ipSpace

Automating Network Security

Ivan Pepelnjak (ip@ipSpace.net)
Network Architect

ipSpace.net AG

Who is Ivan Pepelnjak (@ioshints)

Past

- Kernel programmer, network OS and web developer
- Sysadmin, database admin, network engineer, CCIE
- Trainer, course developer, curriculum architect
- Team lead, CTO, business owner

Present

- Network architect, consultant, blogger, webinar and book author
- Teaching the art of Scalable Web Application Design

Focus

- Large-scale data centers, clouds and network virtualization
- Scalable application design
- Core IP routing/MPLS, IPv6, VPN

CISCO CCIE EMERITUS

vmware vEXPERT

Sounds Familiar?

- Increase flexibility while reducing costs
- Faster application deployments
- Compete with public cloud offerings

Does Your Data Center Look Like This?

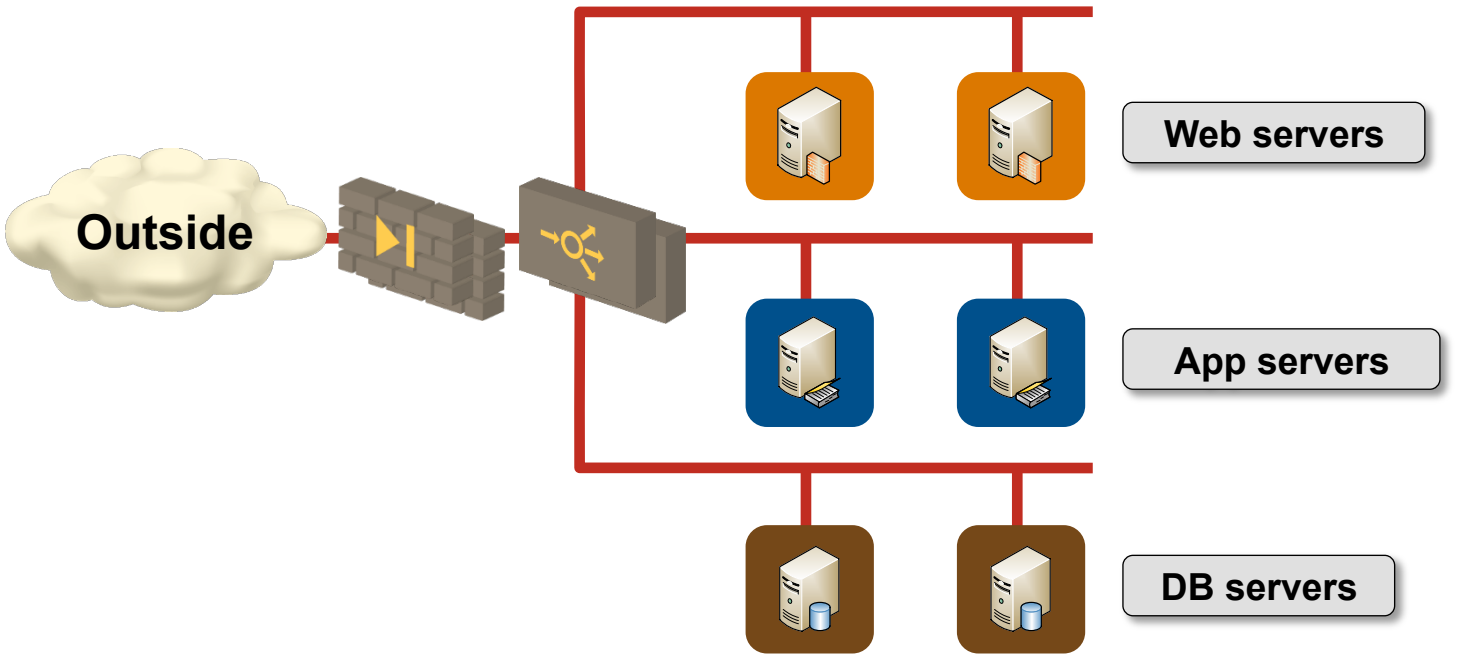

- Single pair of central firewalls and load balancers
- Large VLANs spanning the whole data center
- All inter-subnet traffic goes through an appliance ➔ chokepoint

How Long Does It Take To Create a New Firewall Rule?

