



Data Center Trends And Network Security Impact



Data Center Trends And Network Security Impact

Table of Contents

Introduction	3
Perimeter Firewall	3
Core Network Segmentation	4
Virtualization	5
Cloud Computing	6
Software-Defined Network (SDN)	8
Network Function Virtualization	9
Additional Considerations for Service Providers	10
Summary	11




Introduction

The data center is evolving rapidly with new technologies such as virtualization and cloud-computing, and software-defined networks. These have a fundamental effect on how network security is designed and deployed.

This paper gives a high-level overview of key trends shaping the data center and their impact on network security. The paper is divided into the following topic areas:

- Perimeter firewall
- Core network segmentation
- Virtualization
- Cloud computing (infrastructure-as-a-service)
- Software-defined networking (SDN)
- Network Function Virtualization (NFV)

Considerations for enterprises and service providers to select and deploy network security is discussed, as well as Fortinet's approach to delivering solutions in this new era.



Perimeter Firewall

The Perimeter Is Dead...Long Live the Perimeter!

The perimeter is porous. The enterprise is under siege. Web and e-mail are fat pipes for malware. Advanced threats are already inside the network. Users are mobile and bypassing the enterprise network. The perimeter is an M&M - a thin hard shell with a soft chewy interior. The perimeter is dead.

With all the talk of the demise of the perimeter, one would think that the notion of perimeter security is long gone. But to the contrary - in an interconnected world where there are no longer clear boundaries, a solid perimeter firewall is more important than ever. Rather than thinking of the perimeter firewall only as castle wall that must keep all the bad guys out with no defenses inside, today the perimeter firewall is more like a baseball field - a set of boundaries that establish how and where the game will be played. Without a clear set of bases and markings, of outfield and stands, a baseball game would be chaos. The field lets the players establish where they play offense and defense, while keeping unruly fans out on the sidelines.

The firewall thus establishes that clear deny-by-default boundary and the limited paths into the data center, keeping riffraff out while controlling the chaos of what enters. It anchors where additional layers of protection are then applied, whether at ingress/egress

points or deeper within the network. The perimeter firewall has not been made obsolete, it has become the baseline (quite literally derived from the paths between the bases of the baseball diamond) that shapes how other security layers are deployed.

Mobile Devices and the Internet of Things

With the proliferation of wireless productivity devices such as smartphones and tablets, the number of devices connecting to and accessing applications within the data center is exploding. This is increasing the burden of perimeter security as services are being accessed from anywhere and with greater traffic volume.

Mobile device traffic also may require more emphasis on *small packet performance*, as data center applications are geared more towards smaller screens and smaller bites of information. Some network security solutions achieve their performance specs with larger packet sizes, but can degrade significantly when the traffic shift towards a larger number of users and smaller packet sizes.

Ensuring Availability in a Service-Centric World

Web-based services accessible from the broader Internet will also increasingly become a target of competitors, activists, and others with a negative or political agenda. Widespread denial-of-service attacks are a highly visible means of disrupting business, and motivated interest groups no longer need to have technical sophistication themselves. Armies of botnets are readily available for rent out for *distributed denial-of-service* (DDoS) attacks from organized hacking groups, as long as those special interests have the means to pay.

As the data center becomes more user-centric, employees and customers will rely increasingly on services to be available on-demand. Enterprises therefore need to ensure their business-critical data center services can maintain accessibility from not just technical contingencies but also from motivated opposition as well.

Takeaways

- Baseline perimeter security
- Small packet performance
- DDoS protection

Product Options

- FortiGate
- FortiDDoS

Core Network Segmentation

Moore's Law and Increasing Speed

Network speeds continue to increase in a relentless Moore's Law fashion due to the pace of technological innovation. Always-connected mobile devices are accelerating this trend, as are virtualization and cloud computing. While it wouldn't immediately seem that consolidating servers more efficiently should have any net impact on the amount of network traffic, these technologies have made it easier for IT teams to provision new servers and quickened business team ability to roll out new projects - leading to real phenomenon such as *VM sprawl* and server containment. Cloud computing further empowers new services to "go viral" seemingly without regard to IT constraints on compute or network bandwidth.

