All You Need Are Two Switches

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

Data Center Network (cca 2010)

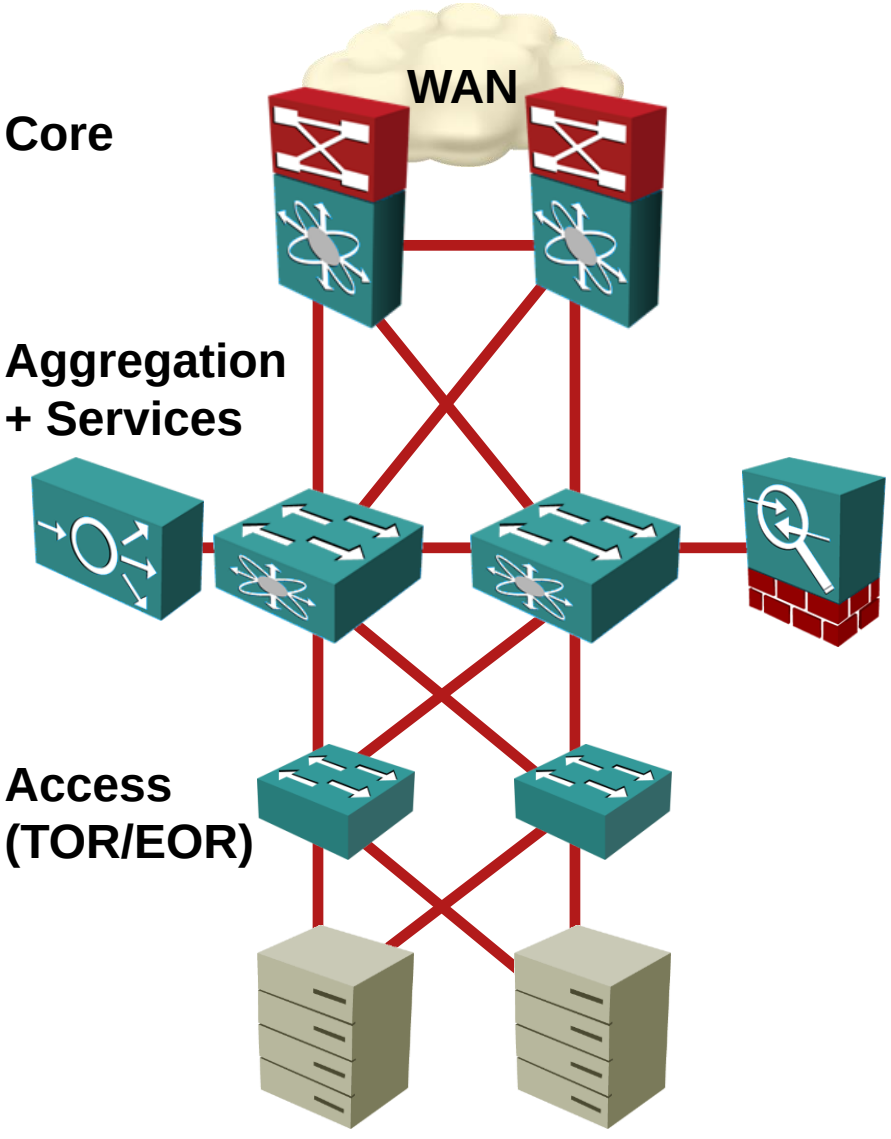

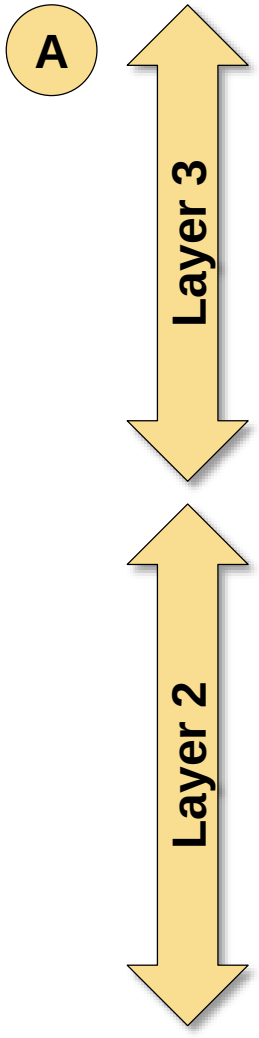

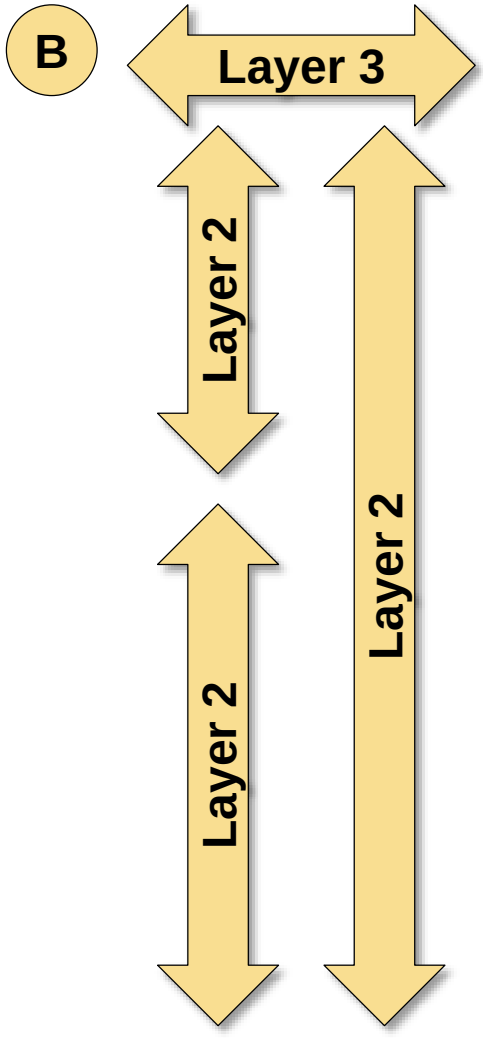

Virtualize the Servers

Rightsizing the Servers

Modern servers support large number of CPU cores

- Find a server model with optimum price-per-core
- Compute virtualization ratio (number of VMs per server) – 50+ VMs per server is reasonable
- Use RAM-per-core ratio (based on VM utilization data) to compute server RAM requirements
- Maximum RAM or # of cores might be the limiting factor
  ➔ repeat as required

