cloud security alliance℠

SecaaS Implementation Guidance

Category 10 // Network Security

September 2012

Contents

Foreword

Cloud Computing represents one of the most significant shifts in information technology many of us are likely to see in our lifetimes. We are reaching the point where computing functions as a utility, promising innovations yet unimagined. The major roadblock to full adoption of Cloud Computing has been concern regarding the security and privacy of information.

Much work has been done regarding the security of the cloud and data within it, but until now, there have been no best practices to follow when developing or assessing security services in an elastic cloud model—a model that scales as client requirements change.

One mission of the Cloud Security Alliance is to provide education on the uses of Cloud Computing to help secure all other forms of computing. To aid both cloud customers and cloud providers, the CSA SecaaS Working Group is providing Implementation Guidance for each category of Security as a Service, as delineated in the CSA's SecaaS Defined Categories of Service. Security as a Service was added, as Domain 14, to version 3 of the CSA Guidance.

Cloud Security Alliance SecaaS Implementation Guidance documents are available at https://cloudsecurityalliance.org/research/working-groups/security-as-a-service/.

We encourage you to download and review all of our flagship research at http://www.cloudsecurityalliance.org.

Best regards,

Jerry Archer	Alan Boehme	Dave Cullinane
Nils Puhlmann	Paul Kurtz	Jim Reavis

The Cloud Security Alliance Board of Directors

Letter from the Co-Chairs

Security as a Service is a specialized area categorized two years ago as growing rapidly and in unbound patterns. Vendors were struggling. Consumers were struggling. Each offering had its own path. We felt it was urgent to address the needs and concerns common to the implementation of Security as a Service in its many forms.

The Defined Categories of Service helped clarify the functionalities expected from each Category. In this series, we hope to better define best practices in the design, development, assessment and implementation of today's offerings.

We want to thank all of the many contributors worldwide who have worked so hard to produce these papers providing guidance for best practices in Cloud Computing Security. Many have been with the Security as a Service Working Group since the beginning; many others joined in this effort. Each has spent countless hours considering, clarifying, writing and/or editing these papers. We hope they help move forward toward those unimagined innovations.

Sincerely,

Kevin Fielder and Cameron Smith
SecaaS Working Group Co-Chairs

Acknowledgments

Co-Chairs

Ken Owens, Savvis

Contributors

Kevin Fielder, Canada Life
Emilio Casbas, Incita
Bernd Jäger, Colt

Peer Reviewers

Vaidy Chandramouli, TCS
Ted DeArmon, Arrowpoint Solutions, Inc.
Michael Roza, Bridgestone
Ganesh Unavane, Wipro Technologies
Henry St. Andre, inContact
Marc Appelbaum, Vonage
Jens Laundrup, Emagined Security

CSA Global Staff

Aaron Alva, Research Intern
Vicki Hahn, Technical Writer/Editor
Luciano JR Santos, Research Director
Kendall Scoboria, Graphic Designer
Evan Scoboria, Webmaster
John Yeoh, Research Analyst

1.0 Introduction

Network security, as applicable to a cloud environment (IaaS, PaaS, and SaaS), consists of the security of the underlying physical environment and the logical security controls that are inherent in the service or available to be consumed as a service. Physical environment security ensures access to the cloud service is adequately distributed, monitored, and protected by underlying physical resources within which the service is built. Logical network security controls consists of link, protocol, and application layer services.

In a cloud environment, a major part of network security is likely to be provided by virtual security devices and services, alongside traditional physical network devices. Tight integration with the underlying cloud software layer to ensure full visibility of all traffic on the virtual network layer is important.

In the cloud network, the classic definition of network perimeter takes on different meanings. For many cloud networks, the perimeter is clearly the demarcation point. For other cloud networks, the perimeter transforms into highly dynamic "micro-borders" around individual customer solutions (to the level of certain data sets/flows within a solution) within the same cloud, consisting of virtual network components. In other cloud networks, there is not clear perimeter at all.

This causes a challenge within a cloud environment. Typically, the inspection and control of network traffic do not pass through physical interfaces where classical control devices can analyze or block them. This happens when cloud servers use a physical server's internal memory pipe (software switch or even direct APIs). This is another reason why effective controls require the integration with the cloud software layer.

This Implementation Guidance addresses cloud environment network security architecture, security gateways (firewalls, WAF, SOA/API), Security Products (IDS/IPS, Sub Tier Firewall, Security Monitoring and Reporting, Denial of Service (DoS) protection/mitigation, and secure "base services" like DNSSEC and NTP.

1.1 Intended Audience

This document is a reference architecture that identifies scenarios and application of network security. It can be used as guidance to those who need and intend to apply network security to their cloud implementation or service provider. If the reader approaches this document with a particular scenario in mind, s/he should be able to find that scenario in the document with the accompanying guidance for his/her desired situation.

1.2 Scope

The scope of this reference architecture is network security considerations and implementation guidance which addresses:

- How to segment networks
- Network security controls
- Ingress and egress controls such as Firewalls (Stateful), Content Inspection and Control (Network-based), Intrusion Detection System/Intrusion Prevention Systems (IDS/IPS), and Web Application Firewalls

- Secure routing and time
- Denial of Service (DoS) and Distributed Denial of Service (DDoS) Protection/Mitigation
- Virtual Private Network (VPN) and Multiprotocol Label Switching (MPLS) Connectivity (Secure Sockets Layer – SSL; Internet Protocol Security – IPsec; Virtual Private LAN Service – VPLS; Ethernet Virtual Private Line - EVPL)
- Threat Management
- Forensic Support
- Privileged User/Use Monitoring

Out of Scope:

- DLP – Refer to CSA SecaaS Implementation Guidance Category 2: Data Loss Prevention
- Encryption – Refer to CSA SecaaS Implementation Guidance Category 8: Encryption
- Internet Access Proxy Servers – Refer to CSA SecaaS Implementation Guidance Category 3: Web Security

2.0 Requirements Addressed

Network security addresses risks relating to the use of, and access to, businesses networks. Network security encompasses protecting data as it traverses the network, protecting data as it traverses public networks such as the internet, protecting systems and data from network-based attacks, and protecting the networking components themselves. By offering services from the cloud that can provide traffic encryption, network monitoring, traffic analysis and controls, Virtual Private Networks (VPN), firewalling and secure networking services, the Cloud Services Provider (CSP) can ensure the security of the customer's network environment. Cloud-based services can be used to secure traditional, non-cloud networks, pure cloud networks, and hybrid internal + cloud networks.

As with most cloud-based services implemented, network security as a service will reduce overhead for both staff and infrastructure, while allowing the customer to leverage the dedicated expertise of the CSP's staff and resources.

The rest of this section provides a high level overview of the concepts and components involved in providing network security from the cloud to both cloud and on premises networks.

2.1 Networking Models

There are many ways to implement a cloud network. Most networks fall into one of the following categories, though each implementation will have a unique architecture: traditional, converged, and cloud-only.

2.1.1 Traditional

Traditional cloud networking will utilize multiple layers, with the hypervisor layered on top of the physical servers that connect to the Access switch layer, and all Virtual Local Area Networks (VLANs) extended to all hosts to enable logical separation within the hypervisor environment.

This model leverages traditional network security components at the distribution level using physical security controls between layer 2 boundaries. Traffic control and security are well understood; limitations involve visibility into the virtualization layer and the threats that affect this layer.

2.1.2 Converged

Cloud networks can leverage the convergence of IP networks and storage networks along with physical and logical networks to create a new cloud network model. This model typically will maintain a physical perimeter switch and security control points, but the underlying architecture is optimized for cloud workloads.

This approach has several network security advantages, as the virtualization network layer becomes the access network, and virtual security appliances can be integrated to provide visibility to virtual machine traffic and secure the virtual network. This approach provides a blend of physical and virtual controls in the cloud environment.

2.1.3 Cloud-Only

A cloud-only network provides direct access to virtual machines over the public Internet or via a routed layer 3 Virtual Private Network (VPN). This model means the security controls relating to the "internal" cloud components must be virtual or fully integrated with the virtual network.

2.2 Network Access Controls

Network access to a public cloud environment is the fundamental security control point that ensures basic attack vectors are mitigated by traditional controls. Controls can be implemented in physical, converged, or virtual appliances.