With all this increased connectivity and access from anywhere, it is even more urgent that the internal network be properly segmented to ensure that external threats or improper access does not permeate the data center. At the same time core firewall segmentation must keep up with ever increasing speeds at the network core.


Next-Generation Interfaces - 40GbE and 100GbE

Wasn't it only just a few years ago that everyone was talking about getting networks ready for 10 Gigabit Ethernet? But things move quickly, and today 40 is the new 10. Indeed, in 2014 already 10GbE will be commonplace with 77% of organizations will be utilizing it in their networks, with 21% adopting 40GbE as well¹, according to a recent study by Network Instruments.

As core network speeds increase, the need for high-speed 40GbE and 100GbE network interfaces and high port density becomes critical, and network security appliances with higher throughput must also efficiently interconnect with high speed network fabric. Infonetics found that with typical firewall throughput requirements in the 100-200Gbps range and increasing, some businesses are even looking at skipping 40GbE and going straight to 100GbE security appliances, as more core network infrastructure becomes available with 100GbE ports in 2014 and 2015².

¹ "Sixth annual state of the network study", Network Instruments, 2013

² "High End Firewall Strategies, Infonetics Research



IPv6 Support

The inevitable march to IPv6 support is already underway in enterprise planning. While the proliferation of mobile devices is not the sole or even primary contributor, certainly it is a stark visual reminder that the world is running out of IP addresses. While enterprises are preparing networks for IPv6 support, not all are scrutinizing *IPv6 forwarding performance* carefully. As traffic migrates to IPv6, there is potential risk that network equipment may not keep at an equivalent rate to IPv4 speeds, causing network bottlenecks. It is therefore important when evaluating new network security devices to ensure that they not only support IPv6 but will also not degrade throughput substantially from IPv4.

Takeaways

- Moore's law increase in network speeds
- High-speed 40/100 GbE interface ports
- IPv6 forwarding performance

Product Options

- FortiGate

Virtualization

It's a Virtual-First World

Virtualization, more specifically x86 server virtualization as popularized by VMware and others, has dramatically transformed the data center in the last decade. What started as workstation technology primarily for testing, development and labs evolved into data center infrastructure for server consolidation - high utilization with efficient capital and operating expenses - and now into a core foundation for cloud computing.

² "High End Firewall Strategies, Infonetics Research, October 2013

Today the number of virtual servers in the world has long surpassed the number of physical servers, with virtualization not only acceptable in production environments but mission-critical. Enterprises are not just consolidating servers and racks, but often re- architecting entire sites and facilities with *data center consolidation* and transformation in mind and "virtual- first" policies – i.e. the planning assumption that any new workloads will be deployed in a virtual machine, and that justification has to be provided for exceptions that need a physical machine.

Mixed Trust Zones

As soon as virtualization moved from test/development into production environments, the issues and concerns on security started early on. Some asserted that there was no change at all in security solutions and security posture when existing workloads went "P2V" (*physical-to-virtual*). Others encountered both architectural concerns and operational issues.

Some of the earliest virtual security discussions were around "mixed trust zones", referring to the risk of hosting virtual servers of different data sensitivity or Internet exposure on the same hypervisor instance (physical server host)³. Sensitive data ran the risk of being breached should a more exposed virtual server be compromised and the underlying hypervisor VM isolation (thus far exceedingly unlikely in practice) as well.

The PCI Council was heavily involved in these debates, as different servers, such as those storing credit card numbers or other payment card industry data and those without, would normally be kept physically separate by function and segmented by network firewalls, per the PCI Council's Data Security Standards (DSS). Fortunately the PCI Council virtualization Special Interest Group (SIG) working group, in providing guidance for revision 2.0 of the DSS, specifically did not put restrictions on the use of virtualization technology nor mixed trust zones in 2010⁴, with the guidelines for the next 3.0 revision maintaining the neutrality of the standards with respect to new technologies, e.g. cloud computing.

However, the use of mixed trust zones can extend the *scope of compliance* audit to additional non-DSS virtual servers, which can increase regulatory and audit costs and efforts.

