

Virtual Firewalls

Ivan Pepelnjak (ip@ioshints.info)
NIL Data Communications



ipSpace

Who is Ivan Pepelnjak (@ioshints)

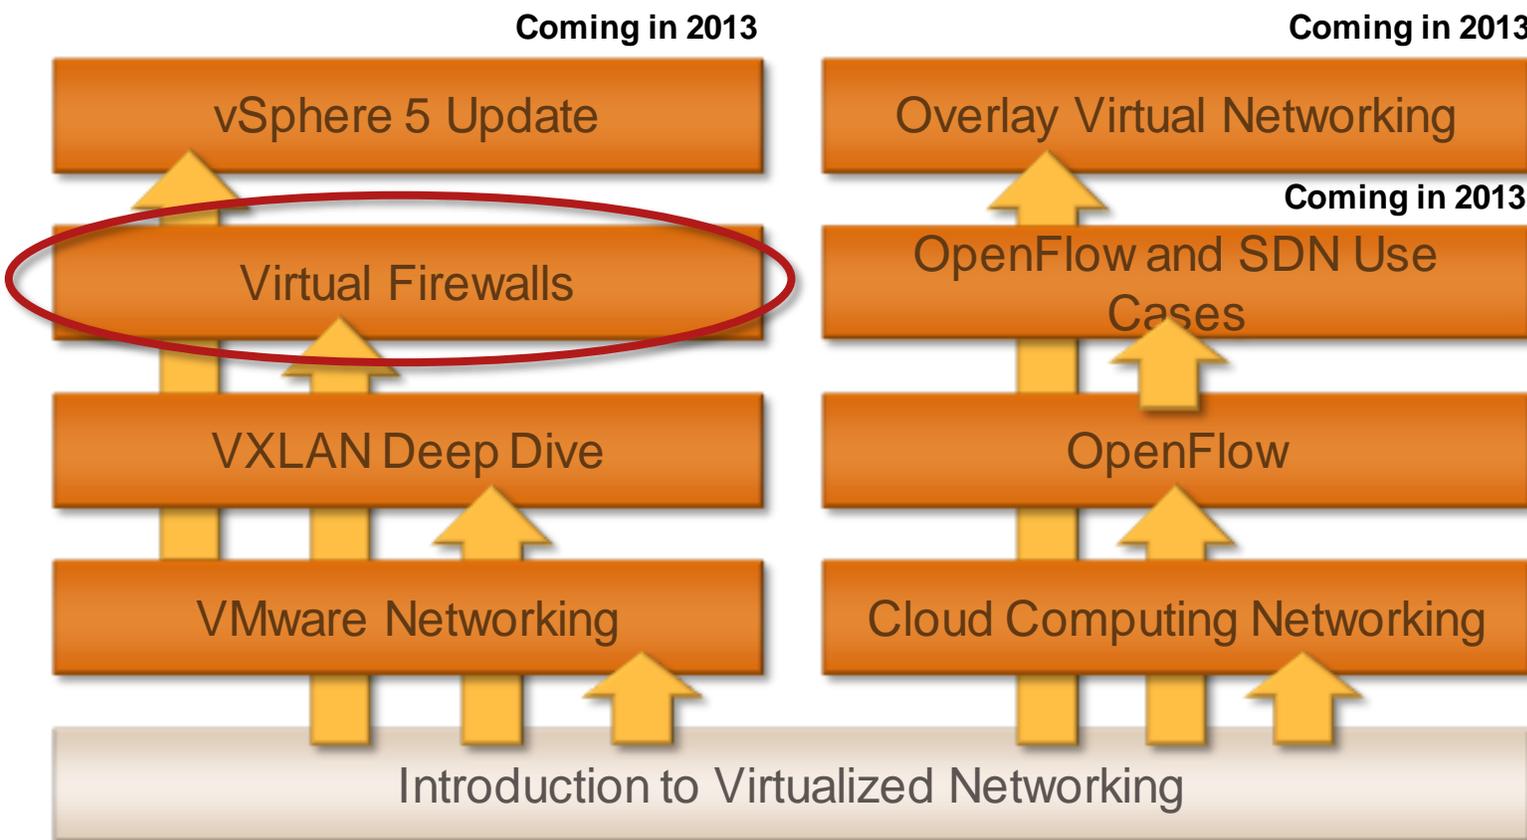
- Networking engineer since 1985
- Focus: real-life deployment of advanced technologies
- Chief Technology Advisor @ NIL Data Communications
- Consultant, blogger (blog.ioshints.info), book and webinar author
- Teaching “Scalable Web Application Design” at University of Ljubljana



Current interests:

- Large-scale data centers and network virtualization
- Networking solutions for cloud computing
- Scalable application design
- Core IP routing/MPLS, IPv6, VPN

Virtualization Webinars on ipSpace.net



Availability

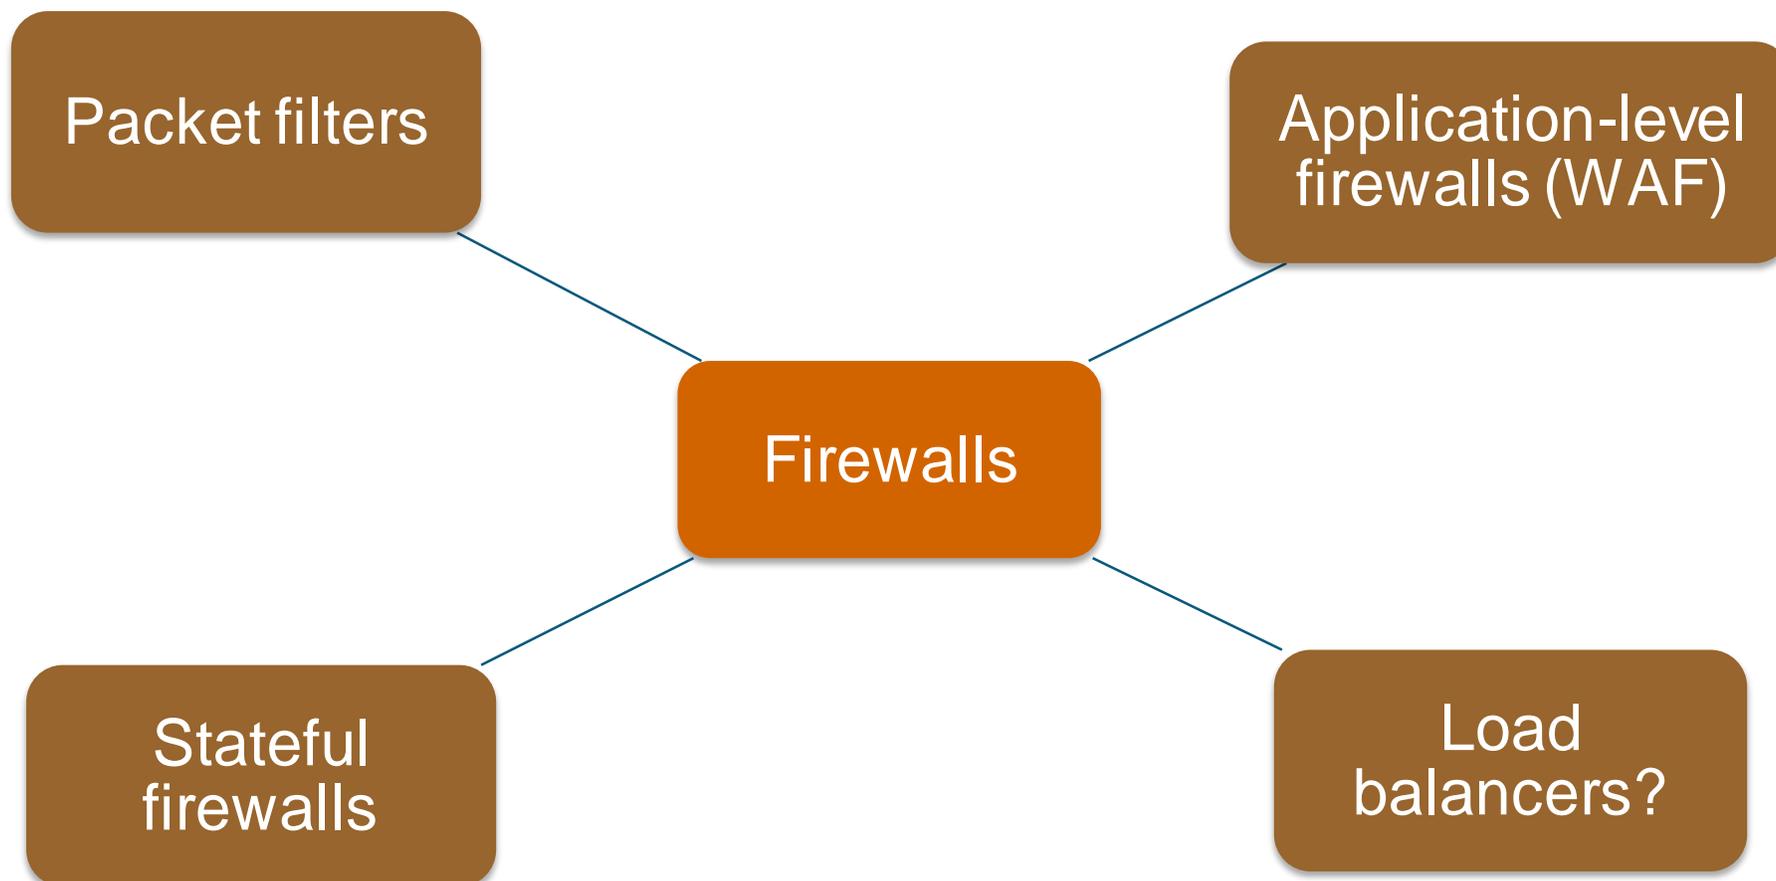
- Live sessions
- Recordings of individual webinars
- Yearly subscription

Other options

- Customized webinars
- ExpertExpress
- On-site workshops

More information @ <http://www.ipSpace.net/Webinars>

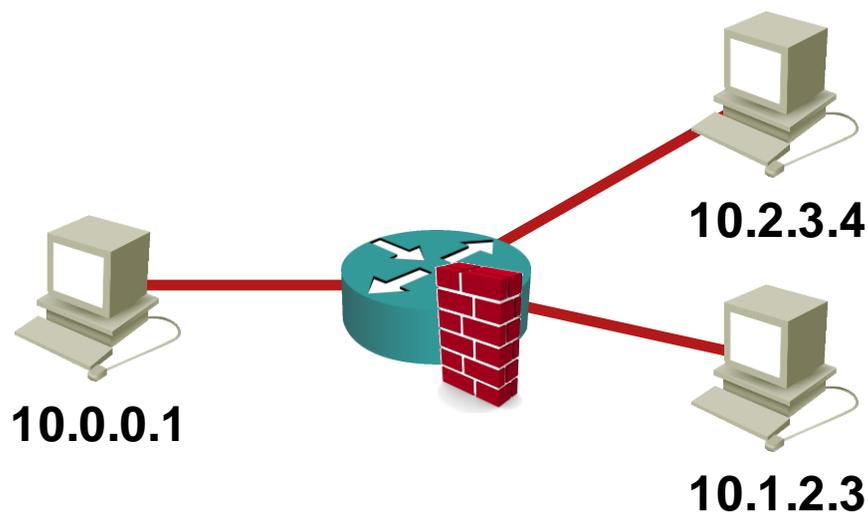
Firewalls Used To Be Easy



Routed or Bridged?

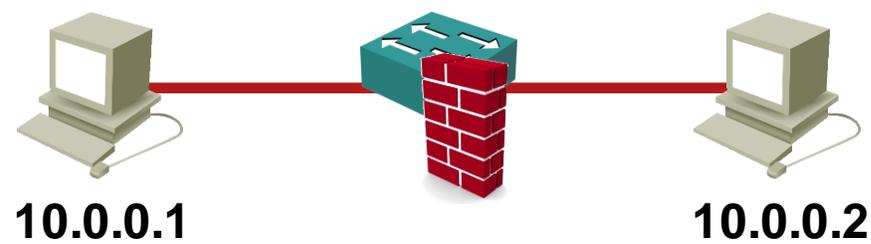
Routed (inter-subnet)

- Packet filtering and IP routing
- *Inside* and *Outside* subnets
- Static routing or routing protocols
- Easy to implement multiple zones



Transparent (bridged)

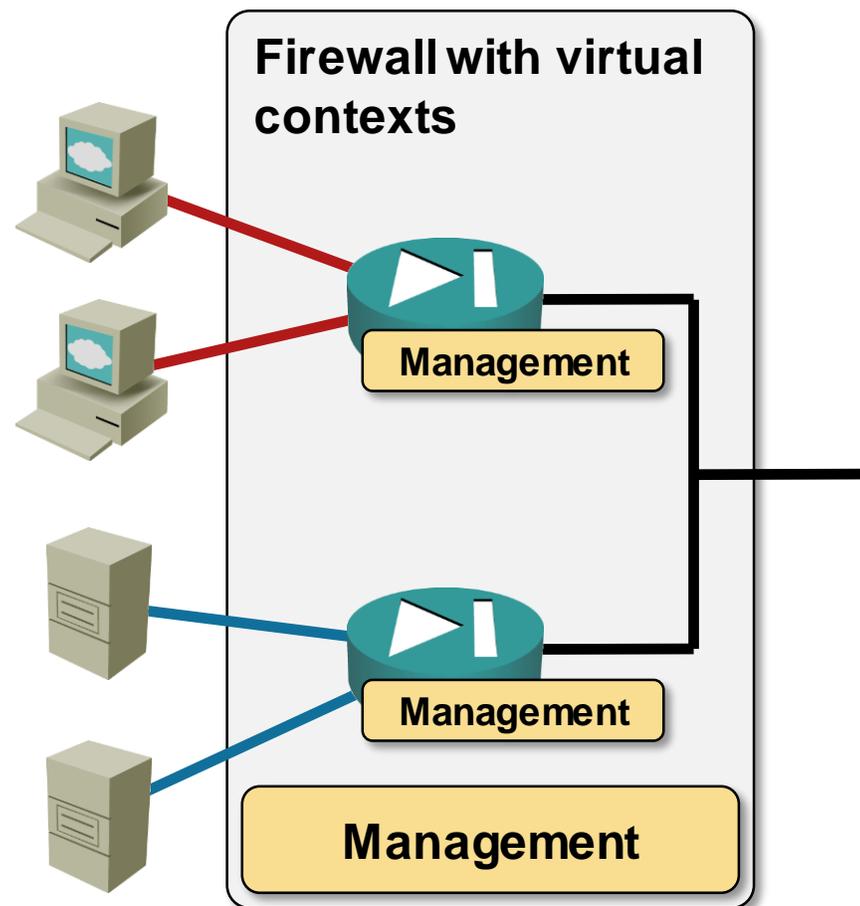
- Packet filtering and bridging
- Simple to insert
- No interaction with routing
- Typically only two interfaces



Anything Is Virtual These Days

Single physical device, multiple virtual contexts

- Separate management plane(s)
- Shared resources (code, CPU, interface bandwidth ...)
- Tied to a physical device



This is not the virtual firewall we're looking for

Virtual Contexts Versus Virtual Firewalls

Transport network independence

- Virtual firewalls run on any transport provided by hypervisor (VLAN, VXLAN, NVGRE ...)
- Virtual contexts support the encapsulations of underlying firewalls software