Ditch Legacy Technologies

1GE or 10GE uplinks? What I Recommended in 2012

Gigabit Ethernet	10 Gigabit Ethernet
+ Well known, field tested + Copper cabling + NIC on motherboard	+ Much faster vMotion + Converged storage+network (FCoE or lossless iSCSI/NFS) + Reduce the number of NICs per server + Built-in QoS with ETS and PFC + Uses fiber cabling (lower energy consumption and error rate)
– Numerous NICs per hypervisor host (user data, vMotion, storage) – No storage/networking convergence – No lossless transport	– More expensive – Usually requires fiber cabling – NICs/CNAs still sold @ price premium

1GE or 10GE uplinks? What I Recommend in 2014

10 GE uplinks (don't even think about 1GE anymore)

10GBASE-T in small deployments

- ToR switches (don't forget inter-switch links)
- Servers with 10GBASE-T LOM or Intel X540-T2 adapter

SFP+/QSFP deployments

- Twinax SFP+ for short distances (when you need SFP+ connectivity)
- Fiber SFP+ for longer (inter-rack or inter-row) distances

Images: Arista 7050T-64, Cisco 5548P, Cisco Twinax SFP+

Minimize Uplinks

Read the Latest vSphere Design Guide

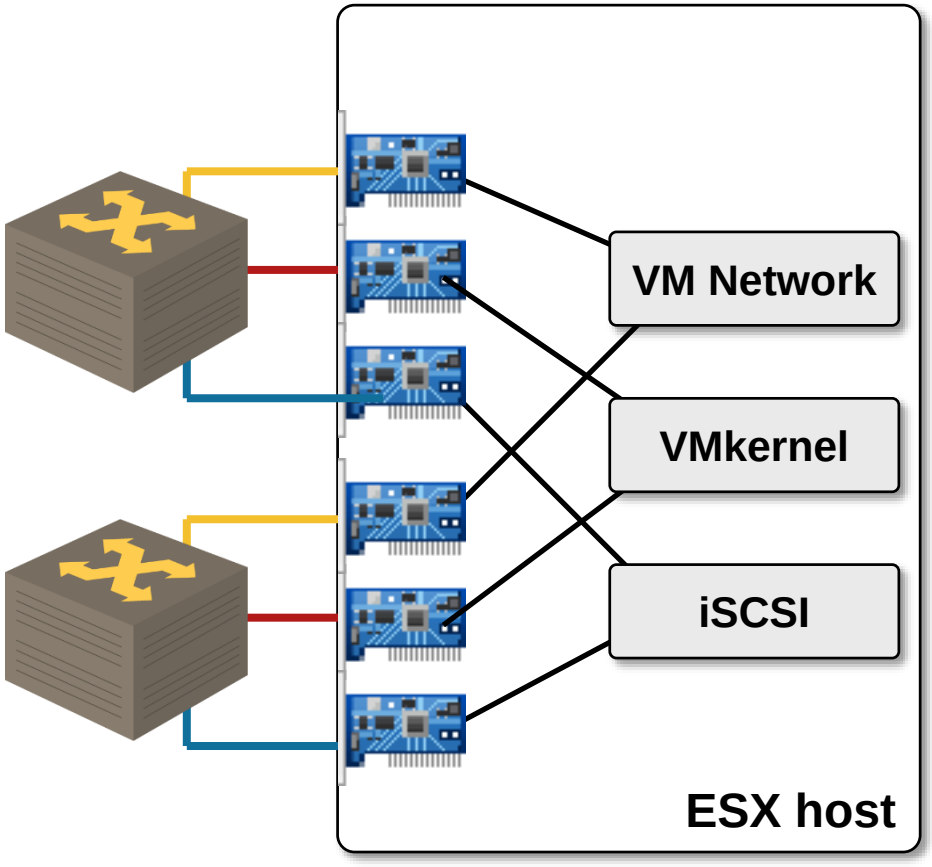

vSphere 4

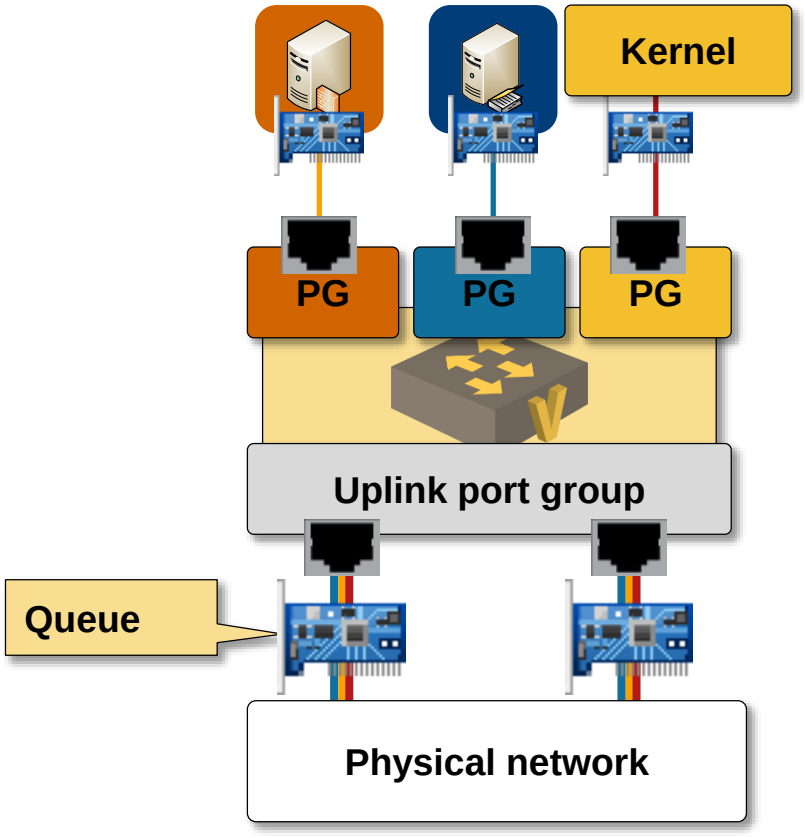

vSphere 5+ (since 2011)

