cohesivenetworks

White Paper

Cloud Security Best Practices

Secure enterprise applications through a cloud-based approach to defense in depth

Ryan Koop, Director of Products and Marketing at Cohesive Networks

September 2015

Securing Enterprise Applications Through A Cloud-Based Approach To Defense In Depth

Executive Summary

Regardless of cloud deployment model (public, private or hybrid cloud), all organizations need security for critical applications and sensitive data. Defense in depth has never been more prevalent than with cloud computing and third party interactions with critical data.

Layers of security can bolster defenses for any application, database, or critical data. In a data center, application owners could physically segment networks and built walls around data. As data centers virtualized, application owners could add logical segmentation at the virtualization layer.

Networking in the application layer is about bringing the network close to the application, and giving control over the network and its configuration to the application owners. Finally, cloud application owners can now dictate specific security rules for each application in network layers 3-7.

First, application owners should selection cloud providers that have published security policies, industry certifications, and recognition. Next, cloud users should take advantage of cloud provider settings that provide additional isolation and network controls for traffic to and from their applications. Finally, application owners must use application-layer security and segmentation they alone own and control.

By adding VNS3, application owners can create an overlay network over the top of a provider's network. This network depends on the native layer, but is fully owned and controlled by application owners. Likewise, IPsec (Internet Protocol Security) tunnels are a vital addition because users can control encryption keys and verify traffic as it travels across the public internet, cloud regions and to third party environments.

Because it offers completely unique application layer security features, VNS3 allows application owners to manage their own authenticated, encrypted SSL tunnels. When used in combination with cloud provider security features, VNS3 security make applications more effective.

Layers of security in cloud computing

Defense in Depth

As more enterprises use cloud computing, circling the wagons around critical data is just impossible. With applications and data spread across the internet, an organization just cannot build a single perimeter around all of their resources.

Including all SaaS-based accounting, CRM, or email services in the cloud category, most organizations have been using cloud computing for years. In early 2015 RightScale reported 93% of enterprise respondents are adopting cloud (88% public cloud, 63% private cloud, and 58% hybrid). Business units, not corporate IT, have been the driving force behind data moves outside of the protected data center, and are now realizing the need for security and control.

Regardless of cloud deployment model (public, private or hybrid cloud), all organizations need security for critical applications and sensitive data. Defense in depth has never been more prevalent than with cloud computing and third party interactions with critical data.

Adding layers of security to base compute and network services can bolster security for any application, database, or critical enterprise resource. With data centers, application owners could physically segment networks and built walls around data. As data centers virtualized, application owners could add logical segmentation at the virtualization layer. Finally, cloud application owners can now dictate specific security rules for each application in network layers 3-7.

Leveraging a defense in depth strategy requires application owners to orchestrate their cloud provider security features, application and software requirements and the application owner's organizational security controls. Tim Phillips describes it as "virtual application networking" or a feature of the network that allows the application owner to define the requirements of each server and applications.

Shared Responsibility in cloud environments

As Gartner analyst Lydia Leong wrote, "IT managers purchasing cloud [Infrastructure as a Service] IaaS should remain aware that many aspects of security operations remain their responsibility, not the cloud provider's. Critically, the customer often retains security responsibility for everything above the hypervisor."

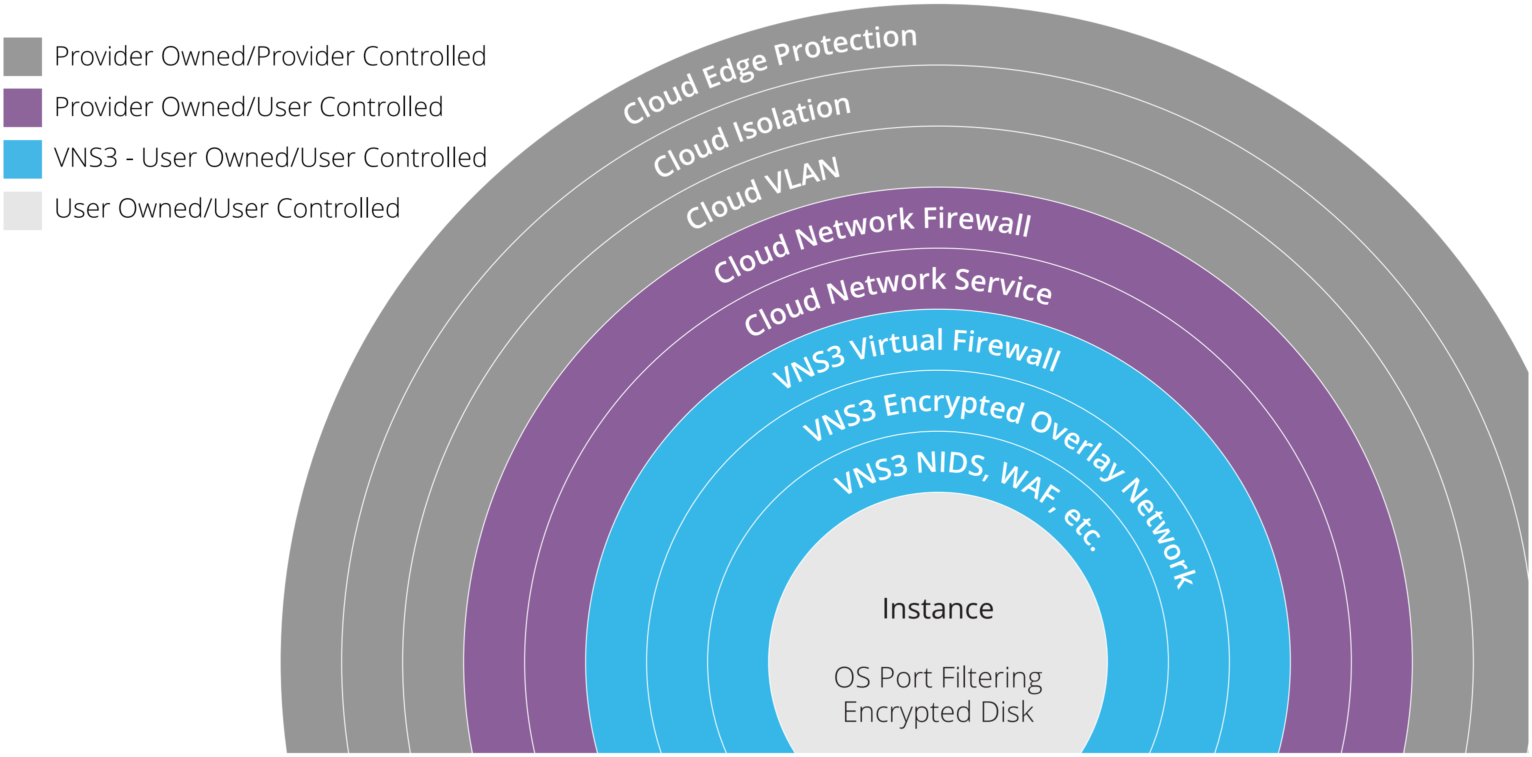

Provider-Owned/Provider-Controlled Security

Enterprise executives are now publicly touting public cloud data security, claiming no business can build, maintain, or secure data centers better than cloud providers. Cloud providers publish regular security white papers detailing certifications, security practices, and principal. Provider-owned, provider-controlled features (such as the physical data center security, the cloud edge, and cloud isolation) provide strong foundations for a layered defense in depth strategy.