Virtual networking support in physical devices

- VLANs (802.1Q)
- Rarely: Q-in-Q (802.1ad)

Exceptions:

- VXLAN supported by F5 (LB), Brocade (LB) and Arista (switch)
- NVGRE supported by F5 (LB)

Virtual Contexts Versus Virtual Firewalls

Transport network independence

Configuration management

- Virtual context configuration tied to physical device
- Virtual firewall configuration moves with it
 - ➔ Stored in virtual disk attached to a VM
 - ➔ Central management software

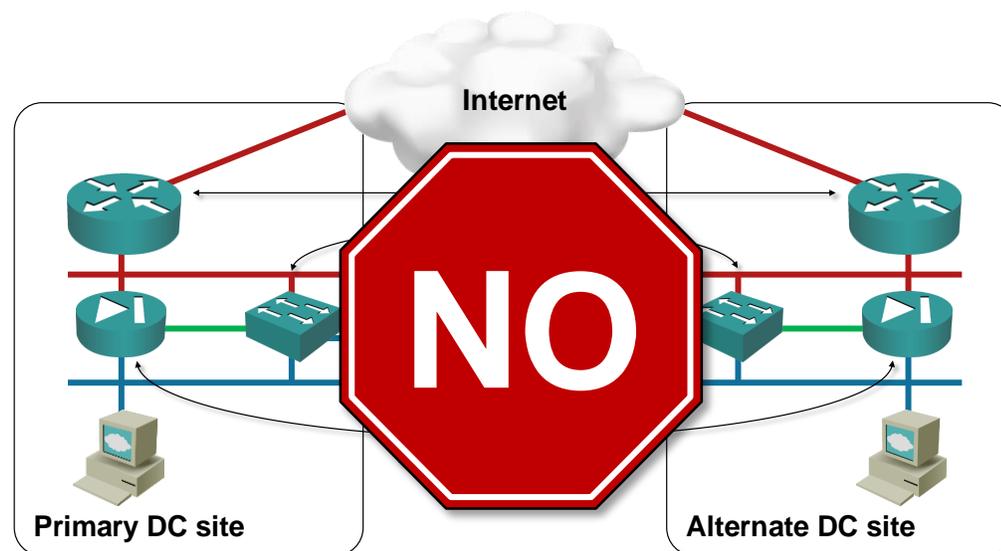
Virtual Contexts Versus Virtual Firewalls

Transport network independence

Configuration management

Workload mobility

- Impossible to move physical device (don't even mention stretched firewalls)
- Virtual firewall migrates with the workload
- Move application stack + L4-7 components in disaster recovery/avoidance procedure



Virtual Contexts Versus Virtual Firewalls

The good news:

- Transport network independence
- Configuration management
- Workload mobility

And now for some bad news:

- Performance
- Attacks on hypervisors, multi-tenant attacks

Virtual Contexts Versus Virtual Firewalls

The good news:

- Transport network independence
- Configuration management
- Workload mobility

And now for some bad news:

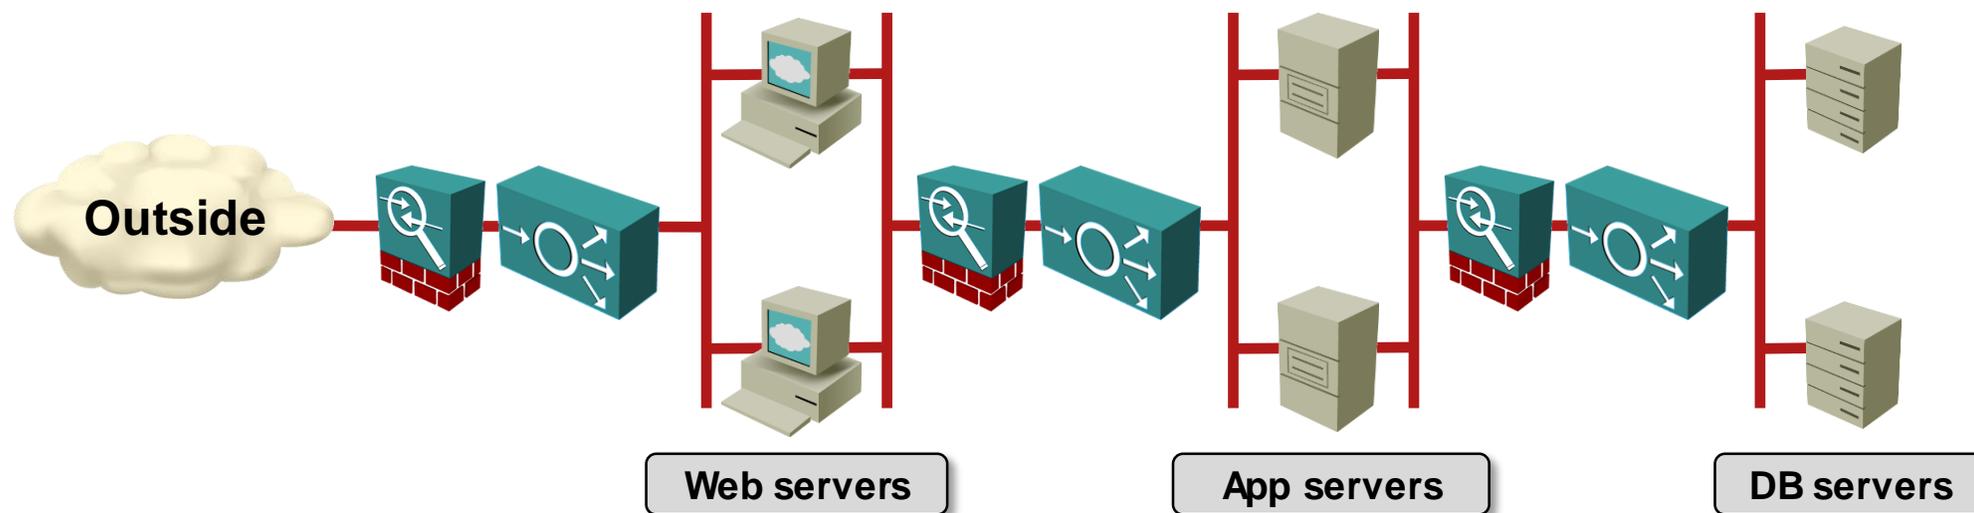
- Performance
- Attacks on hypervisors, multi-tenant attacks

Real question: How secure does your auditor think you have to be?



Virtual Firewalls

Virtual Networking Requirements



User requirements

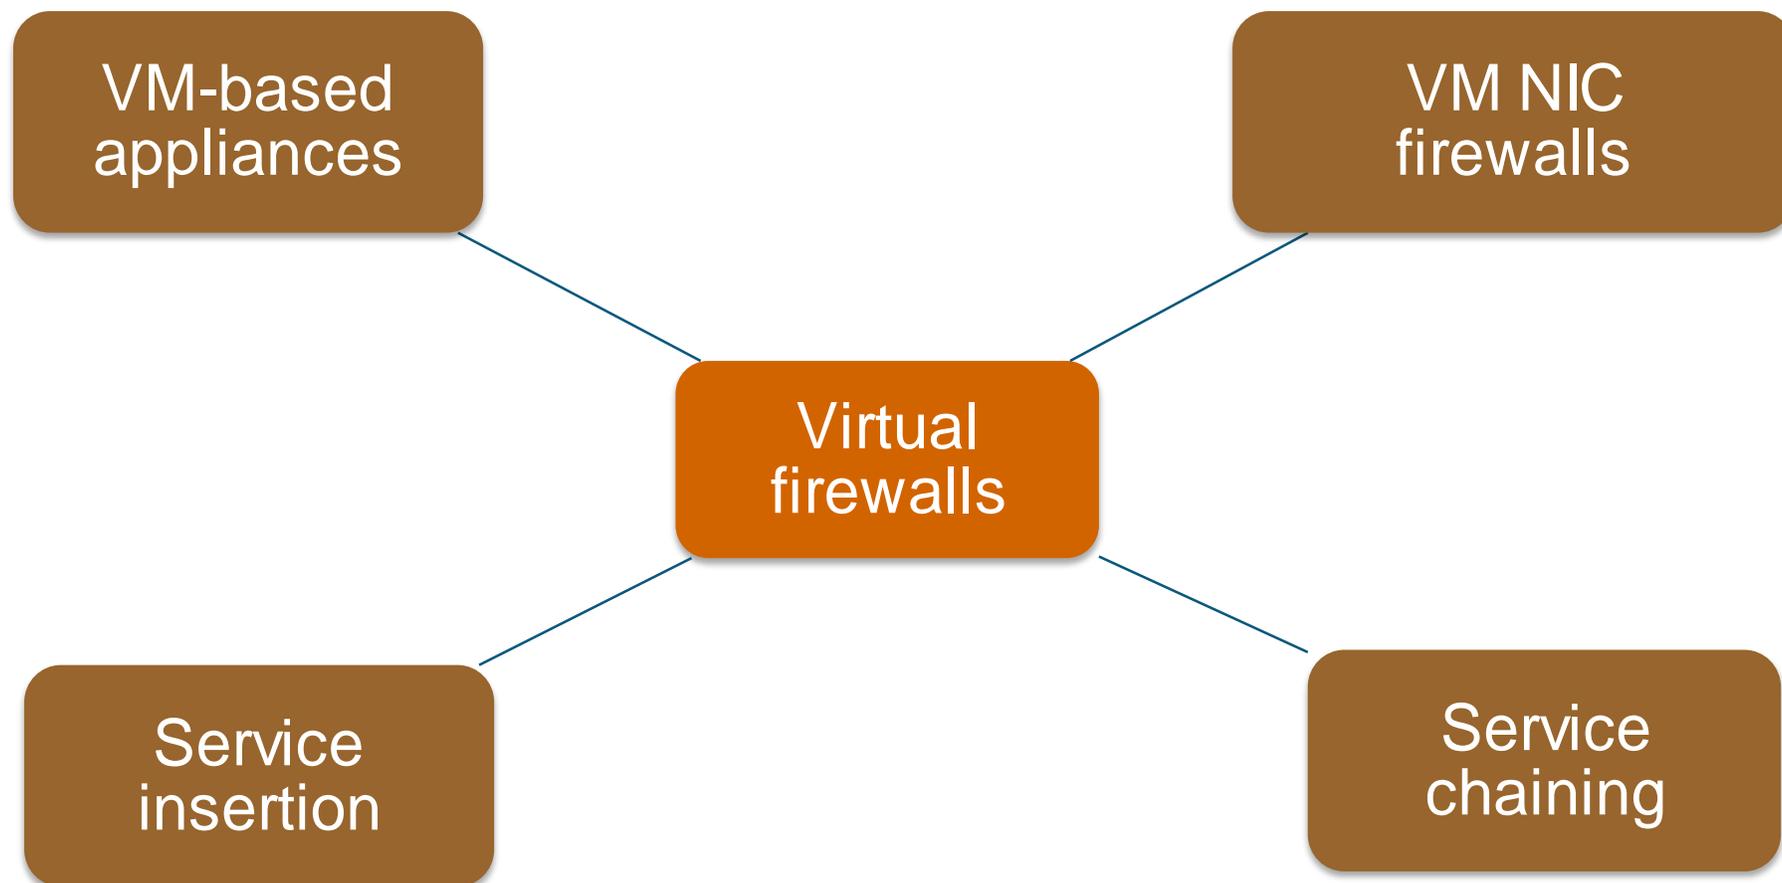
- Use virtual machines like physical hosts
- Deploy and move VMs at will
- Build virtual LANs
- Retain existing application stack
- Retain existing security paradigm

Hypervisor requirements

- Decouple physical hardware from VM NIC (VM mobility)
- Enable inter-VM traffic (intra-hypervisor and across the network)
- Provide inter-VM isolation

Design decision: physical or virtual firewalls and load balancers?

Virtual Firewall Taxonomy



Virtual Appliances

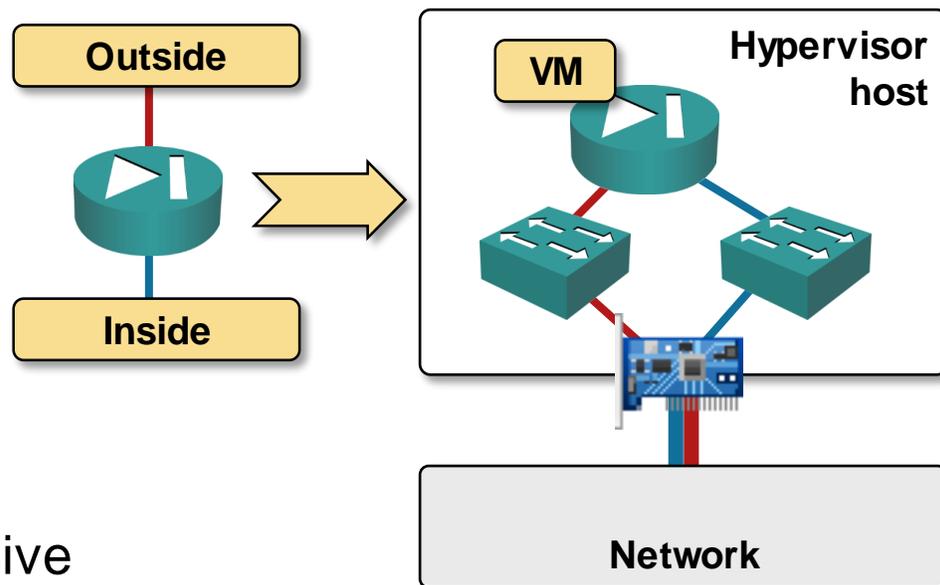
- Most L4-7 devices run on x86 CPU
- Some of them are also offered in VM format
- VM appliances work with all network virtualization technologies (incl. vCDNI and VXLAN)

Drawbacks

- CPU-based packet processing is expensive
- High hypervisor overhead with I/O intensive workload
- Traffic trombones

Sample products

- Firewall: Vyatta, vShield Edge (VMware)
- Load balancer: BIG-IP VTM (F5), Zeus Traffic Manager (now Riverbed), vShield Edge (VMware), Embrane, LineRate Systems (now F5)



Virtual Appliance Performance Issues

Typical performance:

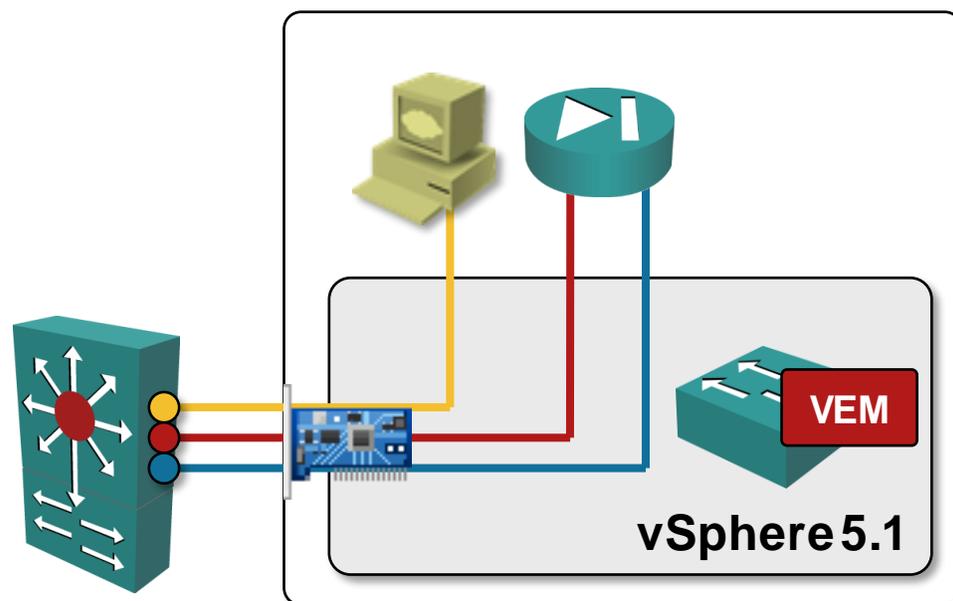
- 40+ Gbps through a Xeon-based server
- A10 load balancer VM on a single core: up to ~ 8 Gbps
- vShield Edge small instance (1 vCPU): ~ 1 Gbps

