cloud security alliance

Security Considerations for Private vs. Public Clouds

Larry Hughes

May 2015

Sponsored by

paloalto NETWORKS®

© 2015 Cloud Security Alliance – All Rights Reserved

All rights reserved. You may download, store, display on your computer, view, print, and link to the Cloud Security Alliance "Cloud Adoptions Practices & Priorities Survey Report" at https://cloudsecurityalliance.org/research/surveys/, subject to the following: (a) the Report may be used solely for your personal, informational, non-commercial use; (b) the Report may not be modified or altered in any way; (c) the Report may not be redistributed; and (d) the trademark, copyright or other notices may not be removed. You may quote portions of the Report as permitted by the Fair Use provisions of the United States Copyright Act, provided that you attribute the portions to the Cloud Security Alliance "Cloud Adoptions Practices & Priorities Survey Report" (2015).

Table of Contents

© 2015 Cloud Security Alliance – All Rights Reserved.

1. Introduction

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services). It is a disruptive technology that has the potential to enhance collaboration, agility, scaling and availability, and provides opportunities for cost reduction through optimized and efficient computing. The cloud model envisages a world where components can be rapidly orchestrated, provisioned, implemented and decommissioned, and scaled up or down to provide an on-demand utility-like model of allocation and consumption.

The cloud movement presents a rare and momentous opportunity to revisit not only how we think about computing, but also how we think about information security. Wise businesses will take the long view and invest in security accordingly. As Thomas Edison once said, "Opportunity is missed by most people because it is dressed in overalls and looks like work."

1.1. Purpose and Scope

The purpose of this paper is to present security-related information which companies should consider when deciding whether to pursue public or private cloud deployment. It is not intended to debate which is more secure because neither inherently is (see the sections on "Virtual Attack Surfaces" and "Compliance and Audit" for explanation why). Besides, as all information security professionals know, nothing about security is ever black or white.

Note that for conciseness, a simplifying assumption is made that public and private cloud models are non-overlapping. In other words, it is assumed that a private cloud's entire infrastructure (e.g., servers, storage, network) is owned, operated and physically housed by the tenant business itself, generally its own IT infrastructure organization. Likewise it is assumed that a public cloud's entire infrastructure is owned, operated and physically housed by an independent Cloud Service Provider (CSP).

1.2. Audience

The intended audience of this document is senior IT management and IT security management. Others can benefit from reading it as well.

1.3. Definitions

Cloud Security Alliance's Guidance features definitions and related subject matter that are based on published work of the scientists at the U.S. National Institute of Standards and Technology (NIST) and their efforts around defining cloud computing. NIST defines cloud computing by describing five essential characteristics, three cloud service models, and four cloud deployment models. They are summarized below.

© 2015 Cloud Security Alliance – All Rights Reserved.

1.3.1. Essential Characteristics

- **On-demand self-service**. The tenant is able to self-provision computing resources, (e.g., server, networking, storage) without requiring human interaction
- **Broad network access**. The tenant is able to access the computing resources over the network through the use of laptops, desktops, mobile devices, servers, etc.
- **Resource pooling**. Computing resources are pooled to serve one or more tenants
- **Rapid elasticity**. Computing resources are elastically provisioned and released commensurate with demand
- **Measured service**. Use of computing resources is monitored, controlled and reported

1.3.2. Deployment Models

- **Private cloud**. The resources are open to a single tenant. They may exist on or off premises, and be owned, managed, and/or operated by the tenant, a CSP, or a combination thereof.
- **Community cloud**. The resources are open to a community of tenants with shared concerns (e.g., mission, security requirements, policy, compliance considerations). They may exist on or off premises and be owned, managed and/or operated by any or all of the tenants, a CSP or a combination thereof.
- **Public cloud**. The resources are open to the general public through a CSP. They exist on the premises of the CSP. They may be owned, managed and/or operated by a business, academic or government institution, or a combination thereof.