Provider-Owned/User-Controlled Security

One of the first techniques to emerge in virtualized infrastructure was port filtering on the host operating system of the hypervisor / operating system itself (like AWS Security Groups). Port filtering prevents packets from ever reaching a virtual adapter. Public cloud providers allow users to control this hypervisor firewall through network mechanisms such as security groups or configuration files. Users can limit rules to only allow ports needed for each application.

VNS3 User-Owned/User-Controlled Security

VNS3 network virtualization allows application owners to control addressing, protocol, topology and security. Network virtualization solves the problem of unencrypted data in motion traveling over public internet or shared regions, beyond a public cloud providers' virtual network protections. VNS3 provides unique cryptographic keys for each host on the network, as well as additional network firewalls on the virtual network adapter. VNS3 encrypted virtual networks allow application owners to lock down applications independently of cloud provider settings.

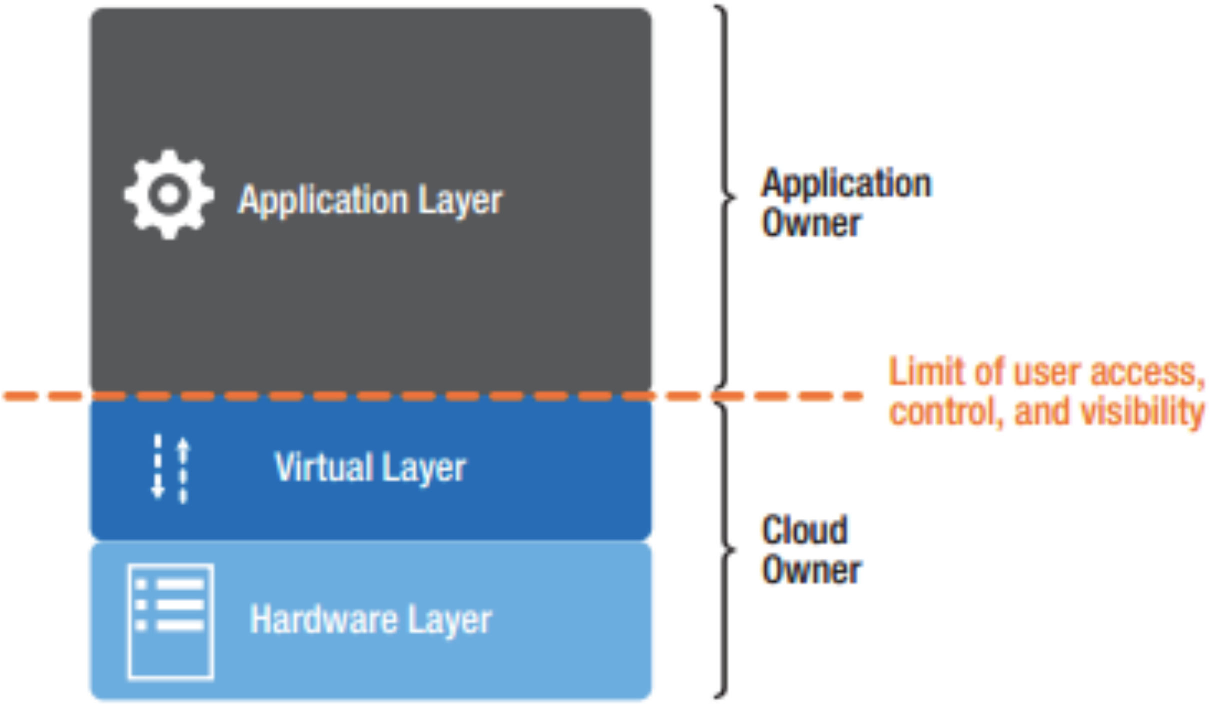

Virtual Application Networking

Network security settings at the hardware and virtualization layers are important when selecting cloud providers. Once application owners set up resources in the cloud (regardless of cloud model), networking capabilities should focus on making applications more effective.

At the application layer, which sits above the line of user access, control and visibility, application owners are now fully responsible for security. Networking in the application layer is about bringing the network close to the application, and giving control over the network and its configuration to the application owners.

Cloud providers offer the ability to run virtual machines (VMs) of various shapes and sizes. Building on the layers of control and security, If users are able to bundle network functions inside of VMs, cloud users can better control security without relying too heavily on the cloud provider. Cloud providers may give some control over underlying network, such as virtual private cloud (VPC), but application layer networking provides capabilities above and beyond the limit of user access, control, and visibility.

Core Network Functions at the Application Layer

The core network functions are the ways to build, contain, and connect networks: through switching, routing, firewalls and virtual private networks (VPNs). By bundling core network functions in the application layer, application owners can effectively build an intranet for an application. This application-centric approach can be very powerful in terms of security, compliance and risk management.

Network Application Services

Applications need services like SSL/TLS termination, load balancing, caching, proxies and reverse proxies. All of these functions were previously delivered as a physical devices plugged into the data center network. Without a data center and a network team inside the cages, application owners can still use these services in software.

Using network application services inside virtual machines in the application layer allows application owners to size, scale and tailor security to a single application. Rather than applying complex, blanket security policies across applications within a typical enterprise network, cloud network operators can simplify management.

Performance is less of an issue when security policy can be tightly matched to the use case at hand. Core network functions at the application layer give cloud users a huge amount of control over their networks, and allows much more streamlined security policy to target specific risks, applications and data.

Lock down the underlay

Public cloud data centers are better equipped to to defend against sophisticated cybersecurity threats, from walls to guard against physical attacks all the way to massive teams of experts to combat hackers. The very public and very damaging data breaches - the US Government's Office of Personnel Management (OPM), Ashley Madison, Sony, TalkTalk, and more - all were in private data centers.

While there is more to cloud providers' offerings than just security, cloud users should be wary of any provider that does not publicize security policies. Major cloud providers regularly publish security white papers detailing certifications, security practices, and cybersecurity trends. FortyCloud also published a security matrix, comparing security offerings across cloud providers.

Provider-owned, provider-controlled features (such as the cloud edge, cloud isolation) provide strong foundations for a layered defense in depth strategy. These critical security offerings include policies for data center physical security, server and software stack security, regular updates and patches, and proper data access and disposal. AWS publishes a Risk and Compliance white paper detailing their external certifications. Likewise, the Google Compute Engine team, Rackspace, and Microsoft Azure list their security policies and compliance on their websites.