³ "Attacking and Defending Virtual Environments, Burton Group, Pete Lindstrom, 2008

⁴ "Securing Virtual Payment Systems", Version 1.0, PCI Security Standards Council, Virtualization Special Interest Group, January 2010

Inter-VM visibility and “Collapsing the DMZ”

With mixing trust zones come practical security problems as well, namely *inter-VM traffic visibility*. The canonical illustration is “*collapsing the DMZ*” of a typical Internet-facing three-tier web application onto a single physical host. With distinct virtual servers for the web, application and database layers all put on the same hypervisor and virtual switch, all the web-to-app inter-VM traffic flows through a virtual switch without leaving the box, effectively losing ability for physical firewalls and appliances to gain visibility to enforce network segmentation.

Security virtual appliances are one logical (no pun intended) solution - putting network security engines themselves into VM's that can now be re-inserted inline into the virtual switch traffic.



North-South vs East-West

Another option would be to string multiple VLAN's, one for each zone or application tier, from each the virtual switch out of the physical host and all the way up the physical network to a more central aggregation layer, where the more traditional firewall appliance would be able to inspect and enforce network zones - topologically, the inter-VM traffic is directed more “north-south” versus the natural “east-west” traffic within the virtual switch.

This is not necessarily much of an issue for VM traffic that spans different hypervisor instances, as the traffic would leave the physical host anyway. But for say a three-tier app that is on a shared host, this can lead

to “*hairpinning*” where traffic exits a physical host only to end up turning right around at the firewall and back down to another VM on the same host. And because live migration (e.g. VMware vMotion) and dynamic resource pooling may move VM's around frequently, it cannot necessarily be predictable when and how much inter-VM traffic will occur.

A network I/O latency study by VMware found that server-server traffic exiting the physical host could add about 10-20 μ s (or 40% more) latency per roundtrip versus pure virtual switch traffic⁵, on top of any latency introduced by the physical switch fabric or security appliances. The added latency can be exacerbated to 100 μ s or more when highly utilized hosts have many VM's queuing network traffic on the physical NIC (and note hairpinning involves doubles the effect with both sending and receiving VM's needing to pass through the physical NIC. This doesn't mean that physical security appliances are not suitable, however maintaining as low latency added as possible from physical security appliances, preferably under 10 μ s, is critical when there is heavy inter-VM traffic.

Takeaways

- Inter-VM traffic visibility
- Low latency physical appliances
- • Virtual security appliances

Product Options

- FortiGate
- FortiGate – VM virtual appliance
- Other Fortinet virtual appliances

Cloud Computing

Virtualization has long achieved more than just static server consolidation. As workloads were encapsulated in VM's, really as VMDK or VHD files on the physical storage, they could be manipulated like files, and analogous to document and content management, to provide other enterprise capabilities such as high availability, backup and redundancy, revision management

⁵ “Network I/O Latency on VMware vSphere 5”, Technical White Paper, VMware, 2012

(snapshotting), workflow and automation, etc. Such agility forms the basis of both private cloud and public cloud computing, but also add more dynamism to the data center and with it additional security considerations.

With security consistently being the #1 concern by IT managers when surveyed about the public cloud, the scope of cloud security is too broad to be fully discussed here, but a few key areas are highlighted here.

Multi-Tenancy


First, in public clouds the entire physical server and storage infrastructure is abstracted by virtualization and generally not visible or auditable below the VM container by tenants. It is further expected that storage, server, and network layers are inherently *multi-tenant* – that in order to maximize hardware utilization and efficiency, the provider will put workloads and traffic from multiple tenants on shared physical infrastructure. Tenant isolation is assured by generally not auditable to the tenants. The *shared responsibility* model is held up by Amazon Web Services and other cloud service providers as the framework for complete security – providers are responsible for tenant isolation and platform security; tenants are still responsible for host, OS, and application security within the VM container or instance, as well as within their virtual network.

The Cloud Security Alliance provides guidance to both enterprise tenants and service providers on how to map security and compliance regulations written originally with internal organization security in mind to provider clouds. The CSA's Cloud Controls Matrix⁶ maps controls in compliance frameworks such as PCI, NIST, ISO 27001/27002 and HIPAA-HITECH to cloud security control responsibilities for tenants and providers.