Think Again …
From Initial Request to Production

It Works for Other People

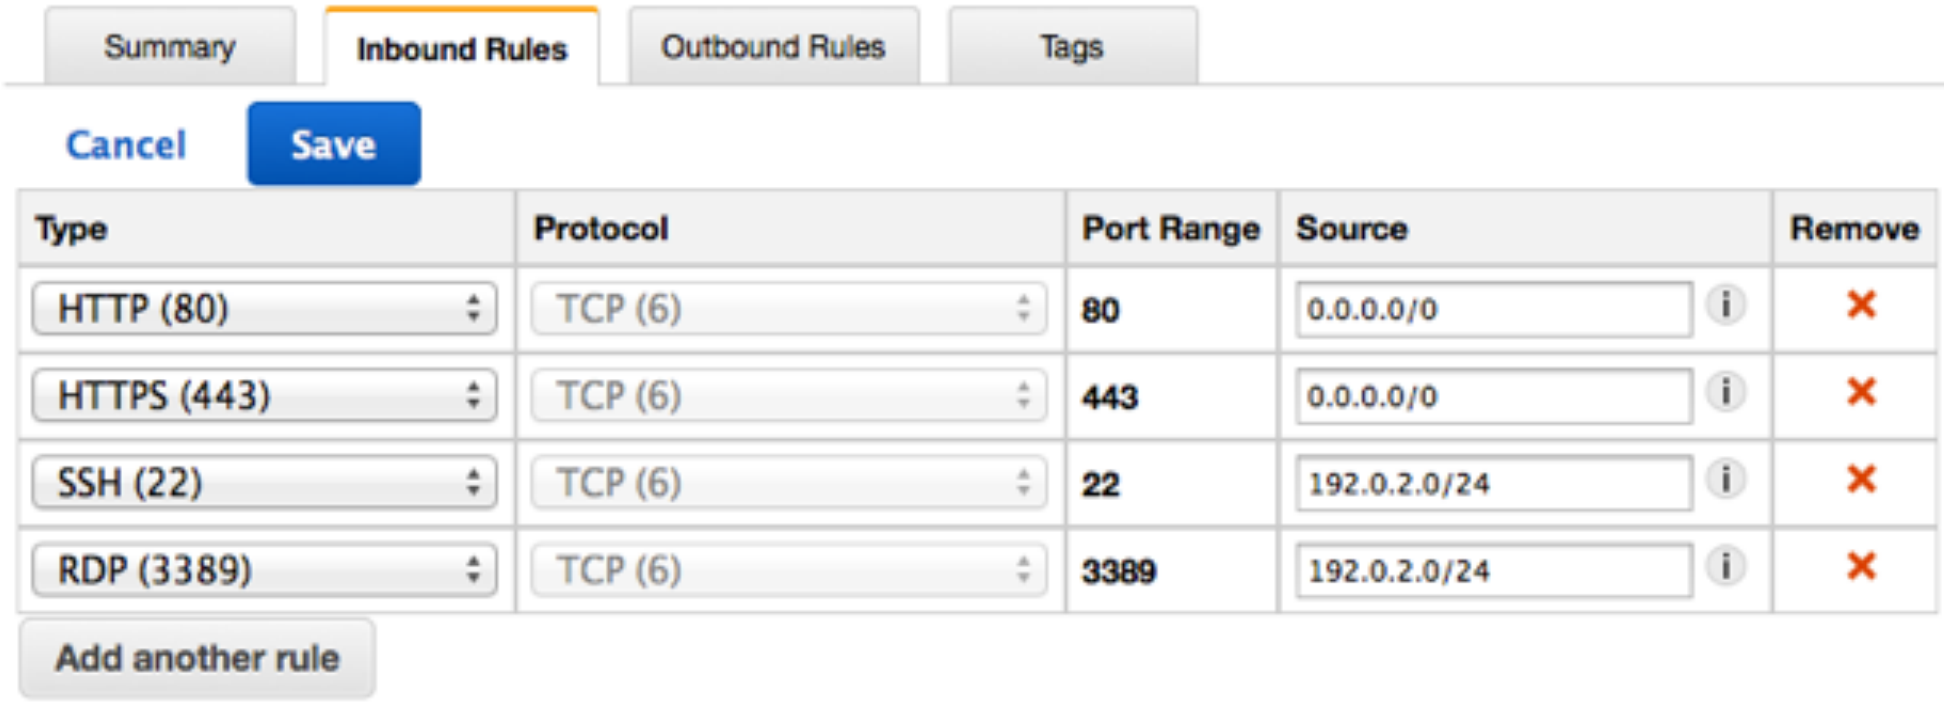

Why does it work?

- Simple rules, no ALG
- Configured by tenants (application owners)
- Deployed in real time

Source: Amazon VPC documentation

Automating Network Security

We Want to Be Here

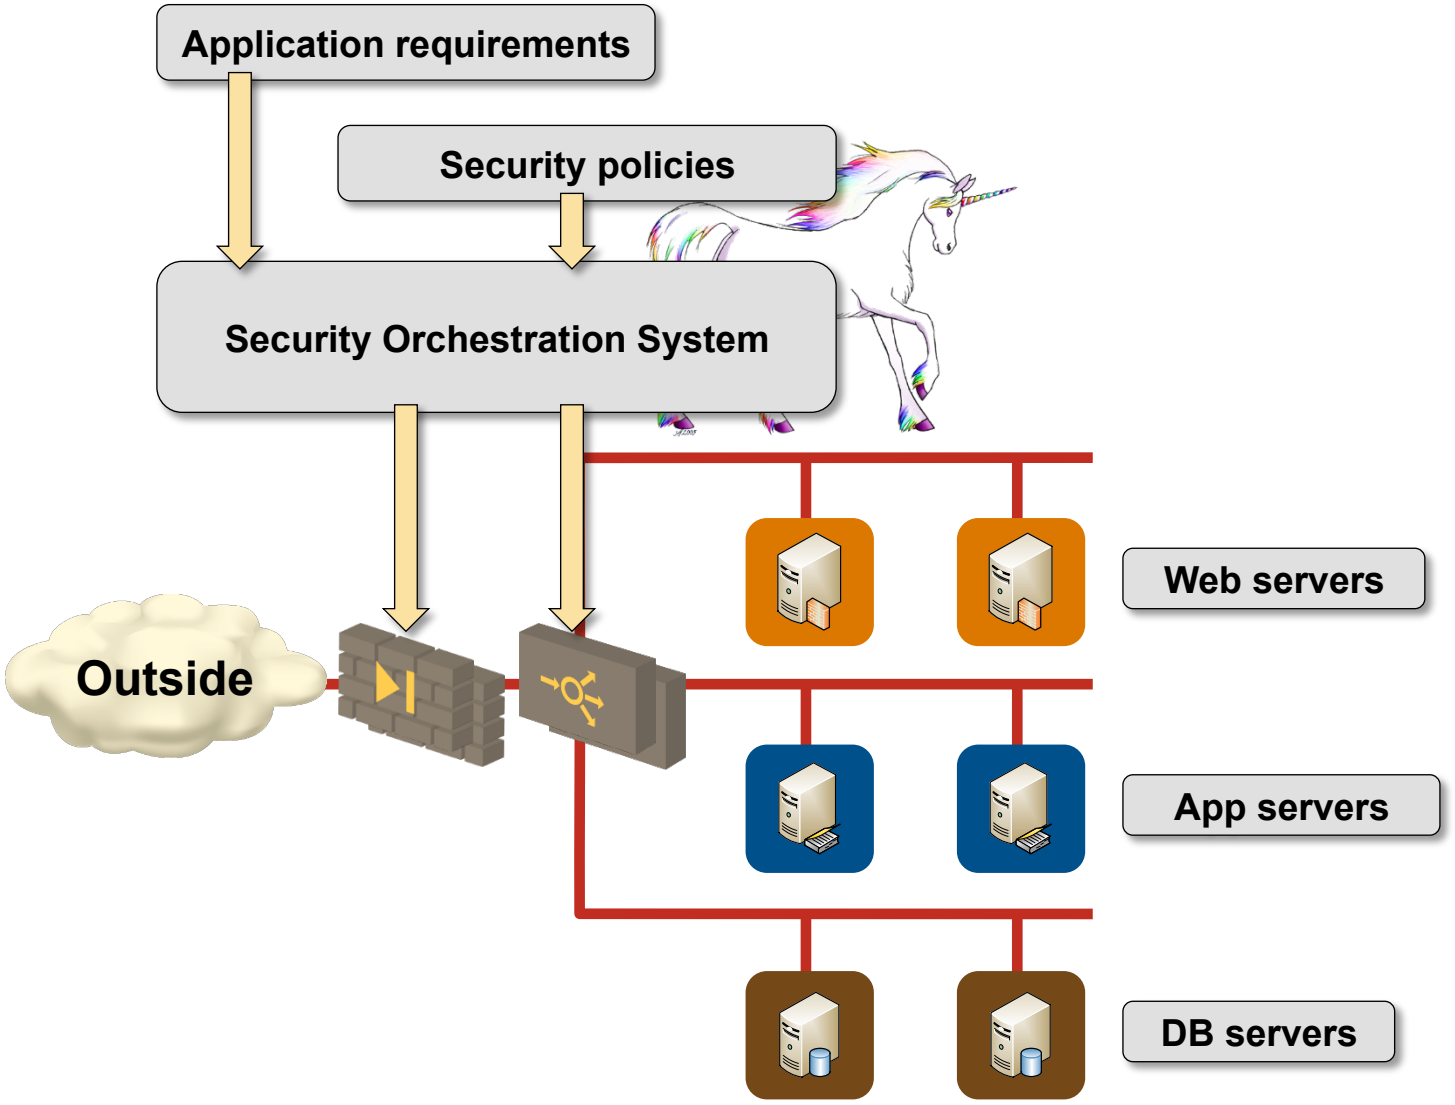

How Do We Get There?