Two performance roadblocks:

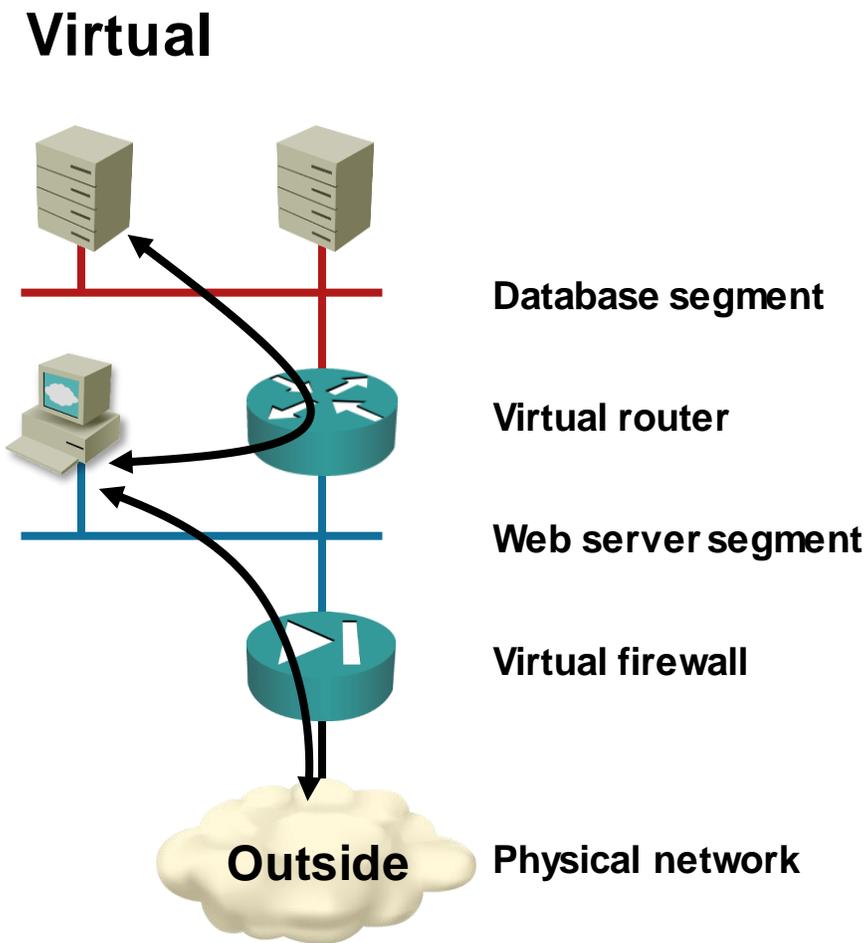
- Linux TCP/IP stack in appliance
- Hypervisor virtual switch

Enhancements:

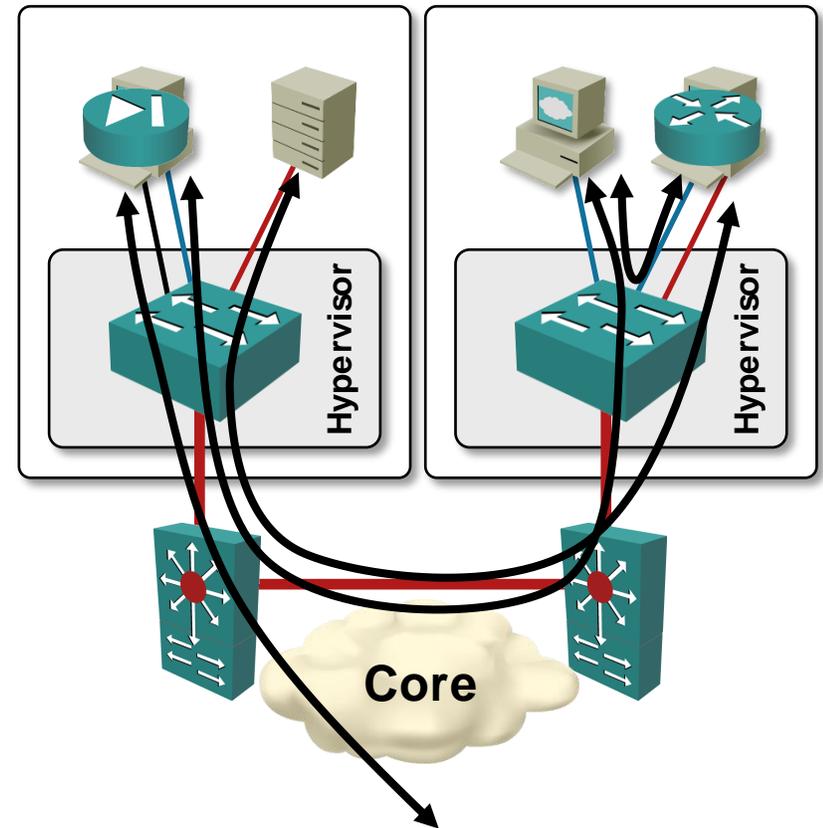
- TCP offload (not on VXLAN)
- Hypervisor bypass (Cisco VM-FEX)
- Third-party TCP stacks (Intel DPDK, 6Wind)



Virtual Appliance-Induced Traffic Trombones



Physical

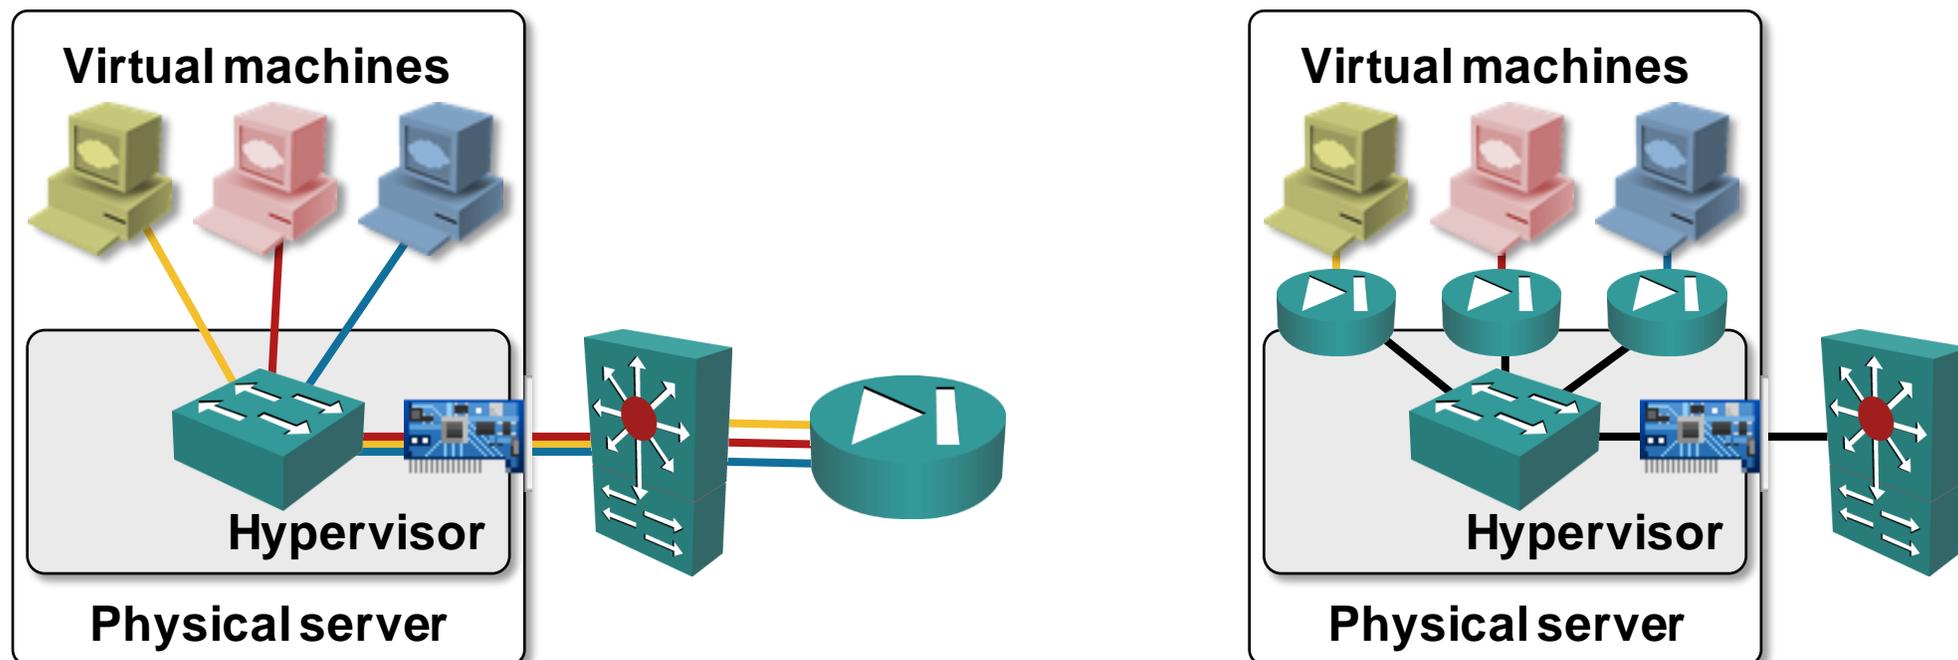


Requires DC design with equidistant end points (Clos architecture)



VM NIC Firewalls

What Is a VM NIC Firewall



- Firewall (or packet filter) inserted between VM Network Interface Card (NIC) and hypervisor virtual switch
- Central management/configuration for scalability
- Firewall rules (or ACLs) and state move with VM

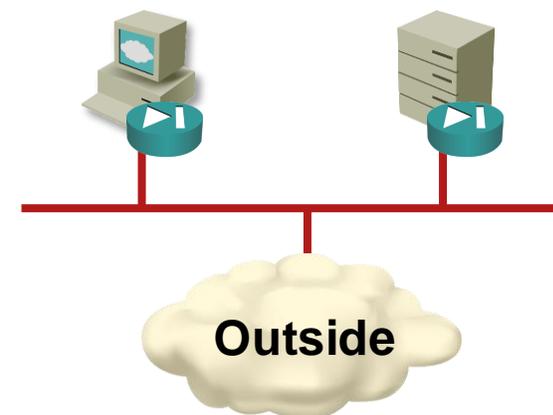
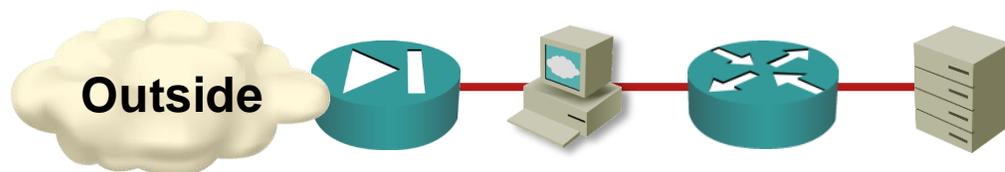
VM NIC Firewalls: Changing the Security Paradigm

Old world security

- Security zones = IP subnets = VLANs
- Add VXLAN/NVGRE ... for scalability
- Subnets segregated with firewalls or virtual appliance firewalls
- Traffic trombones
- Firewalls are choke points

Brave new world

- Firewall rules attached to virtual NICs
- Everything else is “outside”
- Optimal any-to-any traffic flow
- “Infinitely” scalable



VM NIC Firewalls: Sample Solutions

VMware VMsafe Network API

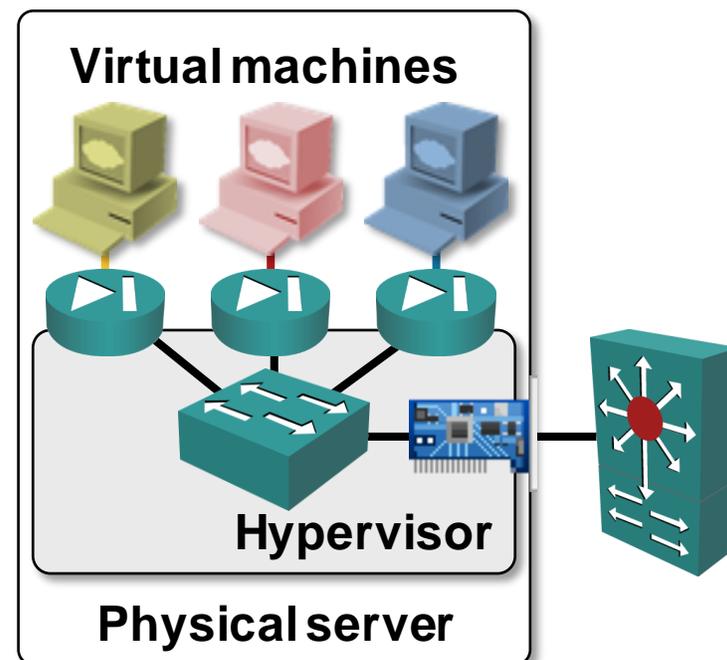
- vShield App/Zones (VMware)
- vGW (Juniper)

Linux (KVM, Xen)

- *iptables, ip6tables, ebtables*
- Open vSwitch with OpenFlow controller
- Midokura Midonet

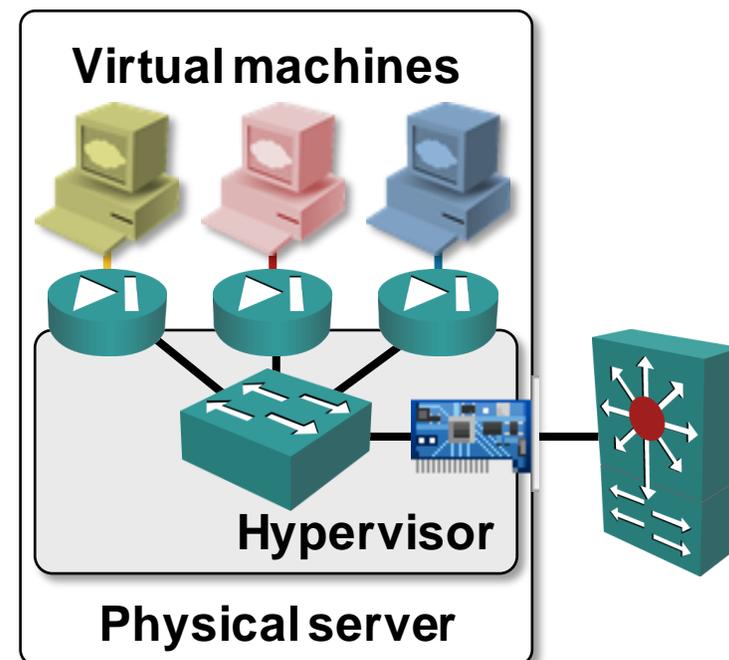
Hyper-V 3.0 Extensible Switch

- 5nine virtual firewall
- NEC ProgrammableFlow virtual switch



Questions to Ask

- Stateless or stateful?
- Filtering in kernel module or userland?
- Per-hypervisor control VM?
- Is control VM involved in flow setup?
- What happens when control VM fails?