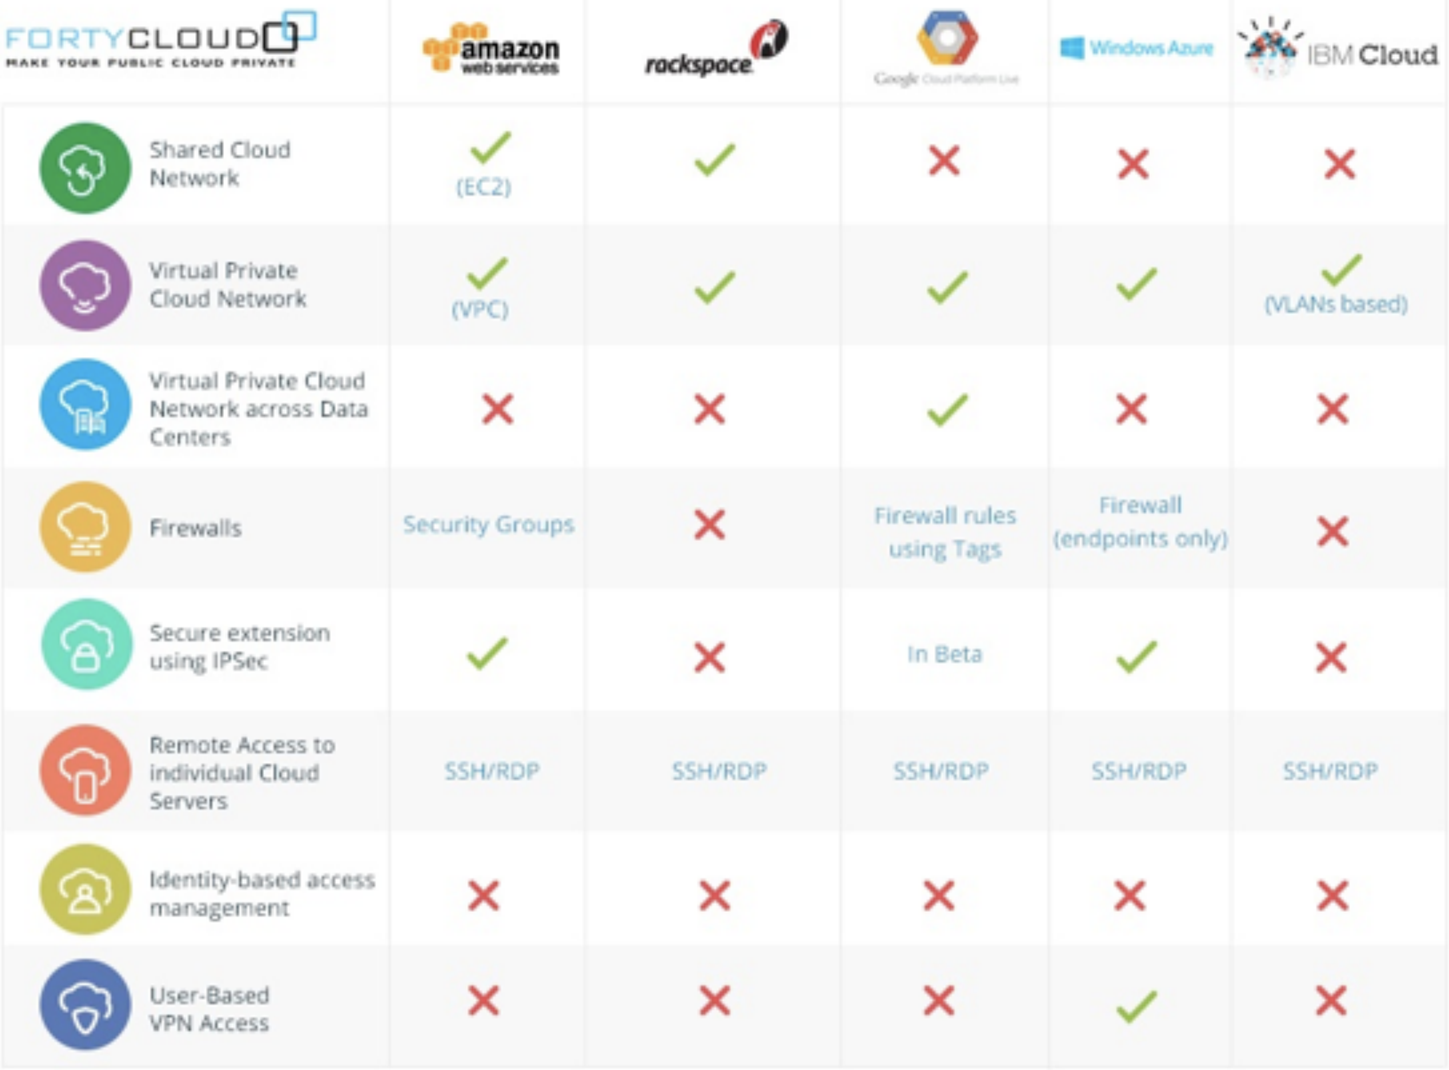

FORTYCLOUD MAKE YOUR PUBLIC CLOUD PRIVATE	amazon web services	rackspace	Google Cloud Platform Live	Windows Azure	IBM Cloud
Shared Cloud Network	✓ (EC2)	✓	✗	✗	✗
Virtual Private Cloud Network	✓ (VPC)	✓	✓	✓	✓ (VLANs based)
Virtual Private Cloud Network across Data Centers	✗	✗	✓	✗	✗
Firewalls	Security Groups	✗	Firewall rules using Tags	Firewall (endpoints only)	✗
Secure extension using IPSec	✓	✗	In Beta	✓	✗
Remote Access to individual Cloud Servers	SSH/RDP	SSH/RDP	SSH/RDP	SSH/RDP	SSH/RDP
Identity-based access management	✗	✗	✗	✗	✗
User-Based VPN Access	✗	✗	✗	✓	✗

Cloud Provider Network Security Musts

The next layer up from cloud provider data center and native network layer security offers are owned by providers but controlled by cloud users. Here the shared responsibility model begins to shift toward application owners' responsibility. These features are available for all users, but must be maintained and customized by each account owner.

Cloud users should take advantage cloud provider settings that offer additional isolation and network controls for traffic to and from their applications. Three key ways to control network traffic in shared environments are virtual private clouds (VPCs) or VLAN isolation, port filtering, and static assignable public IP addresses.

VPCs and VLAN Isolation

Both VPCs and VLAN isolation are ways to limit the ports available to communicate in a network. Limiting ports and interactions allows application owners to simplify network management and prevent any eavesdropping.