2.2.1 Perimeter Firewall Controls

Perimeter firewall security controls provide real-time protocol inspection and detection of known attacks. Place your deployment within a perimeter of security provided by the firewall or Unified Threat Management (UTM) solution, to ensure known attacks anomalies are detected and blocked. This provides the first layer of defense. Base policies for a perimeter firewall limit the source and destination ports and protocol to a limited set required for the service being offered.

In a cloud environment, the perimeter potentially will look more porous, as there will be management connections and potentially traffic going via the CSP. Any deployment likely will share key infrastructure across multiple clients. Keeping data and networks logically separated is critical. The emphasis on greater internal controls, versus the historical focus purely on the perimeter, is in line with recent research regarding the perimeter-less network.

2.2.2 Sub-Tier Firewall Controls

The goal of the sub-tier firewall is to provide a separate security boundary within the virtualization layer of the cloud, to secure the virtual machines and tiers of network created within the cloud network. The base policies for sub-tier firewall limit tier-to-tier network traffic.

2.2.3 Access Control Lists

Access Control Lists (ACLs) provide a basic security control layer to support securing virtual machines from standard layer 2 security threats like flooding and scanning.

2.3 Content Inspection and Control

At a network level, various content inspection control technologies exist to protect the network, business systems and business data from both external attacks and internal data theft. These include IDS/IPS, DLP, and Proxy servers.

2.3.1 Intrusion Detection and Intrusion Prevention

Network Intrusion Detection Systems (NIDS) and Network Intrusion Prevention Systems (NIPS) are usually placed at ingress and egress points of the network in order to detect and prevent anomalous traffic, usually based on a combination of signatures, heuristic behavioral analysis, and statistic protocol anomaly detection.

Any Service Level Agreement (SLA) should define the locations to be monitored, specify service and performance levels, and how rules are added and managed. Secure management, transport (to which locations), segregation and analysis of collected data must be considered and defined in any contracts with the CSP. Any IDS/IPS device must be capable of handling the volume of traffic that is expected to pass through it in order to be effective.

When determining whether and how to use a Network Security vendor to provide Security as a Service, be certain you weigh the benefits of on-site log correlation against the higher volumes of data that must be transferred to the CSP.

2.4 Secure Time

Time synchronization is vitally important as effective prosecution of security incidents requires accurate matching of timestamps on all log files. Discrepancies will complicate and delay incident response. When deploying from the cloud to the on-premises network, ensure time is synchronized between the CSP's systems and the customer's, to enable investigation and log correlation across the entire environment (cloud and non-cloud).

2.5 Network Routing and IP Management

Network routing is one of the most common attack vectors. Limit the risk by managing the IP space and access methods using TLS/SSL or IPSec VPN with IP whitelisting. Ensure that the CSP uses secure implementations for routing devices that are publicly accessible, including permitting them to connect to and receive update suggestions only from trusted routers, and implementing secure update policies.

2.6 DDOS Protection/Mitigation

The most prominent form of a DoS attack is probably the Distributed Denial of Service (DDoS) attack, utilizing thousands of compromised hosts sending malicious traffic to exhaust networking resources or those of the servers hosting the application. The attack results in denial of service to legitimate users because their infrastructure is overwhelmed with illegitimate requests.

Volume based-DDoS attacks can only be mitigated at the CSP backbone or other network that has significantly more bandwidth than the sum of the bandwidth of all attack traffic. Once an attack condition is identified, monitoring entities trigger a re-route of suspicious traffic through a cleansing instance that attempts to filter out the attack traffic while allowing legitimate packets to pass through to the destination network.

A recently recognized trend is the implementation of application level DoS attacks, which trigger resource overload by sending malformed calls to web-applications, causing high utilization on the web-, application- or database server, rather than flooding the network. The appropriate response to an application level DDoS attack is a reverse proxy or filter grid, sometimes called "Web-Application-Firewall" (WAF).

Cloud services depend heavily on management portals and APIs to control almost every function delivered to cloud customer. Therefore, DDoS mitigation is essential to any cloud-based service.

2.6.1.1 DDoS Mitigation Recommendations and Guidance

- Implement volume-based DDoS mitigation with enough bandwidth to exceed the volume of the attack.
- Implement volume and application level protection.
- In a multi-homes environment, ensure all traffic passes through a DDoS mitigation grid. Ensure the CSP does filtering, or engage a SecaaS provider for all traffic.
- Accurate filtering requires thorough profiling of legitimate traffic.
- Test all profiles.
- Have a response procedure in case legitimate traffic gets blocked.
- Determine how fast the configuration of the mitigation filter can be changed.
- Determine if the profiling algorithm can handle all applications on the same subnet or if they must be split.
- Ensure that the CSP allows different deployment models, including constantly protected, on-demand and probably "emergency" mode with almost no time for profiling upfront. This will enable services to be brought up in alternate data centers if required to ensure service availability under prolonged DDoS attack.
- Implement active notification.
- Ensure access to relevant monitoring, alerting, and network performance reports and metrics.
- Consider a "real-time" optimization interface.
- Ensure the solution the mitigation requirements, including maximum mitigation bandwidth, sessions, packets per second, and attacker and victim IP address pair.
- Ensure the solution supports IPv6.

2.7 VPN and MPLS Connectivity (SSL, IPSec, VPLS)

2.7.1 SSL VPN

An SSL Virtual Private Network (VPN) gives remote users secure access to Web applications, client/server applications and internal network connections. TLS-based SSL VPNs are more secure, as TLS is a more secure protocol that is gradually replacing SSLv3.

2.7.2 IPSec VPN

IPsec provides two choices of security service: Authentication Header (AH), which provides authentication of the sender of data, and Encapsulating Security Payload (ESP), which provides both authentication of the sender and encryption of data.

2.7.3 Border Gateway Protocol

The Border Gateway Protocol (BGP) allows PE routers to communicate with each other about their customer connections. Each router connects to a central cloud, using BGP. When new customers are added, the existing routers will communicate with each other via BGP, and automatically add the new customers to the service.

2.7.4 Label Distribution Protocol

Also known as a Layer 2 circuit, this method uses LDP to communicate between PE routers. In this case, every router connects to every other router in the VPN.

2.7.5 VPLS

Businesses considering MPLS VPN services should factor differences between Layers 2 and 3, and opt for the one that best fits their own needs. Layer 2 MPLS technology is limited because it does not scale as well as Layer 3, but Layer 2 services are simpler in architecture and allow customers to retain control of their own routing tables. Layer 3 MPLS VPNs are characterized by fully meshed architectures that enable multicast conferencing for projects involving a dispersed work group. However, outsourcing of routing tables sometimes is seen as a weakness of Layer 3 VPN services, because corporations may not be willing to relinquish control or share their routing schemes.

2.8 Risk Management

Risk management is now one of the principal areas of focus for CIOs, CISOs, and CFOs. Due in large part to regulatory compliance mandates, corporations must understand, manage and report risks to the confidentiality, integrity, and accessibility of their critical data with more granularity and reliability.

It is impossible to eliminate all risks. The only reasonable goal is to manage risks to an acceptable level. Best practices dictate prioritizing threats, and focusing on the most important of them. Providers should offer recommendations as to the proper management of risks within the scope of the existing (or proposed) infrastructure.

Networks will be vulnerable. Keep the window of vulnerability as small as possible, until mitigation measures can be applied. Regular, non-disruptive vulnerability scans are recommended, but require special consideration, as disrupting network services of a cloud infrastructure may impact many tenants.

Patch management of virtual network components may require new processes in the cloud. Some traditional network technology vendors are providing virtual network components which might smoothly integrate into the existing network management infrastructure. Confirm with both application and SecaaS vendors.

2.9 Forensic Support

Network forensics provides information useful in aiding an investigation, or incident response, addressing "hacked" (compromised) systems. As network components are being virtualized, forensic monitoring and collection of logging data that is forensically sound is becoming more challenging. This is due to the fact that many virtual devices rely on the same hypervisor. If the hypervisor becomes in any way compromised, there is a question regarding how trustworthy data from any virtual device is.

2.9.1 Logging

Configure proper logging for all relevant network components that might be virtualized. Capture log content that is required to support network forensic investigations, and forward all logs to a central, hardened log-server and protect logs and server. Protect logs in transit, implement monitoring, apply automatic analysis, correlation and visualization. Determine the retention period and log backup from a compliance and commercial perspective. SIEM products are evolving to become better integrated in and with the cloud. Consider using a hosted log service.

