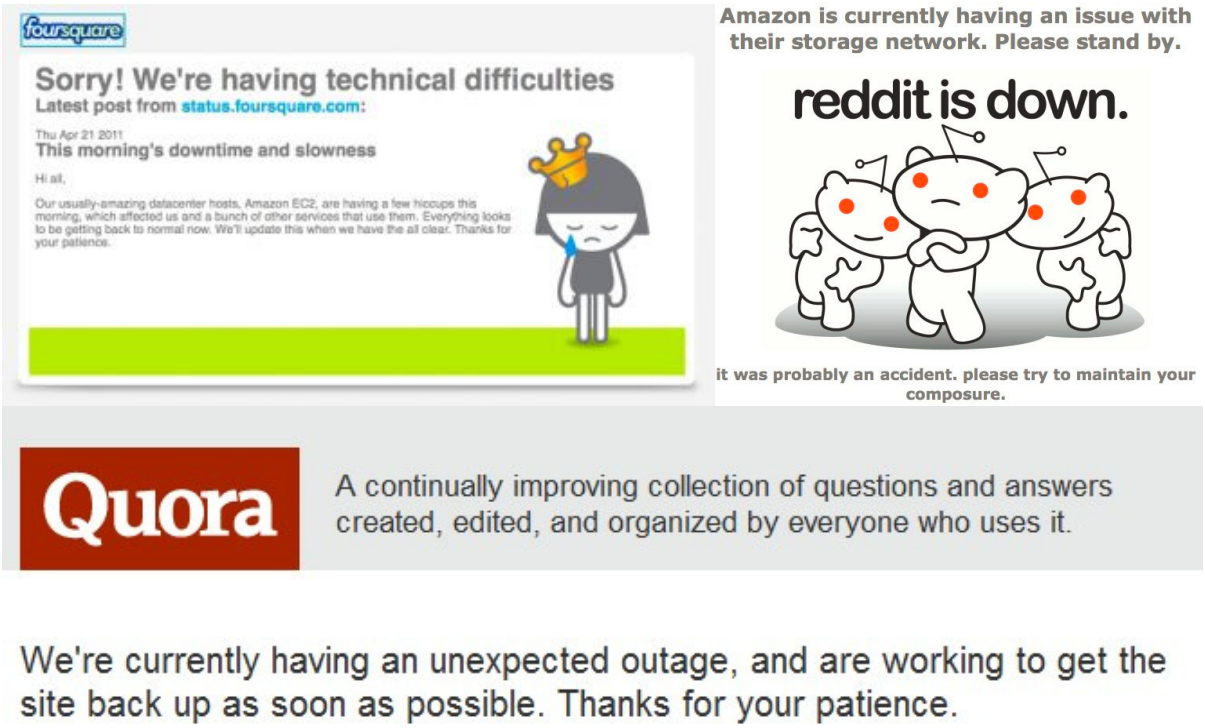

Figure 1: Amazon's EC2 outage took out FourSquare, Quora, Reddit, and others [8-10]

More damaging, however, are the events of the likes of GoGrid cloud hosting and hybrid hosting company advising customers on 30th March this year of a security breach exposing customer details including payment cards [11].

There has also been active use of cloud services as a tool to perform Internet attacks, such as the much publicised use of Amazon's EC2 service to base a penetration of Sony Corporation's online entertainment systems [12]. These are not attacks on "cloud security" as such, but simply security attacks that use cloud resources for the same reasons as any other cloud user. Sony Corporation's PlayStation Network and Qriocity entertainment service exposed 77 million registered users' "names, email addresses, phone numbers, home addresses and user IDs" [13], all of whom were effected "for nearly a month" [14] during the shutdown and recovery; and the Sony Entertainment Online network was also disconnected on 24.6 million users [14], after the discovery that "it is believed credit and debit card details of 24,000 users" [13] were stolen. It has previously been shown [15] that using the cloud to brute force passwords and encryption keys is far more cost-effective than using private resources.

So why use the cloud? The usual answer is that the benefits by far outweigh the risks. But is this really the case? Before you present your business case to your board, you had better make sure you have all the facts for an informed assessment.

Drivers Licence Theory

Before we proceed down the Highway to the Cloud, we first need some basic cartography, including a common terminology and understanding of the concepts.

Over the past 40 years we have moved from a centralised model of computing with single central processors executing programs for multiple users connected by remote terminals over multiplexed serial lines terminal server nodes, through a distributed model where the dumb terminals have been replaced by increasingly more powerful personal computers doing their own processing and exchanging data through the worldwide Internet; and now, as we advance technology further, we move back to the centralised processing model – in the cloud – albeit now a distributed, redundant, centralised processing model, where increasingly powerful personal computers become over-resourced graphics cards to the centralised cloud.

The true origins of the actual term 'cloud computing' are somewhat vague, with clouds being used to abstract telephony and packet networks up to 40 years ago [16-18], the Internet commonly being referred to as 'the cloud' for at least the last ten years or more, and the term 'cloud computing' having appeared in a patent application in 1997 [19]. In more recent times the term was rejuvenated by Google's Executive Chairman (then CEO) Eric Schmidt [19], shortly before Amazon announced it's 'Elastic Compute Cloud' [20].