VMsafe Network (dvFilter) API

VMsafe Network API

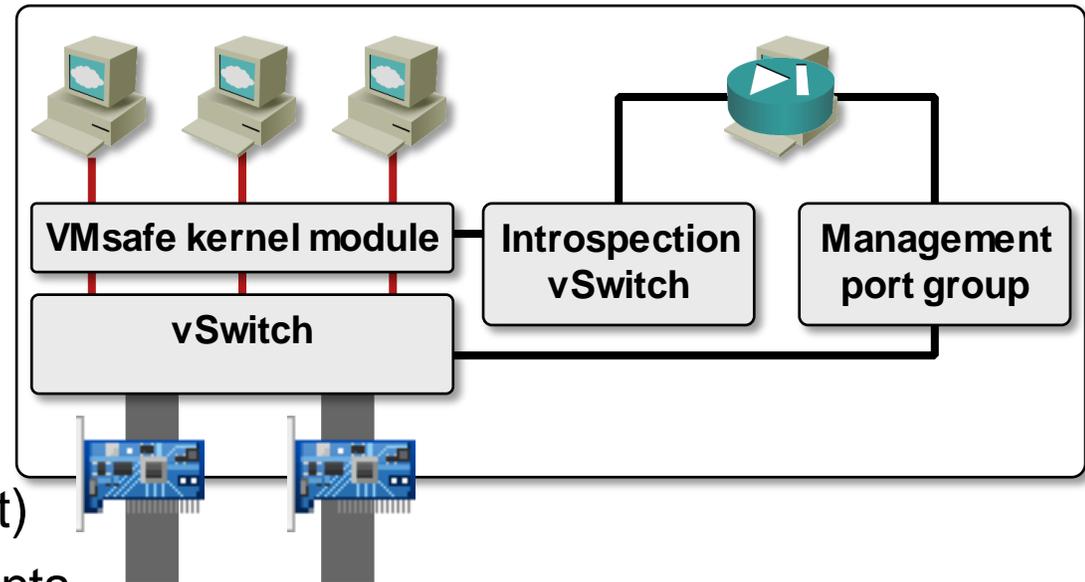
- Allows a security appliance VM to intercept traffic to/from other VMs
- Internal name: dvFilter

Each dvFilter-based product has:

- Data-path kernel module
- Control-path VM (on the same host)
- Communication between components through a hidden vSwitch
- Kernel module or control-path VM can permit, drop or modify VM traffic

Sample products: vShield Zones/App, Virtual Gateway (Juniper), TippingPoint vController (HP)

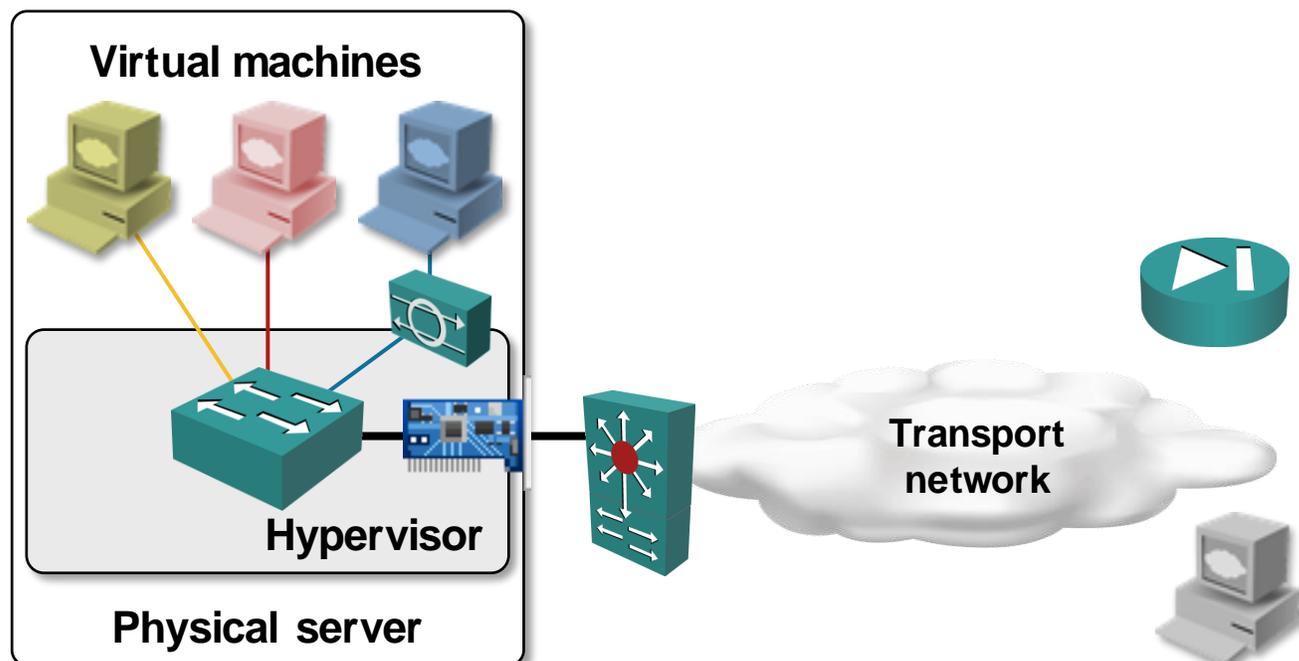
- Significant performance differences based on forwarding path





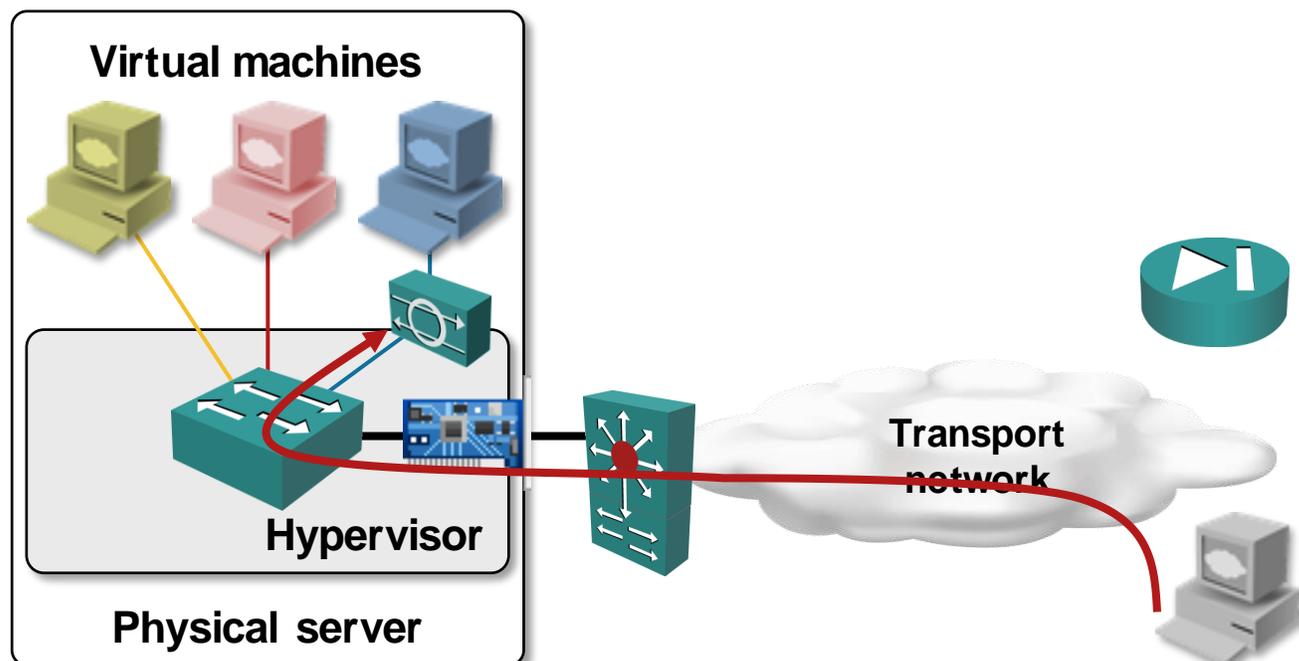
Service Insertion

Service Insertion 101



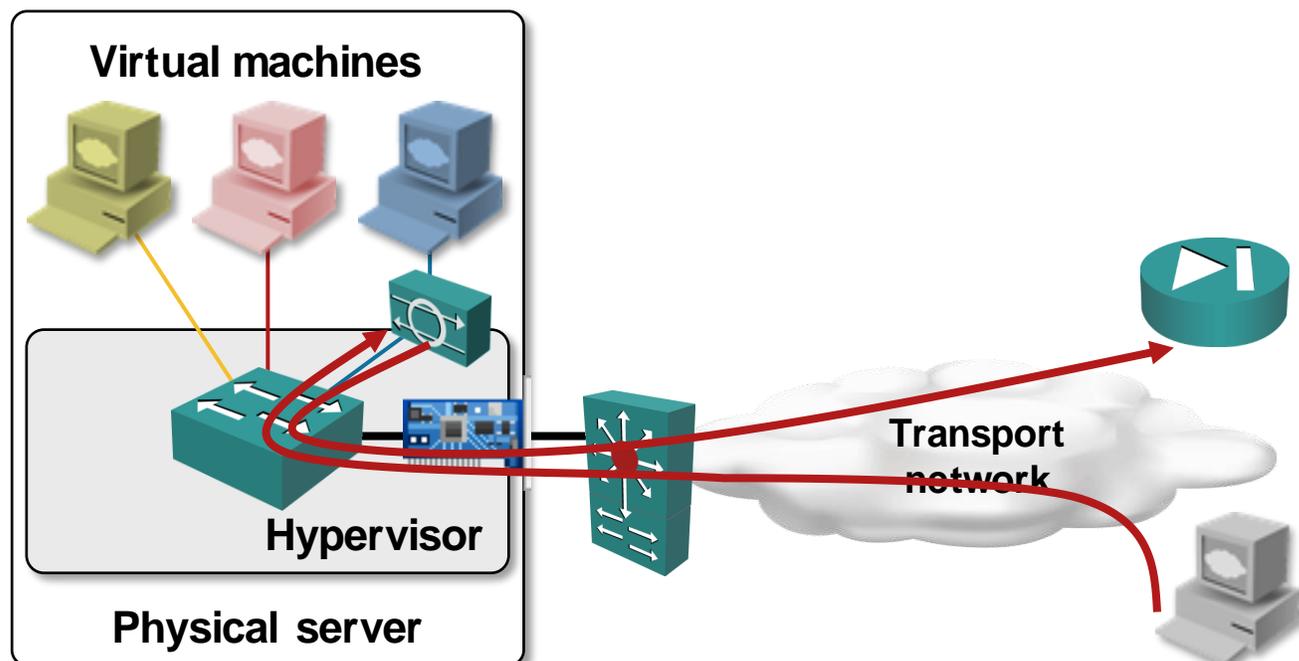
- Hypervisor switch redirects traffic traversing VM NIC
- L4-7 functionality in external device or VM appliance
- Filtered/modified traffic is reinserted at NIC-to-vSwitch boundary
- Optional: approved 5-tuple inserted in hypervisor switch

Service Insertion 101



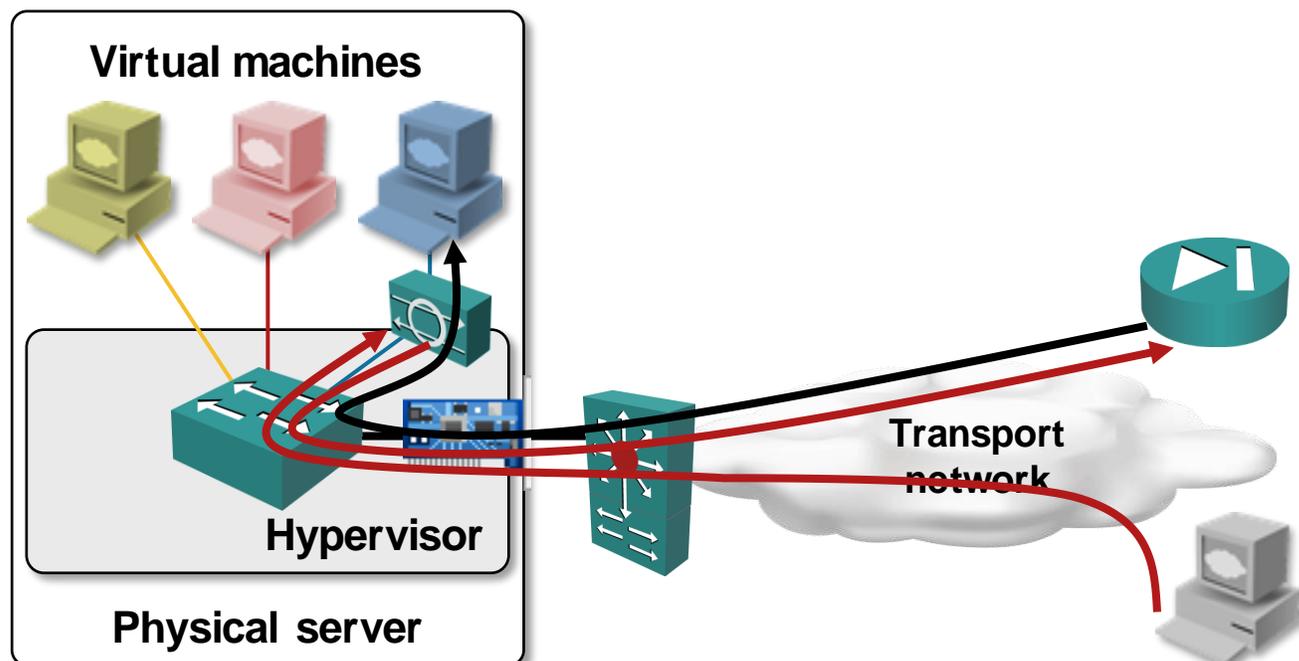
- Hypervisor switch redirects traffic traversing VM NIC
- L4-7 functionality in external device or VM appliance
- Filtered/modified traffic is reinserted at NIC-to-vSwitch boundary
- Optional: approved 5-tuple inserted in hypervisor switch