2.9.2 Capturing Network Traffic

In most cases, capturing real-time network traffic will be required. Ensure the virtualized infrastructure can support this. Virtualization vendors may offer mirror port capabilities, aiding the capture of network traffic by monitoring devices. Ensure virtual switches support sniffing or the topology and solution design allows trunking out relevant traffic.

3.0 Implementation Considerations and Concerns

3.1 Considerations

Like any technology, there are industry standards and best practices that should be followed to ensure that the data processed and stored by that technology is secure. All ingress and egress points to the cloud environment need to inspect traffic monitor, and log network activity at specified periods of time. Any suspicious activity and alerts should be addressed in a defined manner. This section discusses implementation considerations and concerns that should be part of any discussion of Network Security in the cloud.

3.1.1 Isolate Networks

The first network security consideration in a cloud environment is to provide high levels of network isolation between all of the different networks within the environment. These networks include management networks, cloud/virtual server migration networks, IP storage networks, and individual customer networks; which may in turn be further broken down into segregated networks such as databases, file servers, virtual desktops etc. Each network should be segmented from the others.

Isolation can be achieved using various methods, such as the use of separate virtual switches for each of the networks, which also requires the use of separate physical NICs to uplink the virtual switches to the physical network. Another option is the use of 802.1Q VLANs, which will allow for much greater scaling of the virtual environment and allow for the most flexibility. A third option would be a combination of the two methods: using a virtual switch, with the use of 802.1Q VLANs for the management, migration, and IP Storage network; and a virtual switch and 802.1Q VLANs for customer networks.

In addition to switching and routing segregation, these networks should be firewalled from each other to prevent any potential for traffic being accidentally routed among them. Auditability of this segregation can be provided by recording the firewall logs and/or collecting network data.

3.1.1.1 Isolation of Management Networks

Management networks allow the cloud provider to access the environment and manage the different components within that environment. Only authorized administrators should have access to this network.

Control of the management interfaces of the individual cloud hosts allows for complete control of all of the cloud servers on that host. This does not mean they have access to log onto the virtual servers or access the data if it is encrypted; however they can do things like restart, clone, and attempt console level access to the virtual machines. Root access on this interface is analogous to having the keys to a physical rack of servers within a datacenter. Administrator access to the central management console that manages all of the different cloud hosts is analogous to having the keys to the datacenter and every rack within that data center.

Protection of these interfaces is paramount, and a customer should never need direct access to any of the systems within this network.

3.1.1.2 Isolation of Cloud/virtual Server Migration and IP Storage Networks

Both Cloud Migration and IP Storage networks should be on isolated and non-routable networks. No outside connectivity should be necessary.

The reason for isolating this traffic is twofold: performance, as both Cloud Migration and IP Storage traffic need very fast data rates; and the fact that this traffic may travel over the network in clear text, and thus may be susceptible to an attacker sniffing sensitive information. By fully isolating this network, an attacker would require physical access to the network to successfully compromise this data.

3.1.1.3 Isolation of Customer Data Networks

Customer data networks should be isolated from each other and from any management networks. This can be accomplished in a secure and scalable way via the use of 802.1Q VLANs and firewalling between the networks to ensure that no traffic is routed between networks. The use of either physical or virtual appliance firewalls, along with IDS/IPS, can be used to provide a very strong level of security between these networks.

3.1.2 Secure Customer Access to Cloud-Based Resources

Customers need a way to access their resources located within the cloud, and to manage those resources in a secure manner. The Cloud Service Provider (CSP) should supply the customer with a management portal that is encrypted. As the majority of access to CSPs' systems is via the Internet using TLS (preferred) or SSLv3, encryption via a web browser would be the most common approach to securing customer access.

3.1.3 Secure, Consistent Backups and Restoration of Cloud-Based Resources

The cloud environment should supply the customer with a transparent and secure backup mechanism to allow the customer's cloud-based resources to be backed up on a consistent basis. Should there be a loss of any of the customer's cloud-based resources, they should be able to easily and quickly restore those resources. With the capabilities of the cloud technologies that act as the backbone of the cloud, it is not only possible to backup and restore data, but also to restore complete operating systems and applications running within those operating systems.

3.1.4 Strong Authentication, Authorization, and Auditing Mechanisms

It is very important in any shared environment to ensure that users and administrators of the system are properly and securely authenticated, are only able to access the resources they need to do their jobs or the resources that they own within the system, *and nothing more*. It also is very important in cloud to know who is doing what within the system, and when their actions occurred.

The needs to provide separation of duties and enforce least privilege apply to both the cloud environment and the customer. The CSP should ensure that its administrators have access only to what they need and nothing

more. They also should provide the customer with a mechanism to ensure that the customer's own administrative staff has required access to needed resources. Any access to cloud resources by either the customer or the cloud provider should be logged for auditing purposes.

Auditing and authorization are key points where there are strong links between different components of the overall CSA SecaaS Guidance series. A key part of any ability to audit across multiple systems is a method to consolidate and analyze the logs and monitoring data relating to those systems. Best practice for multi-component audit systems is the installation of a Security Information and Event Monitor (SIEM). Further information on the use of a SIEM can be found in the *CSA SecaaS Category 7: Security Information and Event Management* whitepaper.

3.1.5 Resource Management to Prevent DDoS

Many think of the need for resource management as a way to fairly separate the cloud-based resources between customers within the cloud. Resource management also has a very important security function: to prevent the potential for denial of service attacks. If resource management is not in place, a compromised cloud server could allow an attacker to starve all of the other cloud servers within that cloud of needed resources. By using resource management, compromised cloud servers can neither access nor adversely affect other servers within the cloud. Ensure that the CSP has enough resources available to deal with spikes in usage, and that they have resource management in place at the hypervisor level to prevent any individual guests from impacting the performance of other guests sharing the same host.

3.1.6 Bandwidth Availability and Management to Prevent DDoS

The CSP should offer DDoS protection for systems it hosts. This may be accomplished via a combination of large shared bandwidth and DDoS protection tools. DDoS protection tools typically use a combination of signature and heuristic-based detection techniques to drop packets when attacks are instigated.

The tools the CSP deploys may also be deployed in front of the customer's on premises systems to provide some of the technical protection provided to the CSP's own cloud-based system. These tools can instigate responses such as dropping packets or redirecting detected DDoS attempts. However, an on premises deployment does not offer the advantages inherent in the shared bandwidth volumes available in the CSP's data center.

Another technique for protecting on premises systems against DDoS attacks involves automatically routing Internet traffic to fail over systems hosted with the CSP, should a potentially successful DDoS attack be detected by the CSP's monitoring systems.

3.1.7 Encrypting Critical Data

Data encryption adds a layer of protection that remains even if a system is compromised. It is especially important to encrypt data in transit, as that data will be traversing a shared network, and could potentially be intercepted if an attacker can gain access at a critical point in the network. By encrypting the data as it traverses the network, it makes it much more difficult for an attacker to do anything with the data if they are able to intercept it.

Encrypting the critical data "at rest" on the virtual machine, within the virtual disk file, is also very important. This will protect critical data from being accessed by any unauthorized individuals, and will make it much more difficult for an attacker to compromise data even if they are able to compromise an endpoint.

3.1.8 Application Programming Interfaces (APIs)

3.1.8.1 Monitoring APIs

It is critical to understand what monitoring APIs are available from the CSP, and if they match risk and compliance requirements. Network security auditors are challenged by the need to track a server and its identity from creation to deletion. Audit tracking is challenging in even the most mature cloud environments. The challenges are greatly complicated by cloud server sprawl, the situation where the number of cloud servers being created is growing more quickly than a cloud environments ability to manage them. Leveraging the monitoring APIs for audit tracking is recommended.