Simplify

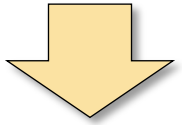

Standardize

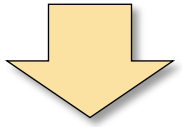

Automate

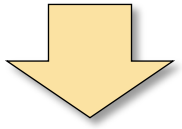

Abstract

Before We Start: Reduce the Blast Radius

Virtualize Everything

NFV Performance Is Acceptable

Some performance maximums

- 50+ Gbps through a Xeon-based server (Multipath TCP)
- 50 Mpps on Open vSwitch with DPDK (~ 130 Gbps with IMIX traffic load)
- 200 Gbps on a Xeon server (Snabb switch)

Commercial products performance

- A10 load balancer VM on a single core: up to 4 to 8 Gbps
- F5 load balancer VM: 3 Gbps @ 2 vCPU
- Vmware NSX Edge Services Router: 10 Gbps firewall, 4 – 10 Gbps load balancer
- Juniper Firefly Perimeter: 1,1 – 4,1 Gbps (IMIX/UDP)
- Palo Alto firewall: 1 Gbps @ 4 vCPU
- Vyatta 5600 series routers: 10 Gbps @ 1 core

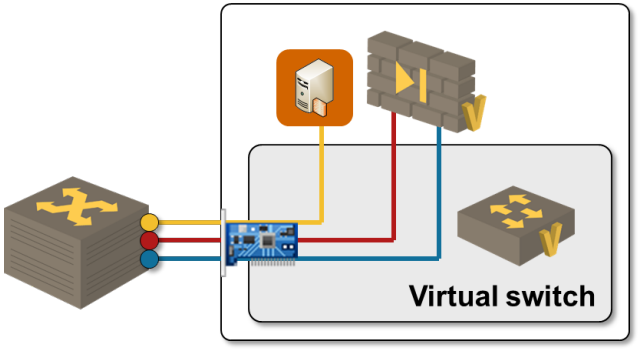

Potential showstopper: SSL tps (key generation)

Minimize Complexity & Standardize

Minimize Complexity & Standardize

Per-application firewall ruleset

- Easier to manage, understand and audit
- Ruleset lives and dies with the application
- Stale rules in central firewall are gone

Standardize rulesets

- Applications are not as special as their owners think
- Create a standard ruleset for each application class
- Changes to ruleset void warranty

Rethink the security policy

- Packet filters are often good enough
- Stateful inspection is required only on links to external networks or for outgoing sessions

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

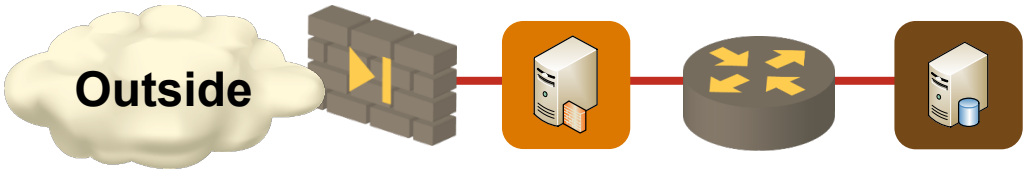

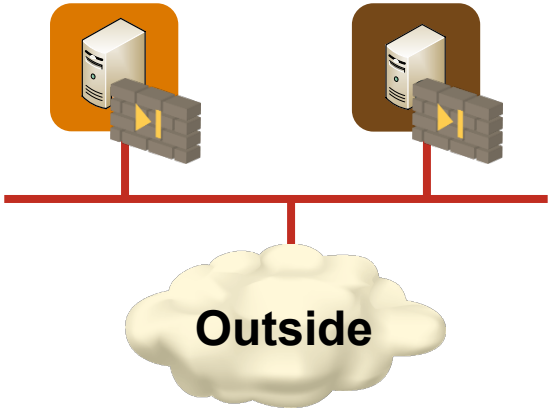

Decouple Virtual and Physical Worlds

Decouple Virtual and Physical Worlds

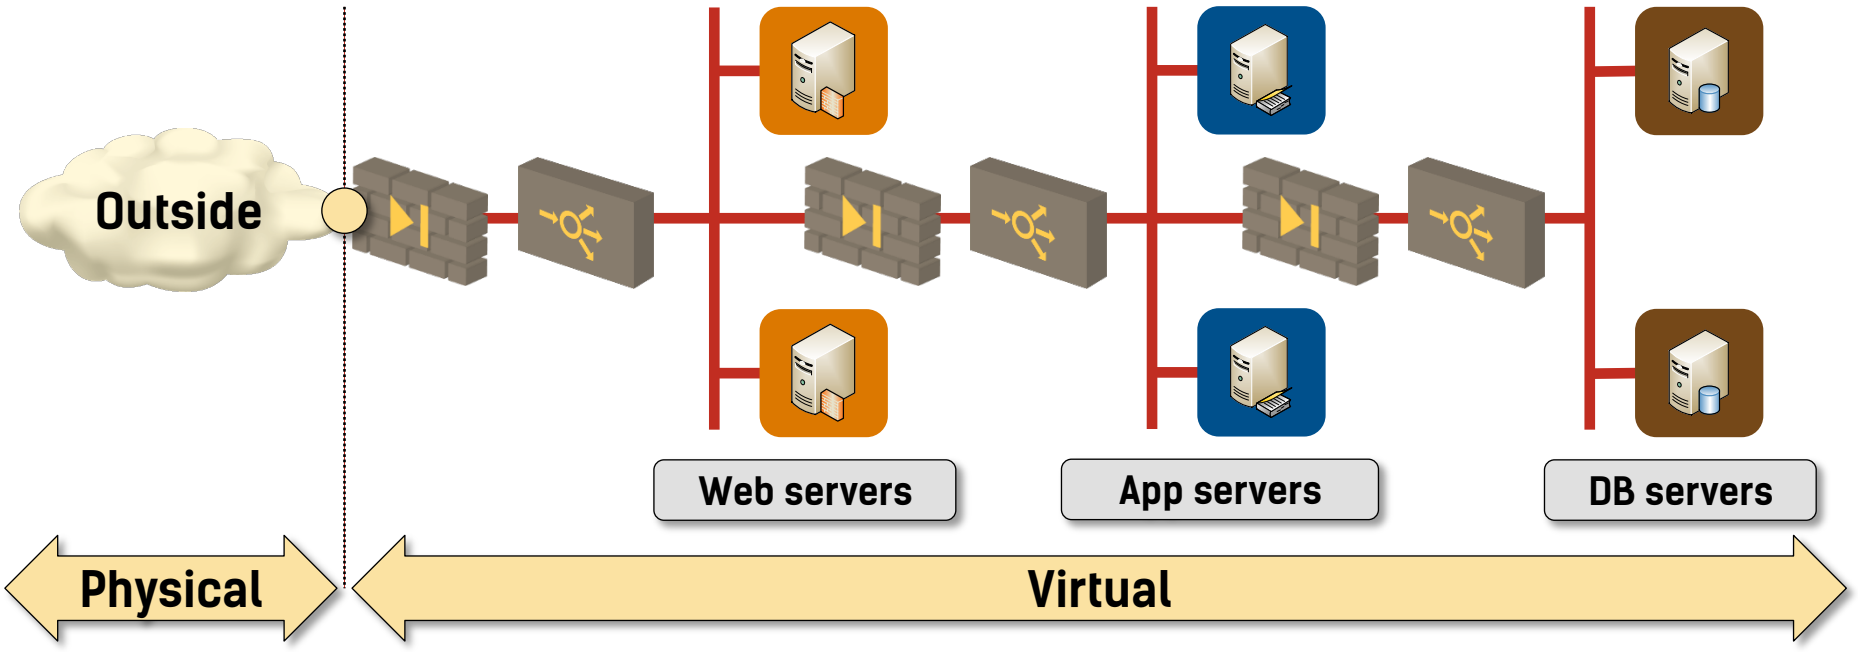

- Simplified workload migration
- Automated deployment
- No interaction with the physical gear ➔ no maintenance windows