1.3.3. Service Models

- **Infrastructure as a Service (IaaS)**. The capability is to provision processing, storage, networks and other resources where the tenant is able to deploy and run arbitrary software, which can include operating systems and applications. The tenant does not manage or control the underlying cloud infrastructure but has control over operating systems, storage and deployed applications, and possibly limited control of select networking components (e.g., host firewalls).
- **Platform as a Service (PaaS)**. The capability is to deploy onto the cloud infrastructure tenant-created or acquired applications created using programming languages, libraries, services and tools supported by the CSP. The tenant does not manage or control the underlying cloud infrastructure, but has control over the deployed applications, and possibly configuration settings for the application-hosting environment.
- **Software as a Service (SaaS)**. The capability is to use the CSP's applications running on the cloud infrastructure. The applications are accessible from various tenant devices through either a thin tenant interface, such as a web browser (e.g., web-based email) or a program interface. The tenant does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

© 2015 Cloud Security Alliance – All Rights Reserved.

2. Business and Legal

When evaluating private vs. public cloud deployment models, there is an abundance of business, legal and technical considerations to be given. Early on it is wise to establish principal business and legal feasibility before investing too heavily in technical requirements gathering. Business and legal topics to consider include contracts, service level agreements, roles and responsibilities, and compliance and auditing. Those topics are discussed in this section, and the technically oriented topics in subsequent sections.

2.1. Contracts

Legally enforceable contracts are generally not warranted for private clouds since, in effect, the tenant is the CSP. They are, however, a necessary ingredient with a public CSP. To the extent a tenant wants its public CSP to be lawfully obligated to fulfill the security terms of a contract, it will pursue crystal clarity around those obligations. To accomplish this it must painstakingly enumerate those security obligations and eradicate assumptions about them.

One method is to have each individual stakeholder — business, technology and legal — think long and hard to make a list of everything pertaining to security that can have a detrimental impact to success. Downtime due to a DoS attack? Compromise of personal information due to a breach? Responding to a court order for e-discovery? Allowing data subject to privacy regulations to leak across jurisdictional boundaries? Compiling the lists is likely to identify most if not all security matters that need to be explicitly addressed in the contract.

2.2. Formal Service Level Agreements (SLAs)

A contractual security SLA is not going to apply in a private cloud as the CSP is the tenant itself. It is naturally required with a public CSP as there are many nuances to consider in a public cloud environment. They include:

- Frequency of vulnerability scanning
- Change management procedures
- Gathering of forensic evidence
- Notices of breach
- Applying patches and updates
- Anti-virus
- Data backup
- Data encryption

These and any others should be clearly enumerated, with roles and responsibilities explicitly assigned across the tenant and the CSP.

© 2015 Cloud Security Alliance – All Rights Reserved.

2.3. Roles and Responsibilities

In cloud computing there is a wide range of security-specific roles and responsibilities that must be instituted. There are too many to enumerate here, but suffice it to say that in a private cloud, the tenant will shoulder them all.

In a public cloud, roles and responsibilities will of course be divided between the tenant and the CSP. One way to begin thinking about the division is to apply a RACI matrix to the six services of the IaaS stack. RACI stands for Responsible ("the doer"), Accountable ("the buck stops here"), Consulted ("in the loop"), and Informed ("keep in the picture"). From an operational perspective, a security RACI matrix might look something like this:

IaaS Service	Tenant	CSP
Physical	C, I	R, A
Compute	C, I	R, A
Data	R, A	C, I
Network	C, I	R, A
Storage	C, I	R, A
Application	R, A	C, I

Bottom line, however, it must be understood that from a compliance perspective, the tenant is the party ultimately accountable for keeping its data secure.

In both the private and public deployment models, once roles and responsibilities are determined they must be meticulously documented, and all of the actors involved must have a thorough understanding of them.

2.4. Compliance and Audit

At this point in time there are few laws and regulations that specifically address cloud computing. They are and will remain for some time works-in-progress. What is clear, however, is that enterprises utilizing clouds are not "off the hook" — they must still demonstrate compliance with standards and statutes. In a public cloud, the tenant and CSP must agree ahead of time which common certification assurance frameworks will be used.