Use IP-based Storage

Replace FC/FCoE with iSCSI or NFS

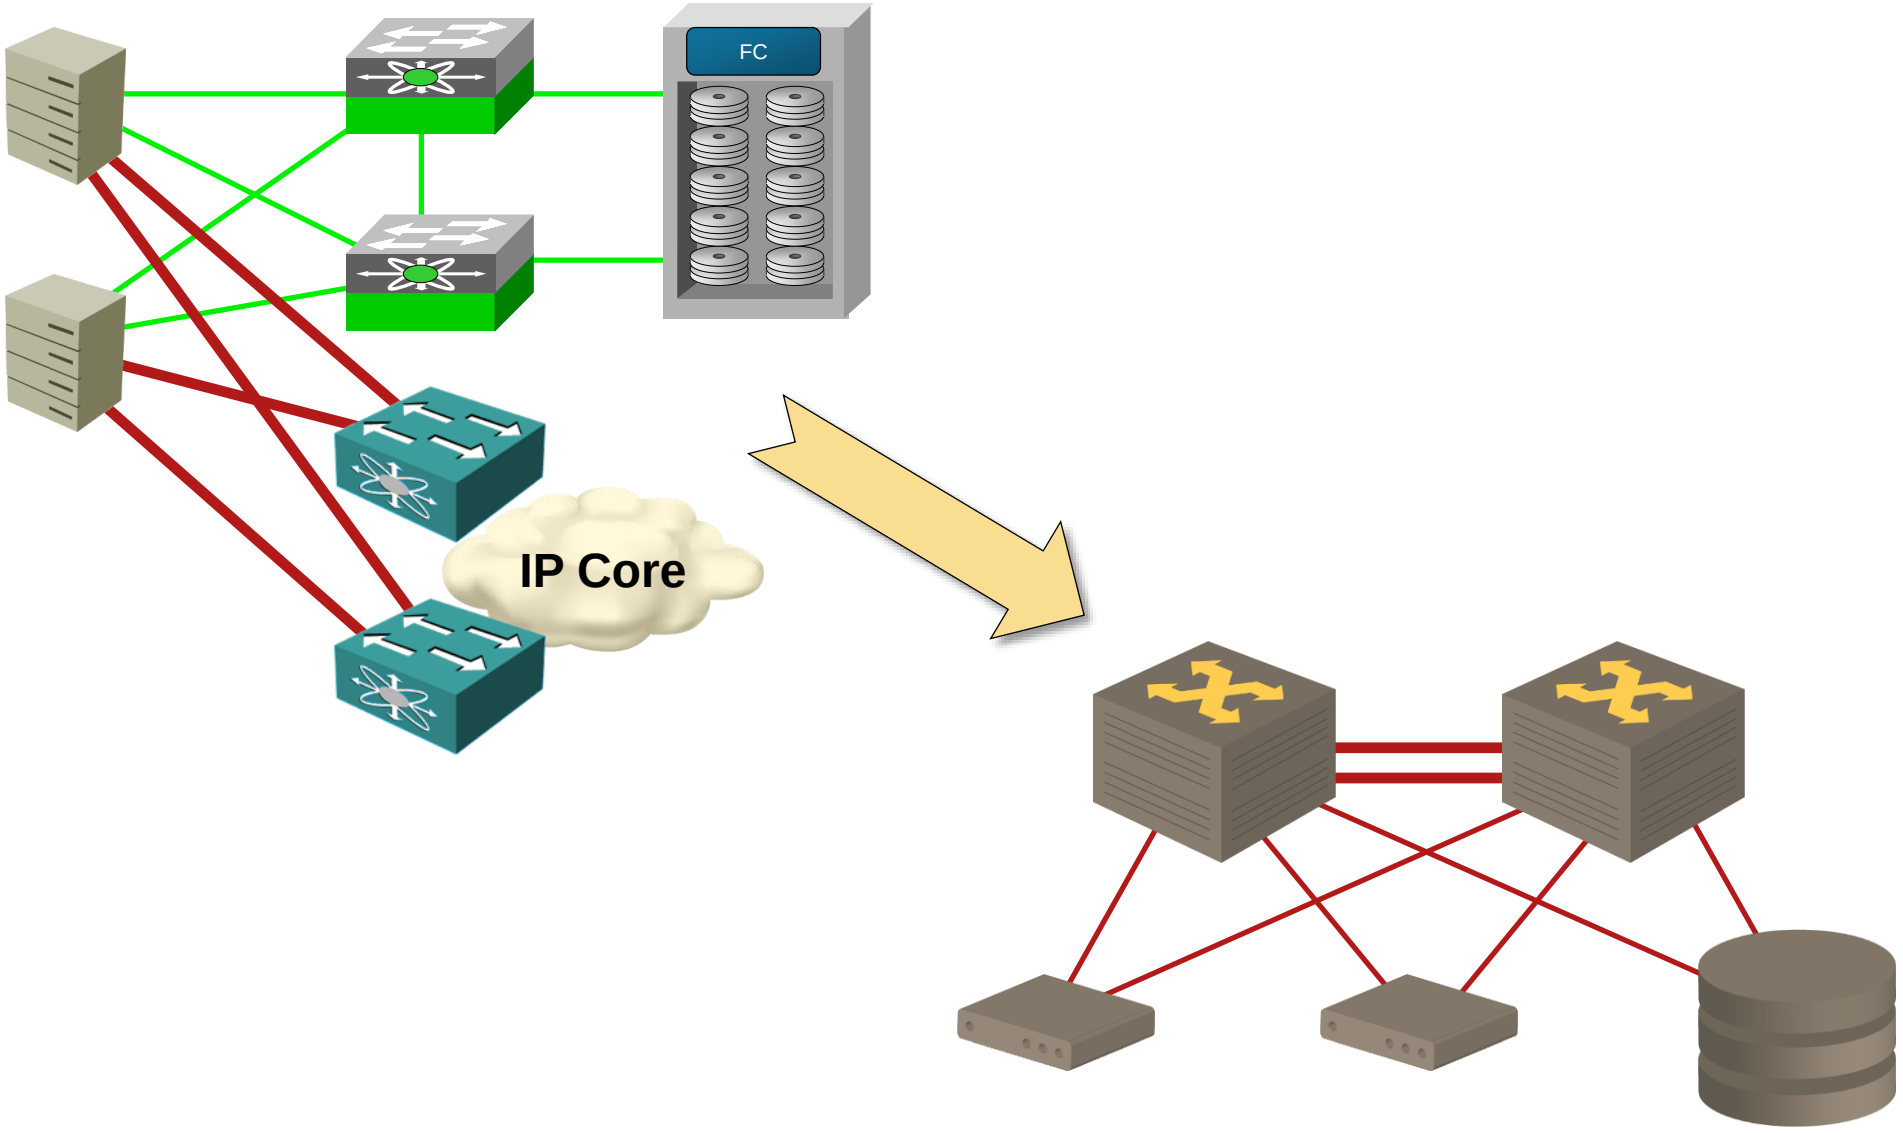

Distributed File Systems

Distributed File Systems with DAS

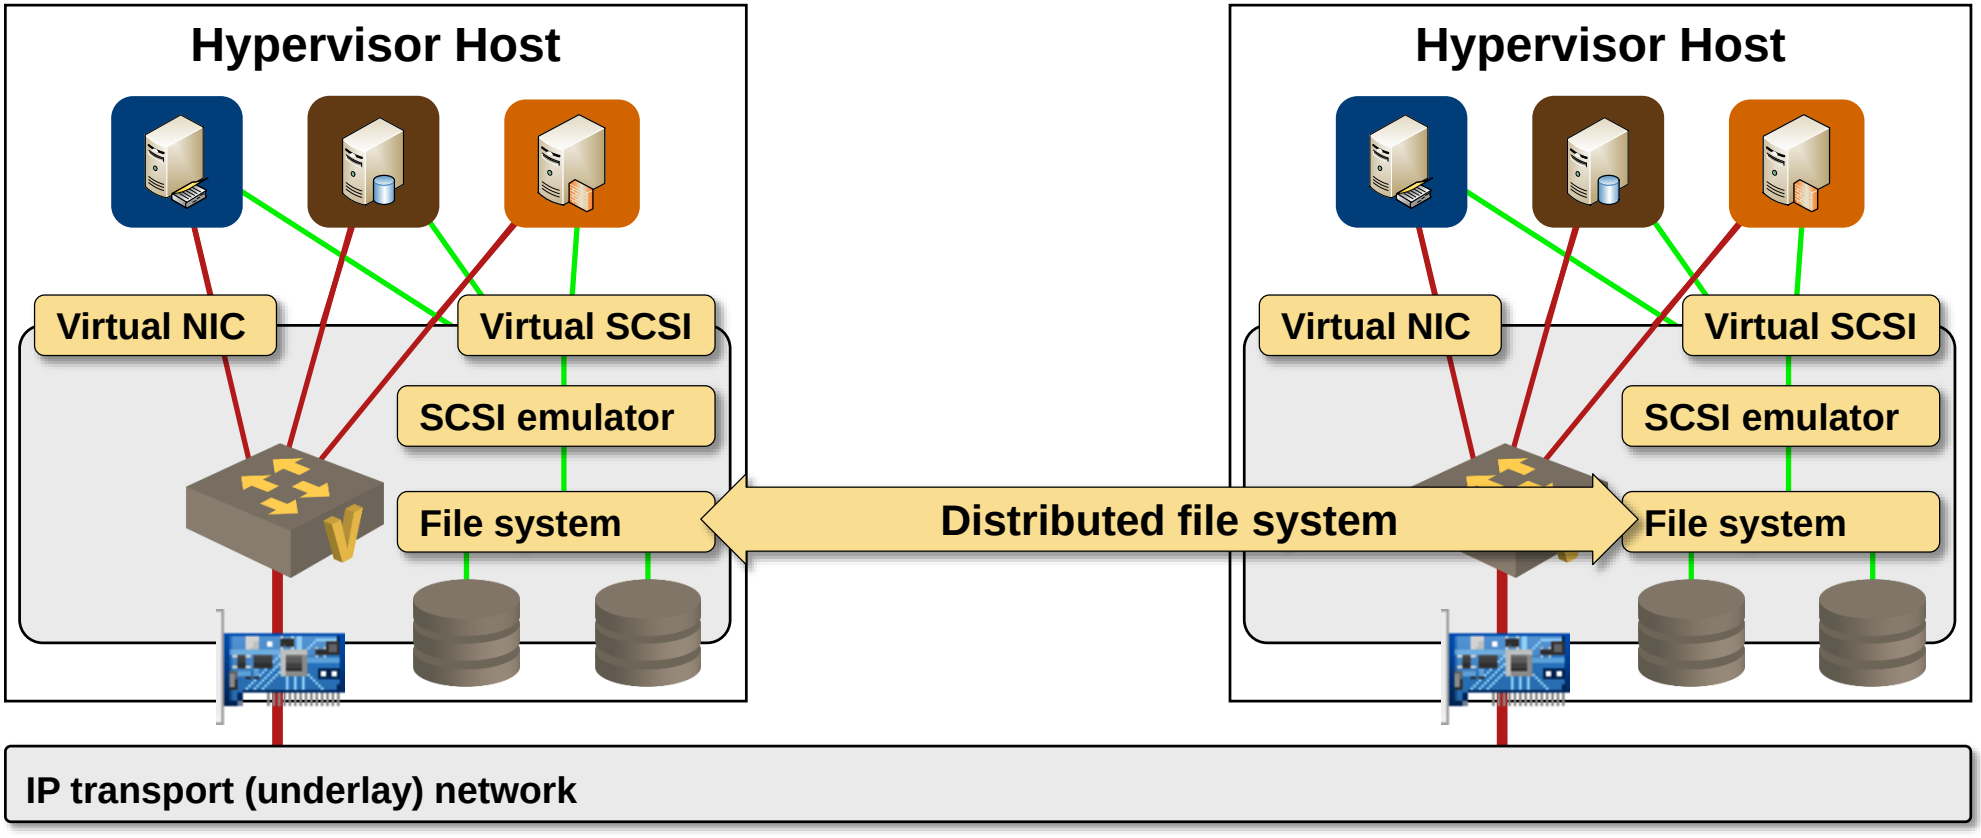

- Each hypervisor uses local disks (DAS) and/or SSDs
- Global file system or object store with file replication between hypervisors nodes
- Examples: Ceph, GlusterFS, VMware VSAN, OpenStack Swift