VLAN Spaghetti and Flatter, Scalable Network

From the provider standpoint, multi-tenant network security exacerbates virtualization issues with much more massive scale. North-south network topologies raise the VLAN spaghetti concern of having a mess of VLAN's broadcast across massive networks and running into the upper limit of 4096 limits. Yet carving up the network into more manageable chunks can go against the premise of horizontally scale-out clouds and flatter Layer 2 network for maximum agility, elasticity and scale.

⁶ "Cloud Controls Matrix", version 3.0, Cloud Security Alliance, September 2013



Delivering Elastic IT-as-a-Service

Many internal IT teams are trying to be like internal service providers and deliver more responsive *IT-as-a-service* from internal data centers, now private clouds. Some may go as far as to adopt multi-tenant paradigms for managing different departments or business units, and in some cases even extend IaaS and PaaS services to actual external partners as well.

In either case it means delivering infrastructure with greater *elasticity* to business units. Many of these marketing or user-based web services are designed around next-generation application stacks designed to scale out horizontally with new server instances. Such internal organizations thus expect server and network capacity to be available on-demand and generally without upper limits.

Agility through Automation and Orchestration

These new types of data center applications are increasingly agile to respond dynamically to changing demand and conditions via automation and orchestration engines. VM instances themselves can be assembled from OS, runtime application stacks and site content on the fly, and the number of instances and tiers ramped up and down automatically based on load. Bursting to run on hybrid or multiple clouds could also be automated based on changing Internet latency and availability.

Furthermore both internal and cloud service providers are pressured to meet scalability, reliability, and security in these agile environments with contractual service levels enforced via service level agreements (SLA's), potentially with even financial or other remunerative penalties.

Takeaways

- Multi-Tenancy
- Elasticity
- Agility, Automation & Orchestration
- Service-Level Agreements

Product Options

- FortiGate
- FortiManager & FortiAnalyzer
- FortiCloud

Software-Defined Networking (SDN)

As if virtualization and cloud computing weren't already disruptive enough, another huge paradigm shift has arisen with the advent of Software-Defined Networking (or the broader context of Software Defined Data Center). SDN, SDDC, network virtualization, network function virtualization (NFV) – what does it all mean, and what impact do they have on security?

Network Virtualization: The Evolution from Virtual Networking

One good start is to approach from the recent history of server virtualization. Server virtualization or more specifically x86 virtualization, abstracted physical compute hardware, namely the x86 CPU, chipset and RAM, from the OS and applications with a hypervisor layer that presented virtualized equivalents – vCPU, vRAM, etc. This abstraction enabled encapsulation of the workload into a VM container and isolation from other containers, and generated savings through hardware consolidation. Along the way, *virtual networking* introduced virtual switches as a convenient mechanism logical mechanism for how multiple vNIC's share a physical network interface card. But this was no longer just an abstraction of the x86 server - after all, the network hardware onboard an x86 server is the NIC or Ethernet adapter, not a switch, as the earliest virtualization products like VMware Workstation didn't have (and still don't have) a virtual switch.

Vendors and customers quickly discovered that virtual networking could provide other network benefits besides consolidating hardware – bandwidth resource pooling, redundancy, NIC redundancy, etc., and virtual networking features quickly expanded.

However, virtual networking was still a byproduct of server virtualization, with the vswitch not a true software switch independent of the hypervisor vmkernel. The virtual network was still dependent on the physical network, rather than the other way around (for example, relying on the physical network to define 802.1Q VLAN's).

Network virtualization promotes the virtual network to a first class citizen and is centered in the physical network fabric, i.e. the switches and routers.


Analogously, just like x86 servers, network ports are abstracted into virtual ports, which can then be combined logically into virtual switches across the entire network fabric. The network hypervisor can even exist independently of x86 hypervisor platform, or even without server virtualization altogether – although most likely any data center today adopting network virtualization will be using server virtualization as well.

Two key topic areas in network virtualization are OpenFlow and overlay networks.