A virtual private cloud (VPC) is a more secure network within the shared infrastructure. Two cloud providers offer VPC features: Amazon VPC and VMware vCloud Virtual Private Cloud OnDemand. IBM Softlayer and Microsoft Azure do not explicitly offer VPCs, but users can configure virtual networks to create VPCs.

A VLAN (virtual local area network) is a way to imitate a limited LAN network with software-only networks. Virtual LANs isolate traffic by restricting port access to a set of "private ports". VLANs isolate a network so that traffic must flow through a trusted router, rather than directly between networks. Microsoft Azure users can create their own VLAN isolation by locking down network settings.

Port Filtering

Most major cloud providers offer a form of port filtering. In Amazon, ports are controlled using AWS Security Groups. Microsoft Azure users Network Security Groups. IBM uses parameters.xml. Google Compute Engine lists the settings under GCE Firewalls. As an example, Cohesive Networks recommends VNS3 users limit overlay network traffic to only UDP port 1194.

AWS, Google cloud, and Azure also allow users to manage their own network access control list (ACL). Network ACLs are lists of access rules for each subnet, evaluated in order. Network ACLs, like AWS security groups, have separate inbound and outbound rules to either allow or deny traffic.

Static and Assignable IP Addresses

Static and assigned IP addresses allow for better disaster recovery (DR) with little human intervention. In Amazon, static IP addresses are called Elastic IP addresses (EIP). Cloud users can speed time to connect and reconfigure settings by using EIPs rather than manually updating each connected device with new IP addresses. VNS3 users can quickly peer or re-peer VNS3 Controllers by assigning EIPs to VNS3 instances. If a VNS3 Controller reboots, the EIP will reconnect automatically.

Deep Dive: Amazon AWS Security Best Practices

In order to examine the specific settings of cloud provided/user controlled security best practices, this paper will examine the Amazon Web Services (AWS) public cloud security features. Our partner and AWS security expert Josh von Schaumberg, the Senior Solutions Architect at Trek10, recommends five Amazon AWS Security offerings all application owners can capitalize on immediately:

1. Disable and delete the root access key

When an enterprise creates an AWS account, the initial account created is the "root account." As the name suggests, this account has full access to do anything with the AWS account. Today, there is no reason a root account should be used for anything outside of a few administrative activities in the console, and the access key should certainly not be activated or in use. Instead, AWS' identity and access management (IAM) accounts should be created for all AWS access. Rather than deleting the access key, AWS users should first deactivate it to test if any issues occur in cloud application.

2. Secure admin ports to VPN-connected users only

Access to administrative ports on public facing servers should not be open to any public IP address. For example, ports 22 (SSH) on Linux instances and port 3389 (RDP) for Windows instances should be locked down to a private subnet — over a VPN connection only.

3. AWS Security Groups

When using overlay networks (with VNS3 or other networking solutions), cloud users can further limit non-public facing ports to an on-premise, private subnet only. This way malicious actors on the public Internet cannot attempt a brute force attack. Alternatively, some cloud users opt for full cloud migration to avoid any on-premise network attacks. For these users, VNS3 can allow for remote access VPN with secure, certificate-based authentication. Users on the road can connect straight into an instance using the AWS private IP address.

4. Force multi-factor authentication (MFA) for all users

After ensuring that MFA is enabled on the root account and there is no outstanding access key, all AWS users should configure MFA. When the "force MFA" IAM policy is attached to a user, it denies all other granted permissions until the user sets up MFA and then logs in using MFA.

4. Enable CloudTrail across all regions and deny access to logs

CloudTrail is an AWS service that used for audit logging; it records virtually every click in the web console, as well as each programmatic API call to AWS. CloudTrail is not enabled by default. To ensure no malicious actors can tamper with audit logs, AWS account owners can attach a policy to all IAM users to explicitly deny access to the CloudTrail bucket.

5. Use "roles" everywhere you can

AWS roles have many different uses. In general, roles are a way to give users or AWS infrastructure the necessary permissions to access other AWS services. With EC2 roles, users can launch an instance with the appropriate S3 permissions, and the application will search for an access key in the instance metadata, which is transparently provided by the EC2 role. This key is then rotated every few hours, making the code much more secure and easier to manage.

Security and Control at Layer 3-7

Overlay networks add security layers to cloud resources that application owners directly own and maintain. Unlike the underlying cloud provider network and security features, cloud users can own, manage, and control security over the top of cloud provider networks.

In the cloud data center, virtual instances (also called virtual machines, or VMs) run in a physical host connected to a switch. Each switch is connected to a firewall, then to a router, and this edge router provides the cloud's connectivity to the internet.

Each sites has a host connected to the public internet by switches, firewalls and routers. Each topology is made up of a combination of physical and virtual devices, multiple layers, and different types of virtualization.