Virtualize Network Services

Virtual Appliances = Network Services in VM Format

Advantages

- Avoid hardware-dependent delays
- Increase flexibility and mobility
- Simplify disaster recovery
- Minimal time-to-deploy
- Minimize sparing

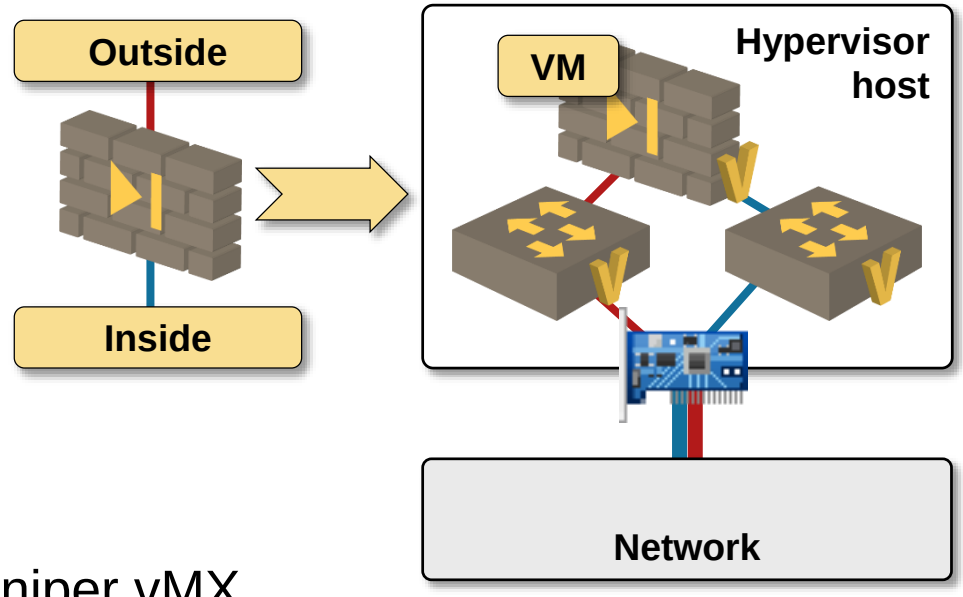

Sample products

- Routers: Brocade Vyatta, Cisco CSR, Juniper vMX
- Firewall: pfSense, Juniper vSRX, Palo Alto, Vyatta, vShield Edge (VMware), vASA (Cisco)
- Load balancer: BIG-IP VTM (F5), Zeus Traffic Manager (now Riverbed), vShield Edge (VMware), Embrane, LineRate Systems (now F5), Citrix NetScaler