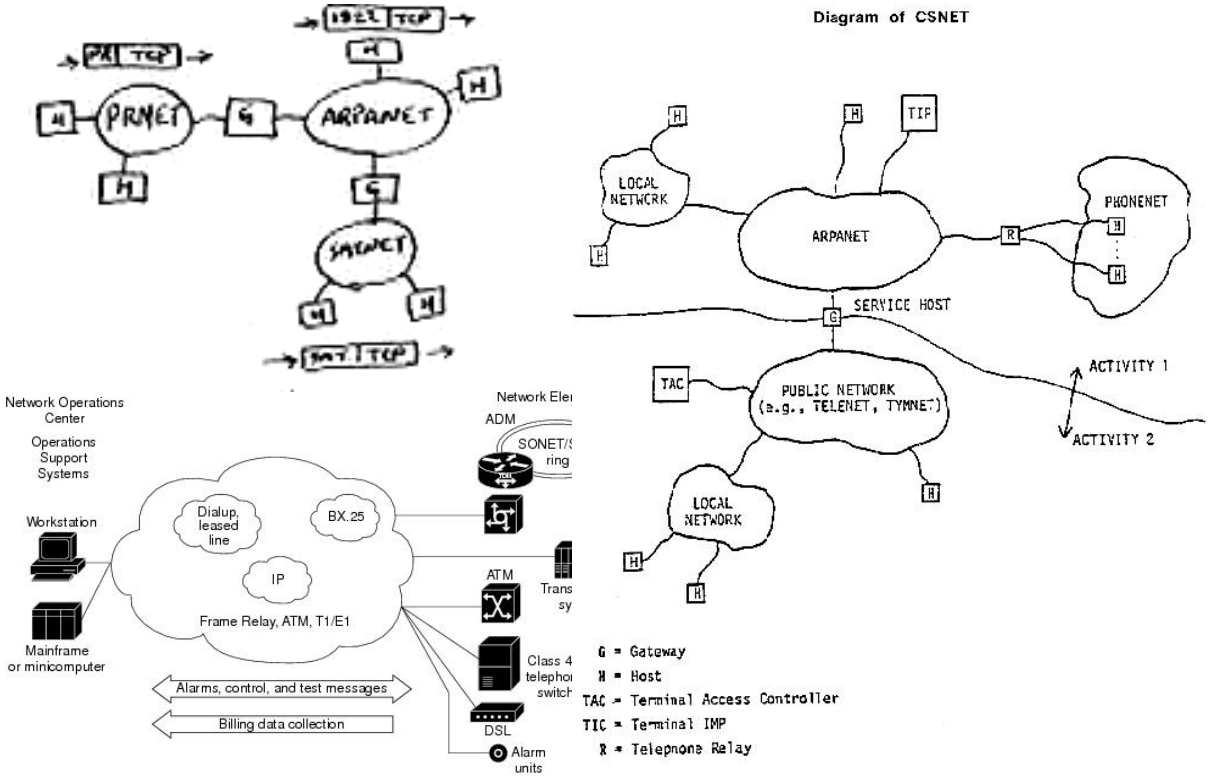

Figure 2: Early Network "Clouds" [16-18]

The Roadmap

The Cloud Security Alliance (CSA) in its *Security Guidance for Critical Areas of Focus in Cloud Computing Version 2.1*, provides a "description of Cloud Computing that is specifically tailored to the unique perspective of IT network and security professionals" [21, p. 13]. It states cloud computing "describes the use of a collection of services, applications, information, and infrastructure comprised of pools of compute, network, information, and storage resources" [21, p. 13], that can be rapidly commissioned, expanded, reallocated, deallocated, and torn down; providing an "on-demand utility-like model of allocation and consumption" [21, p. 13] of resources.

The CSA guidance aligns with the U.S. Department of Commerce's National Institute of Standards and Technology (NIST) Definition of Cloud Computing. On 30th March this year, NIST presented the first version of the NIST Cloud Computing Reference Architecture consistent with the NIST Definition of Cloud Computing [22].

The NIST cloud reference architecture supports **three service models** [22, p. 14]:

- **Software as a Service (SaaS)** – "Deployed applications targeted towards end-user software clients or other programs, and made available via the cloud." The consumer has no control over the network, servers, operating systems, storage, or the applications themselves;
- **Platform as a Service (PaaS)** – "Services for consumers to develop and deploy applications onto the cloud infrastructure, including application containers, application development tools, database management systems, etc." The consumer has no control over the network, servers, operating systems, or storage, but dictates and controls the applications; and
- **Infrastructure as a Service (IaaS)** – "The provisioning of processing, storage, networks, and other fundamental computing resources upon which cloud consumers can deploy and run applications on the cloud infrastructure." The consumer has no control over the hardware infrastructure, but dictates and controls the operating systems, virtual storage, applications, and possibly also some of the virtual networking functions.

The NIST cloud reference architecture also describes **four deployment models** [22, p. 13]:

- **Private Cloud** – The cloud infrastructure (internal, or external by a third party) is solely for one organisation;
- **Community Cloud** – The cloud infrastructure (may also be a third party provider) is shared by several organisations for a specific community of interest, with shared "security requirements, policy, and compliance considerations";
- **Public Cloud** – The cloud infrastructure tenancy is not restricted to any particular organisation(s); and
- **Hybrid Cloud** – The cloud infrastructure is a composition of private, community, or public, that enables data and application portability "(e.g., cloud bursting for load balancing between clouds)."

The NIST cloud reference architecture provides **five essential characteristics** [21, p. 15]:

- **On-demand self-service** – consumers can provision services as required without requiring human interaction with a provider;
- **Broad network access** – services allow access by thin or thick clients or other cloud services;
- **Resource pooling** – provider's resources are pooled and dynamically assigned to serve multiple consumers in a multi-tenant model;
- **Rapid elasticity** – resources can be quickly provisioned and released, possibly appearing unlimited; and
- **Measured service** – resource usage can be "monitored, controlled, and reported — providing transparency for both the provider and consumer of the service."