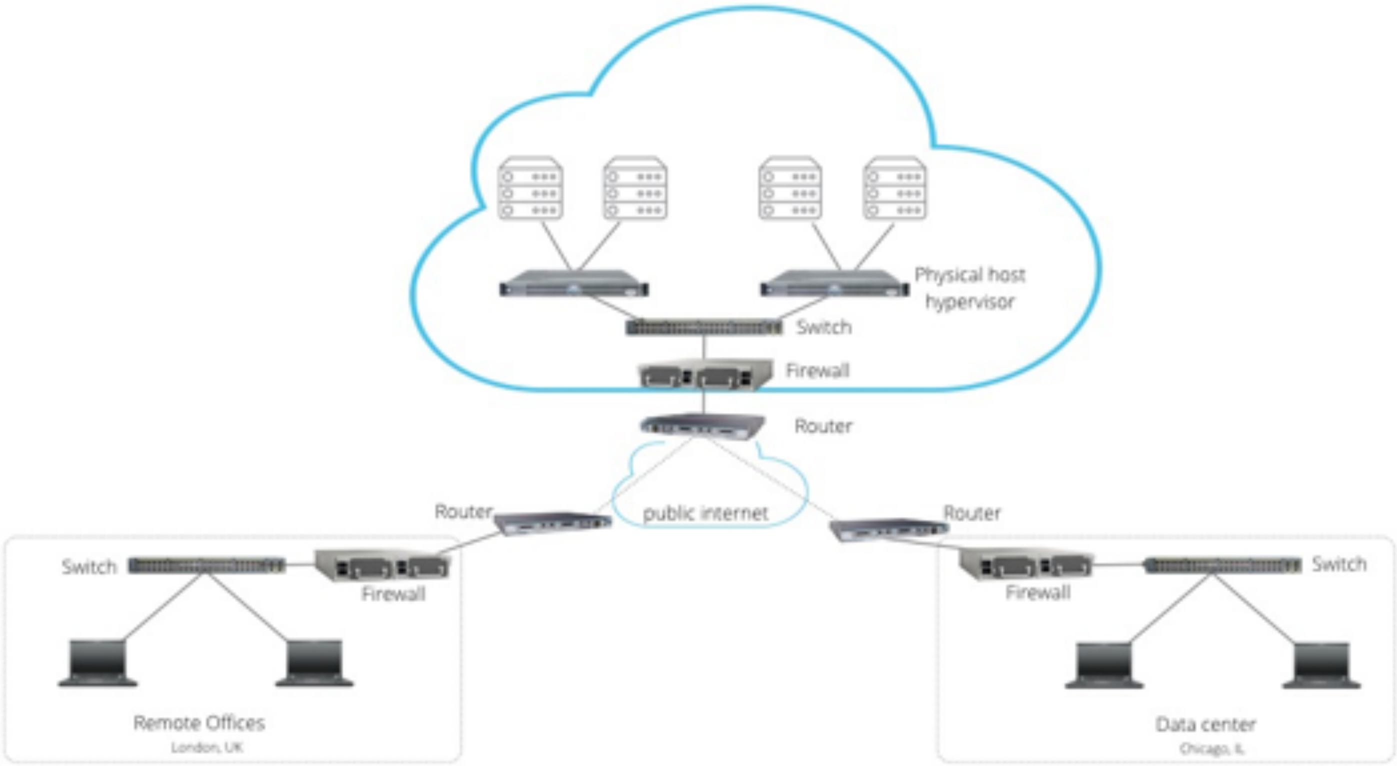

© 2016

VNS3 Overlay Networks

By adding VNS3 to the cloud network, an application owner can create a second layer 3-7 network over the top of the provider's network. Native devices are still present in the topology, and provide the critical physical and virtual connections. The overlay network depends on the native layer, but is fully owned and controlled by the VNS3 application owners.

VNS3 acts as a virtual switch, providing connectivity to each of the compute hosts and switches traffic between hosts. The data center hosts will see the VNS3 Controller as the next logical hop in the topology. The VNS3 Controller is configured with a subnet, giving it a range of IP address available for the user to configure the static allocation to client servers, allowing clients to join overlay network.

In this example, the traffic still travels on the native network, on the same native devices. Logically VNS3 is layered over these devices to provide an overlay network function. The application owner can select the IP address for each client server in the public cloud, and these overlay IPs are addressable from both London and Chicago.

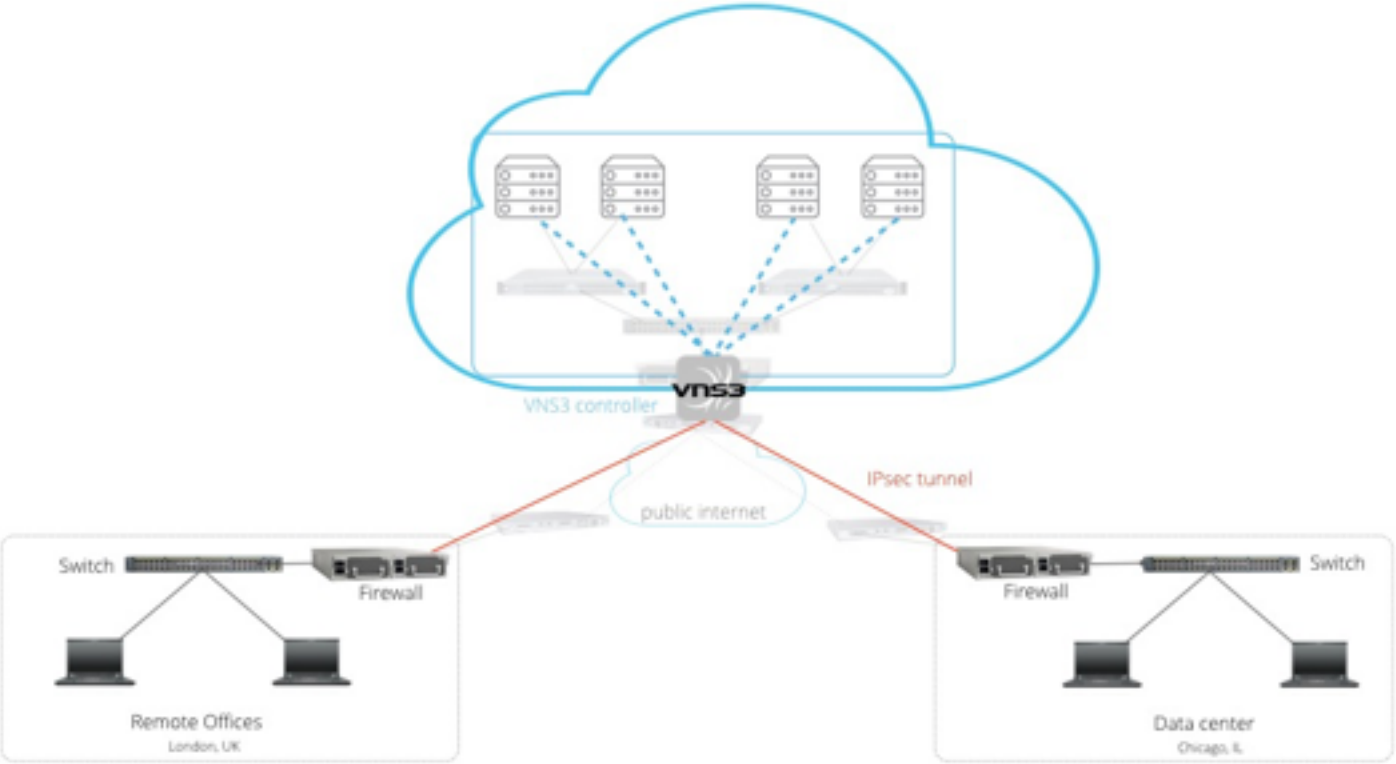

© 2016

Seal Off a VNS3 Overlay Network

When VNS3 users configure a cloud instance in the VNS3 topology, a secondary interface is created on the server called “tun0” which is layered over the existing native interface and all overlay network traffic goes via tun0. This is in addition to the existing primary interface usually labeled “eth0”. Usually eth0 can be shut off to all traffic other than traffic via tun0 and this creates a sealed network.

High Availability with Multiple, Peered VNS3 Controllers

Building on the VNS3 network topology example, application owners can create high availability by peering two VNS3 Controllers. The peered VNS3 devices exchange the topology's routing information and share client server credentials. Any connected client servers can connect to either Controller, and in turn both Controllers will still be able to access the clients, no matter which Controller the client server connects to.

Similarly, all IPsec tunnels connected to either Controller are accessible via both Controllers and all client servers. The two VNS3 Controllers function like a pair of high availability switches or routers running HSRP (Hot Standby Routing Protocol). The diagram shows active tunnels from each endpoint in continuous red lines, while dashed red lines are passive tunnels ready to take over.