Virtualize the Appliances

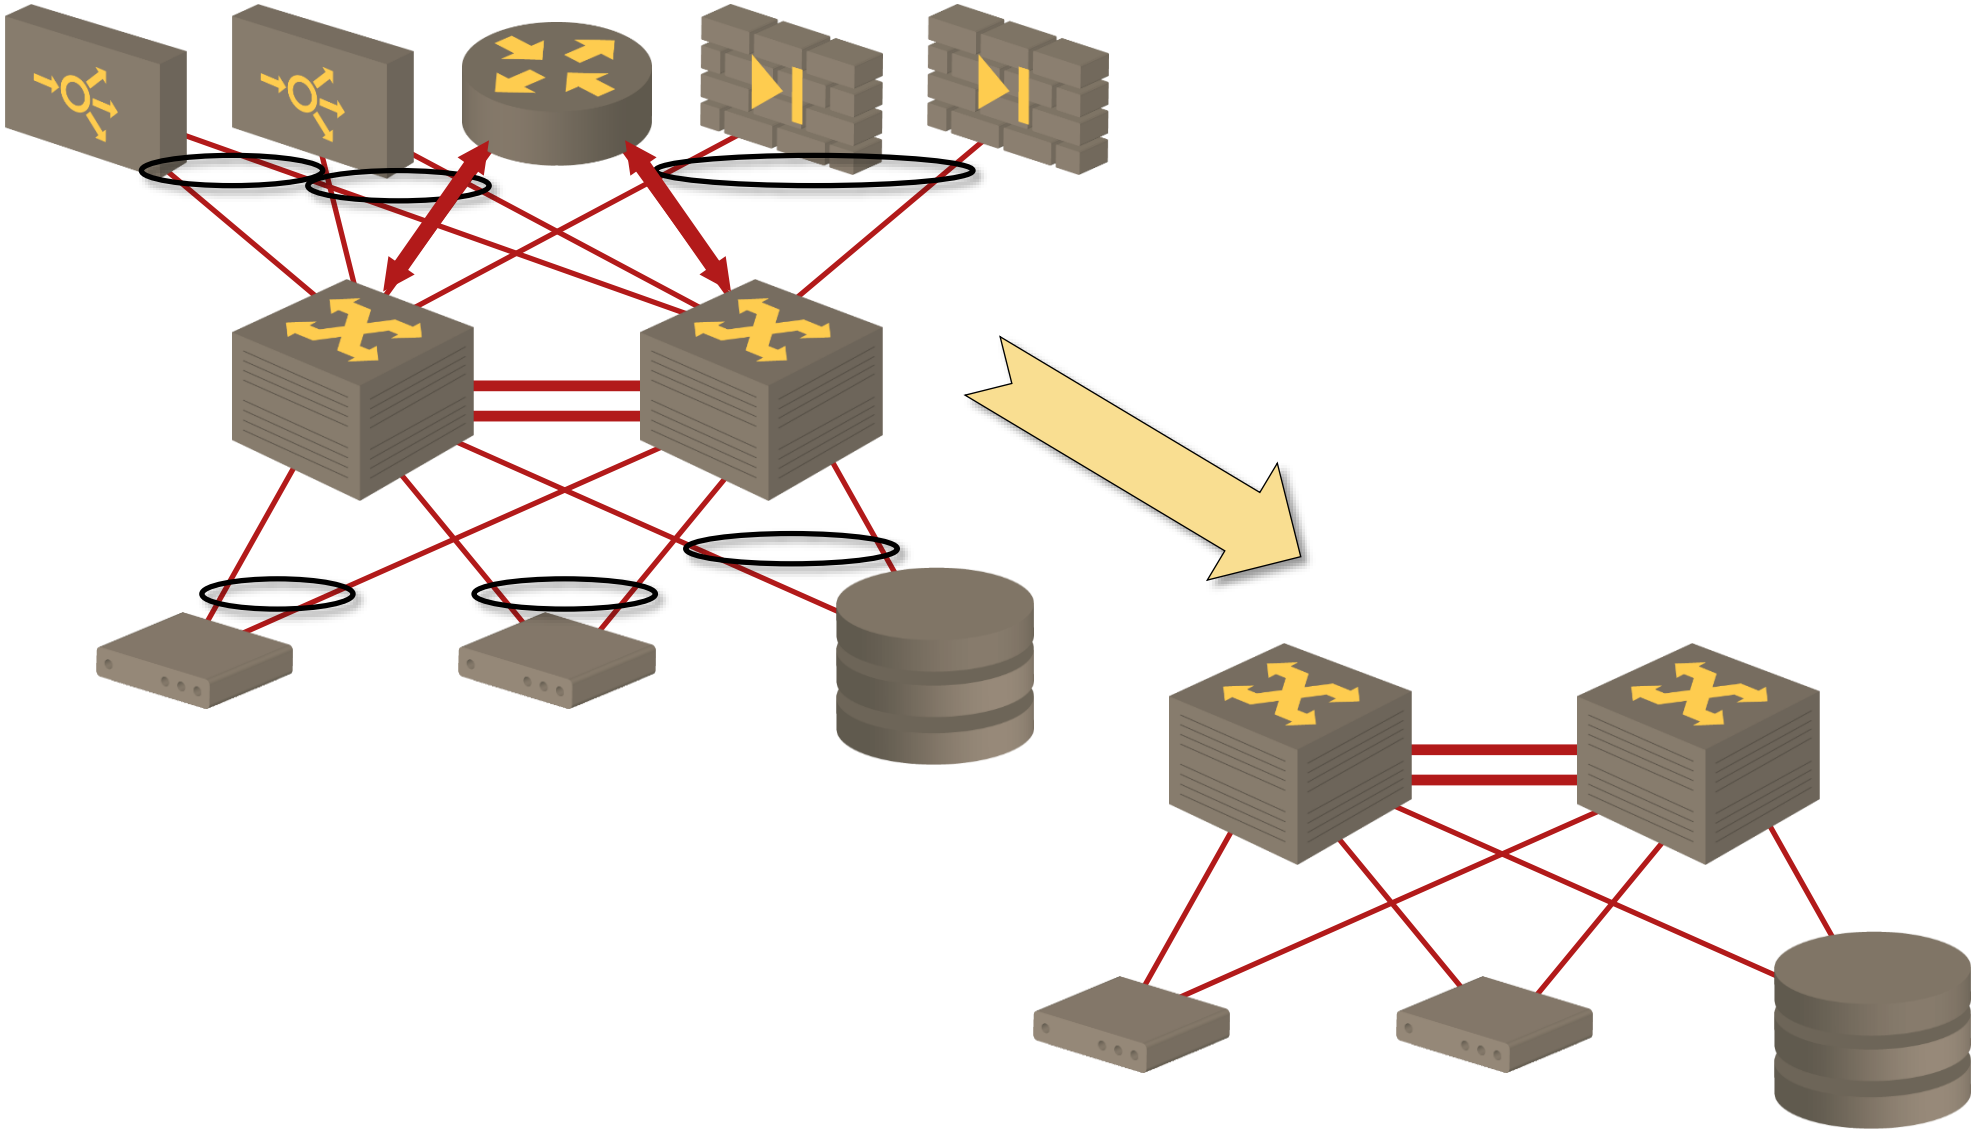

Final Result

Going All Virtual

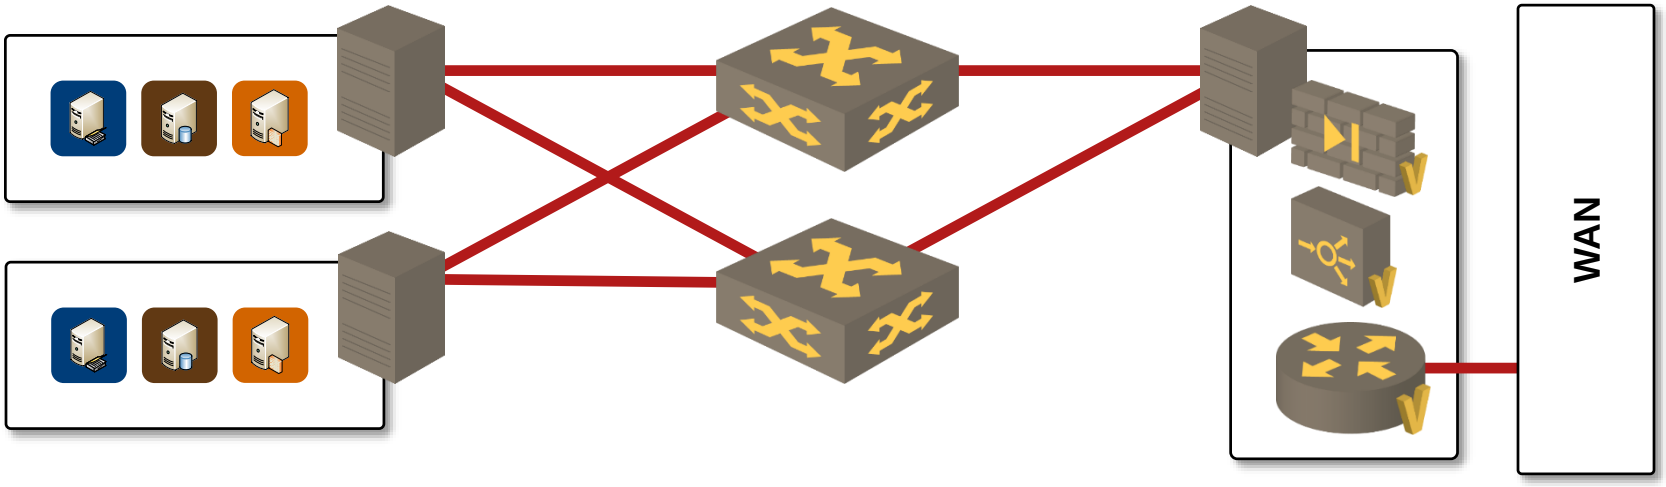

- All servers are virtualized ➔ small number of hypervisor hosts
- Distributed file system ➔ no extra storage components
- Virtualized appliances ➔ only two hardware components

Will It Work?