Roadworthiness

Cloud computing relies at its very heart on virtualisation technologies. It is the virtualisation of hardware, networks, servers, operating systems, platforms, and applications, that form the basis of all cloud computing endeavours. Modern advances in virtualisation technologies have enabled cloud computing as a viable business process. The security of virtualisation technologies is intertwined with the security of cloud computing services at its lowest level.

In its most basic form, virtualisation is in use on almost every computer today, such as in disk partitioning to create virtual disks. There are three main types of virtualisation involved in the provision of cloud services:

- **Storage virtualisation** – Storage resources are pooled and centrally managed to appear as a single elastic set of storage blocks;
- **Network virtualisation** – The resources and functionality of network components are pooled and centrally managed to provide flexible networking options such as quality of service and virtual local area networks (VLANs) over virtual switches, routers, and firewalls; and
- **Server virtualisation** – Server hardware resources are pooled and centrally managed so that they can be allocated to virtual machines to consume resources on an as-needs basis without limiting the virtual machines to a single set of hardware.

Server virtualisation of the lowest layers, the hardware, occurs in various forms relevant to cloud computing:

Full virtualisation, where the virtualisation layer completely emulates a set of hardware for each guest operating system (OS), using binary translation to trap all system calls for the guest (virtual) machine and translate to the appropriate calls to host (native) machine, such that guests that do not require to be aware of the virtualisation and do not require any modification to run as virtual machines; or

Para virtualisation, where the virtualisation layer emulates some of the hardware functions, but each guest is aware of the virtualisation and requires modifications to drivers and system calls to call the virtualisation layer directly to handle such requests rather than require the virtualisation layer to emulate all the hardware at a significant processing overhead and trap and translate such calls.

In both cases, the advances in hardware assisted virtualisation, where the physical processors automatically trap guest calls to the hypervisor Virtual Machine Monitors (VMM) running in a special Root Mode privilege level and provide data structures to store VM states, can be used to supplement full virtualisation or para virtualisation functions.

The virtualisation layer can be provided by two primary types of hypervisor:

Type 1 Hypervisor or **Bare Metal Hypervisor**, which runs on the host (native) machine as the base OS and controls all native activity on the host, either as the sole footprint (e.g. VMware ESXi) or with a hardened general purpose OS providing management capabilities which can create new domains and manage virtual devices and physical hardware such as network interfaces and hard disk controllers (e.g. VMware ESX [+Linux], Xen, XenServer [+Linux], or Microsoft Hyper-V [+Windows Server 2008]); and

Type 2 Hypervisor or **Hosted Hypervisor**, which runs on top of an existing general purpose OS on the host hardware gains the advantage of the most versatile host hardware integration at the expense of having to make all calls via the host OS in addition to all the Type 1 hypervisor tasks and being exposed to all of the host's general purpose OS footprint and vulnerabilities (e.g. VMware Server, VMware Workstation, VMware Player [on Windows or Linux], Parallels Workstation [on Windows or OSX], Microsoft VirtualBox [on Windows], or KVM [on Linux]).

Uncharted Waters

In a virtualised environment, virtualised machines have their own independent operating systems running as if they were instantiated on their own individual physical device, however their calls to the physical hardware instead go to the hypervisor that is controlling all of the physical resources on one or a cluster of physical machines, sharing these resources among any number of virtual machines.

The virtualised environment provides virtual networks that may enable the virtual machines to network with each other or with the actual physical network interfaces on the physical machines. These virtual networks include virtual switches and may include also include other network management functionality.

Many current information security policies and procedures, indeed often entire Information Security Management Systems (ISMS), including the information security risk management methodologies, do not account for the use of virtualisation technology in production environments.

Compliance and Audit

Take for example, the Payment Card Industry Data Security Standard (PCI DSS) [23]. This is of particular relevance to those organisations planning to deploy private clouds that may either contain Cardholder Data (CHD) or in any way impinge on their Cardholder Data Environment (CDE).

The PCI DSS is a highly prescriptive list of controls dictated by the PCI Security Standards Council (PCI SSC), founded by the five card brands: Visa Inc., MasterCard Worldwide, American Express Company, Discover Financial Services, and JCB International [24]. Every merchant and service provider dealing with credit cards must be compliant, including conducting quarterly external and internal vulnerability scans and annual external and internal penetration tests on their systems and heavy users must prove it with an annual audit.

Meeting the security requirements of the PCI DSS requires the implementation of a fixed set of controls with very limited exception – where it can be shown additional controls, beyond those already required by the standard, are in place, effective, and maintained.