Peered VNS3 Controllers provide overlay network failover and high availability, but not IPsec failover. The endpoint IPsec device should have its own high availability features (eg. Cisco ASA offers a multi 'Peerlist' feature so the ASA can drive the IPsec failover, and can detect if the IPsec tunnel is down).

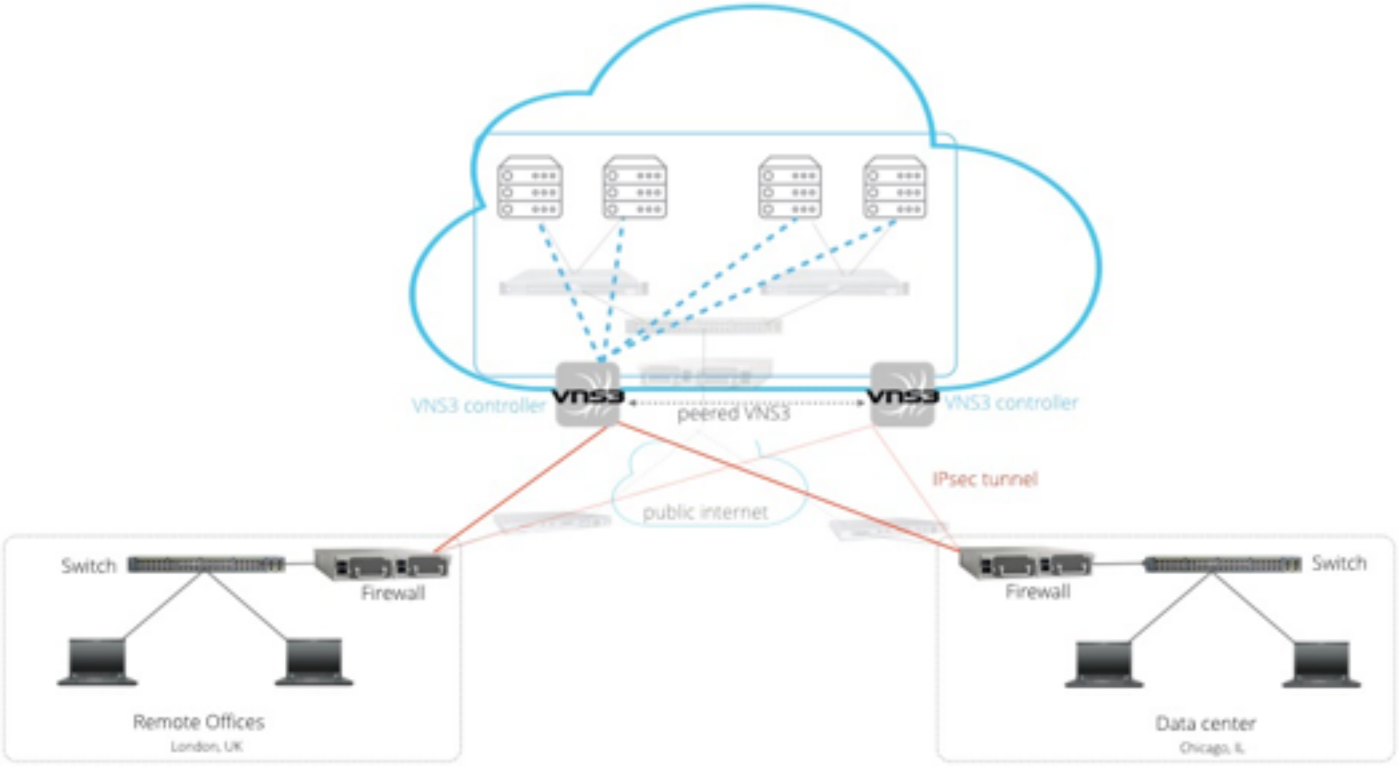

© 2016

IPsec Based Security

IPsec (Internet Protocol Security) is a protocol suite for securing IP (network layer) communications between peers by authenticating and encrypting each packet of communication.

In the VNS3 network diagrams, the two red lines represent IPsec tunnels from the VNS3 Controller to the two remote firewall devices. The London and Chicago locations are two different endpoints.

Site-to-Site Connections

IPsec endpoints are the remote devices that a VNS3 controller instance connects to. IPsec endpoints are typically extranet devices like Cisco ASA, Juniper Netscreen, or Palo Alto. IPsec tunnels are the actual remote-to-local subnet definitions VNS3 users configure for the IPsec Endpoints.

For example, a tunnel can connect an IPsec endpoint in a local subnet (e.g. Overlay or unencrypted VPC VLAN) to a remote subnet (e.g. your data center subnet, partner subnet, customer subnet).

Matching IPsec Tunnel Negotiations

IPsec communication is divided into two phases: Phase 1 initial negotiations and Phase 2 peer exchange. In phase 1, network peers find each other, trade parameters, and create session keys (using the Internet Key Exchange Protocol or IKE). In Phase 2, the peers trade parameters and create an encrypted tunnel using IPsec Protocol Encapsulating Security Payload (ESP) to securely encrypt the network traffic (packets).

In Phase 1, the peers exchange a number of parameters in order to authenticate their connection, including NAT-Traversal (NAT-T), pre-shared key (PSK), Phase 1 Hash, and so on. These interactions, or negotiations, create a secure connection from end to end. VNS3 users can define the endpoint to any device in a network, subnet, cloud or data center. The key is to explicitly match all IPsec parameters, from tunnel definitions to session keys between endpoints.

The VNS3 IPsec Configuration Guide and Troubleshooting Guide outline the step by step instructions to set up IPsec tunnels with VNS3, as well as what types of Phase 1 and Phase 2 settings are allowed.

Encrypt SSL/TLS Tunneled Overlay Network Traffic

IPsec tunnels add direct connections to cloud-based resources. Unlike the network access settings at the cloud provider layer, users can both direct traffic and control the keys to encrypt the traffic as it travels across IPsec tunnels. For example, cloud providers offer data-at-rest encryption (data stored in cloud servers) but do not provide data-in-motion encryption (data traveling across networks, regions, or public cloud).

VNS3 uses Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL) cryptographic protocols to secure network traffic. This means all traffic flowing across the tunnel is private, authenticated and encrypted. A VNS3 IPsec tunnel is a vital addition to cloud security because users can verify there has not been any eavesdropping and tampering in transit from IPsec device to the cloud.