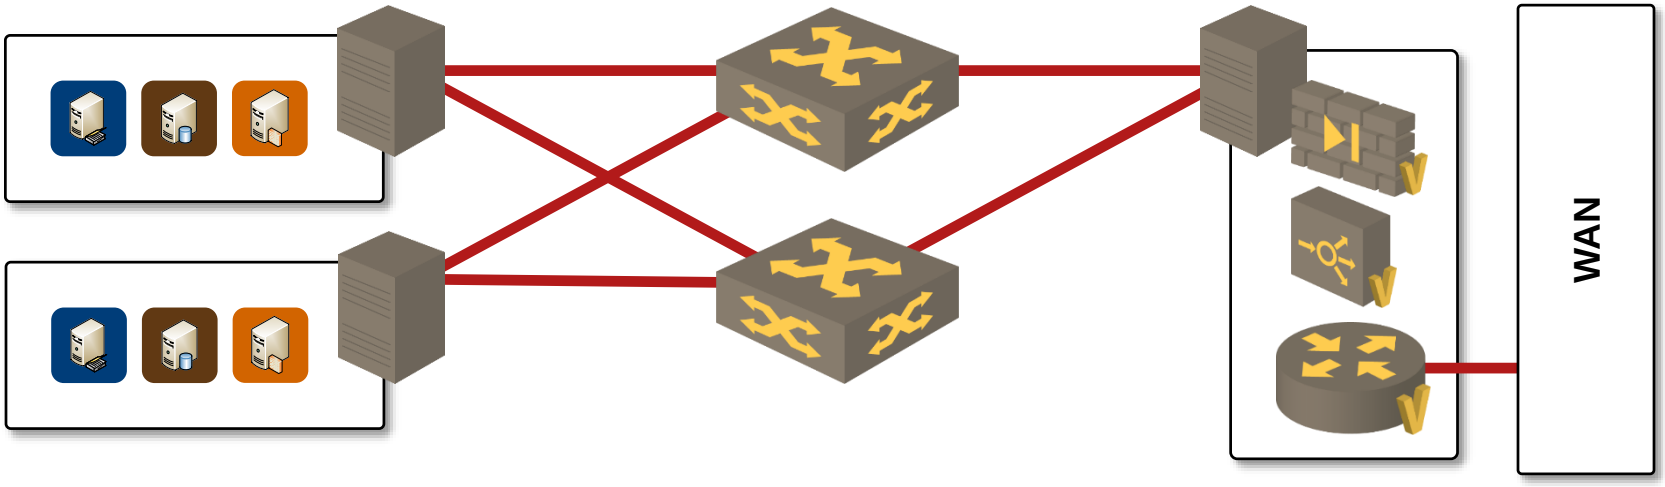

- 1500 VMs in 28 RU (56 blade hosts)
  http://virtual-red-dot.info/1000-vm-per-rack-is-the-new-minimum/
- Cisco Nexus 93120 has 96 10GE ports