Service Insertion 101



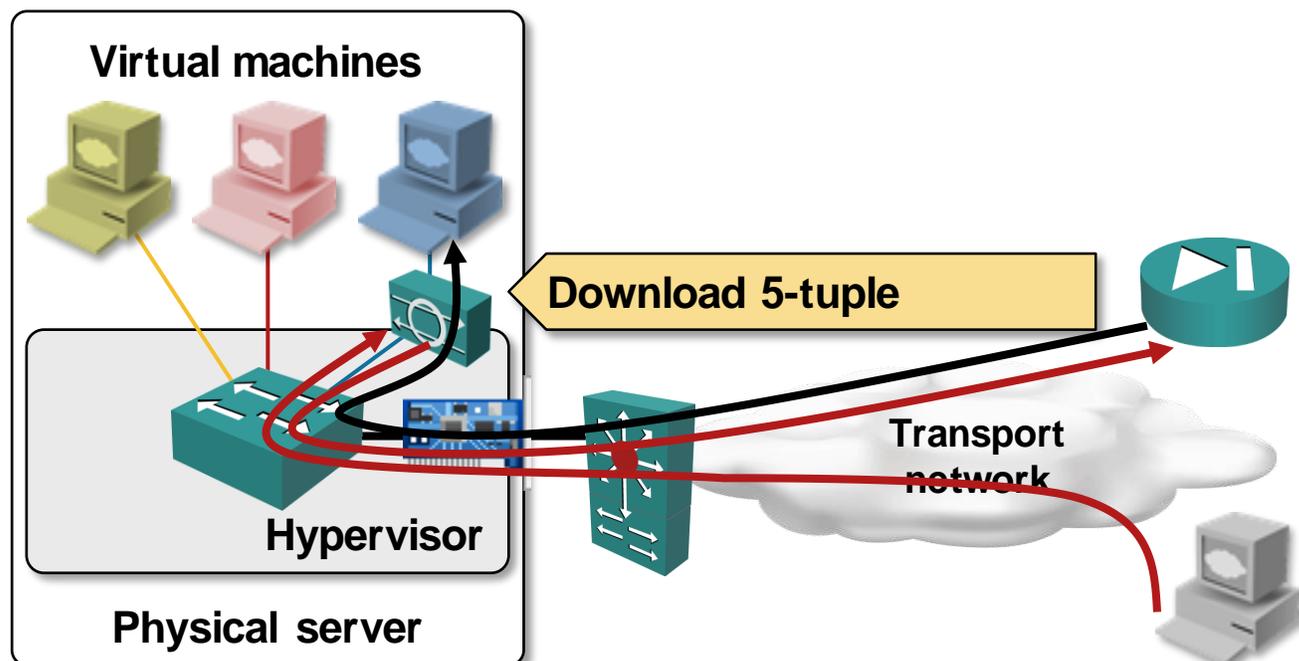
- Hypervisor switch redirects traffic traversing VM NIC
- L4-7 functionality in external device or VM appliance
- Filtered/modified traffic is reinserted at NIC-to-vSwitch boundary
- Optional: approved 5-tuple inserted in hypervisor switch

Service Insertion 101



- Hypervisor switch redirects traffic traversing VM NIC
- L4-7 functionality in external device or VM appliance
- Filtered/modified traffic is reinserted at NIC-to-vSwitch boundary
- Optional: approved 5-tuple inserted in hypervisor switch

Service Insertion 101



- Hypervisor switch redirects traffic traversing VM NIC
- L4-7 functionality in external device or VM appliance
- Filtered/modified traffic is reinserted at NIC-to-vSwitch boundary
- Optional: approved 5-tuple inserted in hypervisor switch

HP TippingPoint vController

TippingPoint = IPS appliance

vController = per-vSphere host VM

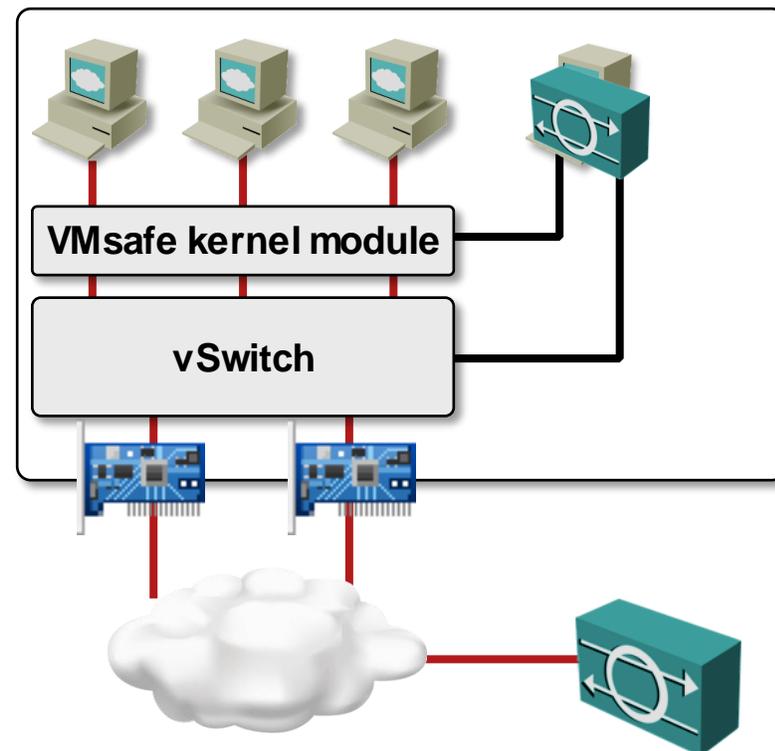
VMsafe Network API used for service insertion

Typical packet flow

- vController intercepts VM traffic
- vController sends VM traffic to IPS
- IPS inspects VM traffic and returns it to vController
- vController forwards the traffic to VM or vDS

Benefits and drawbacks

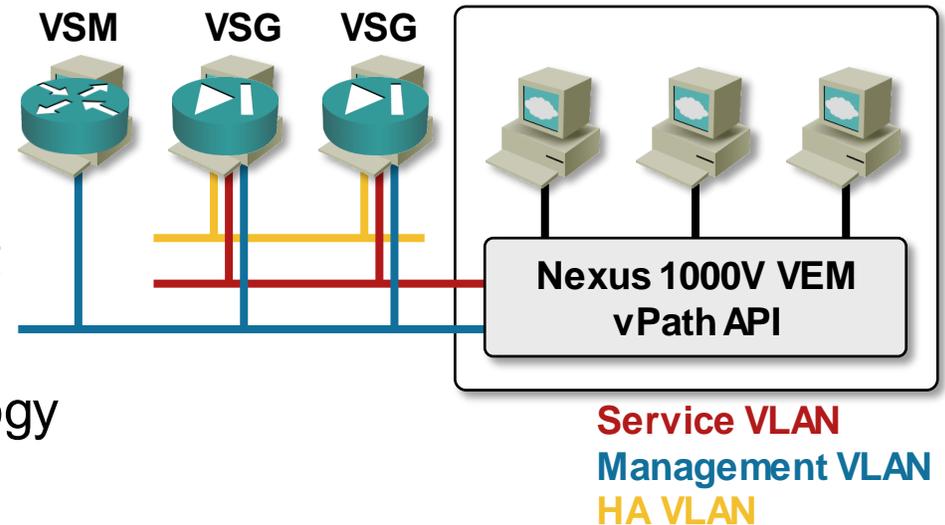
- Leverages existing IPS appliance
- Reduced CPU load on the ESX host
- Still requires a vController VM on each ESX host



Virtual Security Gateway (Cisco)

Some terminology

- Nexus 1000V : vSwitch replacement
- VSM: Nexus 1000V control plane
- VEM: switching element in vSphere host
- VSG: stateful layer-2 firewall
- vPath: Cisco's service insertion technology

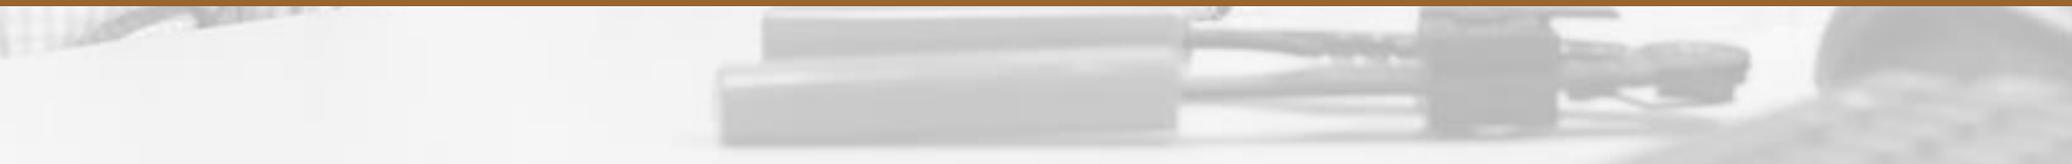


Principles of operation

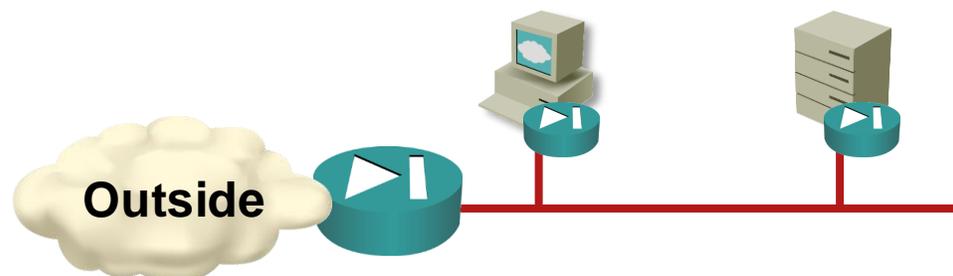
- Service interception done in vSwitch, not in NIC driver
- VN-service defined on port profile in Nexus 1000V
- Traffic forwarded to VSG on service VLAN or encapsulated in IP
- VSG can download 6-tuple (+VLAN) to VEM (fast-path offload)



Service Chaining



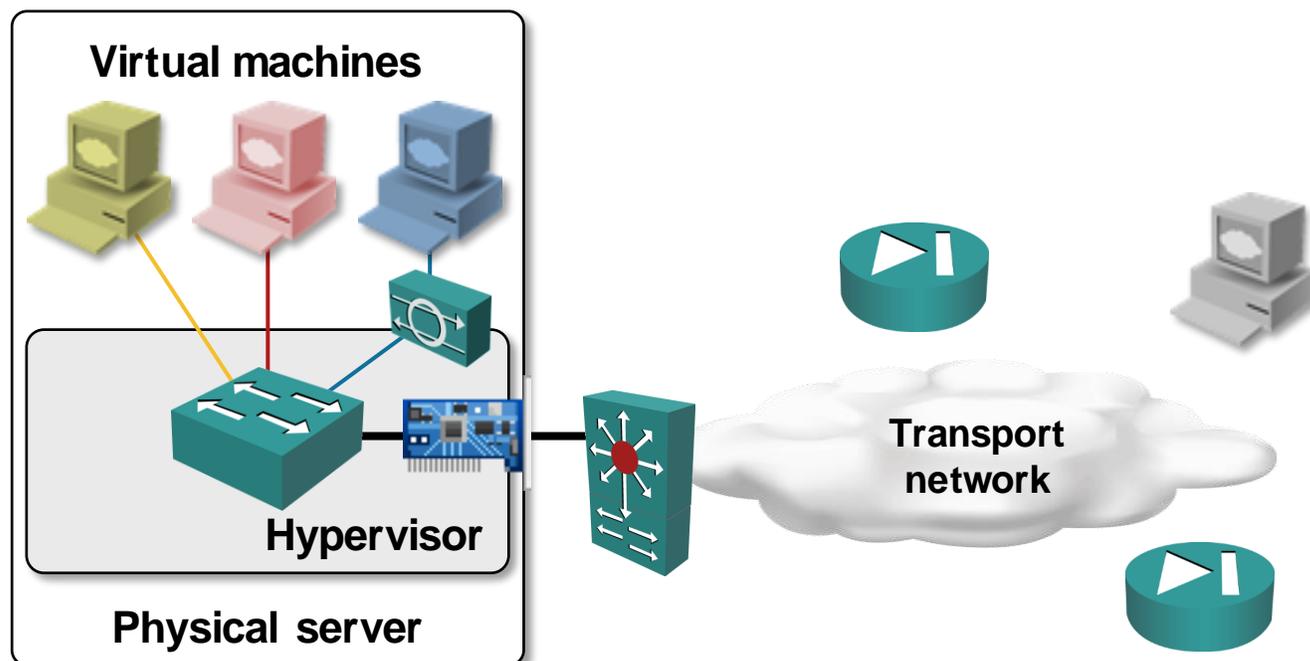
Problem: Combining L3 and L2 Services



NIC-level firewall + routed firewall, load balancer or WAF

- Easy to implement with VM appliances + NIC-level firewalls
- More interesting when used with service insertion

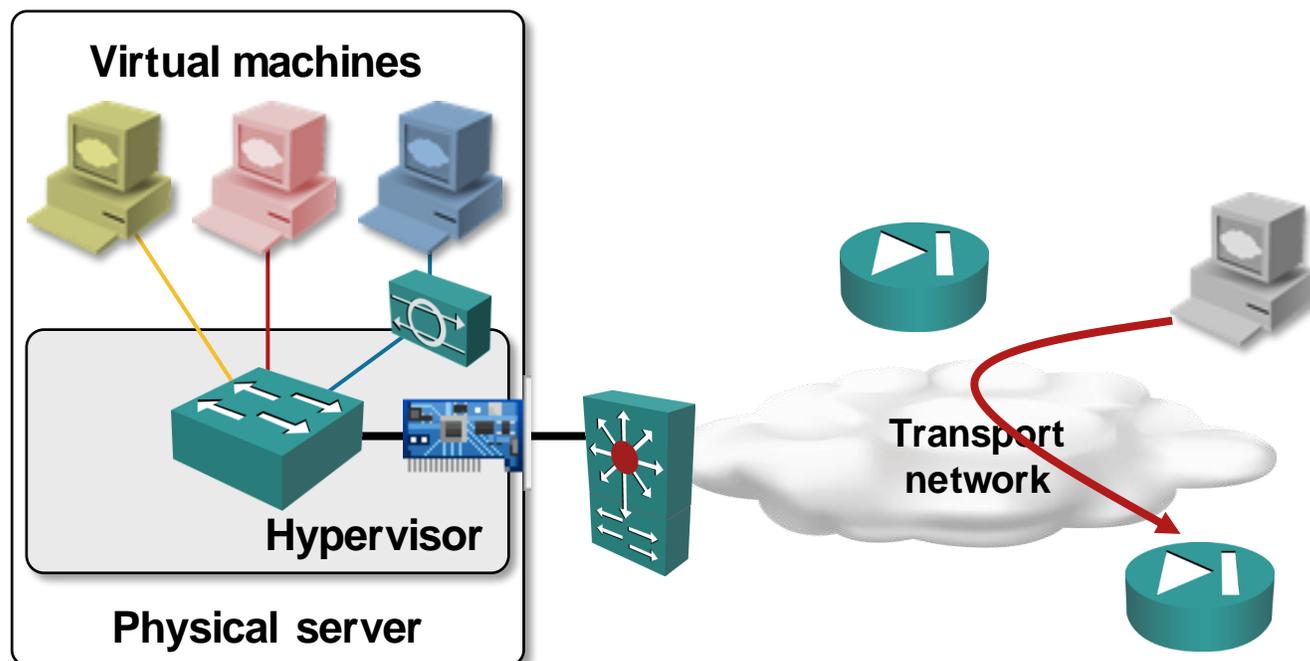
Service Insertion Gone Bad



- External traffic is sent to L3 appliance (based on IP routing)
- L3 appliance forwards traffic toward VM MAC address
- Hypervisor switch (or NIC driver) intercepts the traffic → Traffic is rerouted to IPS/L2 firewall
- VM receives traffic after IPS/L2 firewall inspection