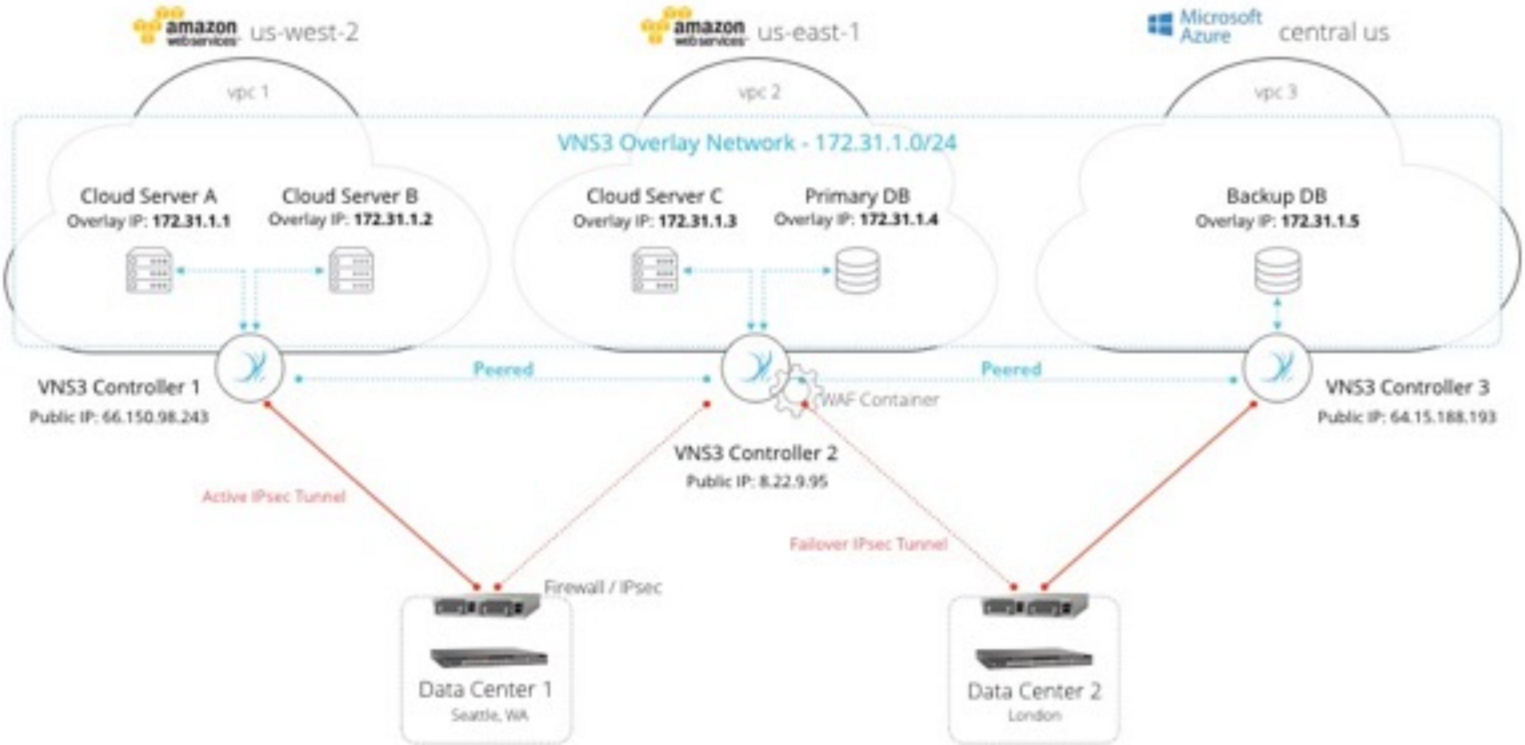

Unique Application-Layer VNS3 Security

Clientpacks

When first configuring a new VNS3 Controller, application owners begin by generating unique X.509 cryptographic keys. These keys, called clientpacks, allow each instance to authenticate connections with each Overlay Network IP. Clientpacks, used along with an SSL client (such as OpenVPN), connect endpoints to the overlay network using a specific IP address over an encrypted SSL tunnel.

Clientpacks are available as a single configuration file optimized for Linux (vnscubed.conf) and Windows (vnscubed.ovpn). Embedded within these files are the certificates and keys needed to connect to the Controller. Only end VNS3 users can view, edit, create,or save Clientpacks and the corresponding security token. Best practices for Clientpack are to tag, disable and regenerate as needed for your topology.

Tagging clientpacks allows application owners to map a clientpack to a particular endpoint or server, reinforcing logical subnet decisions. Tags also allow application owners to group clientpacks by function, and better organize networks as they grow.

Because each Clientpack is tied to a specific Overlay Network Address, application owners have the ability to disable any unused clientpacks. Disabling Clientpacks can prevent unwanted network activity on that IP address. Shutting down or disconnecting an endpoint or client from the topology also allows VNS3 users to reuse the corresponding Clientpack.

Finally, application owners can regenerate an old, lost or compromised Clientpack to increase the usability of the Overlay Network for road warrior VPNs. Regenerating Clientpacks deletes the old record and generates a completely new and unique key associated with the same address. In the VNS3 web UI, application owners can toggle each Clientpack list as available or in use (checked out) to keep track of used credentials and to prevent an API call from fetching a pack already in use.

Secure Account Access

Establish user roles to create the necessary permissions to access VNS3 products. The VNS3:ms product, designed for complex networks and multiple network users, is an easy way for application owners to manage and monitor VNS3 networks, VPN connections and underlying VLAN settings. VNS3:ms offers user authentication and role management for groups. With all VNS3 editions, application owners can secure access with separate UI and API passwords and multi-factor authentication.

Separate UI and API passwords

Each new VNS3 Controller comes with default login settings, and clearly best practices are to immediately changes both usernames and passwords. Because VNS3 answers API calls on the same port 8000 as the web interface, make separate passwords for the API. Cohesive Networks does not have any key access or remote access to any VNS3 Controllers, so the Support team cannot recover usernames or passwords.

Multifactor/multiparty authentication for Remote Support

Each VNS3 Controller runs a restricted SSH daemon, with access limited only to Cohesive for debugging purposes. Only VNS3 users can log in to the web UI to toggle the Remote Support function on. Ensure application owner teams use multi-factor or multiparty authentication for all remote support interactions.

In the event Cohesive Networks needs to observe runtime state of a VNS3 Controller for support requests, application owners will need to open Security Group access to SSH from the Cohesive support IP range and toggle Remote Support via the Web UI. Next, Cohesive sends an encrypted passphrase to generate a private key. Access to the restricted SSH daemon is completely controlled by the application owners. Once the support interaction is complete, the VNS3 users must disable remote support access and invalidate the access key.

VNS3 Snapshots for Configuration and Recovery

Snapshots are a compressed file stored on the VNS3 web interface, in a file that retains all the configuration settings of the VNS3 Controller. Once a VNS3 Controllers is configured and running, application owners can save the configuration with a runtime snapshot, then quickly reconfigure a new Controller with the same SSL Certificates and Keyset.