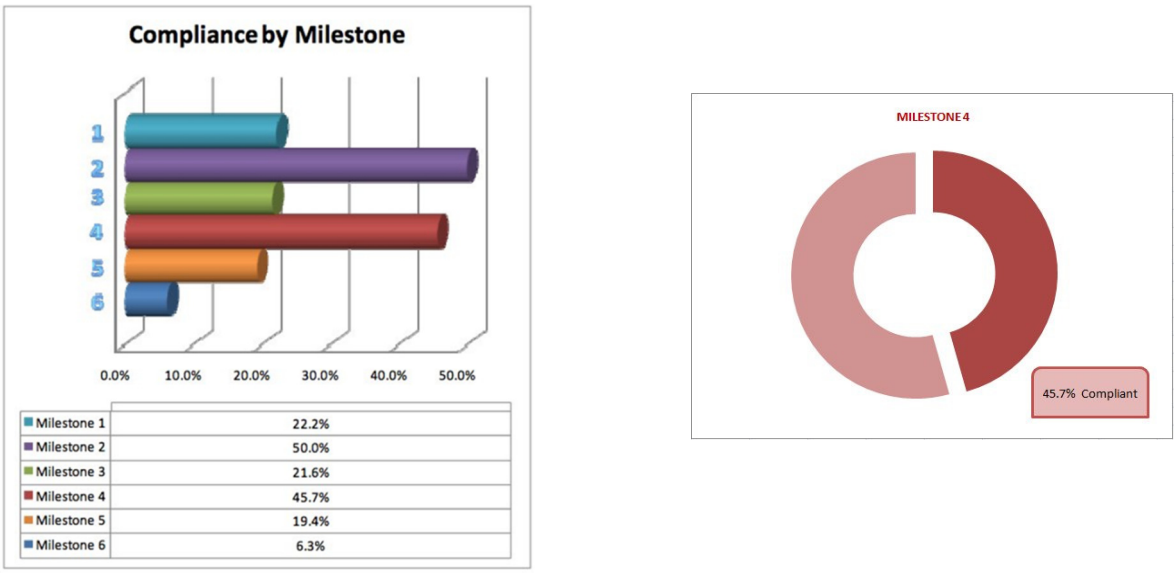

Figure 3: Example Gap Analysis against the PCI DSS v2.0

The PCI community has been long waiting the inclusion of virtualisation in the PCI DSS, with the previous versions leading many PCI QSAs to simply declare virtualised environments **non-compliant**, or at least declare segregation (network segmentation) ineffective and thus make the cardholder data environment (CDE) subject to the most severe restrictions, encompass the entire virtualisation cluster.

It was therefore with great expectation that organisations and QSAs waited for the release of Version 2.0 of the PCI DSS, including virtualisation for the first time.

PCI DSS Version 2.0

The PCI DSS version 2.0 was released on 28 October 2010 and came into effect on 1 January this year. This latest revision provided many clarifications and additional guidance on the previous revision; however, as the standard is becoming more mature, only one of the twelve sections was even moderately rearranged and significantly updated.

The PCI DSS version 2.0 now included references to virtualised environments, as the definition of "system components", previously defined as "any network component, server, or application that is included in or connected to the cardholder data environment" [25], was expanded to "system components also include any virtualisation components such as virtual machines, virtual switches/routers, virtual appliances, virtual applications/desktops, and hypervisors" [23].

The new version of the standard also added that if virtualisation is implemented all components within the virtual environment will need to be identified and considered in scope of the review, including individual virtual hosts or devices, management interfaces, central management consoles, hypervisors, and all intra-host and external communications and data flows [23]. The implementation of a virtualised environment "must meet the intent of all requirements such that the virtualised systems can be effectively regarded as separate hardware" [26, p. 5].

So while the PCI DSS now included the specific burden of securing virtualised environments, unlike all other aspects of the standard, it provided no guidance on the actual controls to achieve this, referring further advice to come from a long awaited report from the *Virtualisation Special Interest Group* (SIG), started at the end of 2008 and initially expected in the last half of 2010, finally released in June 2011.

The *Information Supplement: PCI DSS Virtualization Guidelines* from the *Virtualisation SIG* starts out by defining virtualisation and associated terminology, before proceeding with scoping guidance for PCI DSS impacts on the extent of the CDE, to the effect that if any VM, virtual switch, network, or virtual security device, is in-scope, then the associated hypervisor and all possible physical hardware are automatically in-scope for PCI DSS assessment [27].

Scope guidance on 'cloud computing' mirrors that of any hosted CDE components: "The cloud provider should clearly identify which PCI DSS requirements, system components, and services are covered by the cloud provider's PCI DSS compliance program. Any aspects of the service not covered by the cloud provider should be identified, and it should be clearly documented in the service agreement that these aspects, system components, and PCI DSS requirements are the responsibility of the hosted entity to manage and assess. The cloud provider should provide sufficient evidence and assurance that all processes and components under their control are PCI DSS compliant." [27, p. 9].

The PCI DSS virtualization guidelines make special mention of the new attack surface introduced by a hypervisor. This is of particular relevance for Type 1 hypervisors that use a general purpose OS for management, particularly those hypervisors without a hardened OS;

and especially Type 2 (hosted) hypervisors. "Hypervisors are not created equal, and it is particularly important to choose a solution that supports the required security functions for each environment." [27, p. 10].

The guidelines then detail considerations that must be taken into account regarding mixing VMs of different trust levels, VM-to-VM and VM-to-hypervisor attacks, the loss of separation of network, system administration, and audit duties, dormant VMs and the security of images and snapshots, but leaves it entirely to your selected QSA as to what meets compliance in a virtualised environment.

Some QSAs will require a separate virtualised cluster for the CDE using separate physical networks controlled by separate physical firewall from other internal virtualised clusters and a separate physical DMZ virtualised cluster, again with separate physical network and separate physical perimeter firewall out to the border routers. Other QSAs may well sign off on a single encompassing virtualised cluster including CDE, DMZ, firewalls, and other internal hosts.

The rest of us lie anywhere on the range in between. Note that in view of yet another hypervisor breach, breaking out of KVM [28], presented at the Black Hat USA 2011 and Defcon 19 conference in August, most QSAs will reject Type 2 hypervisors in a mixed environment.

Navigating a Steady Course

An overarching consideration in assessing an organisation for any sort of information security, is just what constitutes the sensitive data. This is a primal concern in determining not only the scope but the controls required and their locations.

The hypervisor

It must be remembered, that since the hypervisor affects the security of its virtualised guests, if any virtualised guest system contains sensitive data, then that automatically means the hypervisor controlling that guest is automatically part of the control environment. As such all of the physical machines constituting that cluster controlled by that hypervisor are subject to all the physical requirements of machines relative to the most sensitive data on any VM.