OpenFlow - Abstracting control and data planes

In the OpenFlow model, the logical abstraction in the decoupling of the management or control plane physically from the actual switches via a “network hypervisor” or “SDN controller”. *OpenFlow* is one of the proposed standards for how the two communicate, thus the formally highlighting this separation – by defining a vendor-agnostic client- server API between “smart” SDN controllers that would define and dictate flow control to arrays of “dumb” physical switches/ports.

OpenFlow has been embraced by most major network hardware companies; however, not all vendors intend to drive all value through open standards. Cisco's ONE (Open Network Environment) embraces OpenFlow but also seeks to extend functionality through proprietary layers.



Service Insertion - Northbound vs Southbound API

Flow control such could be one means to integrate security appliances such as network monitoring or inline firewall appliances into the logical network as well; however, security products should not be using calling the OpenFlow protocol (or Southbound interface) directly to modify flows in the switches, as fundamentally there should be only a single brain or SDN controller as an OpenFlow client. In order to coordinate with the controller, security products should leverage available *northbound Interfaces* in either the controller or associated orchestration frameworks to coordinate with other network services and with the core network flow itself.

However, this raises the challenge that SDN controllers, such as VMware NSX or even open- source Floodlight, do not have a standard Northbound interface. One potential solution are projects like Openstack Quantum, which is an open source project that provides orchestration including northbound API's, but instead of serving as the SDN controller, interfaces with a number of supported controllers.

Even then, redefining flows continually and in realtime is not the ideal way to leverage SDN for security policy enforcement, as it exposes complexity and latency.

Instead.

VXLAN and Overlay Networks

A different aspect of SDN are network overlays (and underlays), including proposed standards such as VXLAN and NVGRE.

VXLAN enables Layer 2 subnets to be tunneled across Layer 3 networks and WAN/Internet, again creating logical network abstractions on top of the physical network. VXLAN can also overcome the earlier discussed VLAN limits of 4096 addressable ID's, expanding that to over 16 million.

Takeaways

- Control vs data plane abstraction
- Overlays and flatter networks
- Service insertion of network security

Product Options

- FortiGate
- FortiGate – VM virtual appliances
- FortiManager

Network Function Virtualization

Beyond SDN to NFV

Related to Software Defined Network is another movement driven by carriers, called Network Function Virtualization (NFV), starting in 2012. Rather than defining the network topology itself as in network virtualization, it is more about virtualizing or abstracting the network services and devices that sit on the network, from switches to firewalls to load balancers. Thus NFV is often considered closely related to, but distinct from, SDN itself.

Some of the key fundamental concepts of NFV are about making network services more agile, elastic and scalable as the compute/network infrastructure itself. With the latter gaining capex, opex, and manageability benefits from x86 and network virtualization, doing the same for network services can gain like efficiencies while also ensuring they do not service as impediments or bottlenecks to the underlying Infrastructure.

If NFV sounds a bit familiar like the virtualized security appliances discussed earlier, it is not far off, and certainly those virtual appliances do exist today from a number of vendors. However, there is additional SDN-like emphasis that it is not just about putting the services into VM's, but also about being able to manage, automate, and orchestrate heterogeneous services to deliver the agility and elasticity of those services.

Commodity vs Proprietary Hardware - A Red Herring

NFV has gained a lot of attention even though the working groups are still dealing more with concepts than even any concrete standards that vendors could implement just yet. There is also another agenda, or at least inherent assumption, in NFV that is attracting debate, that of commodity vs proprietary hardware.

Within the networking industry acceptance of OpenFlow and control plane abstraction, there was never any premise that network switches and routers would move to x86, for example, and various equipment vendors continue to embrace their own strengths in customer ASIC's, merchant silicon, or generalized x86 platforms. NFV takes the further step of championing x86, or really commoditized hardware, over proprietary hardware as the means to the end goals of service agility and elasticity.

This is somewhat of a red herring in the debate. While virtual appliances will generally be x86-based, it is again more than just about running in a VM, and more about orchestrating those services in a better way.

Those same concepts can be applied to proprietary platforms as well with the proper management integration, especially if the services on those non-x86 platforms can be 'virtualized' into allocable logical units and resource pools.