3.1.8.2 Cloud APIs

A valid threat vector for cloud is the API access. The majority of CSPs today support public API interfaces available within their networks and likely over the Internet. Access to these APIs should require certificate-based TLS/SSL encrypted access to ensure that these interfaces do not become a new point of attack for injections, DDoS, and code level penetration. Since this single interface enables all provisioning, monitoring, billing, and real-time auditing of the cloud, this threat vector has several considerations from a network security standpoint. At a minimum, the CSP's network security offering should:

- Require SSL
- Audit Calls
- Offer an IDS
- Protect against DDoS
- Provide Security Penetration Protection for:
    - Code injection
    - Malformed Requests
    - SQL Attacks
- Limit request message size
- Check for XML, and reject DOCTYPE (prevents external XML element definition)
- Manage the depth and complexity of XML trees
- Offer automatic retry on target service
- Provide Authentication and Authorization
- Provide Credential Caching and Expiration
- Allow IP Restrictions (White Listing)
- Provide Rate Limiting
- Provide API Service Level Monitoring
- Monitor overall health

3.2 Concerns

It is likely that any network security monitoring of traditional on premise systems will involve the CSP installing some infrastructure at the local site to perform tasks such as traffic analysis, blocking and dropping connections, and collating log files. These components should be managed by the CSP, and the logs sent to the CSP for analysis and action.

Encrypted traffic will be a challenge for many devices used in network monitoring, including NIDS/NIPS, network-based DLP and firewalls. If the data remains encrypted, the monitoring systems are essentially blind, and unable to perform the roles for which they were designed. In order to perform network monitoring, the CSP may be given the ability to unencrypt and view the traffic for monitoring purposes. The ability to decrypt the data enables a CSP to provide improved monitoring, by enabling their systems to interrogate all data. The monitoring requirements and potential benefits need to be carefully weighed against the increase in risk inherent in unauthorized data access if the data is unencrypted at any point. The general recommendation is that data should be encrypted at all times when transferring to and from the cloud and when in the cloud, thus the recommendation is to keep the data encrypted and perform the best monitoring possible around this fact.

A common concern across many cloud-based solutions is that of the CSP's access to the customer's data and systems. Ensure there is correct and enforced separation of both duties and their various customers' data.

For deployments where the CSP is monitoring on premise systems, bandwidth between the customer and the CSP will be a consideration. Potentially large amounts of logging and analysis data will need to be sent to the CSP.

Bandwidth is unlikely to be a concern when the CSP is monitoring systems already hosted in the cloud, as they are effectively on premises in relation to the CSPs Security systems, and it is assumed that the CSP will manage any scale and bandwidth issues.

3.2.1 DDoS Mitigation

Once the business has transferred information into the cloud, it will be available only via some sort of network (which might well be a public network like the Internet). The consumer cannot simply go to the server where the data is stored and use the local console if the network becomes unavailable; therefore, network integrity becomes crucial.

DDoS attacks can make cloud services and access to data unavailable (e.g., by flooding the network with packets until no remaining bandwidth is available, or by overloading gateways or services within the network path). Mitigation, especially for packet flooding, can be implemented within the core network only if there is sufficient bandwidth to deal with this type of attack. The service provider either must operate a mitigation infrastructure, or buy a service and re-route the customer traffic through a third-party filter grid.

3.2.2 Cost Effectiveness

Ensure the provider's countermeasures and controls address business risks in a cost effective way. Consider the cost of making granular usage of some of the CSP's controls, and implementing others in house. Confirm that the provider offers APIs to assist the customer in performing internal network security monitoring services.

3.2.3 Reports

The CSP should provide a reporting facility to present information about the security concerns in the cloud environment. The reports should be available on-demand, according to a specified schedule, and also should be summarized into monthly reports.

4.0 Architecture Overview and Implementation Steps

This section provides implementation guidance for secure tunneling and encryption in Transit, Access and Authentication controls, Security Gateways (firewalls, WAF, SOA/API), Security Products (IDS/IPS, Sub Tier Firewall, Security Monitoring and IR), DoS protection/mitigation, and secure "base services" like DNSSEC and NTP.

4.1 Architecture Overview

Traditional environments segment physical servers by utilizing different VLANs. In order to maintain segregation of different customers' systems, and separation within customer networks in line with their on premise networks, cloud environments should take the same approach and segment cloud networks and servers. One method is to segregate VLANs through Port Group configurations. This is a well-understood and mature technology, supported by other security building blocks like security gateways, while at the same time supporting the dynamic nature and requirements of cloud network environments.

In the traditional environment, traffic flows are visible to traditional network-based security protection devices, such as the network-based intrusion prevention systems (IPSs). The concern in cloud environments is that the cloud provides limited visibility to inter-virtual machine traffic flows, as these remain entirely within the realm of the virtual infrastructure and hypervisor. By default, these traffic flows are not visible to traditional network-based security protection devices, such as the network-based intrusion prevention systems (IPSs) located in the data center network. The model of a cloud service is most often depicted as an "Internal Cloud" that can distribute workloads to either a "Private Cloud" or an "External Cloud." In concept, this sounds logical. However there are several network security concerns that architecturally need to be addressed.

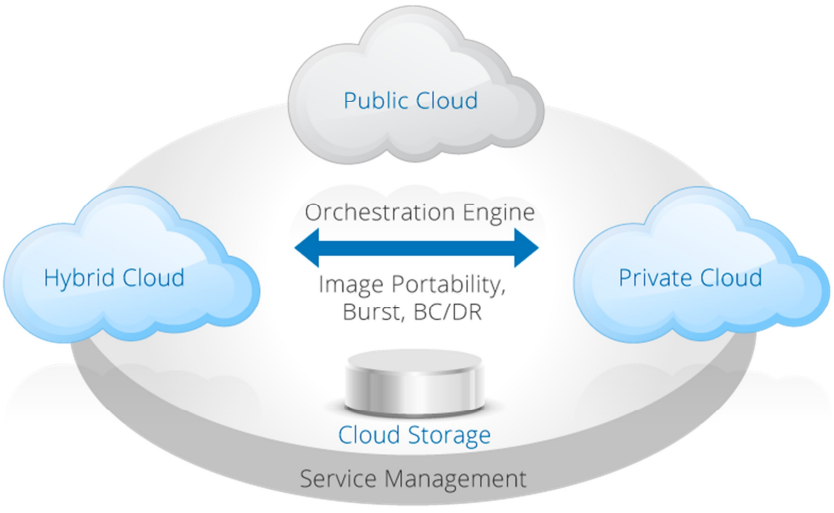

Figure 1: Cloud Environment

The first concern is the network connectivity between the different cloud components. This involves network access controls, content inspection, encryption, API, network routing, and auditing of data flows. The second concern is the network security controls within each cloud environment. Finally, common practices controls like DDOS, NTP, and IP addressing are important to discuss.

There are many ways to implement a cloud network. Possible configuration approaches and their implications are discussed below.

4.1.1 Traditional Approach

The traditional networking design will utilize multiple layers comprising Core switching and routing, with distribution/server layer 2/3 switching connecting to Core along with layer 2 edge switching for the access layer as shown in Figure 2.

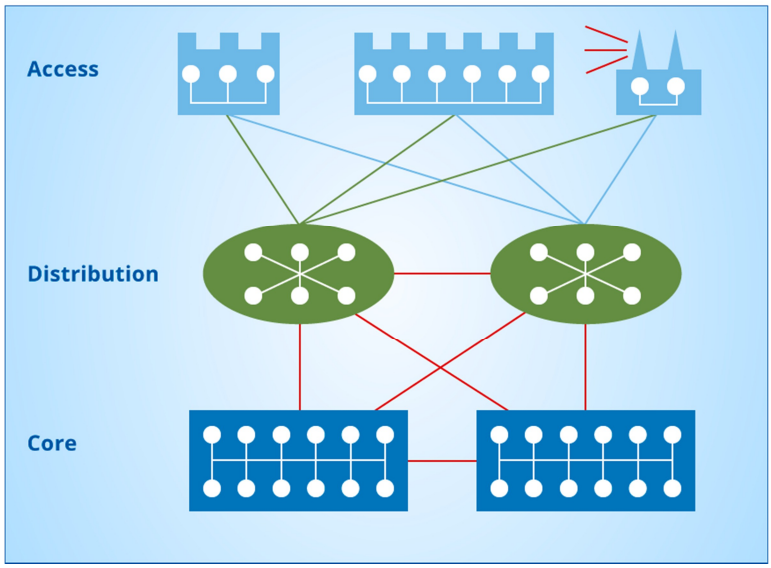

Figure 2: Traditional Three-Tier Network

In this model, the hypervisor is layered on top of the physical servers that connect to the Access switch layer, and all VLANs are extended to all hosts to enable logical separation within the hypervisor environment. This model leverages traditional network security components at the distribution level using physical security controls between layer 2 boundaries. This approach will have at least two physical Ethernet interfaces per server (1GB or 10GB). Traffic control and security are well understood; limitations involve visibility into the virtualization layer and the threats that affect this layer.