Service chaining: remove extra hops

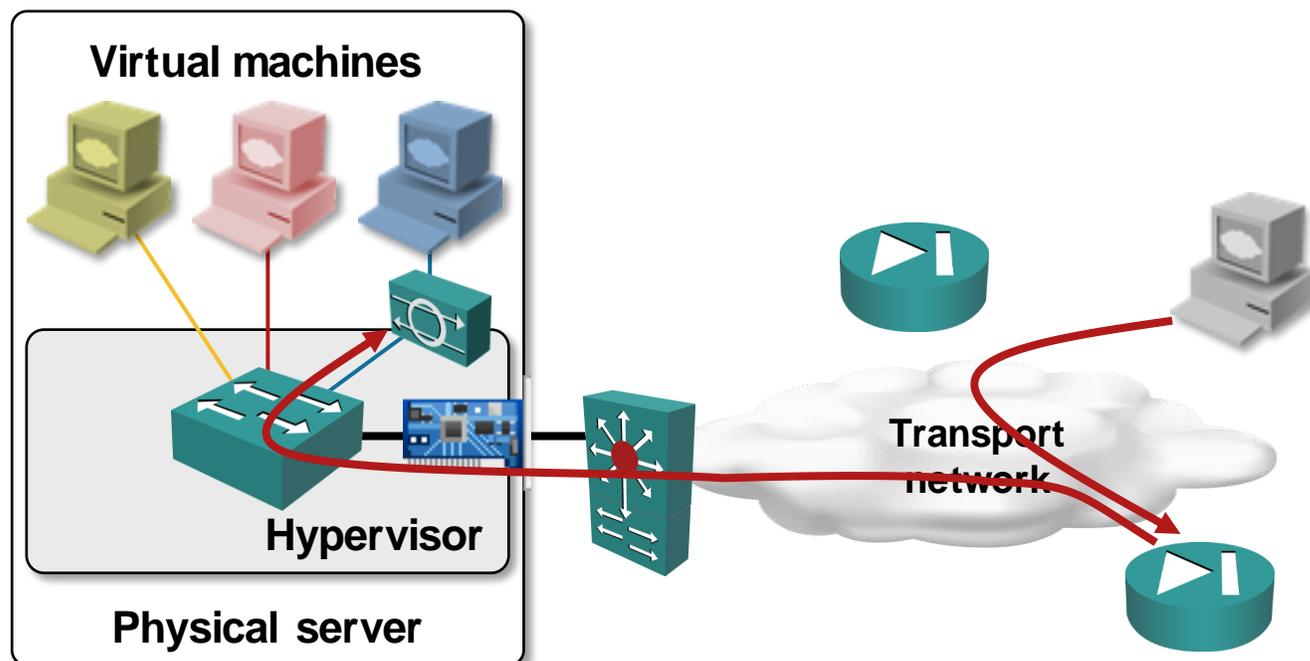
Service Insertion Gone Bad



- External traffic is sent to L3 appliance (based on IP routing)
- L3 appliance forwards traffic toward VM MAC address
- Hypervisor switch (or NIC driver) intercepts the traffic → Traffic is rerouted to IPS/L2 firewall
- VM receives traffic after IPS/L2 firewall inspection

Service chaining: remove extra hops

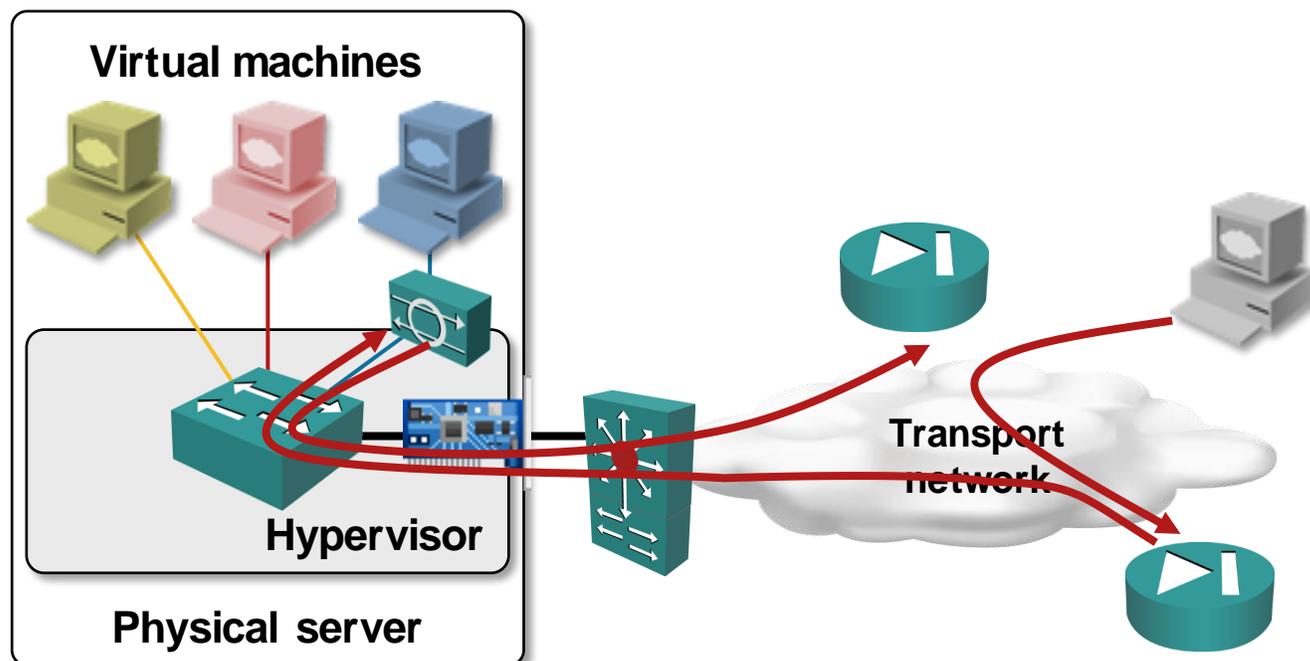
Service Insertion Gone Bad



- External traffic is sent to L3 appliance (based on IP routing)
- L3 appliance forwards traffic toward VM MAC address
- Hypervisor switch (or NIC driver) intercepts the traffic → Traffic is rerouted to IPS/L2 firewall
- VM receives traffic after IPS/L2 firewall inspection

Service chaining: remove extra hops

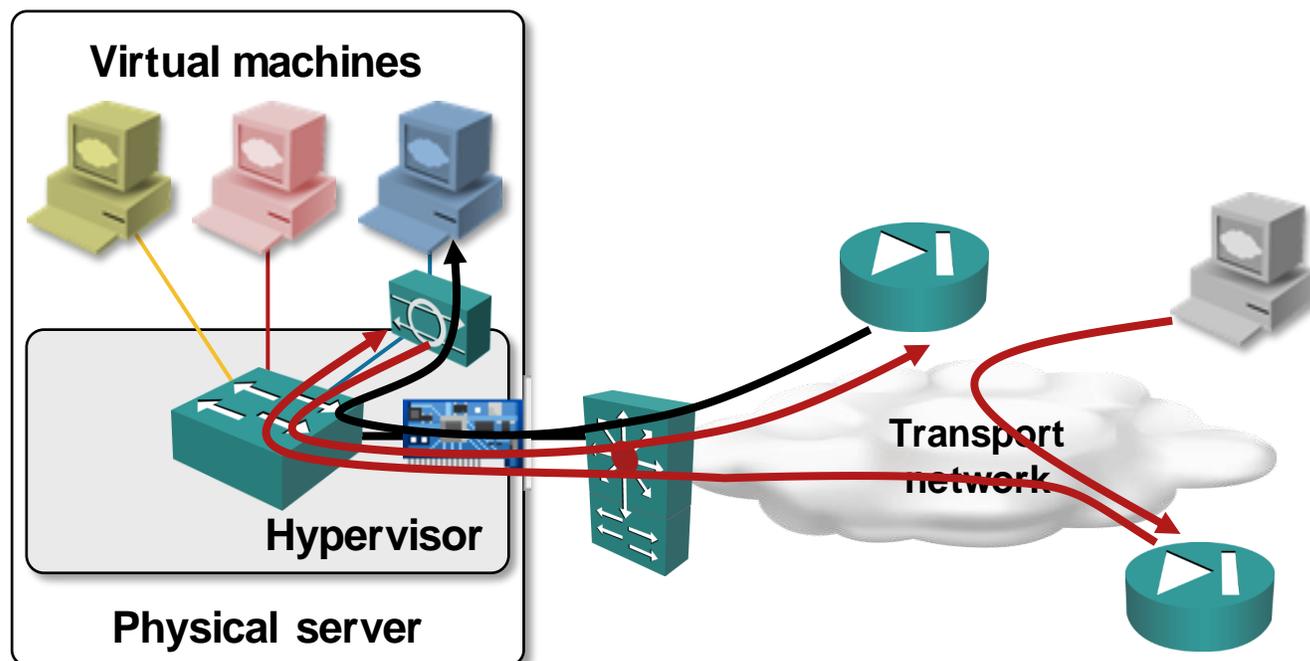
Service Insertion Gone Bad



- External traffic is sent to L3 appliance (based on IP routing)
- L3 appliance forwards traffic toward VM MAC address
- Hypervisor switch (or NIC driver) intercepts the traffic → Traffic is rerouted to IPS/L2 firewall
- VM receives traffic after IPS/L2 firewall inspection

Service chaining: remove extra hops

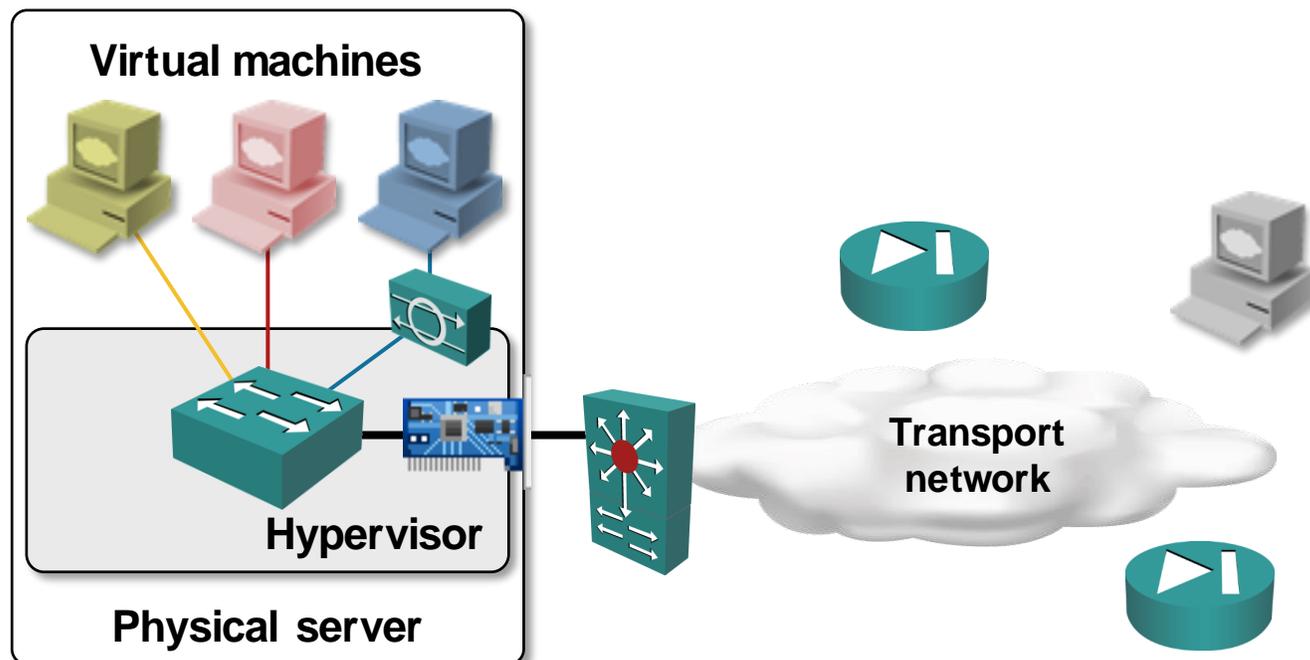
Service Insertion Gone Bad



- External traffic is sent to L3 appliance (based on IP routing)
- L3 appliance forwards traffic toward VM MAC address
- Hypervisor switch (or NIC driver) intercepts the traffic → Traffic is rerouted to IPS/L2 firewall
- VM receives traffic after IPS/L2 firewall inspection

Service chaining: remove extra hops

Service Chaining 101

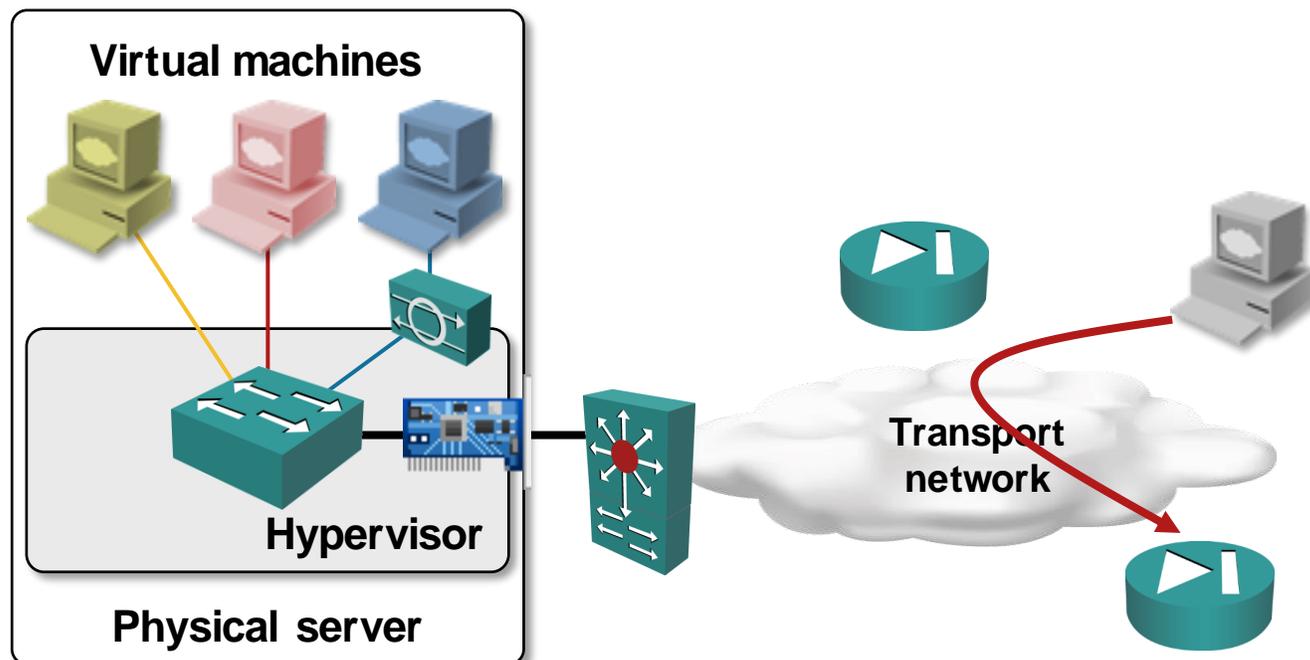


- Hypervisor switch redirects L3 appliance traffic directly to L2 appliance
- An extra hop through the hypervisor is eliminated

Sample commercial implementation: vPath 2.0 (Cisco)