About VNS3

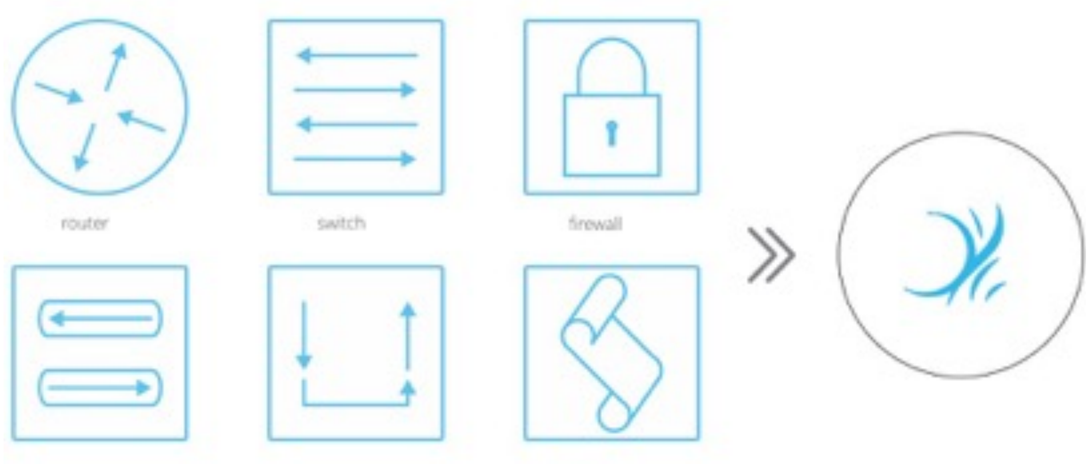

VNS3 products are priced based on network complexity. Unlike other networking products, Cohesive Networks measures network complexity by endpoints and secure IPsec tunnels. VNS3 editions are separated out by the number of IPsec endpoints, VNS3 Controllers, IPsec tunnels, clientpacks, or containers. See the pricing page for the specifics.

VNS3 is available immediately in the AWS Marketplace (click here for VNS3:vpn, and here for VNS3:net Lite Edition). and the Microsoft Azure Marketplace (click here for VNS3:vpn, and here for VNS3:net Lite Edition). For Google Compute Cloud, IBM Softlayer, CenturyLink Cloud and others, please contact Cohesive Networks support to deliver an image directly to a cloud account.

Summary

Provider-Owned/Provider-Controlled Security
Public cloud data centers have more advanced cybersecurity in place than the average organization. Major cloud providers should publicize their security innovations, certifications, and best practices. Provider-owned, provider-controlled features are a strong foundation for a solid defense in depth strategy.

Provider-Owned/User-Controlled Security
The next layer up from data center networks offers security measures owned by providers but controlled by cloud users, including VPC and VLANs, port filtering, static IP addresses, user role tracking, identity and access management, and multi-factor identification. Here the shared responsibility shifts toward cloud users. Layer 3-4 network security features are available for all cloud users, but must be maintained and customized by each account owner.

VNS3 User-Owned/User-Controlled Security
At the application layer, which sits above the line of user access, control and visibility, application owners are fully responsible for security.
With VNS3, cloud users can virtualize critical network security functions. and application owners can better control addressing, protocol, topology and security. By locking down all ports to and from an application, cloud users can verify that all traffic in and out of the cloud passes through VNS3 and in private, encrypted and authentic. No other cloud security device offers the same level of control and flexibility across so many cloud provider environments.

When used in combination with cloud provider security features, VNS3 networking capabilities make applications more effective.

About the Author

Ryan Koop, Director of Products and Marketing

Mr. Koop is responsible for product planning, customer engagement and marketing operations at Cohesive Networks. He is responsible for product development and manages teams for public relations, international events, and content marketing. His role spans the technical product development, customer support, business development and thought leadership needs of a growing company. Previously, he worked at a trading platform software company in the US Derivative Markets.

References

1. Weins, Kim. "Cloud Computing Trends: 2015 State of the Cloud Survey." Cloud Computing Trends: 2015 State of the Cloud Survey. 18 Feb 2015. Web. 15 Sep. 2015. <http://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2015-state-cloud-survey>.
2. Phillips, Tim. "Let Software Take the Strain off Your Data Centre." The Register. 28 Feb 2013. Web. 15 Sep. 2015. <http://www.theregister.co.uk/2013/02/28/datacentre_sdn/>.
3. "Lydia Leong - A Member of the Gartner Blog Network." Web. 15 Sep. 2015. <http://blogs.gartner.com/lydia_leong/>.
4. Asay, Matt. "The Mega-clouds Are Coming for Your Data Center." InfoWorld. InfoWorld. Web. 15 Dec. 2015. <http://www.infoworld.com/article/3014358/cloud-computing/mega-clouds-coming-for-your-data-center.html>.
5. Cheng, Ganti, Lubsey, Shekha, Swan, et al. "Open Data Center Alliance Master Usage Model: Software-Defined Networking Rev. 2.0." Open Data Center Alliance. Web 15 Sep 2015. < http://www.opendatacenteralliance.org/docs/software_defined_networking_master_usage_model_rev2.pdf>.
6. Linthicum, David. "The Public Cloud Is More Secure than Your Data Center." InfoWorld. InfoWorld. Web. 1 Dec 2015. <http://www.infoworld.com/article/3010006/data-security/sorry-it-the-public-cloud-is-more-secure-than-your-data-center.html>.
7. Naftali, Amir. IaaS Security State of the Industry - Comparing IaaS Providers. FortyCloud. 26 May 2015. Web 15 Sep 2015. < http://fortycloud.com/iaas-security-state-of-the-industry/>.
8. "Amazon Web Services: Risk and Compliance." Amazon Web Services. Dec 2105. < https://d0.awsstatic.com/whitepapers/compliance/AWS_Risk_and_Compliance_Whitepaper.pdf>.
9. "Google Security Whitepaper." Google Cloud. 26 May 2015. Web 15 Sep 2015. < https://cloud.google.com/security/whitepaper>.
10. "Rackspace security management." Web 15 Sep 2015. <https://www.rackspace.com/en-us/security/management>.
11. "Microsoft Azure Trust Center." Web 15 Sep 2015. < https://azure.microsoft.com/en-us/support/trust-center/>.
12. Palekar, Ashwin. "Network Isolation Options for Machines in Windows Azure Virtual Networks." Microsoft Azure blog. 28 March 2014. Web 15 Sep 2015. < https://azure.microsoft.com/en-us/blog/network-isolation-options-for-machines-in-windows-azure-virtual-networks/>.

Images:

- "Open Data Center Alliance Master Usage Model: Software-Defined Networking Rev. 2.0." Open Data Center Alliance. <http://www.opendatacenteralliance.org/docs>.
- IaaS Security State of the Industry - Comparing IaaS Providers. FortyCloud. <http://fortycloud.com/iaas-security-state-of-the-industry/>.
- All Other Images copyright Cohesive Networks