Questions that arise as to the issues invoked by moving virtual machines from one physical machine to another and whether or not other virtual machines are then also part of the sensitive data environment – or can such machines be effectively segregated?

The PCI DSS requirements

The PCI DSS has six goals that drive 12 requirements. These 12 requirements dictate 196 control components, with another six for managed service providers, which in an audit requires 989 items of evidence to be verified and documented.

Requirement 1.1.3 requires a firewall at each Internet connection and between any DMZ and the internal networks.

In a virtualised environment, providing a firewall between the DMZ and the internal networks can be achieved by deploying a separate physical virtualised cluster for the DMZ separated by physical firewall from the internal virtualised cluster. This splits the physical hardware and thus reduces the efficiency of the virtualisation.

There is a further restriction under requirement 1.2.1 to use firewall and router configurations that restrict inbound and outbound traffic to and from the cardholder data environment. Where the cardholder data environment is segregated from the rest of the internal networks, this requires segregation within the internal networks. Under this scheme, this would dictate yet another separate physical cluster so that routers and firewalls can provide the segregation from the rest of the internal cluster. This drastically reduces the efficiency and utility of a virtualised environment.

Another option relies on the hypervisor to provide the segregation between hosts by tying these groups of virtual hosts to separate physical network interface cards (NIC) so that a separate physical firewall can be used to guarantee this segregation.

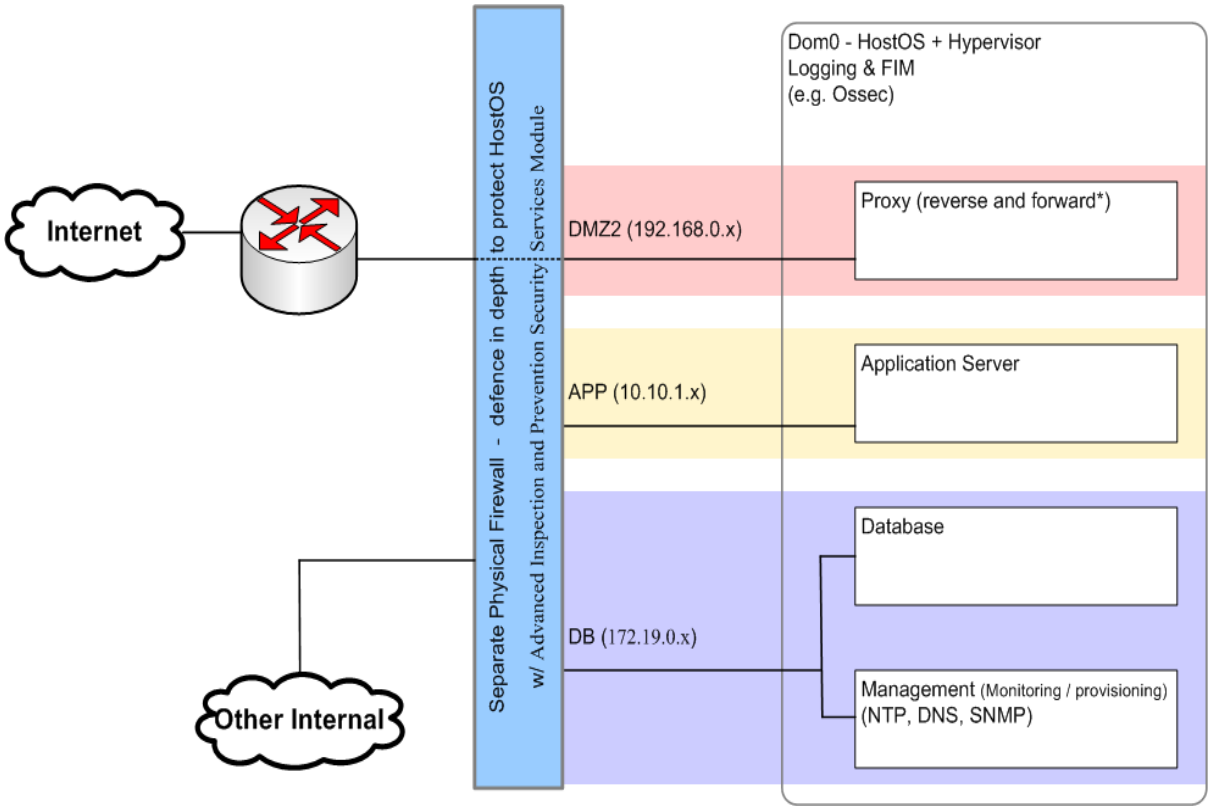

Figure 4 : Example of Using Separate Physical Firewall(s)

While this helps maintain the utility of the internal virtualised cluster, this still limits the number of externally controlled segments and the utility of the network interface cards. Current technology is moving to maximise physical network throughput by teaming all physical interfaces together and providing segregation at the logical level.

Virtual Alternative

An alternative involves using a single enterprise virtualised cluster and deploying virtual firewalls within the cluster two separate environments. However it is up to your QSA, and not you, to declare whether the technology you use meets the PCI DSS requirements for segregation of networks. Many QSA is would not consider this to achieve effective segregation.

If you use para-virtualisation, that is a hypervisor that runs in a general-purpose operating system such as Microsoft windows, UNIX, or UNIX-like operating systems, then the hypervisor is subject to the security vulnerabilities of the underlying operating system and thus can be subverted by exploits on the underlying operating system, in turn subverting the security of the guest virtual machines as well is any segregation between them.

The use of so-called “bare metal” virtualisation technology exposes a vastly reduced attack surface, making such “virtual segregation” more palatable to your QSA.