SAN environments traditionally are set up with dedicated storage switching that is configured and maintained separately from the IP network.

4.1.2 Converged Network Approach

The converged network approach leverages the convergence of IP networks and Storage networks along with physical and logical networks to create a new cloud network model. This model typically will maintain a physical perimeter switch and security control points, but the underlying architecture is optimized for cloud workloads. This optimization typically includes converging the IP and SAN networks into FCoE or 10GE IP with iSCSI for SAN connectivity. Some providers are investigating local disk techniques to lessen the SAN requirements.

In this approach, servers typically have at least two converged network interfaces that are 10GE IP-based. The SAN fabric still has an A and a B side; however, the SAN is directly connected off the access layer switch instead of needing to connect to access, distribution, and core SAN switches. Traffic control in this approach is trickier, as IP traffic, management traffic, and storage traffic all share a common network fabric. Care should be taken to engineer bandwidth requirements to allocate percentages of bandwidth to various traffic types.

This approach has several network security advantages, as the virtualization network layer becomes the access network, and virtual security appliances can be integrated into this network layer to provide visibility to virtual machine traffic and secure the virtual network with firewall, IPS, File Integrity, AV, etc. This approach provides a blend of physical and virtual controls in the cloud environment.

4.1.3 Cloud-Only Network

A cloud-only network moves away from traditional controls, and instead providing direct access to virtual machines over the public Internet or via a routed layer 3 VPN network. This model puts the network security controls completely on virtual routers, load balancers, VPNs, firewalls, and IPSs.

4.2 Guidance and Implementation Steps

4.2.1 Network Access Controls

Network access to a public cloud environment is the fundamental security control point to ensure that basic attack vectors are mitigated by traditional controls. These controls can be implemented in physical, converged, or virtual appliances.

4.2.1.1 Perimeter Firewall Controls

Perimeter Firewall security controls consist of real-time protocol inspection for detection of known attacks.

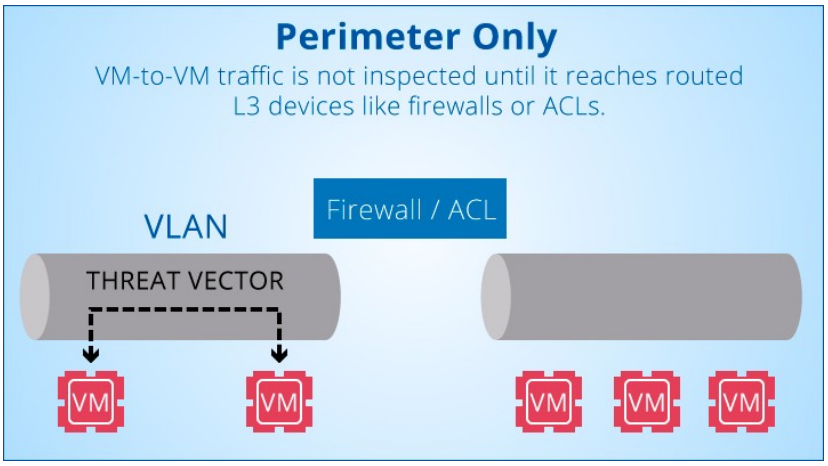

Figure 3: Perimeter Firewall

The goal is to place the entire deployment within a perimeter of security provided by the firewall or UTM solution, to ensure all known attacks anomalies are detected and blocked. This provides the first layer of defense. The base policies for a perimeter firewall limit the source and destination ports and protocol to a limited set required for the service being offered.

When network security is being provided by a CSP, the perimeter potentially will look more porous, as there will be management connections and potentially traffic going via the CSP. Any deployment in the cloud likely will share key infrastructure, such as firewalls and IPSs, across multiple clients at the infrastructure level. This section highlights the concept of logical separation. Keeping data and networks logically separated is critical in a cloud-based deployment. While perimeter controls are still very valid and should form part of any secure network design, the levels of protection around specific "internal" networks and host-based protection of servers are becoming ever more critical. The emphasis on greater internal controls, versus the historical focus purely on the perimeter, is in line with much of the research regarding the "perimeter-less network."

4.2.1.2 Sub-Tier Firewall Controls

The sub-tier firewall provides a second layer of virtualization aware real-time protocol inspection and detection. This layer of security ensures that VM to VM traffic that stays within the virtualization network has security policies to protect against internal threats or compromised machines.

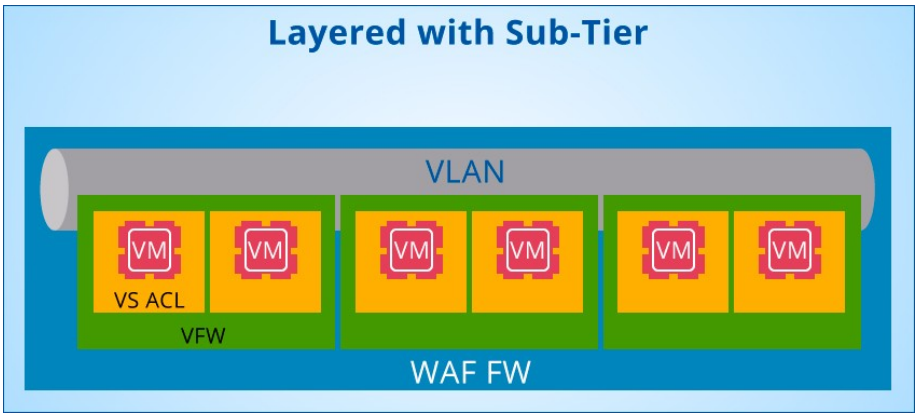

Figure 4: Sub-Tier Firewall

The goal of the sub-tier firewall is to provide a separate security boundary within the virtualization layer of the cloud, to secure the virtual machines and tiers of network created within the cloud network. The base policies for sub-tier firewall limit tier-to-tier network traffic. This traffic is limited based on source, destination port and protocol between tiers or virtual machines. Traffic will be limited to allow only the ports and destinations required for application/service functions; traffic from all non-required ports and IP addresses will be dropped by default.

4.2.1.3 Access Control Lists

The final level of access security is the Access Control List (ACL) on the virtual switch. ACLs are a basic security control layer to support securing virtual machines from standard layer 2 security threats like flooding and scanning.

The final architecture:

Figure 5: ACLs

This architecture provides three levels of network access security: from perimeter, to server sub tiers, to Virtual Network Interface Cards' (vNICs) Access Control Lists (ACLs).

4.2.2 Web Application Firewall

The Web Application Firewall service will provide application layer protection for web servers from known attacks via signature-based detection, along with some protection from unknown web-based application attacks via heuristic detection techniques. Examples of the types of attack these firewalls can prevent are the web application vulnerabilities listed by organizations such as SANs and OWASP with their constantly evolving list of top 10 attacks. Following the principle of "defense in depth," web application firewalling services should be deployed in addition to the traditional boundary firewalls that provide connection state-based network protection.

4.2.3 Secure Time

Time synchronization is vitally important to your organization. From a security perspective, effective prosecution of security incidents requires accurate matching of timestamps on all log files. Any discrepancies will complicate and delay incident response. There also is a reasonably large body of security software which requires accurate time information to work effectively. If you are a software development organization, correct time information across your NFS servers and clients can make or break your development–particularly if you use a parallel/distributed build product.

The CSP must provide accurate, cloud-wide time synchronization to ensure the security of your organization, both in terms of internal analysis of security data from event logs and other information sources, and when you wish to prosecute computer crime cases. CSPs need to support NTP across their entire cloud.

When deploying network security solutions from the cloud to the on-premises network, it is also important to ensure time is synchronized between the CSP's systems and the customer's, to enable investigation and log correlation across the customer's entire environment (cloud and non-cloud).