Network Appliance Implementation: Decision Points

- Per-application appliances ➔ use virtual appliances, not contexts
- Firewall & load balancer self-service ➔ use virtual appliances
- Rigid security rules ➔use physical devices (but you'll get the cloud you deserve)

Compromise: Combine the virtual and the physical appliances

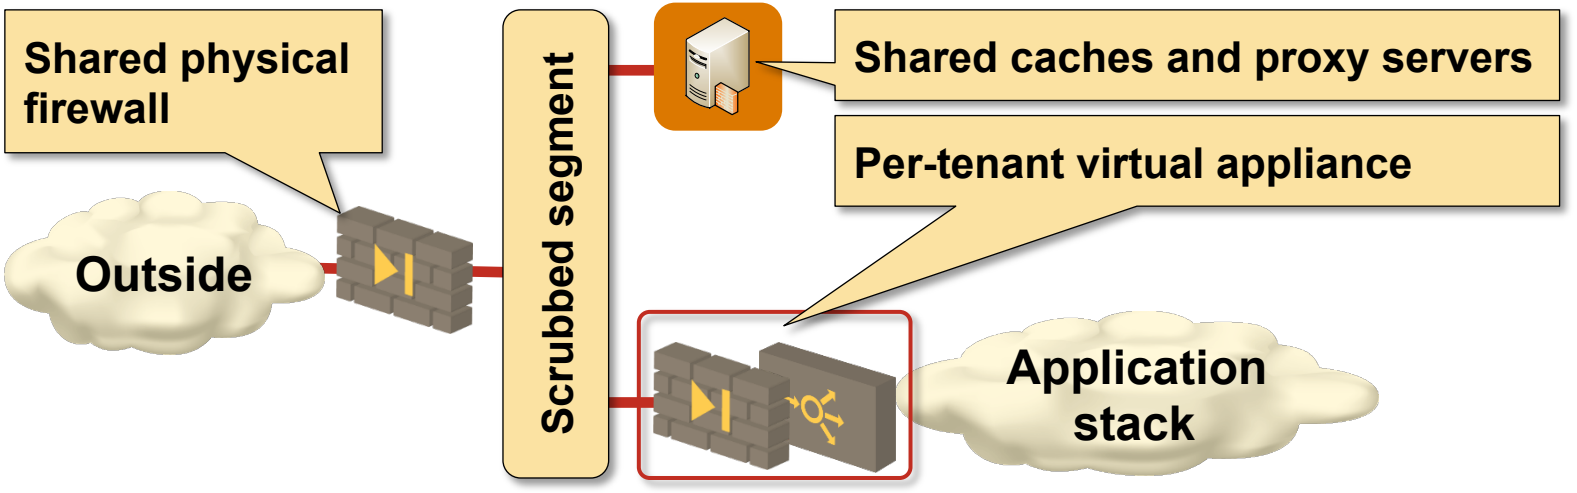

What Others Are Doing

Multi-Tenant Isolation With Firewalls

Typical IaaS cloud provider approach

- Customer VMs attached to “random” L3 subnets (or per-customer subnets)
- VM IP addresses allocated by the IaaS provider (example: DHCP)
- Predefined configurations or user-controlled firewalls

Implementations

- OpenStack (*iptables* or vendor plugin)
- CloudStack (*iptables*)
- VMware NSX
- Juniper Contrail
- Hyper-V

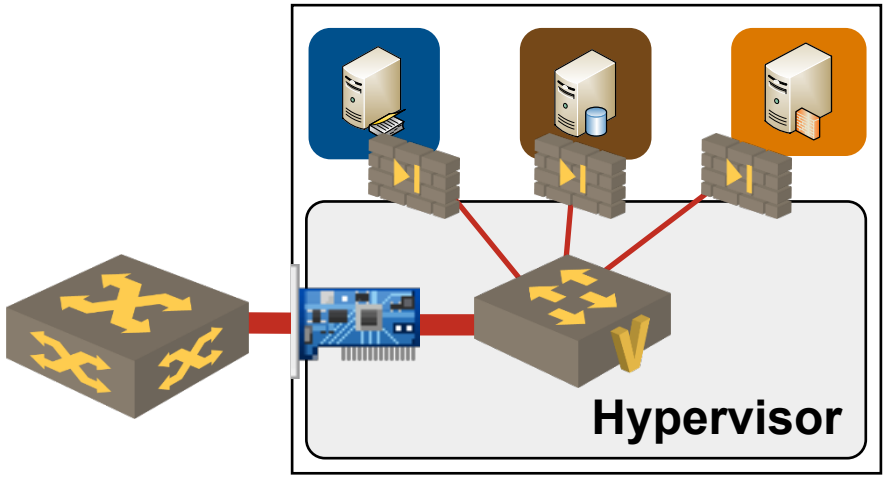

Most implementations use functionality equivalent to reflexive ACLs

Integration with Orchestration Systems

Fact: Most traffic filtering products include L3/L4 information in filter rules

Challenge: Integrating filter rules in dynamic VM environment

Traditional solutions: use subnets to indicate security group membership

Tighter integration with orchestration systems

- Defined security group membership based on port groups, VM names or other VM attributes
- Collect VM information from the orchestration system (example: vCenter)
- Automatically build L3-7 filter rules from orchestration system information

Sample products:

- Security groups in OpenStack and CloudStack
- VMware vShield Edge and NSX Distributed Firewall
- Palo Alto Panorama
- Cisco Prime Network Services Controller