It is tempting to assume that it is easier to demonstrate compliance when using a private cloud. This is not necessarily the case. For example, a PCI audit encompasses all networks and systems that handle cardholder data. In many legacy environments this is known to be an extremely onerous exercise. Deploying cardholder environments in a public cloud can dramatically reduce the surface area and volume that needs to be audited. The tenant would rely on attestations from the CSP that its cloud infrastructure meets the tenant's auditing requirements. The CSP needs to undergo its own audits in order to make its attestations honorably.

© 2015 Cloud Security Alliance – All Rights Reserved.

3. Physical Attack Surfaces

An attack surface is the sum of the vulnerabilities that are accessible to would-be attackers. From the IaaS perspective, attack surfaces can be viewed as either physical or virtual. Physical attack surfaces are discussed in this section, virtual attack surfaces in Section 4.

Surface	Physical	Virtual
Humans	✓	
Datacenters	✓	
Servers and Storage	✓	✓
Networks	✓	✓
Operating Systems		✓
Hypervisor / hosting operating systems		✓
Cloud management consoles, tools & APIs		✓

3.1. Humans

A 2015 survey by Vormetric found that senior management at 89 percent of companies polled across the globe feel their organization is vulnerable to an insider attack, one-third of them "very" or "extremely" so[1].

In both private and public cloud deployments, it behooves tenants to implement best practices to ameliorate this risk. They include conducting background checks on new hires, requiring them to sign NDAs, providing thorough and periodic security training and awareness programs, promptly modifying privileges when an employee undertakes a new role, and promptly terminating physical and digital access when an employee departs. In public cloud deployments, it also behooves tenants to demand extra diligence from the CSP since the CSP's employees execute duties in an environment with a great many tenants; any errors they commit could have widespread impact.

3.2. Datacenters

In a private cloud, for physical access to datacenters, it is wise for tenants to meet or exceed the standards set forth in ISO/IEC 27002 or similar. In a public cloud, tenants should require this of their CSPs. It will encompass the topics of:

- Physical security perimeter (e.g., fences, concrete walls)
- Physical entry controls (e.g., mantraps, biometrics)
- Security offices, rooms and facilities (e.g., physical keys, cardkeys)
- External and environmental threats (e.g., earthquake proofing)
- Working in secure areas (e.g., surveillance, supervision)
- Delivery and loading areas (e.g., security personnel)

[1] http://www.vormetric.com/campaigns/insiderthreat/2015/

© 2015 Cloud Security Alliance – All Rights Reserved.

3.3. Server, Storage and Network Devices

NIST has published well-established procedures for decommissioning hardware in a legacy infrastructure[2]. They include wiping of long-lived storage media.[3] In a cloud environment, be it private or public, where storage is virtual, a CSP will often achieve redundancy by stripping data (in some cases down to the virtual block level) across datacenters. This can become problematic if even the slightest amount of data crosses a jurisdictional boundary since many countries have vigorously-enforced laws about the exportation of personally identifying information (PII), even if it is encrypted. A tenant of a public cloud must require that its CSP have adequate controls in place to prevent this from occurring. Since it will not be in a position to investigate those controls by itself, the tenant will need to rely on an attestation from the CSP that they are being correctly implemented.

4. Virtual Attack Surfaces

Cloud networks face the same security challenges as legacy enterprise networks, and new ones as well. That said, it would be a fallacy to say that private clouds are inherently more secure than public clouds. That would be the same as saying that all CSPs provide inherently weaker controls to tenants than tenants provide to themselves. Judging by all-too-common headlines about security breaches, this is not a safe conclusion to draw. Some cloud providers actually offer better security features than some tenants implement for themselves. Furthermore, it is (and always will be) tempting for even the most security-conscious enterprises to take shortcuts on security when they physically possess servers in their private datacenters. The best course of action for any enterprise is to continually pursue rock-solid security in their cloud, independent of whether the cloud is a private or a public one.