4.2.4 Network Routing and IP Management

Network routing is one of the most common attack vectors, as everyone uses standard routing protocols such as BGP and OSPF. These protocols can be fairly easily manipulated, and traffic then can be diverted and inspected. The recommended approach to limit the risk is to manage the IP space and access methods. TLS/SSL or IPSec VPN with IP whitelisting is the preferred access and IP management method for protecting traffic and data in the network of the CSP, or before entering the CSP network.

Ensure that the CSP uses secure implementations for any of its routing devices that are publicly accessible, including permitting them to connect to and receive update suggestions only from trusted routers, and implementing secure update policies.

4.2.5 DDOS Protection/Mitigation

A Denial of Service (DoS) attack is an explicit attempt by attackers to prevent legitimate users from using that service. Examples include attempts to:

- Flood a network
- Use all available processing power or resources (e.g., memory) on the end system
- Disrupt connections, preventing access to a service or disrupting functionality of the service
- Prevent a particular individual from accessing a service

Focusing on network security, the most prominent form of a DoS attack is probably the Distributed Denial of Service (DDoS), which in the past had been a volume-based attack, utilizing thousands of compromised hosts sending malicious traffic to the target either to exhaust networking resources or those of the servers hosting the application.

These attacks typically involve (sometimes secretly) installed software on a master computer and a collection of compromised zombie computers (bigger bot networks may utilize many control servers to improve their resilience to being taken down). The attacker hides its true identity and location by using these zombie machines to launch the attack. The attack results in denial of service to legitimate users because their infrastructure is overwhelmed with illegitimate requests.

Areas vulnerable to attack include:

- Routers
- Firewalls
- Web Servers
- DNS Servers
- Mail Servers
- Voice Over IP (VoIP) gateways
- Indirect Victims: elements that share the victims' network (for example, other VMs in a Cloud environment)

The following diagram outlines the primary areas on a network that are vulnerable to DDoS Attacks.

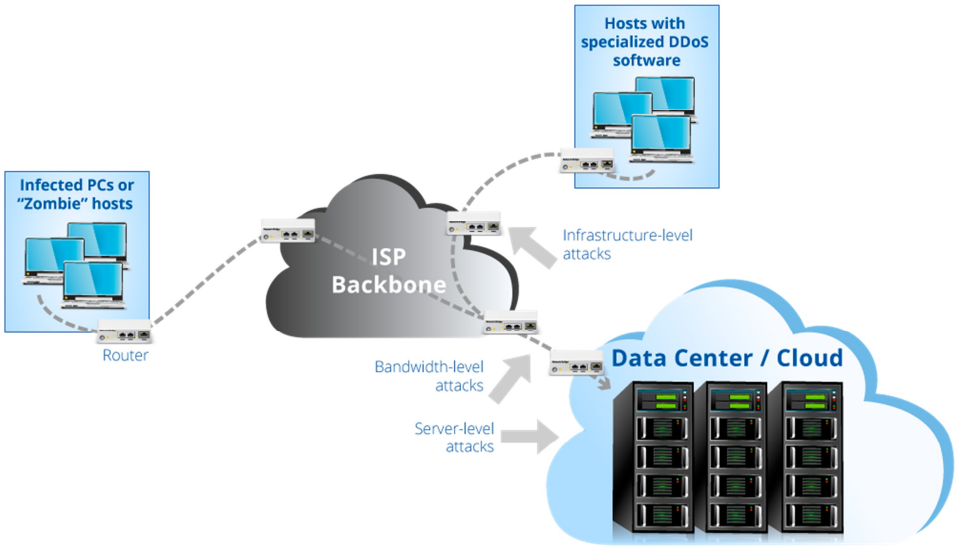

Figure 6: Different Target Infrastructure Layer

Typical (but fairly simplified) mitigation architectures are shown in the diagram below:

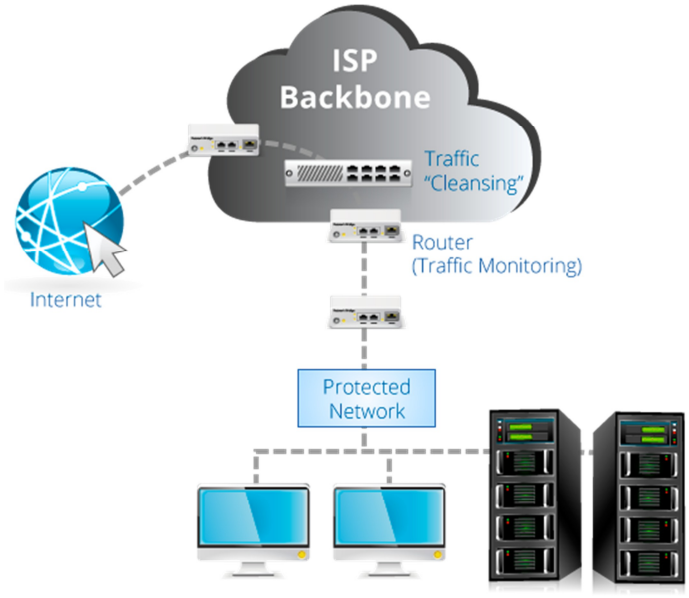

Figure 7: DDoS Mitigation – Traffic Monitoring

Volume-based DDoS attacks, filling the uplink pipe of a target, can only be mitigated at the CSP backbone or other network that has significantly more bandwidth than the sum of the bandwidth of all attack traffic.

During normal operations (figure 6), special traffic monitoring entities constantly try to identify traffic anomalies. Once an attack condition is identified, the monitoring entities trigger the network to re-route suspicious traffic through a cleansing instance that attempts to filter out the attack traffic while allowing legitimate packets to pass through to the destination network, as shown in Figure 7, below.

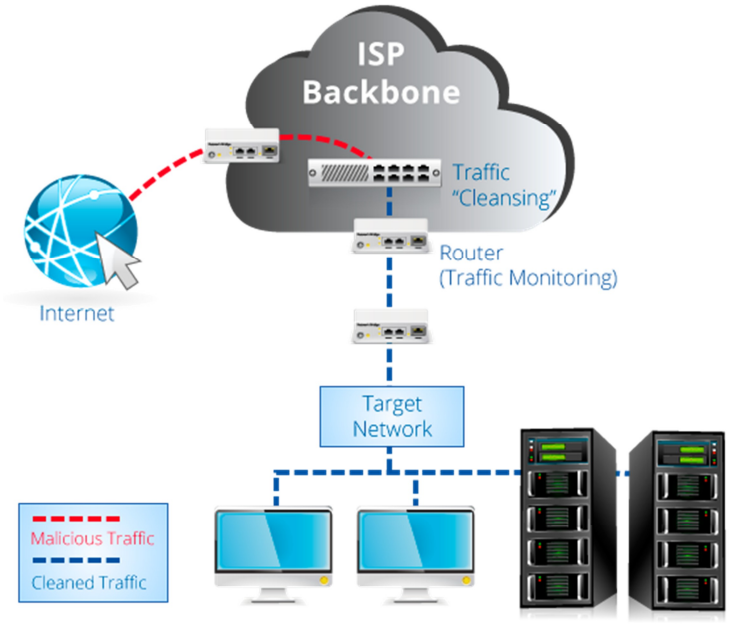

Figure 8: Traffic Cleansing

After the attack traffic has stopped, the routing reverts to normal.

4.2.5.1 Availability

Once your data is uploaded into the cloud, you cannot just walk over to a server to access it. In the case of any network issue, availability becomes crucial. Cloud services heavily depend on management portals and APIs to control almost every function delivered to cloud customer. Therefore, DDoS mitigation is essential to any cloud-based service.

4.2.5.2 Recommendations

Although DDoS is not new, the resurgence of hacking groups and the readily available tools to conduct such attacks have increased the number of attacks significantly. This increases the risk of becoming a target.

The cloud in particular depends on the availability of interfaces and data. This, along with the fact that many customers are hosted on the same infrastructure, means that robust DDoS protection and mitigation must be a key component of a CSP's networking services.

Recent observations and analysis of DoS attacks seems to show a trend that attacker move from volume-based attacks towards application level DoS attacks, triggering resource overload by sending malformed calls to web-application that will cause high utilization on the web-, application- or database server, rather than flooding the network.

The appropriate response to an application level DDoS attack would be a reverse proxy or filter grid, sometimes called Web-Application-Firewall (WAF).