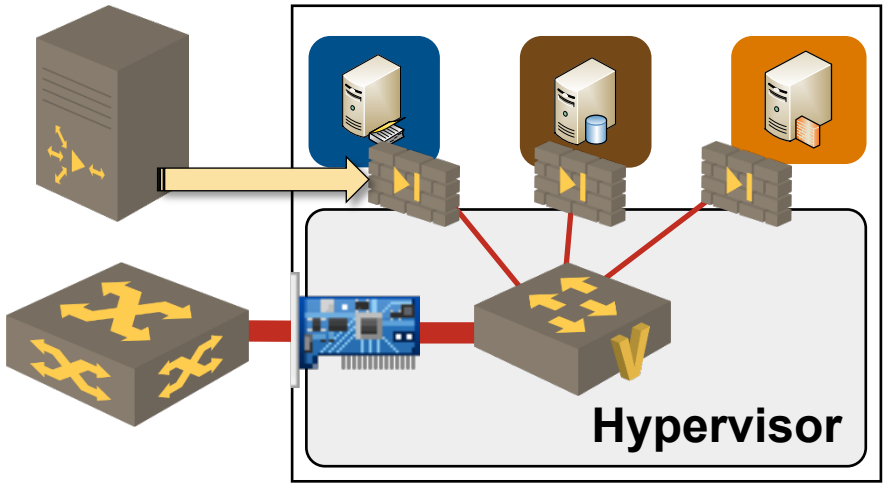

Tap Aggregation on Arista 7150

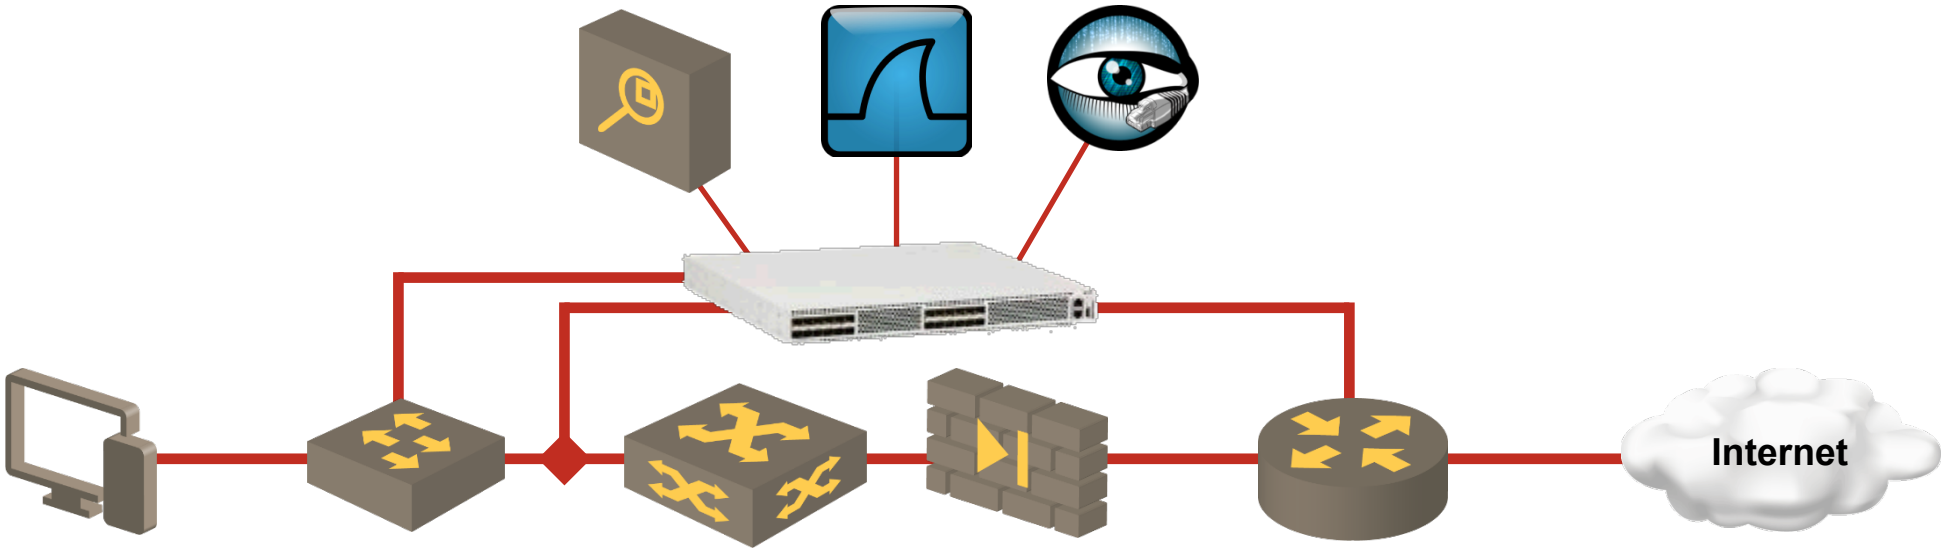

- Regular, tap and tool ports
- Tap tools are input-only ports, tool ports are output-only ports
- MAC learning and STP is disabled on tap and tool ports
- Switch operates in *normal* or *tap aggregation* mode
- Tap and tool ports disabled in normal mode, regular ports in aggregation mode
- Programmed through XMPP, REST API, Puppet, Chef …