Virtual Firewalls

In addition, virtual firewalls also come in two forms: a firewall running as a virtualised guest and so competing with all the other virtualised guest machines for resources physical resources; or for a new breed of virtualised systems where firewalls, intrusion detection, antivirus, and the like, are provided with hooks into the hypervisor itself guaranteeing true access to the data flows in and out of the virtualised guest machines.

Intrusion detection

Requirements for intrusion detection state that all traffic in the CDE must be monitored. In a virtualised environment this is a requires either requires host IDS (HIDS) on every single virtualised guest machine (effectively resulting in multiple instances of the same code running simultaneously on the physical hardware) or network IDS (NIDS) patched into all of the virtual networks within the virtual cluster, monitoring the traffic between the virtual machines and to and from the outside world. Where such a system runs as a guest in the virtualisation environment, it must compete with the other virtual machines for physical resources.

Again, the latest virtualisation technology allows IDS vendors to provide solutions that hook into the actual hypervisor to guarantee access to all of the data flows to and from the virtual machines.

Antivirus

The PCI DSS requires antivirus software be deployed on all systems commonly affected by viruses, including servers. In a separate physical environment, this requires an instance of the antivirus software on every machine. Similarly, in a virtualised environment every machine must have antivirus. Where an instance of the antivirus software is running on each virtual guest machine, this results in many instances of the same software running simultaneously within the physical cluster.

By using the latest virtualisation technologies and deploying an antivirus solution that hooks in to the hypervisor itself, thus monitoring all data flowing into all of the virtual machines, to detect and action viruses malicious code in transit to the virtual machine, as well as scanning the virtual machine image as it is loaded or swapped, may meet these PCI requirements and substantially reduce the computing overheads involved with antivirus in a virtualised environment.

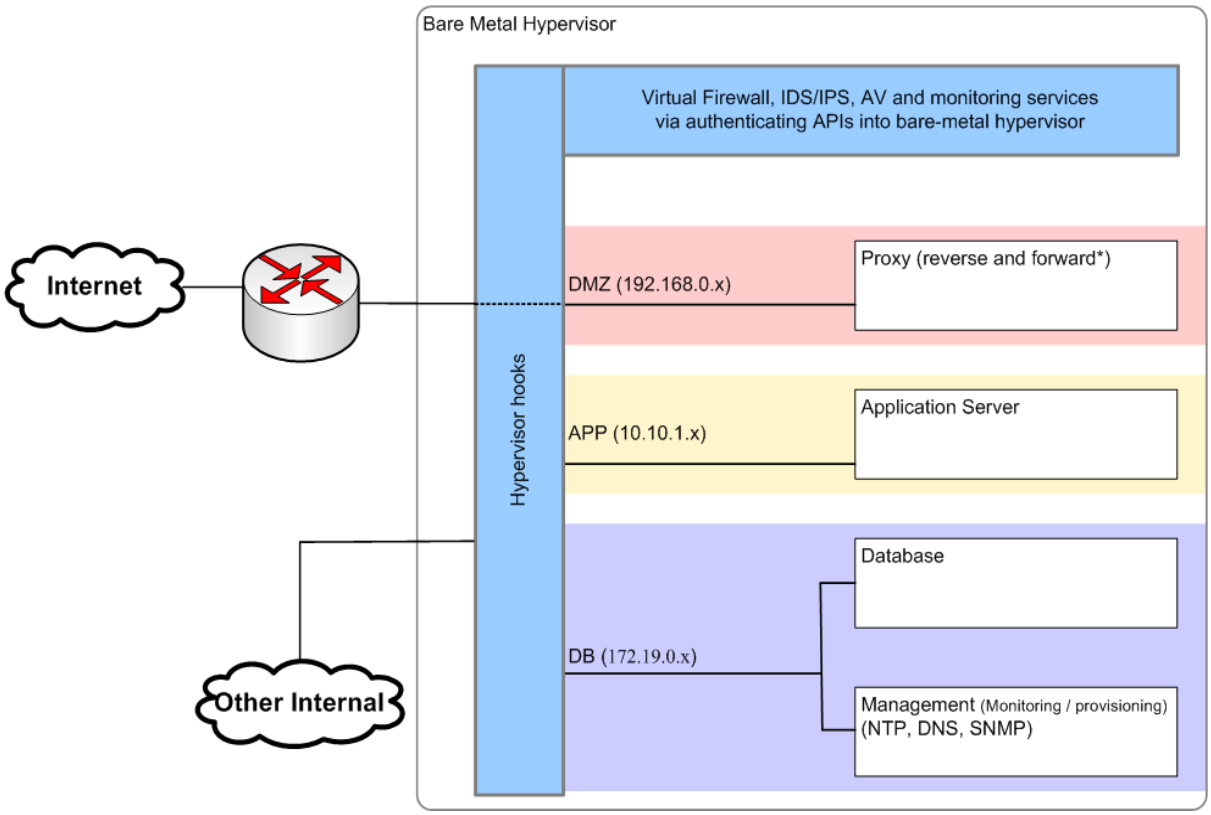

Figure 5: Example of a Bare-Metal Hypervisor providing hooks for Firewalls, IDS & AV

Conclusion

The security and auditing of virtualisation environments is now a pressing need in both government and commercial environments. In the commercial world, and in particular the Payment Card Industry, virtualisation is a leading topic of research and debate.

New and emerging virtualisation technologies, although untested from a compliance and audit perspective, now have the potential to provide tangible benefits not only in power, space, and cost efficiencies, but now provide increased abilities to secure and monitor Information Systems.

For those who are subject to the PCI DSS, you should engage and work with your QSA at the earliest opportunity to explore the opportunities and advantages you can leverage from virtualisation – *before* you commit your resources and design and build your systems.