DoS Mitigation Recommendations and Guidance:

- Implement volume-based DDoS mitigation at a point in your network with enough bandwidth to exceed the volume of the attack.
- Implement volume and application level protection.
- If your cloud is multi-homed (multiple ISP gateways via BGP), ensure all traffic passes through a DDoS mitigation grid. Ensure the CSP does filtering, or engage a SecaaS provider for all traffic.
- Detecting an attack pattern is sometimes not an easy task, and accurate filtering requires a thorough profiling of legitimate traffic. Ad-hoc, emergency filters may not work well. False positives may occur.
- Test all profiles.
- Have a response procedure in case legitimate traffic gets blocked. Determine how fast the configuration of the mitigation filter can be changed.
- Different applications may have quite different, sometimes very dynamic usage/traffic profiles. Determine if the profiling algorithm can handle all applications on the same subnet or if they must be split.
- Ensure that the CSP allows different deployment models, including constantly protected, on-demand and probably “emergency” mode with almost no time for profiling upfront, in order to enable services to be brought up in alternate data centers if this is required to ensure service availability if a specific data center is under prolonged DDoS attack.
- Implement active notification in case suspicious activities are detected and re-routing/filtering kicks in.
- A customer should have access to relevant monitoring, alerting, and network performance reports and metrics.
- Some solutions offer a “real-time” optimization interface.
- The solution must fit the customer’s mitigation requirements with regard to:
    - Max mitigation bandwidth
    - Sessions
    - Packets per second
    - Attacker- and victim IP address pair
- The solution should support IPv6.

4.2.6 VPN and MPLS Connectivity (SSL, IPSec, VPLS)

4.2.6.1 SSL VPN

An SSL VPN is a form of VPN that can be used with a standard Web browser. It does not require the installation of any client software on the end user's computer. It is used to give remote users access to Web applications, client/server applications and internal network connections.

A virtual private network provides a secure communications mechanism for data and other information transmitted between two endpoints. An SSL VPN consists of one or more VPN devices to which the user connects by using a Web browser. The traffic between the Web browser and the SSL VPN device is encrypted with the SSL protocol. TLS-based SSL VPNs are more secure, as TLS is a more secure protocol that is gradually replacing SSLv3.

4.2.6.2 IPSec VPN

IPsec is a framework for a set of protocols for security at the network layer that is useful for implementing virtual private networks and for remote user access through private networks.

IPsec provides two choices of security service: Authentication Header (AH), which provides authentication of the sender of data, and Encapsulating Security Payload (ESP), which provides both authentication of the sender and encryption of data as well.

4.2.6.3 Border Gateway Protocol

The Border Gateway Protocol (BGP) allows PE routers to communicate with each other about their customer connections. Each router connects to a central cloud, using BGP. When new customers are added, the existing routers will communicate with each other via BGP, and automatically add the new customers to the service.

4.2.6.4 Label Distribution Protocol

Also known as a Layer 2 circuit, this method uses LDP to communicate between PE routers. In this case, every router connects to every other router in the VPN.

4.2.6.5 VPLS

Businesses considering MPLS VPN services should factor differences between Layers 2 and 3, and opt for the one that best fits their own needs.

Layer 2 MPLS technology is limited because it does not scale as well as Layer 3. Layer 2 network services are simpler in architecture and allow customers to retain control of their own routing tables.

Layer 3 MPLS VPNs are characterized by fully meshed architectures that enable multicast conferencing for projects involving a dispersed work group. In some Layer 3 offerings, the service provider takes over all WAN

routing. Outsourcing of routing tables is sometimes seen as a weakness of Layer 3 VPN services, because corporations may not be willing to relinquish control or share their routing schemes.

4.2.7 IDS/IPS

Network Intrusion Detection Systems (IDS) and Network Intrusion Prevention Systems (IPS) are usually placed at ingress and egress points of the network in order to detect and prevent anomalous traffic, usually based on a combination of signatures, heuristic behavioral analysis, and statistic protocol anomaly detection.

Intrusion Prevention Systems actually prevent attacks, rather than providing just logging and alerts to attacks, which is the function of an Intrusion Detection System. An IPS responds to malicious threats by initiating a defense which may include dropping malicious packets, resetting connections, or blocking traffic from a specific address. An IPS also can be set to monitor only, like an IDS, so new rules and heuristic rules can be tested and embedded prior to the system being allowed to actually respond to malicious use of the network. Where possible, the devices should be placed in-line, as this enables them to drop traffic. Dropping traffic is much more effective than sending reset packets, which is the only response available to devices not placed in line.

The main considerations for any IDS/IPS deployment whether cloud-based or on traditional networks are to:

- Understand where the devices need to monitor. This is likely to be ingress and egress points from the network, and potentially points between the server network or key server networks and the rest of the environment.
- Understand the performance requirements. Any IDS/IPS device must be capable of handling the volume of traffic that is expected to pass through it in order to be effective. These devices usually will "fail open," so as traffic increases beyond their capacity to analyze, they will not stop system functionality, but system effectiveness will rapidly degrade.

Within a CSP's network, the customer should define the locations to be monitored, and agree to service and performance levels, along with how rules are added and managed.

When deploying on-premise to monitor your local systems, whether they are traditional or virtual/private cloud deployments, the role of the customer likely will be more complex, as there will be a greater need to understand data volumes; determine which devices should be deployed; and weigh the benefits of on-site log correlation against the higher volumes of data that must be transferred to the CSP.

Secure management, transport (to which locations), segregation and analysis of collected data must be considered and defined in any contracts with the CSP.

4.2.8 Threat Management

Risk management is now one of the principal areas of focus for CIOs, CISOs, and CFOs. In the past, periodic vulnerability scanning, along with the requisite OS and application patching process, was considered to be sufficient to uncover and manage risks. Today, due in large part to a growing number of regulatory compliance

mandates, corporations are compelled to understand, manage and report risks to the confidentiality, integrity, and accessibility of their critical data with more granularity and reliability.

As companies struggle to stay competitive, additional pressure to grow infrastructure to accommodate partner companies with extranets, provide more Internet access to assets, and increase global connectivity in general, all generate new risks and expose existing vulnerabilities to threats that were previously manageable. From a risk management perspective, such growth, while seemingly positive from a business perspective, can represent a serious challenge for executives who must balance revenue growth with an increased exposure to threats.

Risk = Threats x Vulnerabilities x Asset Values

It is impossible to eliminate all risks. The only reasonable goal is to manage risks to an acceptable level. Determining what is acceptable is an impossible task in itself without the right process and technology. Best practices dictate prioritizing threats, and focusing on the most important of them. Providers should address this problem, in part, by providing customers with the means to manage risks within the scope of the existing (or proposed) infrastructure.

One important part of the threat management life-cycle is vulnerability management. As vendor products will continue to be vulnerable, and human beings will continue to introduce configuration errors from time to time, networks will be vulnerable. The added complexity (technical, management processes and new admin roles) of a virtualized environment might not improve the situation, but increased automation might.

The goal is to keep the window of vulnerability as small as possible, until mitigation measures can be applied. As in non-cloud environments, regular, non-disruptive vulnerability scans are recommended, but require due care, professional skills and mature tools, as disrupting network services of a cloud infrastructure may impact many tenants.

The patch management of virtual network components might require new processes in the cloud, as this likely will be done via the virtualization platform management interface. However, traditional network technology vendors are providing virtual network components which might smoothly integrate into the existing network management infrastructure, tools and network administrator's skill set.

4.2.9 Forensic Support

Network forensics provides information useful in aiding an investigation, or incident response, addressing "hacked" (compromised) systems. Network forensics is a mature area, and virtualization typically does not change its tactics.

However, as network components like routers, switches, firewalls or intrusion detection devices are being virtualized, forensic monitoring and collection of logging data that is forensically sound is becoming more challenging. This is due to the fact that virtual switches, firewalls, and any other virtual devices rely on the same hypervisor. If the hypervisor becomes in any way compromised, there is a question regarding how trustworthy data from any virtual device is. For this reason, the security of the underlying hypervisor is absolutely key to all cloud security efforts and concerns.

A sound forensic process includes, but is not limited to:

- Evidence Acquisition
- Investigation and Analysis
- Reporting

This section focuses on implementation guidance regarding network forensics support within a virtualized environment.