Traffic Tapping with OpenFlow-Enabled Switches

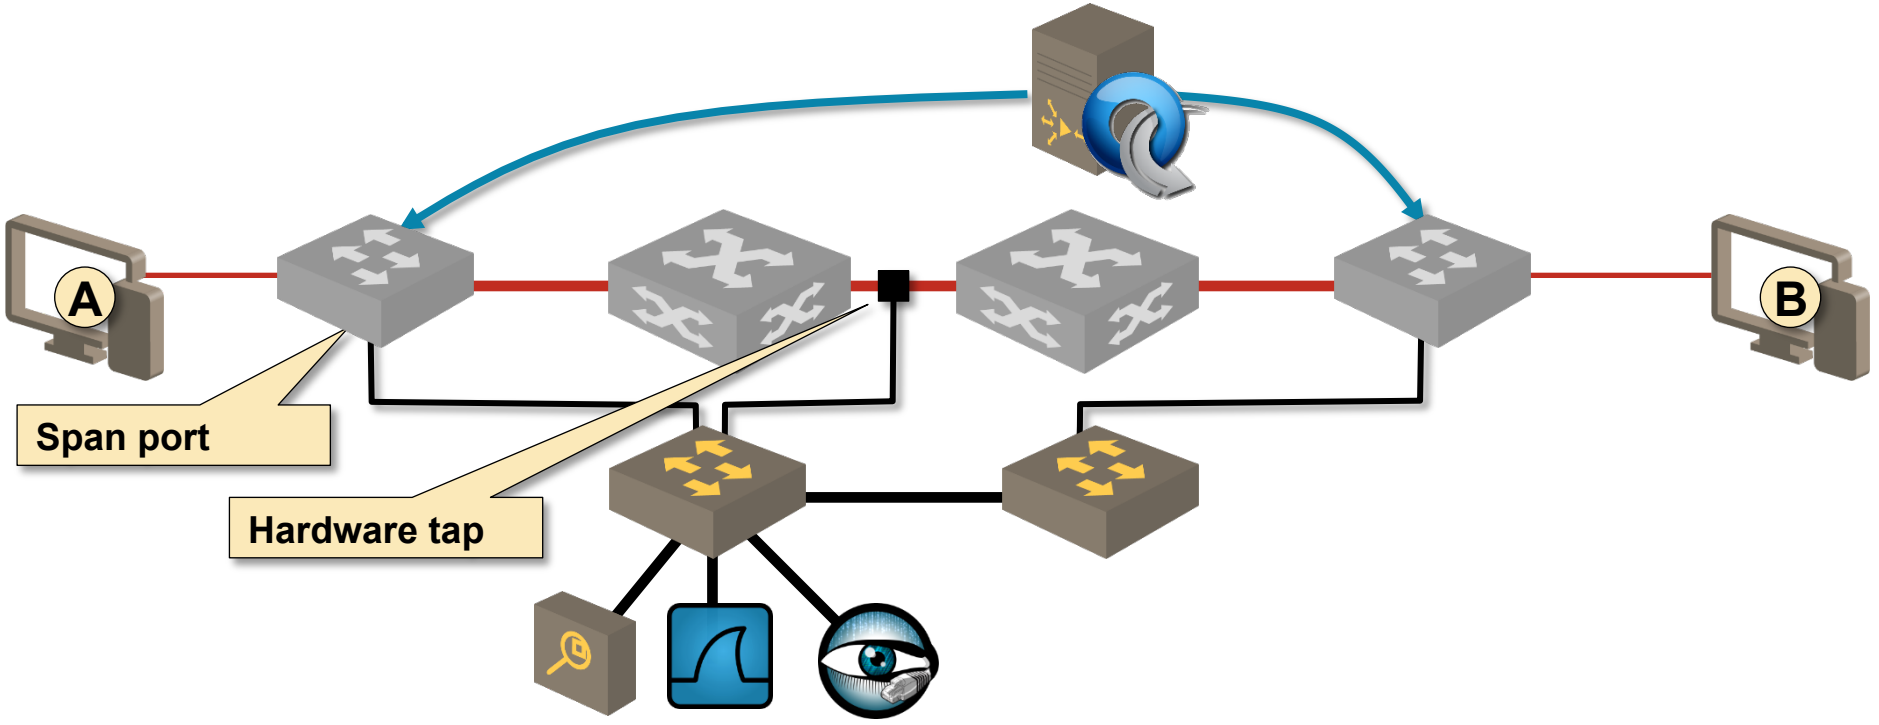

- Use OpenFlow flows to mirror traffic to SPAN ports
- Higher traffic redirection granularity ➔ lower number of SPAN ports required
- Any OpenFlow controller capable of inserting individual flows could be used

Solution from Cisco (on Nexus 3000 switches) and Arista (monitor bind mode)

Firewall Cluster with Deterministic ECMP

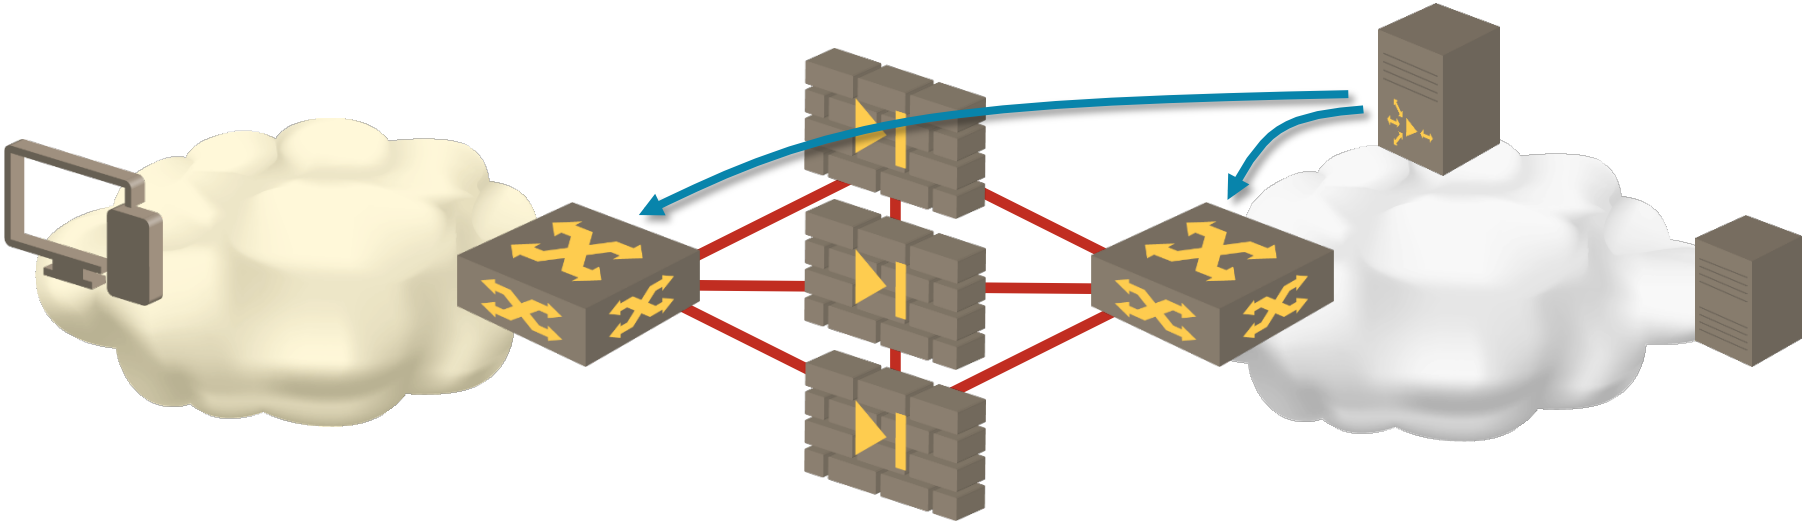

Switches have synchronized ECMP behavior

- Same hashing mechanism
- Coarse-grained ECMP based on source or destination IP addresses

Caveats

- Number of active links must match on both ends (need central control)
- Addition / removal of devices should cause minimal disruption

Demonstrated: Arista + Palo Alto

Remote-Triggered Black Hole: a Decade of SDN

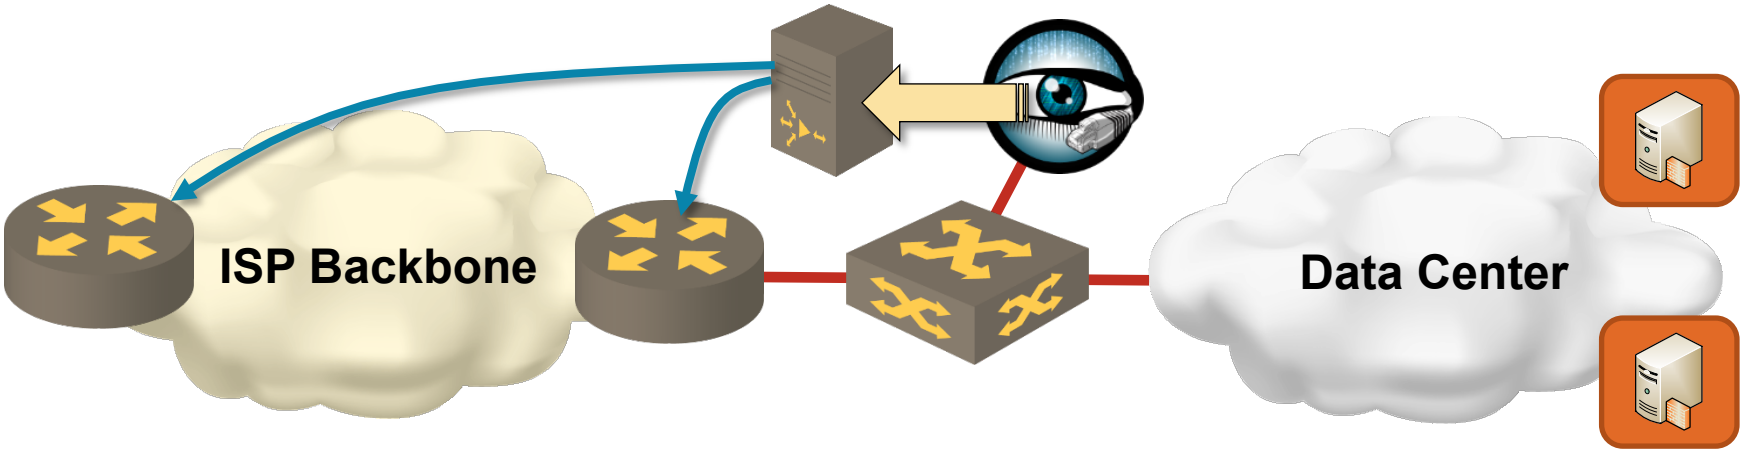

RTBH 101

- Install a host route to a bogus IP address (RTBH address) pointing to *null* interface on all routers
- Use BGP to advertise host routes of attacked hosts with next-hop = RTBH address (alternative: use BGP communities, set next-hop locally)
  ➔ drops DoS traffic at backbone ingress point
- Use uRPF to drop traffic *from* DoS sources