Look no further than the lessons of x86 virtualization itself. Certainly VMware and other hypervisors have hastened the shift away from RISC-based servers to x86 based ones. Yet it is not at all a commodity market. Yes, white-box x86 vendors now account for a huge share of the server host market, with Quanta alone supposedly supplying one out of seven servers shipped worldwide. Yet, in this same timeframe, Cisco UCS has also risen in a few years from zero to about 16% share of the branded x86 server market, offering a vertical story of networking, compute, services, and support that is anything but commodity.

In addition, the x86 server is more than a general purpose CPU, with many of the most I/O intensive network and storage functions supported by add-in cards with proprietary ASIC's. Intel executives themselves have scoffed at the notion that Intel CPU's will replace ASIC's anytime soon for high-performance networking or other functions, and the first NFV working group whitepaper candidly accedes that ASIC's are required for "high-throughput applications". In essence, the right answer today is not a single approach based on ideals, but a very practical matter of using the right tool for the right job.

For these reasons, the real attention on NFV should be better focused on how network services can become agile, and let the market and technology determine what hardware and virtualization technologies can fulfill that, and ultimately deliver the lowest TCO to customers and end-users.

Takeaways

- Agile network and security services
- Commodity vs proprietary hardware

Product Options

- FortiGate
- FortiGate – VM virtual appliances
- FortiManager

Additional Considerations for Service Providers

Security-as-a-service

With cloud computing, the shared responsibility model was introduced to distinguish the dual roles of both providers and tenants to provide a complete security posture. But rather than forcing every tenant to BYO network and host security via service VM's, more cloud providers are starting to integrate and offer security as a virtual, on-demand service on their IaaS/PaaS cloud services. In other words, provisioning security is really no different fundamentally than offering compute, memory, or storage for a VM instance.

"Security-as-a-service" can be deployed by cloud service provider directly into their cloud hosting infrastructure with either hardware or virtual appliances, and offered to tenants as on-demand service options complete with service level agreements. Or it could be more hands-off, such as offering private-label or branded virtual appliances through an integrated cloud marketplace. With either approach, providers are being driven to allow tenants to consume security and network services in the same manner as the VM instances - for example, with pay- as-you-go pricing (say per instance hour), on-demand (deployable at any time), and with service level agreements (SLA's).

Securing From the Cloud, for the Cloud

Cloud security will also have to become more efficient, scalable and easy to use as administrators are tasked to deal with an increasingly complex IT environment. With simplicity in mind, security management can also be delivered as a cloud service, e.g. central, Web- based management that can manage individual or aggregated security devices, and could include hosted log retention, automatically storing valuable log data in the cloud, and categorized by traffic, system events, Web, applications and security events.

Cloud management is conducive to managing security-as-a-service within a cloud provider, or could be delivered as a managed service, either by a network security vendor SaaS or by third party managed security service providers (MSSP). These cloud services could reach into managing security in tenant instances in public clouds or even back into the internal enterprise data center.

Tenant Partitioning, Delegation & Self-Service

Central to provider efficiencies and low TCO at scale is offloading as many IT functions as possible to automation or tenant self-service. This may mean not only delivering security services to tenants, but also logically isolating those security services for each tenant. So each tenant may get a logical firewall service with a separate virtualized runtime; furthermore the more that the administration of that tenant's security policies can be logically isolated from other tenants through administrative domains, the easier to delegate security management to tenants themselves. This further reduces provider costs while empowering tenants with self-service. For example, rather than all tenants having to perhaps choose from a limited set of security policies or profiles with largely on-off control, each tenant could craft very tailored policies per their own threat and regulatory environment, just as they would do within their own data center Infrastructure.

Takeaways

- Security-as-a-service
- Securing from the cloud
- Administrative delegation and self- service

Product Options

- FortiCloud
- FortiCarrier
- FortiGate VDOM virtual domains
- FortiManager ADOM administrative domains

Summary

Fortinet's Approach

With areas like cloud computing and SDN still early in the customer and vendor adoption cycle, there is no one-size-fits-all answer for everyone. That's why Fortinet is investing and innovating in a number of different areas as customers adopt these nascent technologies.