Cloud Security Alliance, Australia Chapter

In February 2011, Gary Blair, Executive General Manager, CBA (Sydney) made enquires with Jim Reavis, Executive Director, Cloud Security Alliance (USA) in regards to an Australian Regional Coordinating Body and State Chapters.

On 7th April 2011, David Ross, CISO, Bridge Point Communications (Brisbane) with 24 Queensland members of the CSA LinkedIn group, one QLD/WA member, and one NSW/ACT member applied to start CSA Australia and expand it throughout Australia. This application required a greater breadth of membership to proceed.

On 29th April 2011, Jim Reavis (CSA) invited all interested parties who had contacted him from around Australia to get together and advance the application.

On 5th May 2011, Gary Blair hosted a meeting via teleconference, chaired by David Ross and attended also by Jim Reavis (CSA) and those aforementioned interested parties. This meeting decided to proceed with a single overall Australian chapter, with a structure to be reviewed after six months. Nominations for roles on the Board of Directors were called.

On 16th May 2011, A CSA Australia BoF session was announced at all AusCERT conference sessions. This meeting concluded the nominations and the Board of Directors was finalised soon after.

David Ross sent a refreshed application supported by 58 Australian members of the CSA LinkedIn group to Jim Reeves and on the 21st June 2011 the ***Cloud Security Alliance, Australia Chapter*** was approved for development.

The Founding Directors of CSA Australia are:

- Ben Chung (HP – NSW),
- Gary Gardiner (Check Point – QLD),
- Craig Lawson (HP – QLD),
- Wipul Jayawickrama (Infoshield – QLD),
- Richard Keirstead (Ernst & Young – VIC),
- Phil Kernick (CQR Consulting – SA),
- Archie Reed (HP – NSW),
- David Ross (Bridge Point – QLD),
- Darren Skidmore (FIS Australasia – VIC),
- Tim Smith (Bridge Point – QLD),
- Marcel Sorouni (BUPA Australia – NSW),
- Michael Trott (Bridge Point – QLD),
- Chad Walker (Infoshield – QLD),
- Marcus Wong (CBA – NSW),
- Jason Wood (CBA – NSW).

The name of the association is the **Cloud Security Alliance, Australian Chapter**, hereinafter called **CSA Australia** and abbreviated as **CSA-AU**.

The purposes for which the association is established are, in line with CSA Global, to promote the use of best practices for providing security assurance within Cloud Computing, and provide education on the uses of Cloud Computing to help secure all other forms of computing.

The geographical boundary of CSA, Australian Chapter, until otherwise amended, is the Commonwealth of Australia. The geographical boundary does not affect membership in CSA Australia, but simply defines the service area.

The membership of CSA Australia shall consist of all current members of the global Cloud Security Alliance LinkedIn group who have elected to join the chapter's official LinkedIn sub-group within CSA's LinkedIn group.

The **CSA** (global) LinkedIn group is at: http://www.linkedin.com/groups?gid=1864210

The **CSA Australia** LinkedIn group is at: http://www.linkedin.com/groups?gid=3966724

The Cloud Security Alliance's *Certificate of Cloud Security Knowledge* (CCSK) is the industry's first user certification program for secure cloud computing. "The CCSK is designed to ensure that a broad range of professionals with responsibility related to cloud computing have a demonstrated awareness of the security threats and best practices for securing the cloud" [29].

The CCSK provides evidence that an individual has successfully completed an examination covering the key concepts of the CSA *Security Guidance for Critical Areas of Focus in Cloud Computing, V2.1* and the European Network and Information Security Agency (ENISA) whitepaper "*Cloud Computing: Benefits, Risks and Recommendations for Information Security*". Jim Reavis, CSA executive director said "The CCSK is a low cost certification that establishes a robust baseline of cloud security knowledge. Combined with existing professional certifications, it helps provide necessary assurance of user competency in this important area of growth" [29].

Bibliography

[1] BBC. *Amazon fault takes down websites*. BBC News Technology, 21st April, 2011. [Online] Available: http://www.bbc.co.uk/news/technology-13160929

[2] Amazon Web Services. *Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region*. 21st April, 2011. Amazon Web Services LLC. [Online] Available: http://aws.amazon.com/message/65648/

[3] Charles Babcock. *Amazon Issues Post Mortem On Cloud Outage*. InformationWeek, 15th August, 2011. UBM TechWeb, UBM LLC. [Online] Available: http://www.informationweek.com/articles/231400281

[4] Lee Pender. *Another Bad Day for the Cloud: Microsoft Office 365, Dynamics CRM Go Down*. Redmond Channel Partner, August 18, 2011. 1105 Media Inc., Chatsworth, CA, USA. Posted August 18, 2011 at 1:35 PM. [Online] Available: http://rcpmag.com/blogs/lee-pender/2011/08/microsoft-office-365-and-dynamics-crm-go-down.aspx

[5] Tim Beyers. *Gmail Fails, Cloud Computing Wins?*. The Motley Fool. 24th February, 2009. [Online] Available: http://www.fool.com/investing/high-growth/2009/02/24/gmail-fails-cloud-computing-wins.aspx

[6] Erick Schonfeld. *Why Gmail Failed Today*. In TechCrunch. 1st September, 2009. AOL Inc. [Online] Available: http://techcrunch.com/2009/09/01/why-gmail-failed-today/

[7] Robert Charette. *Google Gmail's "Lost Emails" Restored to Life*. In The Risk Factor, IEEE Spectrum, 1st March 2011. Institute of Electrical and Electronics Engineers, Inc. [Online] Available: http://spectrum.ieee.org/riskfactor/telecom/internet/google-gmail-lost-emails-restored-to-life