4.1. Virtual Networks

Virtual network security is a complex topic for which there are many subtopics that merit separate discussion. The salient point behind them all, however, is that in cloud computing there exists a corresponding virtual “form factor” for every physical networking component that pertains to security: switches (layer 2), routers (layer 3) and firewall appliances (layers 2 – 7).

4.1.1. Switches

Any network can accrue some security benefit by implementing layer 2 switching, which is used to establish and isolate broadcast domains. In a cloud environment, the risk is that compromising a hypervisor can have the net effect of compromising a multitude of physical layer 2 switches in one fell swoop, one for each virtual machine (VM) executing under the hypervisor’s control.

[2] http://csrc.nist.gov/publications/nistpubs/800-123/SP800-123.pdf
[3] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf

© 2015 Cloud Security Alliance – All Rights Reserved.

If this occurs in a private cloud deployment, the sole tenant alone will suffer the consequences. However, in a public cloud deployment, a tenant that does not sufficiently leverage its CSP security offerings stands to cause other tenants to also suffer the consequences. (Depending on the scope of hypervisor compromise, this may not be an issue for VMs that are encrypting their network communications at layer 3 or above.)

4.1.2. Routers

In a private cloud, the tenant has complete control over routing paths and can therefore keep layer 3 packets confined to the specific trust zones it has established. In a public cloud, the CSP is charged with this responsibility on behalf of the tenant. If an attacker can manipulate the CSP's routing fabric such that a tenant's packets leak outside their intended trust zones, they are subject to being sniffed and having their payloads compromised.

Routers are also frequently used to filter certain types of network traffic from ingressing or egressing certain interfaces. For most enterprises, router filters are prone to accumulate much faster than they dissipate and, over time, it becomes increasingly difficult to comprehend their collective behavior. Leaving traffic filtering to firewalls, which can do much more besides layer 3 filtering, is often a much better decision.

4.1.3. Firewalls

A modern-generation firewall appliance will address a diverse set of security issues, each one complex in itself. It will support myriad permutations of:

- Network filtering (layers 2-4)
- Application layer filtering (layers 5-7)
- URL filtering
- User and user session filtering
- Intrusion Prevention (IPS)
- Antivirus protection
- Malware protection

The challenge of implementing a hardware appliance for these purposes in a cloud is that it is exceptionally inefficient; virtualized traffic must be routed out of the cloud to a legacy infrastructure before the filtering can occur. This problem is allayed when firewalling and threat prevention are implemented within the cloud in a virtual form factor instance that deploys alongside each VM. When knitted together in a framework, they give the VM collective every appearance of being a single firewall appliance.