For example, Fortinet is investing in both physical and virtual security appliance technologies. Fortinet has long been a pioneer in physical FortiGate firewall appliances with proprietary FortiASIC hardware technology that delivers the highest performance, lowest latency to meet elastic cloud requirements, and a 10X price/performance advantage for lowest cloud TCO. Yet even with hardware ASIC strengths, Fortinet has also invested in porting its entire FortiOS security software stack to run with full functionality in x86 virtual appliances, and today has one of the largest portfolios of virtual appliances that run on VMware and other hypervisor and in public clouds to monitor "east-west" virtual traffic. These range from FortiGate-VM to nearly a dozen other security virtual appliances including FortiWeb-VM web application security, FortiMail-VM e-mail gateway security to FortiADC-VM application delivery controller.

Fortinet has also virtualized security management with FortiManager-VM and FortiAnalyzer-VM virtual appliance editions of central management and logging solutions. Fortinet is further moving security management to the cloud with FortiCloud SaaS- based management, logging, and analytics.

Fortinet is working to bring virtualization and abstraction to physical appliances as well, so that they can be equally agile within cloud and SDN orchestration frameworks. For example, with innovative virtual domain (VDOM) technology, a single FortiGate firewall appliance can be divided into hundreds of individually managed and process- isolated logical devices. With FortiManager administrative domain (ADOM) technology, those logical instances can further be isolated and managed separately in multi-tenant and delegated scenarios.

Fortinet is also investing in SDN technologies, exploring integrations with OpenFlow, OpenStack, and VMware's NSX platform.

Deployment Considerations

A few considerations for organizations to evaluate the appropriate mix of physical and virtual security approaches:

Fixed vs variable network capacity – Network bandwidth is continually growing, but that does not mean that it is all the same. How much data center network traffic is steady or predictable, such as the usage patterns of employees? How much is variable, such as customer demand, and does the variability follow cyclical patterns such as seasonality, or are they highly unpredictable, like new online business initiatives or marketing campaigns? At today’s price/ performance levels, hardware approaches can be very compelling for delivering fixed capacity at lowest absolute cost, while virtual appliances can be great for almost unlimited and linear scale-out for highly elastic traffic requirements, with potentially no practical limits.

Network Topology – Is the overall network topology more hierarchical (north-south) or flat (east-west)? Will the answer to that change between today’s infrastructure and tomorrow’s with adoption of SDN and network virtualization? How much network bandwidth is lateral between applications within the data center and how much is external to the Internet or the enterprise campus?

Throughput vs Latency - What is the right balance of throughput vs latency? A high performance physical appliance could deliver cost-effective throughput for north-south traffic as long as the added latency is minimized and acceptable. Meanwhile a virtual appliance enables east-west traffic but can local vCPU/vRAM bottlenecks during high-utilization periods.

Regulatory environment – Do customer agreements, government regulations or industry compliance restrict or dictate use of certain technologies or approaches. Even when they do not, such as with PCI DSS, does the burden of additional auditing such as expanded “in-scope” requirements make some approaches expensive or impractical?

Conclusion

Enterprises need to evaluate their data center initiatives under way today and tomorrow and how they impact network security design.

Fortinet is also working with both industry leaders and more agile smaller players. These range from bigger technology partners such as VMware to smaller innovators like BigSwitch and HyTrust, and to telco and service providers who are customers but also partners in enabling the cloud for enterprises.

Fortinet is committed to investing in new data center technologies to bring benefits to businesses and customers as soon as they are ready.



GLOBAL HEADQUARTERS
Fortinet Inc.
899 Kifer Road
Sunnyvale, CA 94086
United States
Tel: +1.408.235.7700
www.fortinet.com/sales

EMEA SALES OFFICE
120 rue Albert Caquot
06560, Sophia Antipolis,
France
Tel: +33.4.8987.0510

APAC SALES OFFICE
300 Beach Road 20-01
The Concourse
Singapore 199555
Tel: +65.6513.3730

LATIN AMERICA SALES OFFICE
Paseo de la Reforma 412 piso 16
Col. Juarez
C.P. 06600
México D.F.
Tel: 011-52-(55) 5524-8428