- Combines Cisco ASA 1000V Cloud Firewall with VSG

Service Chaining 101

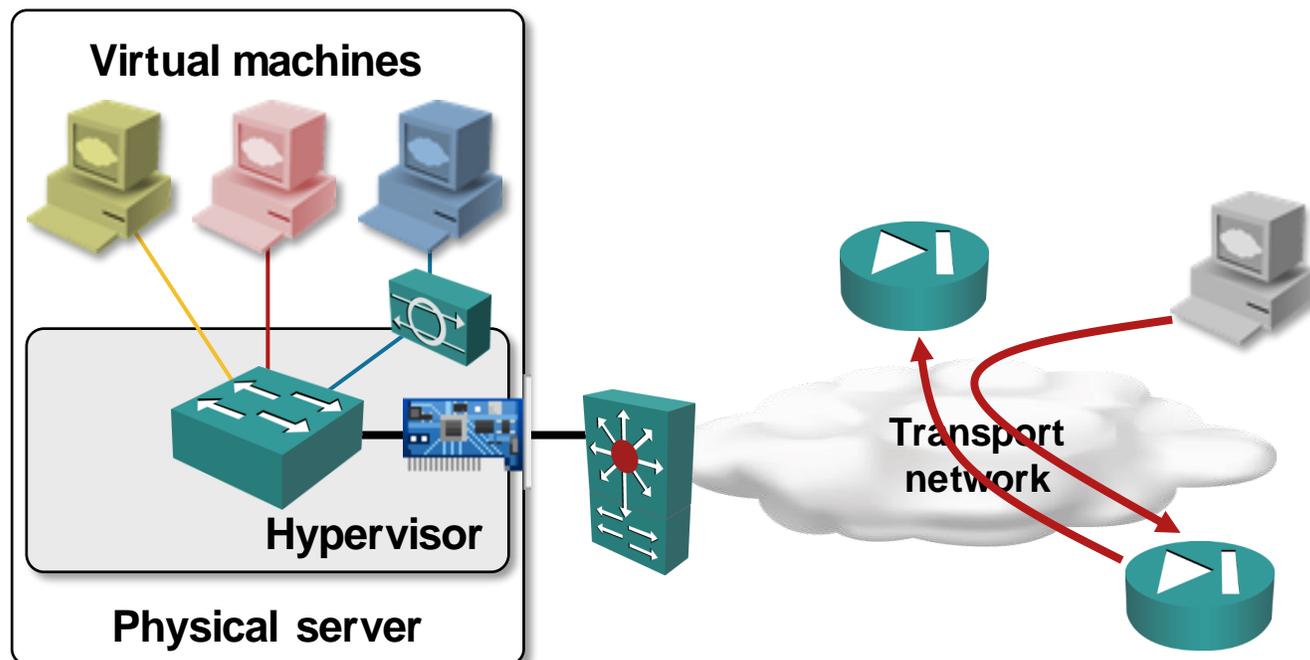


- Hypervisor switch redirects L3 appliance traffic directly to L2 appliance
- An extra hop through the hypervisor is eliminated

Sample commercial implementation: vPath 2.0 (Cisco)

- Combines Cisco ASA 1000V Cloud Firewall with VSG

Service Chaining 101

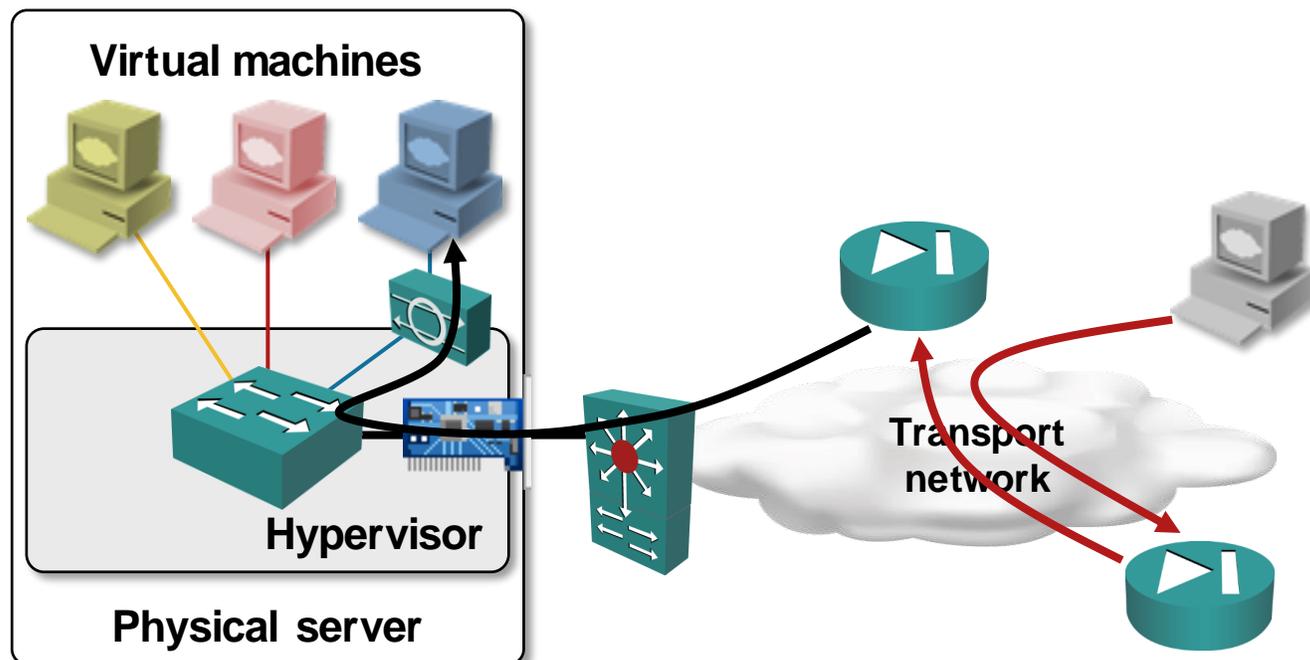


- Hypervisor switch redirects L3 appliance traffic directly to L2 appliance
- An extra hop through the hypervisor is eliminated

Sample commercial implementation: vPath 2.0 (Cisco)

- Combines Cisco ASA 1000V Cloud Firewall with VSG

Service Chaining 101



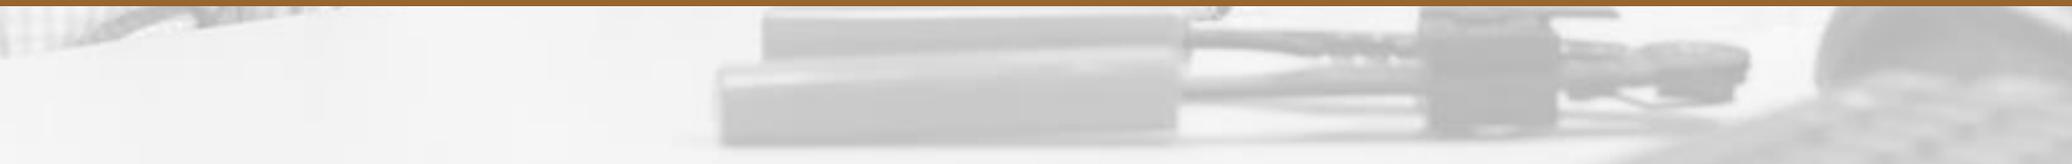
- Hypervisor switch redirects L3 appliance traffic directly to L2 appliance
- An extra hop through the hypervisor is eliminated

Sample commercial implementation: vPath 2.0 (Cisco)

- Combines Cisco ASA 1000V Cloud Firewall with VSG



Conclusions



Conclusions

VM appliances

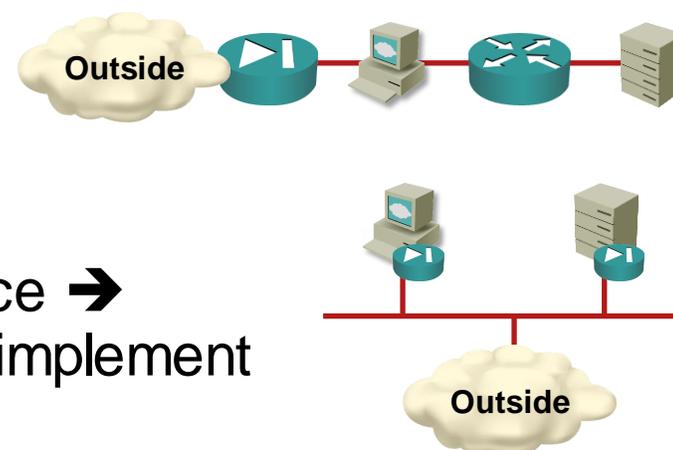
- Almost equivalent to physical devices
- Dedicated servers in high-security environments
- Work best with data center fabrics with equidistant endpoints

NIC-level firewalls

- Linear scale-out performance ... assuming you're ready for new security paradigms

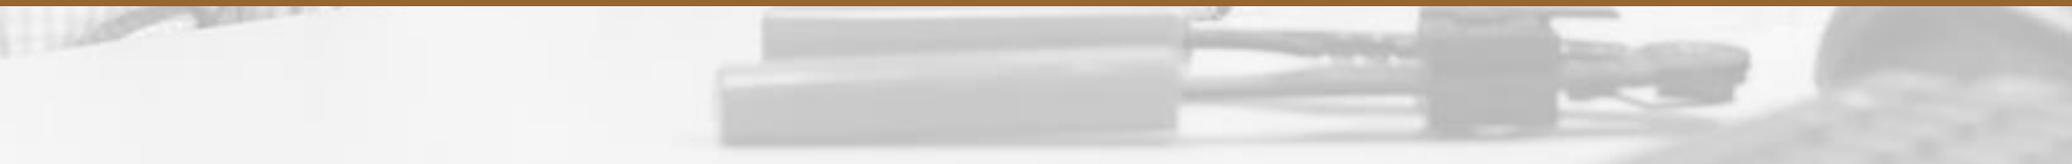
Service insertion and chaining

- Best of both worlds?
- Needs fast-path flow processing for performance → anything beyond smart packet filters is hard to implement





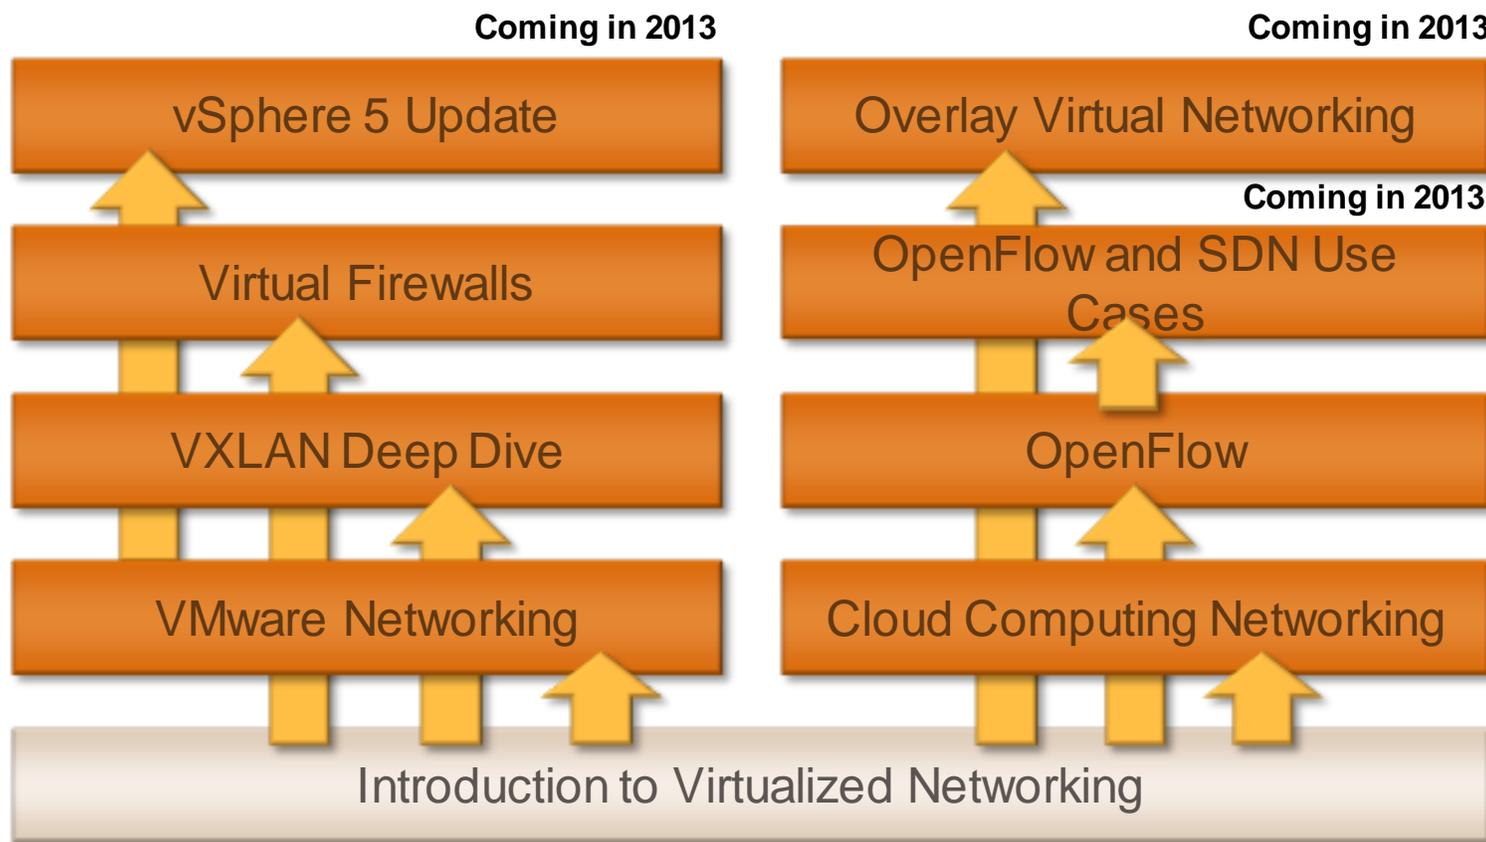
More Information



More Information: Blogs and Podcasts

- Packet Pushers Podcast & blog (packetpushers.net)
- Yellow bricks (Duncan Epping, VMware)
- Frank Denneman's blog
- Scott Lowe's blog
- RationalSurvivability.com (Christopher Hoff, Juniper)
- it20.info (Massimo Re Ferre, VMware)
- ChrisColloti.us (Chris Colloti)
- The Lone Sysadmin (Bob Plankers)
- High Scalability Blog (Todd Hoff)
- Errata Security (Robert Graham)
- Network Heresy (Nicira – dormant)
- Virtualization Security Roundtable
- blog.ioshints.info & ipSpace.net (yours truly)