[8] AP Photo. *foursquare.com*. In FoxNews.com, Downed Amazon Servers Leave Websites Across Internet Struggling. 22nd April, 2011. FOX News Network, LLC. [Online] Available: http://a57.foxnews.com/static/managed/img/Scitech/396/223/Amazon%20Outages%20take%20out%20Foursquare.jpg

[9] Leon Katsnelson. *reddit_down-due-to-AWS.jpg*. In BigDataOnCloud + FreeDB2, When cloud fails. 17th March, 2011. [Online] Available: http://freedb2.com/wp-content/uploads/2011/03/reddit_down-due-to-AWS.jpg

[10] Charles Arthur. *The message greeting visitors to Quora during the outage*. In guardian.co.uk, 21st April, 2011. Guardian News and Media Limited. [Online] Available: http://static.guim.co.uk/sys-images/Technology/Pix/pictures/2011/4/21/1303394021477/Quora-outage-007.jpg

[11] Craig Balding. *GoGrid Security Breach*. 30th March, 2011. [Online] Available: http://cloudsecurity.org/blog/2011/03/30/gogrid-security-breach.html

[12] Pavel Alpeyev, Joseph Galante and Mariko Yasu. *Amazon.com Server Said to Have Been Used in Sony Attack*. In Bloomberg News, 15th May, 2011. Bloomberg L.P. [Online] Available: http://www.bloomberg.com/news/2011-05-13/sony-network-said-to-have-been-invaded-by-hackers-using-amazon-com-server.html

[13] John Kennedy. *Amazon EC2 service used in Sony attack?*. In Siliconrepublic, Stragegy, 14th May, 2011. Siliconrepublic Publishing Ltd. [Online] Available: http://www.siliconrepublic.com/strategy/item/21798-amazon-ec2-service-used-in

[14] Paul Wagenseil. *Amazon's Cloud Servers Possibly Used in Sony Attack*. In SecurityNewsDaily, 16th May, 2011. MSNBC Interactive News LLC. [Online] Available: http://www.msnbc.msn.com/id/43054841/ns/technology_and_science-security/t/amazons-cloud-servers-possibly-used-sony-attack/#.Tk4TQF3pcYo

[15] Dan Goodin. *PlayStation Network hack launched from Amazon EC2 (Cloud economics strikes again)*. In The Register, Cloud Business, 14th May, 2011. Situation Publishing. [Online] Available: http://www.theregister.co.uk/2011/05/14/playstation_network_attack_from_amazon/

[16] Vinton Cerf. Untitled network diagram 1973. In Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins Of The Internet*. Simon&Schuster, 1998. ISBN 0684832674.

[17] Takeshi Utsumi. *Diagram of CSNET*. c.1981. In Electronic Global University System and Services (Unpublished draft July, 1998). [Online] Available: http://www.friends-partners.org/utsumi/bookwriting/part_i/chapter_i/total/Insertions/NSF/CSNET/CSNET.jpg

[18] Cisco Systems, Inc. *Multiple networks are included in the DCN network cloud*. In Cisco Network Solutions for the Telco DCN: Telephone Switch Environments, January, 2008. [Online] Available: http://www.cisco.com/en/US/i/000001-100000/80001-85000/82001-83000/82354.jpg

[19] John M. Willis. *Who Coined The Phrase Cloud Computing?*. 31st December, 2008. [Online] Available: http://www.johnmwillis.com/cloud-computing/who-coined-the-phrase-cloud-computing/

[20] Reuven Cohen. Response to John Willis. *Who Coined The Phrase Cloud Computing?*. 31st December, 2008. [Online] Available: http://www.johnmwillis.com/cloud-computing/who-coined-the-phrase-cloud-computing/

[21] Glenn Brunette and Rich Mogull, editors. *Security Guidance for Critical Areas of Focus in Cloud Computing Version 2.1*. The Cloud Security Alliance. December, 2009. [Online] Available: https://cloudsecurityalliance.org/csaguide.pdf

[22] National Institute of Standards and Technology. *NIST Cloud Computing Reference Architecture Version 1*. U.S. Department of Commerce, 30th March, 2011.

[23] PCI Security Standards Council LLC. *Payment Card Industry (PCI) Data Security Standard, Requirements and Security Assessment Procedures, Version 2.0*. Wakefield, MA, USA, October 2010.

[24] PCI Security Standards Council LLC. *PCI DSS Quick Reference Guide, Understanding the Payment Card Industry Data Security Standard version 2.0*. Wakefield, MA, USA, October 2010.

[25] PCI Security Standards Council LLC. *Payment Card Industry (PCI) Data Security Standard, Requirements and Security Assessment Procedures, Version 2.0*. Wakefield, MA, USA, October 2010.

[26] PCI Security Standards Council LLC. *Payment Card Industry (PCI) Data Security Standard, Navigating PCI DSS, Understanding the Intent of the Requirements, Version 2.0*. Wakefield, MA, USA, October 2010.

[27] PCI Security Standards Council Virtualization SIG. *Information Supplement: PCI DSS Virtualization Guidelines*. PCI Security Standards Council LLC, Wakefield, MA, USA. June 2011.

[28] Nelson Elhage. *Virtunoid: A KVM Guest -> Host privilege escalation exploit*. In Proc. Black Hat USA 2011, August, 2008.

[29] Cloud Security Alliance. *Cloud Security Alliance announces availability of Certificate of Cloud Security Knowledge (CCSK)*. September, 2010. [Online] Available: https://cloudsecurityalliance.org/csa-news/cloud-security-alliance-announces-availability-of-certificate-of-cloud-security-knowledge-ccsk/