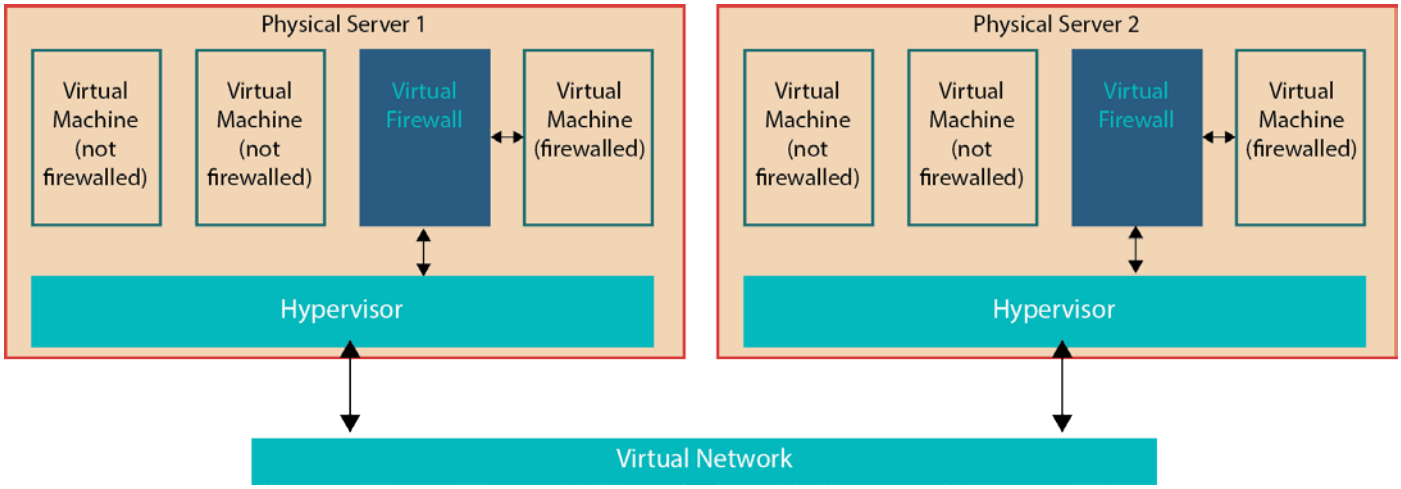

© 2015 Cloud Security Alliance – All Rights Reserved.

This technique introduces some added complexity in implementation, so care must be taken when selecting a firewall vendor. If, for example, there are 200 VMs participating on the firewalled network, and each one happens to be hosted on a different physical server, then the 200 instances will need to deploy in lockstep. Policy changes also need to be implemented across the fleet in lockstep. Done incorrectly, the virtual firewall will not behave with the same degree of determinism as the physical firewall appliance, which, in effect, is a giant step backward in protecting the applications and data — precisely the opposite of what is needed in a cloud.

Cloud tenants should be aware that although this approach makes it fairly straightforward to deploy the firewall in the cloud, it does not make it easier to administer it relative to a physical appliance. In both private and public clouds, the tenant will still need security professionals on staff to interact with the cloud administration team while managing firewall policies.

4.2. Virtual (Guest) Operating Systems

Ideally, OS security is invested equally in both public and private clouds. A compromised OS opens the door to compromising the hypervisor. A disciplined hardening approach includes:

- Removing software known to be unnecessary (e.g., sendmail)
- Keeping pace with patches
- Strong passwords on accounts
- Locking down SSH and RDP remote access protocols
- Least privilege for user accounts
- Disabling unnecessary services
- Utilizing vendor or open source host-based firewalls

NIST publishes more than one hundred OS hardening guidelines at its National Checklist Program Repository.[4]

4.3. Hypervisors

In a private cloud, in addition to selecting its hypervisors, the tenant has the tall order of configuring and operationally managing them (including upgrading and applying security patches) with rigorous attention. In a public cloud, the tenant delegates most of that effort to the CSP. The tenant should be mindful of how the CSP does this, and fully understand everything it takes to manage massive fleets of hypervisors. The CSP should be transparent about how it prevents issues such as:

- Unauthorized access to management interfaces (ports)
- Inter-VM attacks
- “Instant-On” gaps
- Data comingling
- VM data not being destroyed
- VM image tampering
- Unencrypted VMs

[4] https://web.nvd.nist.gov/view/ncp/repository?category=Operating+System&startIndex=0

© 2015 Cloud Security Alliance – All Rights Reserved.

4.4. Management Consoles and APIs

No cloud would actually be a cloud without the management consoles, tools and APIs that power the on-demand self-service aspects of cloud computing. It must be kept in mind that, in the cloud environment, having access to virtual servers, storage and networks through consoles and APIs is tantamount to having access to their physical counterparts in a legacy IT environment, but made orders of magnitude easier. Painstaking care must be given to controls that address:

What	Console Controls	API Controls
Authentication	Multi-factor	API keys
Access control	Rights management	Rights management
Audit	Logging and monitoring	Logging and monitoring
Confidentiality	HTTPS	Transport and/or application encryption
Integrity	HTTPS	Digitally signed requests

In a public cloud, a tenant should implement all of these controls with rigor. In a private cloud, a tenant might opt to implement only a subset of them, although that is a decision that should be made consciously and with great precaution.

5. Operational Differences

Operations management pertains to the administration of practices used to create the highest possible degrees of efficiency within an organization. As investors have seen in recent years, security breaches can deliver striking blows to efficiency and, by extension, profits and even valuation. Breaches at very large companies in recent years are known to have exceeded tens and even one hundred million dollars. An enterprise operating in the cloud with flawed security operations is one operating ultimately with a low degree of efficiency.