Virtualization Webinars on ipSpace.net



Availability

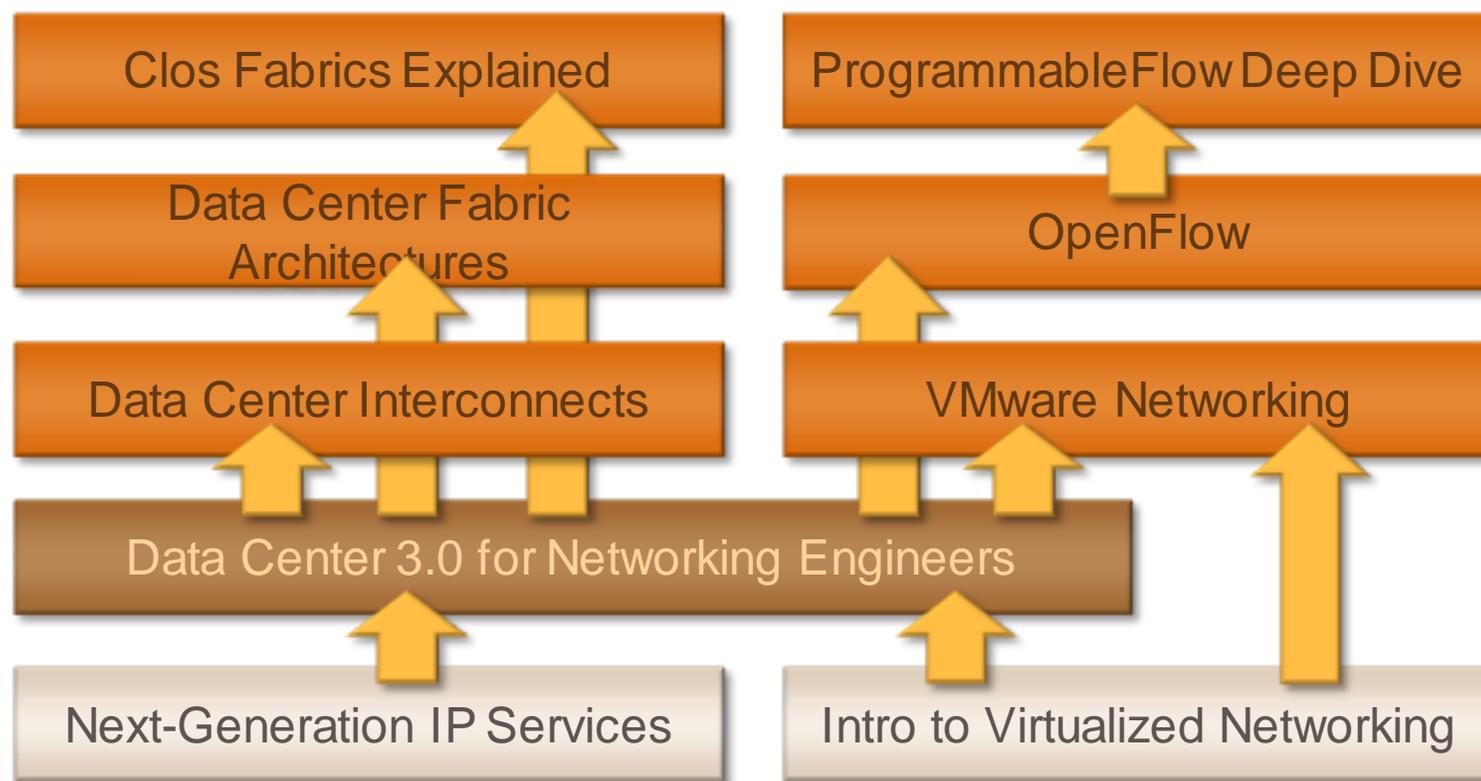
- Live sessions
- Recordings of individual webinars
- **Yearly subscription**

Other options

- Customized webinars
- ExpertExpress
- On-site workshops

More information @ <http://www.ipSpace.net/Webinars>

Data Center Webinars on ipSpace.net



Availability

- Live sessions
- Recordings of individual webinars
- **Yearly subscription**

Other options

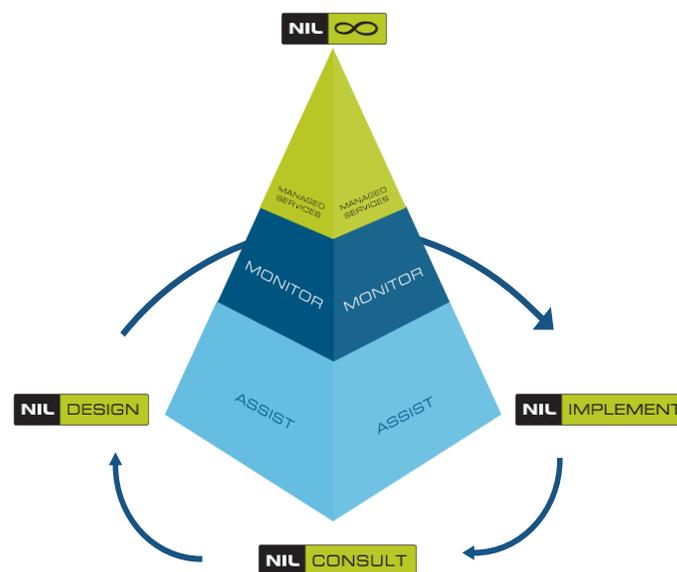
- Customized webinars
- ExpertExpress
- On-site workshops

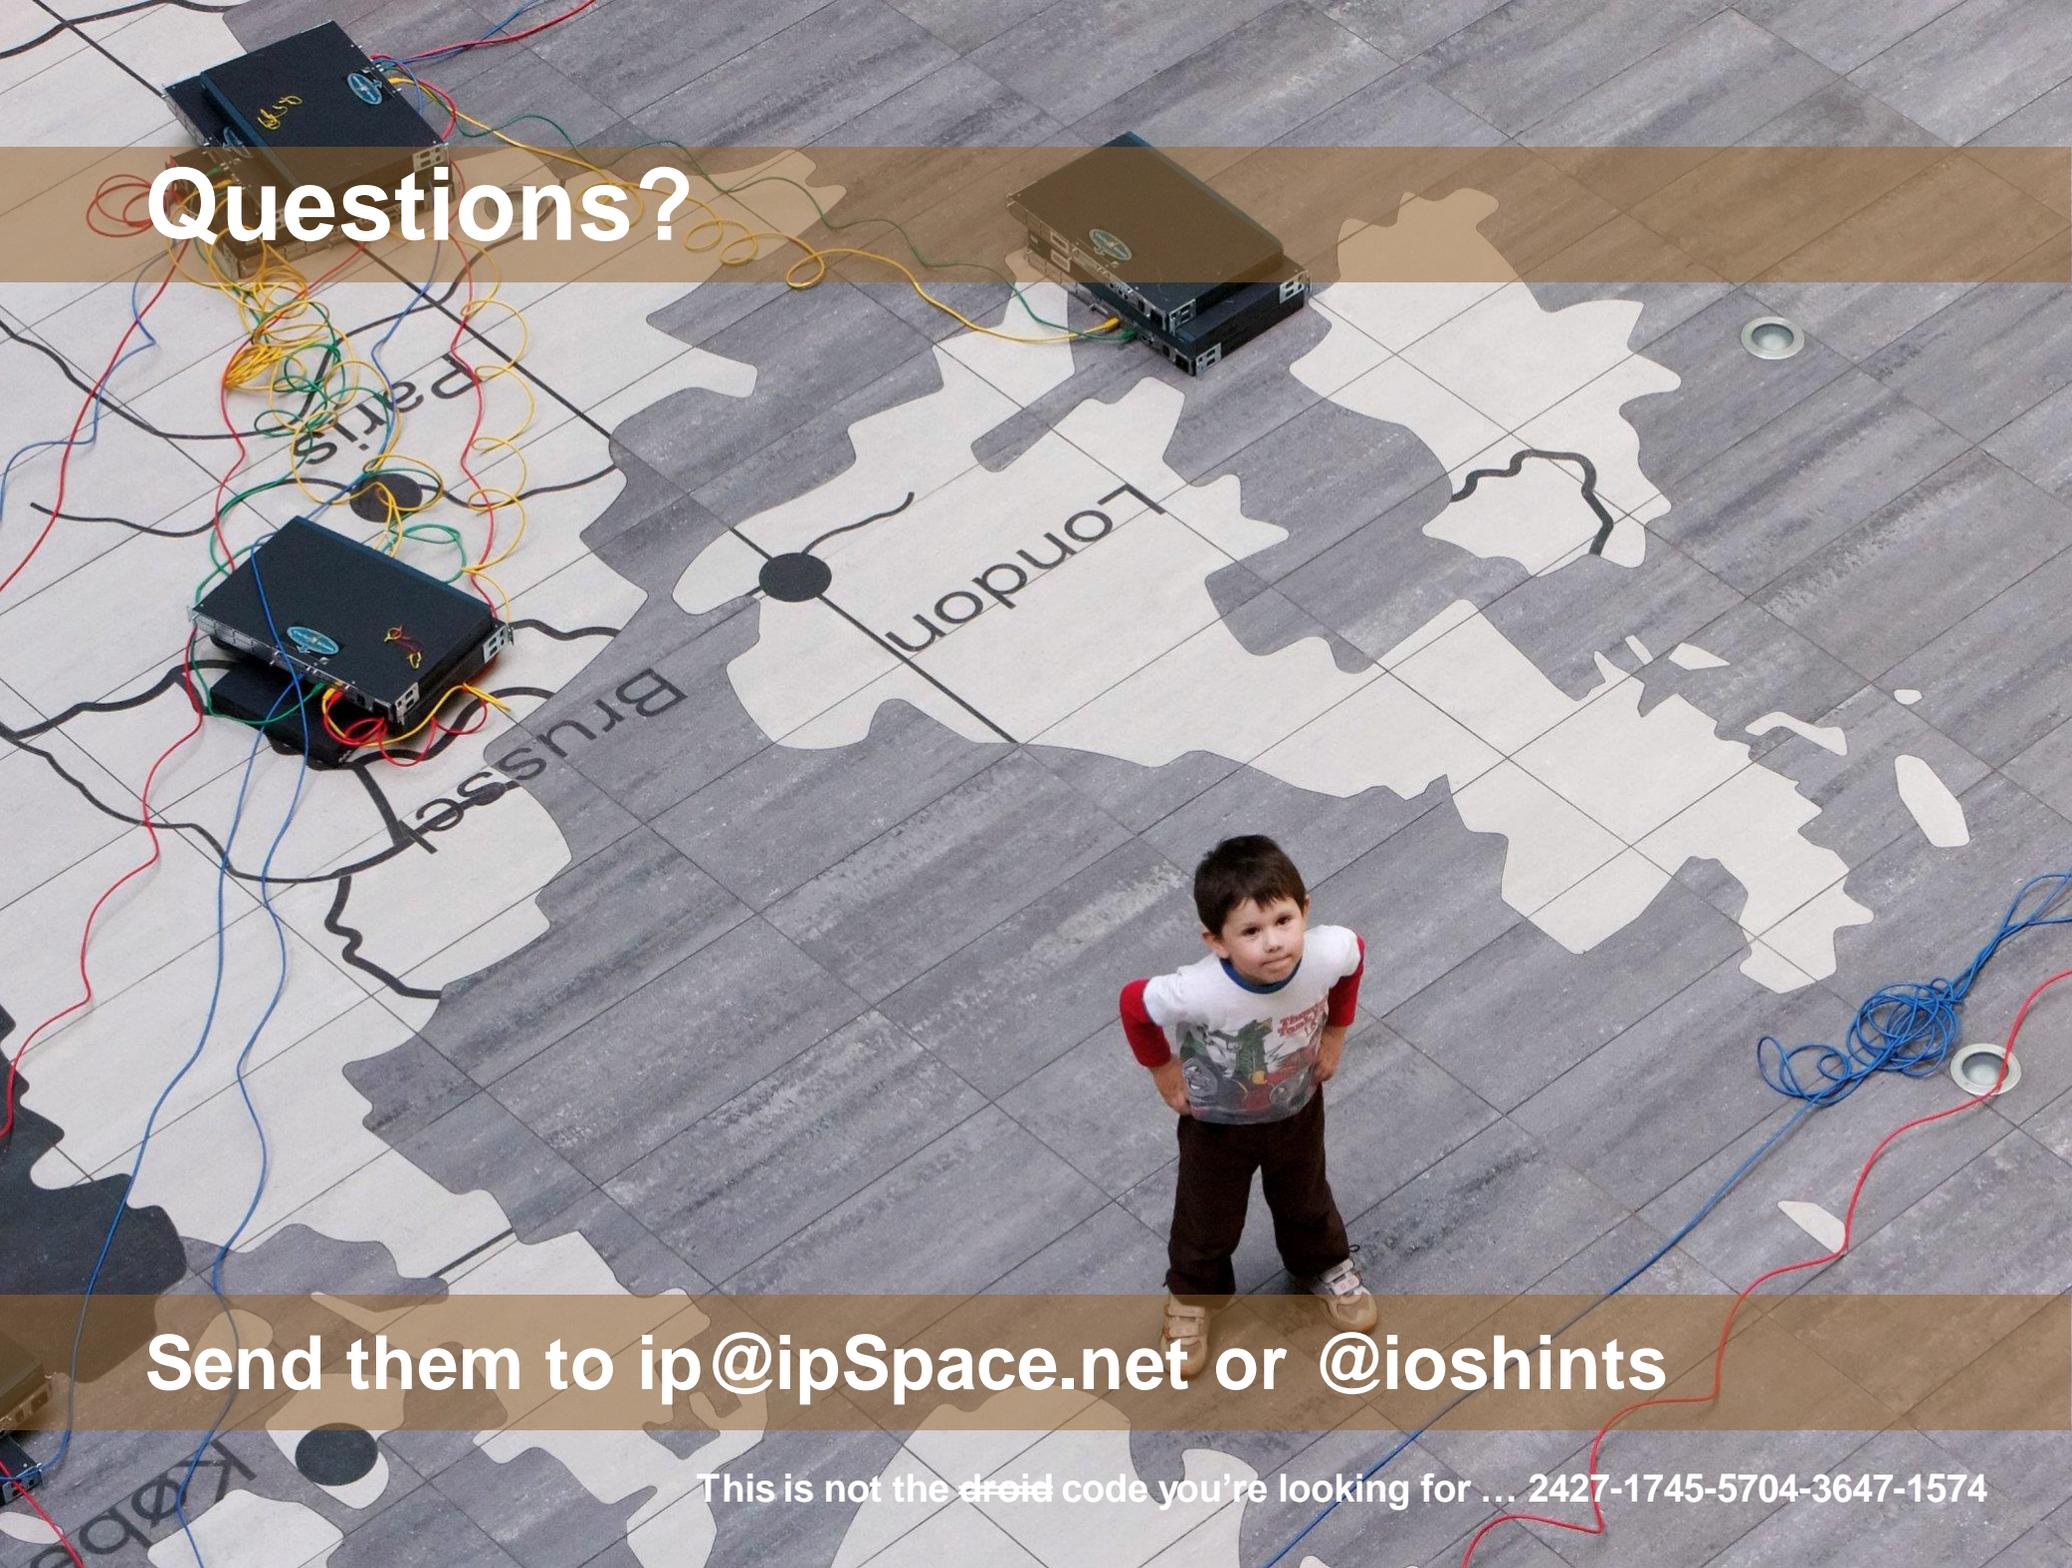
Need help?

[ExpertExpress](#) for quick discussions, reviews or second opinions

NIL's Professional/Learning Services

- In-depth design/deployment projects
- Data Center-, virtualization- and cloud-related training
- Details: www.nil.com, flipit.nil.com



A young child stands on a floor map of Europe. The map is drawn on a grey tiled floor and includes labels for 'Paris', 'London', and 'Brusset'. Several black network switches or routers are placed on the map, with a dense network of colorful cables (red, blue, yellow, green) connected to them. The child is wearing a white t-shirt with red sleeves and dark pants. The scene is set in a room with a grey tiled floor and a circular floor vent.

Questions?

Send them to ip@ipSpace.net or [@ioshints](https://twitter.com/ioshints)

This is not the friend code you're looking for ... 2427-1745-5704-3647-1574