Conclusion: Two switches is all you need to build a data center

Cloud Computing Webinars on ipSpace.net

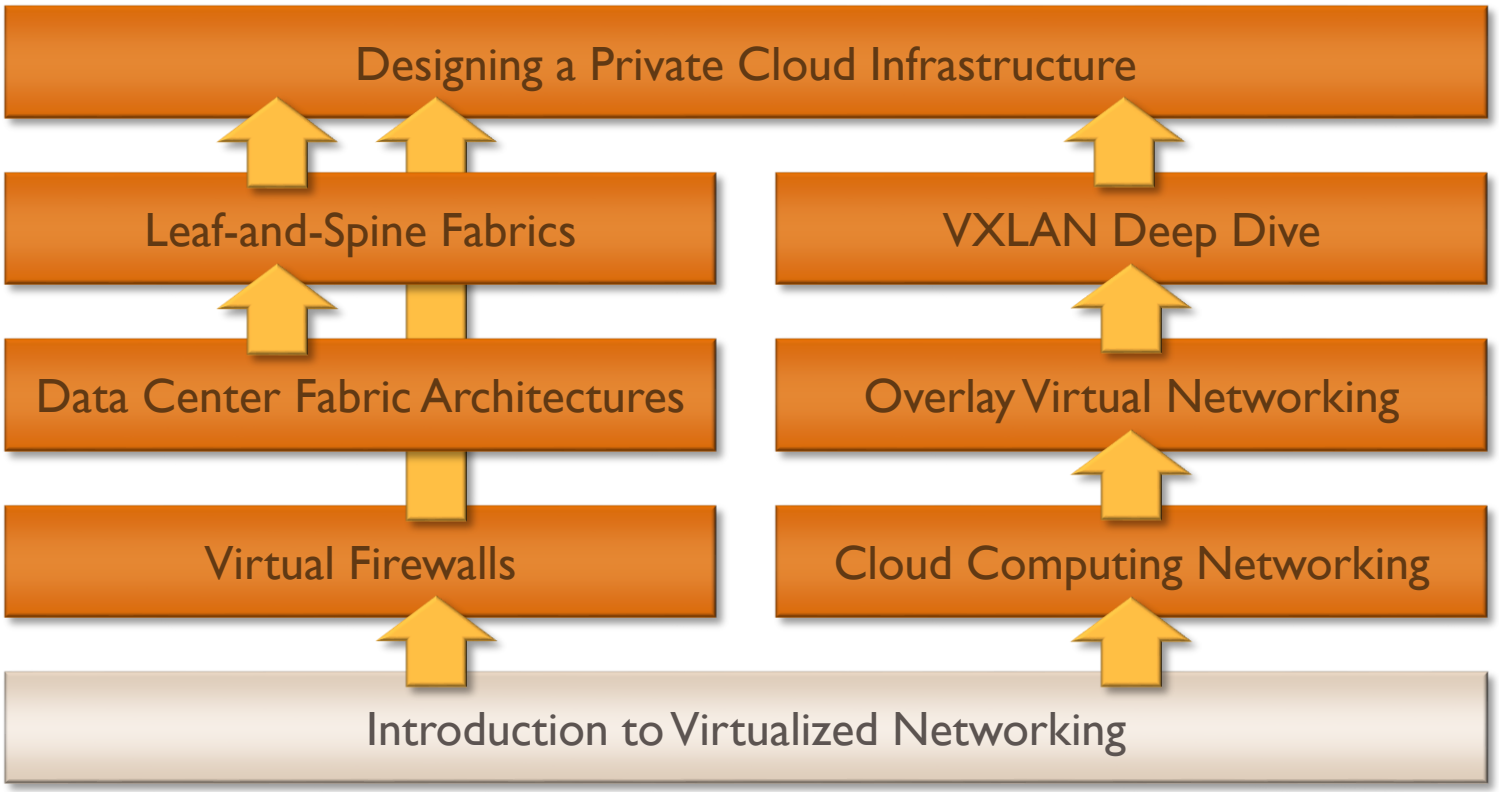

Availability

- Live sessions
- Recordings of individual webinars
- Yearly subscription

Other options

- Customized webinars
- ExpertExpress
- On-site workshops

More information @ http://www.ipSpace.net/Cloud

Virtualization Webinars on ipSpace.net

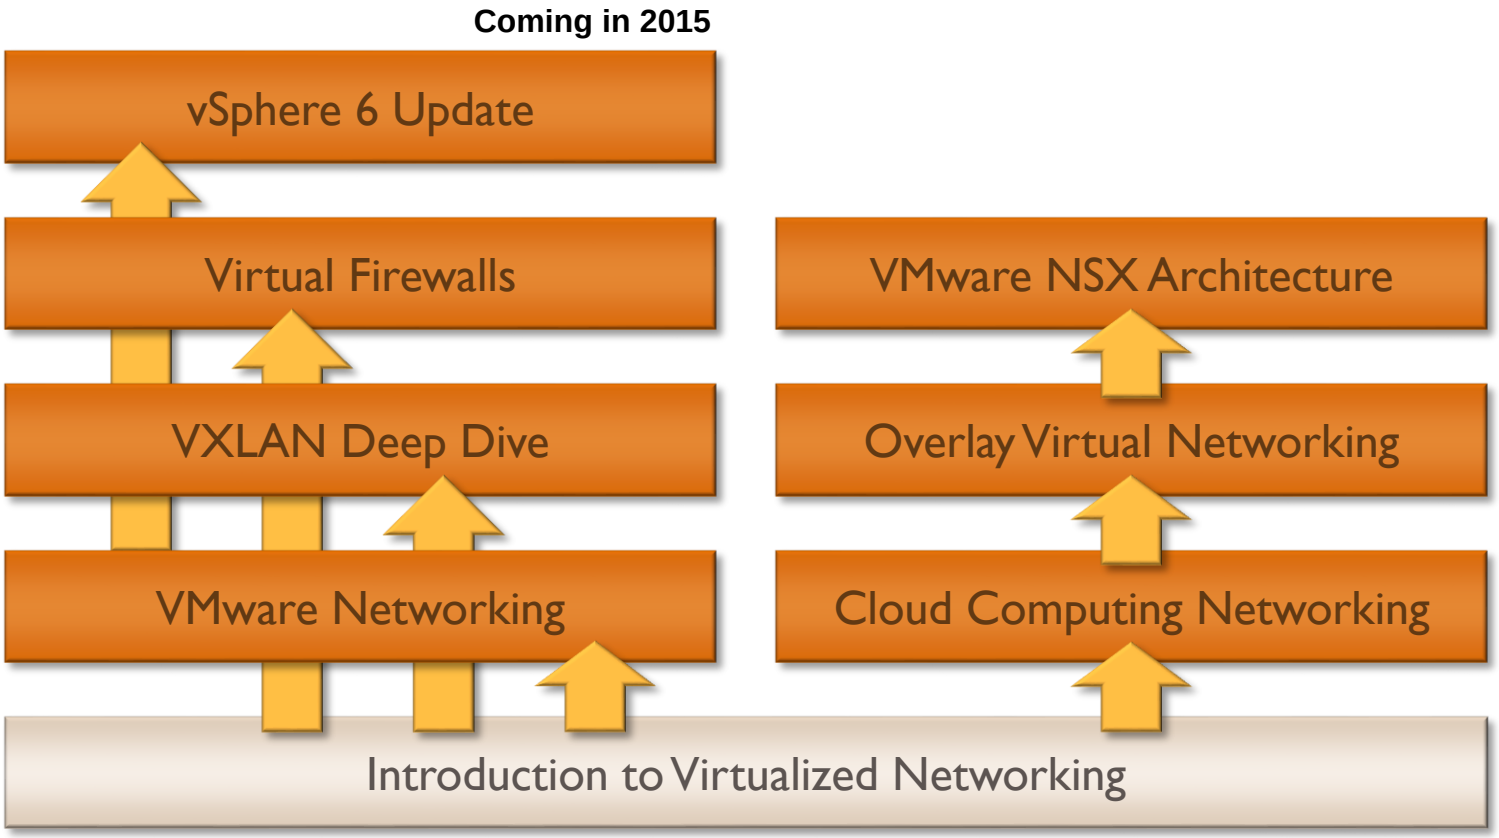

Availability

- Live sessions
- Recordings of individual webinars
- Yearly subscription

Other options

- Customized webinars
- ExpertExpress
- On-site workshops

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

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

Questions?

Send them to ip@ipSpace.net or @ioshints