5.1. Data Migration

The security ramifications of migrating data to a cloud infrastructure, private or public, are not to be taken lightly. Painstaking consideration must be given to:

- **Involved parties**. In a private cloud deployment, a tenant can conceivably implement a data migration on its own. In a public cloud, in all but the most trivial of circumstances the CSP may need some degree of involvement. At the least the CSP should provide the tenant with tools and documentation to ease the migration. If the tenant is subject to compliance mandates, the tenant's compliance auditors should also have a hand in the planning, lest it be discovered after-the-fact that policy violations occurred. Some enterprises might opt to involve specialized migration service providers.
- **Means of transport**. If data is subject to compliance, regardless of how it is moved, it must continue to comply while it is being moved. For example, if the data calls for AES 256 encryption while in-motion over an internal network, it will not do to FTP the data into the cloud over a VPN link that uses AES 128 encryption. Likewise, if the data needs to be physically transported due to sheer volume, the disks on which the data resides must meet the minimum at-rest encryption requirements.

© 2015 Cloud Security Alliance – All Rights Reserved.

- **Authorizations**. Often, actors involved in data migration will require temporary elevated privileges to manipulate cloud data stores. Such authorizations should be carefully documented for potential future reference. They should also be promptly revoked when no longer needed.

- **Temporary files**. In all but the smoothest of migrations, it is common practice to create temporary files for purposes such as debugging or simply to divide the work piecemeal. Documentation should be kept about the purpose and location of these files along with who created them. When the migration is complete they should be deleted using an endorsed cryptographic wiping technique.

5.2. Change Management

In a private cloud, a tenant will already be well versed in change management processes. In a public cloud, it is important the tenant know whether and under what circumstances to engage their CSP in the change process. To begin, the tenant and CSP must first agree on a start-to-finish change process; this should occur when negotiating the SLA. It should be understood by the CSP that, without exception, the process is to be followed under both calm (e.g., routine firewall rule modification) and crisis (e.g., applying an emergency kernel patch). The NIST model calls for these sequential steps[5]:

1. Request the change
2. Record the request
3. Determine if change control is required
4. Analyze for security impact
5. Test for security impact
6. Approve the change
7. Implement the change
8. Verify the implementation
9. Close the request

5.3. Logging, Monitoring and Measuring

It is essential to develop a crisp logging, monitoring and measuring strategy well ahead of cloud deployment. In a private cloud environment, an enterprise most likely has some logging infrastructure in place, although odds are that it is neither uniform nor centralized. It is common to have several if not many logging ecosystems (e.g., Windows Event logs, Unix syslogs). If opting for a public cloud deployment, this is an excellent opportunity to improve upon that. Some CSPs provide API-driven log management services that normalize log messages into a common format, accompanied by policy-driven monitoring and event-driven alarming capabilities.

Care should be taken on at least two fronts. One, it should be explicitly agreed upon what the tenant is logging vs. what the CSP is logging. If logging is divided between the entities, there will be ample circumstances where the tenant's logs and the CSP's logs need to be analyzed side-by-side in order to piece together what happened prior to and during a security incident. Two, it is not unheard of for logs to be tainted by information that managed to leak into them that should not have, for example, debug messages containing credit card numbers or other sensitive information. It is wise to have pre-arrangements made for scrubbing the logs when the need arises.

[5] http://csrc.nist.gov/publications/nistpubs/800-128/sp800-128.pdf

© 2015 Cloud Security Alliance – All Rights Reserved.

5.4. Incident Management and Recovery

Subscribing to the NIST security incident framework mode of thinking about incidents is a healthy practice. It describes four distinct stages for managing security incidents[6]:

1. Preparing for the incident
2. Detecting and analyzing the incident
3. Containing, eradicating and recovering from the incident
4. Performing post-incident activities

In a public cloud, the tenant and CSP should have a formal agreement about the roles and responsibilities of both sides when an incident occurs. (See the earlier sections on Contracts, SLAs, and Roles and Responsibilities.) It must be kept in mind that incidents can take days or even weeks to resolve. It is important the CSP have the same steely determination to resolve incidents as the tenant.