Widely used in ISP environments

Scale-Out IPS with OpenFlow Controller

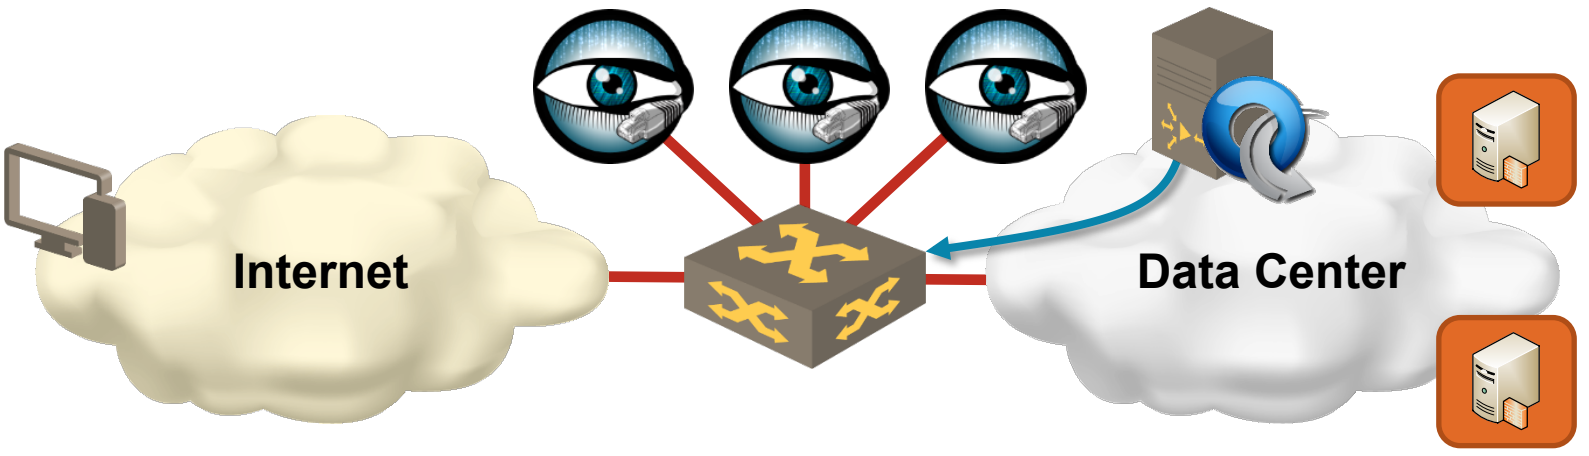

OpenFlow used to distribute the load to multiple IDS appliances

- Coarse-grained flows deployed on the OpenFlow switch
- Flow granularity adjusted in real time to respond to changes in traffic
- Each appliance receives all traffic from a set of endpoints (complete session and endpoint behavior visibility)

Pretty easy to implement, used by (at least) Indiana University

Scale-Out IDS Using OpenFlow to Block Traffic

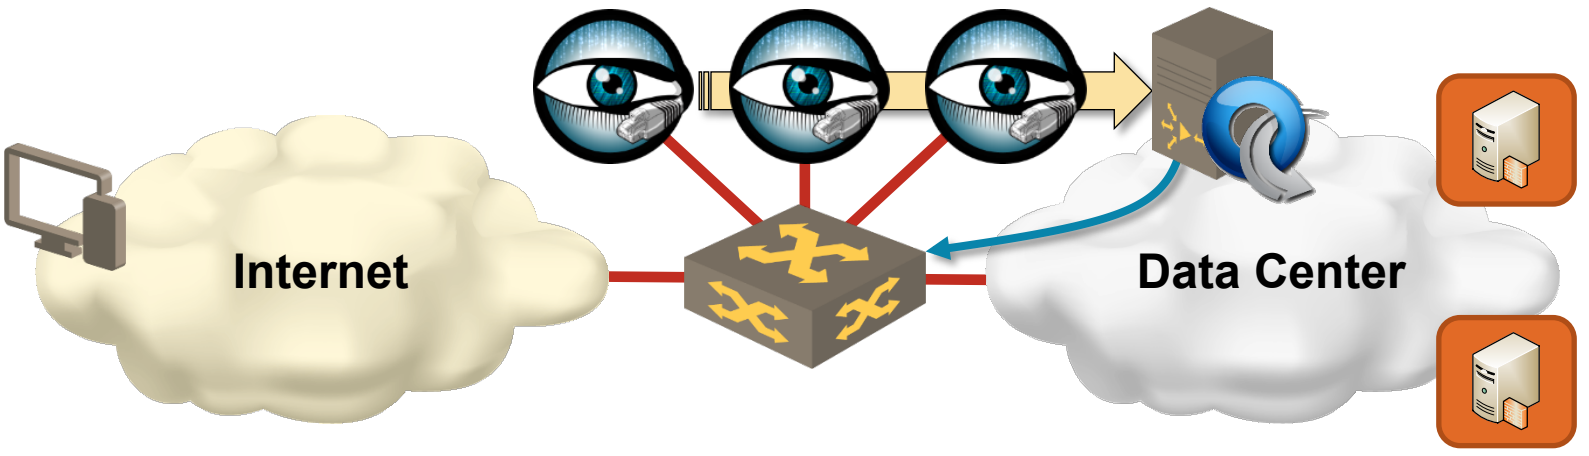

DoS detection system reports offending X-tuples

- Source IP addresses
- Targeted servers
- Applications (port numbers)

OpenFlow controller installs *drop* flows

The Roadblocks

The Roadblocks

Internal

- Mindset
- Processes and procedures
- Rigid security policies

External

- Licensing
- Management tools
- Auditing tools

Does It Work Today?

- Virtual appliance vendors: Palo Alto, Checkpoint, Fortinet, Cisco, VMware
- Distributed virtual firewalls: VMware NSX, Palo Alto, …
- Orchestration tools: VMware vCloud, CloudStack, OpenStack…
- Automation tools: Ansible, Chef, Puppet…

Start NOW

Gartner on Shiny New Object Syndrome

[…] address the following questions before introducing any new technology:

- Can the root issue be addressed via a policy or process change?
- If we wait a year, will this become a commoditized capability from established providers (or my existing providers)?
- Do we have existing network, security, or management capabilities that can address the bulk (i.e., 85%) of the technological requirements?
- Do we have the right process and staff expertise to properly leverage the new technology?