4.2.9.1 Logging

Configure proper logging for all relevant network components that might be virtualized. Capture log content that is required to support network forensic investigations; more is better. Forward all logs to a central, hardened log-server and protect logs and server using hashing logs, file system integrity controls, tight monitoring of access. Protect logs in transit using a secure network path and/or encryption. Implement monitoring that generates alarms and send automatic notifications in case an event feed stops. Apply automatic analysis, correlation and visualization; otherwise, storing a vast amount of data is meaningless. Consider the retention period and log backup from a compliance and commercial perspective. Using a hosted log service might be an option.

SIEM products are evolving to become better integrated in and with the cloud. Cloud providers are fed data from different customers who expect their data to be protected, segmented from other customers, controlled, secured, and monitored. A cloud provider should not access customer data for their use or benefit, unless specifically allowed by the customer contract.

4.2.9.2 Capturing Network Traffic

In most cases, capturing real-time network traffic will be required, ensure the virtualized infrastructure can support this. Recently, virtualization vendors added mirror port capabilities to their products, thus aiding the capture of network traffic by monitoring devices. To support network forensics, make sure virtual switches support sniffing or the topology and solution design allows the trunking out of relevant traffic.

4.2.9.3 Deployment Scenarios

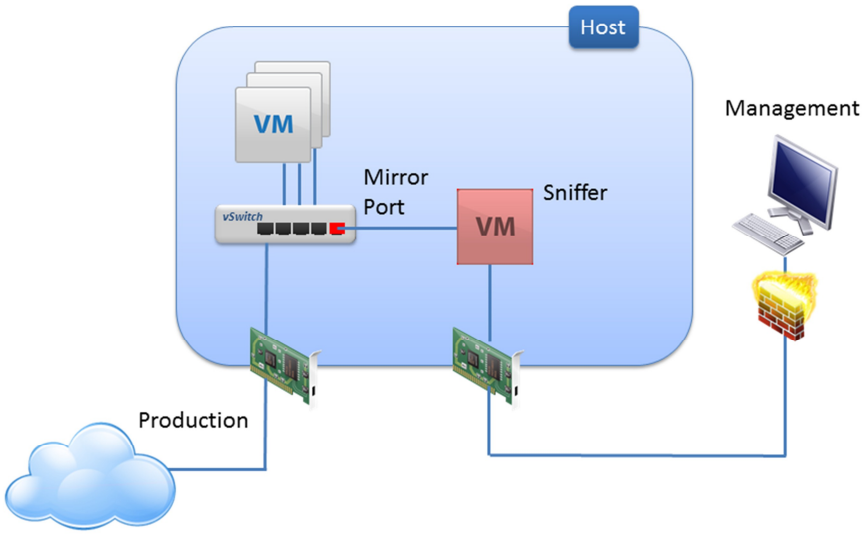

Figure 9: Virtual Deployment

In the figure above, a virtual switch with mirror port capabilities is used to capture traffic and a separated VM acts as a sniffer, managed over a dedicated pNIC.

Advantage:

- Easy deployment

Disadvantages:

- Not all vSwitches support mirror ports.
- The sniffer VM might generate high load on the host (in particular if some analysis, de-duplication or correlation scripts are active).
- If the hypervisor is compromised, the sniffer can no longer be trusted.

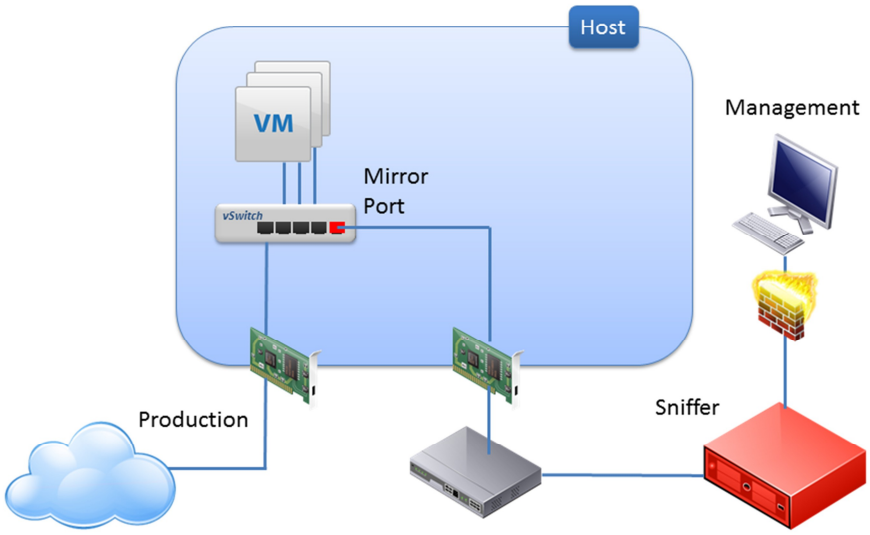

Figure 10: Physical Sniffer Deployment

Within this deployment, the sniffer is a physical appliance outside the hypervisor.

Advantages:

- Physical Sniffers are independent from the hypervisor, and not susceptible to compromise via the hypervisor.
- High load on the sniffer does not impact virtual machines or the host.

Disadvantages:

- Not all vSwitches support mirror ports.
- Dedicated sniffer hardware and sometimes scarce spare physical NIC port required.

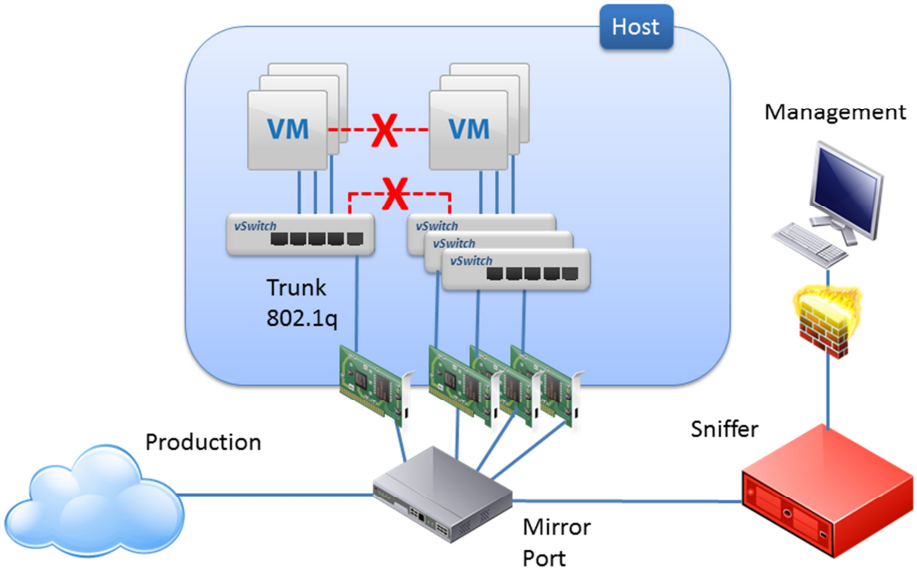

Figure 11: Trunk out all traffic

The deployment shown above trunks out all VM traffic for inspection outside the hypervisor. This works only if direct communication (red dash-lines) is disabled.

Advantages:

- Sniffer is independent from hypervisor
- High load on the sniffer does not impact virtual machines or the host.
- Works with all virtual switches.

Disadvantages:

- Most hardware intensive.
- Dedicated sniffer hardware and sometimes scarce spare physical NIC port required.
- Direct VM-2-VM or vswitch-2-vswitch communication will not be seen (and could easily be enabled by accident or intention).

To tap into "all" communication channels within a virtualized environment, pay attention to specific interfaces that might allow direct communication between entities like VMs, bypassing all network interfaces (virtual and physical). An example of a channel that may remain internal to the physical host and not be obviously visible to

monitoring tools would be the “Virtual Machine Communication Interface” VMCI for the VMware hypervisor. As these interfaces typically increase the attack surface, consider disabling them if not required.

Carefully consider what should be captured, recorded and stored in a virtual environment. Sniffing on a vMotion network might disclose (and store somewhere temporarily for investigation) sensitive information like clear text passwords or keys or credit card numbers, etc., as they are included within the VM’s RAM being transferred in clear text.

5.0 References and Useful Links

5.1 Useful Links

http://www.net-security.org/article.php?id=1509