6. Summary

This paper has presented a broad high level array of security issues that enterprises should consider before deciding whether to pursue a private or public cloud deployment. While thorough, it should not be considered exhaustive. Nor should it be considered absolute. No two enterprises seeking to deploy a cloud infrastructure do so for exactly the same reasons. Regardless, security should be a foremost consideration for all cloud-bound enterprises.

[6] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

© 2015 Cloud Security Alliance – All Rights Reserved.

7. Resources and References

Cloud Security Alliance

Cloud Controls Matrix v3.0.1
https://cloudsecurityalliance.org/download/cloud-controls-matrix-v3/

Security Guidance for Critical Areas of Focus in Cloud Computing V3.0
http://www.cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

Enterprise Architecture V2.0
https://downloads.cloudsecurityalliance.org/initiatives/tci/TCI_Reference_Architecture_v2.0.pdf

National Institute of Standards and Technology (NIST)

The NIST Definition of Cloud Computing
http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Guide to Security for Full Virtualization Technologies
http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf

Guide to General Server Security
http://csrc.nist.gov/publications/nistpubs/800-123/SP800-123.pdf

Guidelines on Firewalls and Firewall Policy
http://csrc.nist.gov/publications/nistpubs/800-41-Rev1/sp800-41-rev1.pdf

Guide for Security-Focused Configuration Management of Information Systems
http://csrc.nist.gov/publications/nistpubs/800-128/sp800-128.pdf

Computer Security Incident Handling Guide
http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

National Checklist Program Repository
http://checklists.nist.gov

International Standards Organization (ISO)

ISO/IEC 27002:2013 Information technology — Security techniques — Code of practice for information security controls
http://www.iso.org/iso/catalogue_detail?csnumber=54533

© 2015 Cloud Security Alliance – All Rights Reserved.

Other

Virtualized Firewalls
Palo Alto Networks
https://www.paloaltonetworks.com/products/platforms/virtualized-firewalls/vm-series/overview.html

Learn more about the Palo Alto Networks VM-Series
Palo Alto Networks
https://www.paloaltonetworks.com/products/platforms/virtualized-firewalls/vm-series/overview.html

AWS security best practices
Amazon.com Web Services
Presented at Palo Alto Networks Ignite 2015, April 4-6, 2015
http://www.slideshare.net/AmazonWebServices/security-on-aws-palo-alto-ignite-conference-2015

Security Concerns During Data Migration
SANS Institute
http://www.giac.org/paper/gsec/3800/security-concerns-data-migrations/106121

© 2015 Cloud Security Alliance – All Rights Reserved.

8. Appendix

Palo Alto Networks VM-Series Next-Generation Firewall

The Palo Alto Networks VM-Series of virtualized next-generation firewalls supports exactly the same security feature set as found in its physical form-factor firewalls. The VM-Series is based on a single-pass software architecture that first identifies the application, regardless of the port or ports it is using, while simultaneously determining the maliciousness of the content and the identity of the user. These three business relevant elements -- the application, the content and the user -- form the basis of the overall virtualization security posture, inclusive of visibility, policy control, forensics and reporting. The VM-Series allows the following types of security policies to be applied to public, private or hybrid cloud computing environments:

- Segment applications and data through whitelisting, while implicitly blocking all else, regardless of port, evasive tactic, or SSL decryption
- Control VM-to-VM communications based on specific applications as a means of reducing threat footprint and risk exposure
- Apply application-specific policies to prevent known and unknown threats from gaining access to, and moving laterally within, the virtualized environment
- Exert granular control over access to applications and data based on user need and credentials
- Minimize the time gap that may exist between workload additions, removals or changes with Panorama and native management features and a rich set of APIs

The VM-Series supports a wide range of leading virtualization and cloud environments. Learn more about the virtualized form factor of our next-generation firewall here.

© 2015 Cloud Security Alliance – All Rights Reserved.