Source: http://blogs.gartner.com/andrew-lerner/2015/01/15/netsecdirtydozen/

Advanced SDN and Network Automation Track

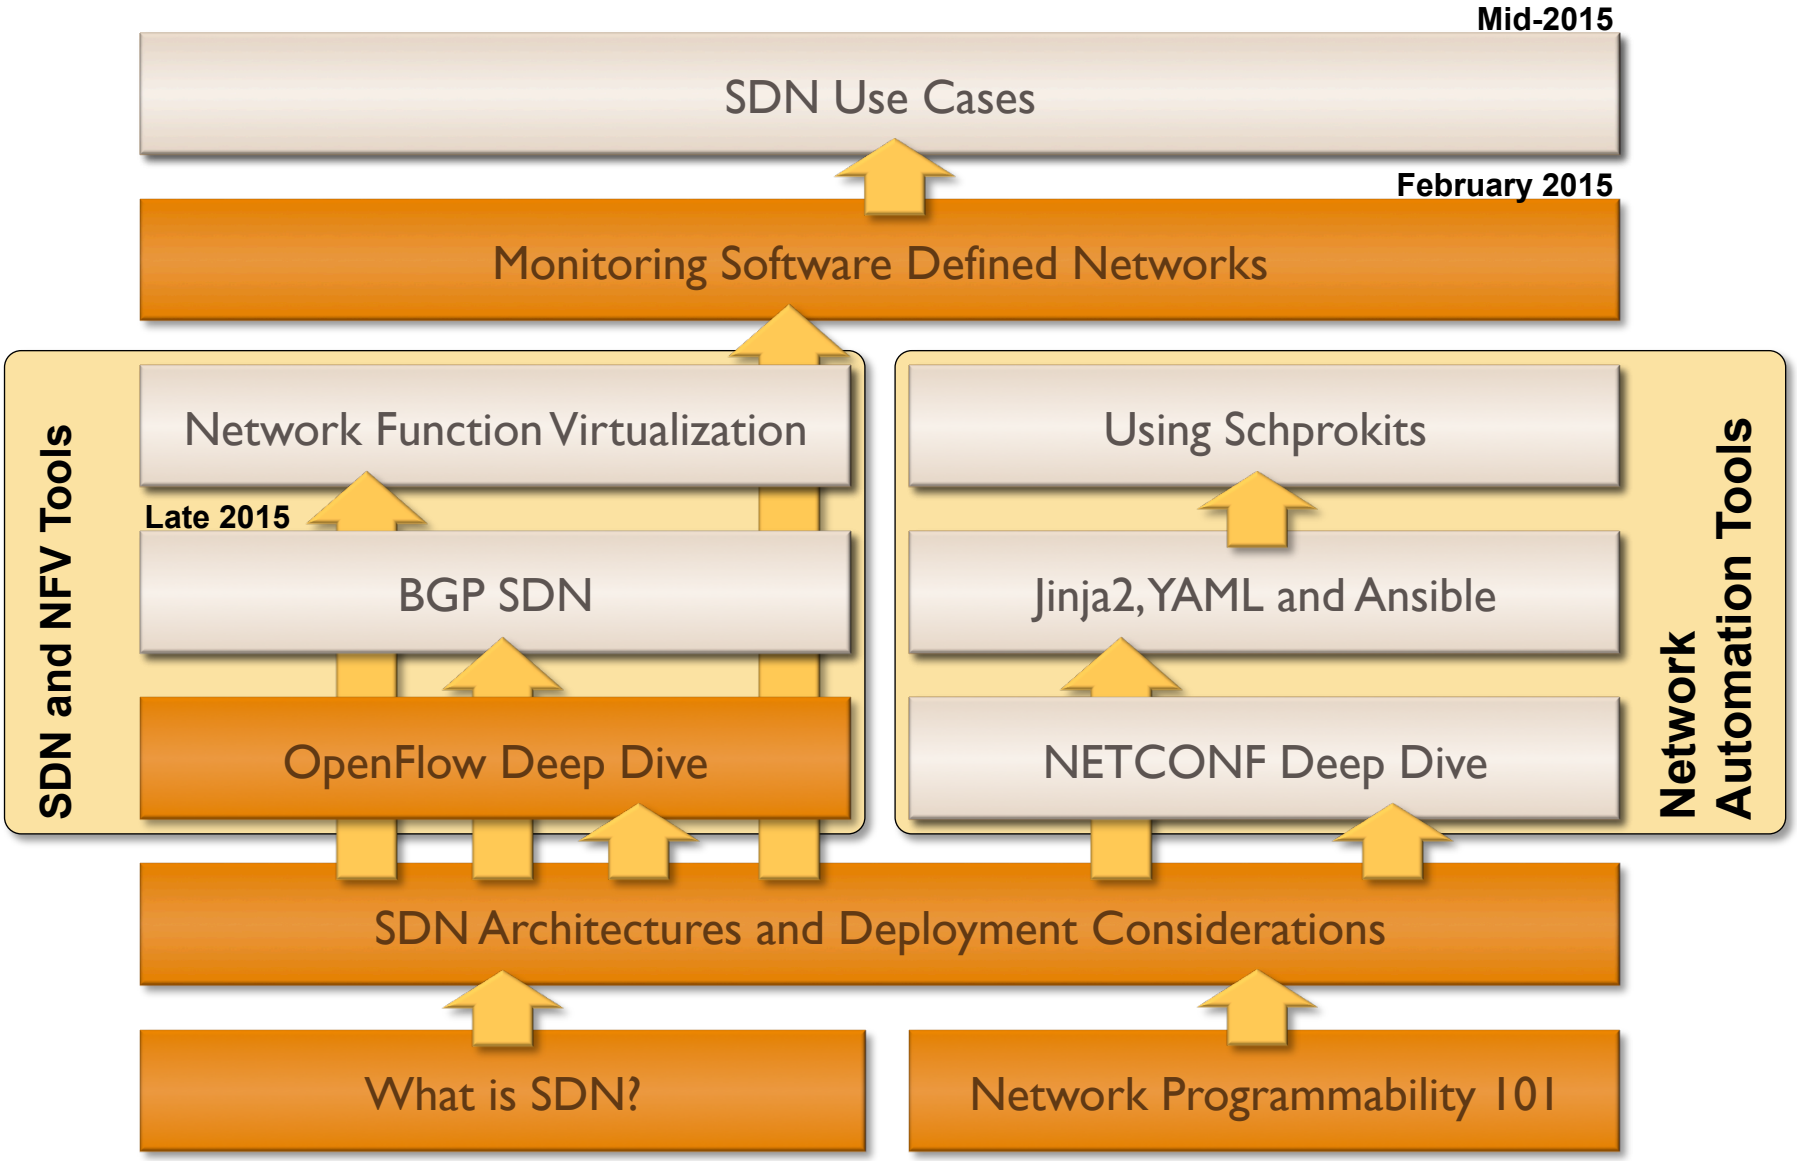

Stay in Touch

Web:	ipSpace.net
Blog:	blog.ipSpace.net
Email:	ip@ipSpace.net
Twitter:	@ioshints

SDN:	ipSpace.net/SDN
Webinars:	ipSpace.net/Webinars
Consulting:	ipSpace.net/Consulting