Virtual Routers

Do they make sense?

Ivan Pepelnjak (ip@ipSpace.net)

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

Typical Router Hardware Architecture

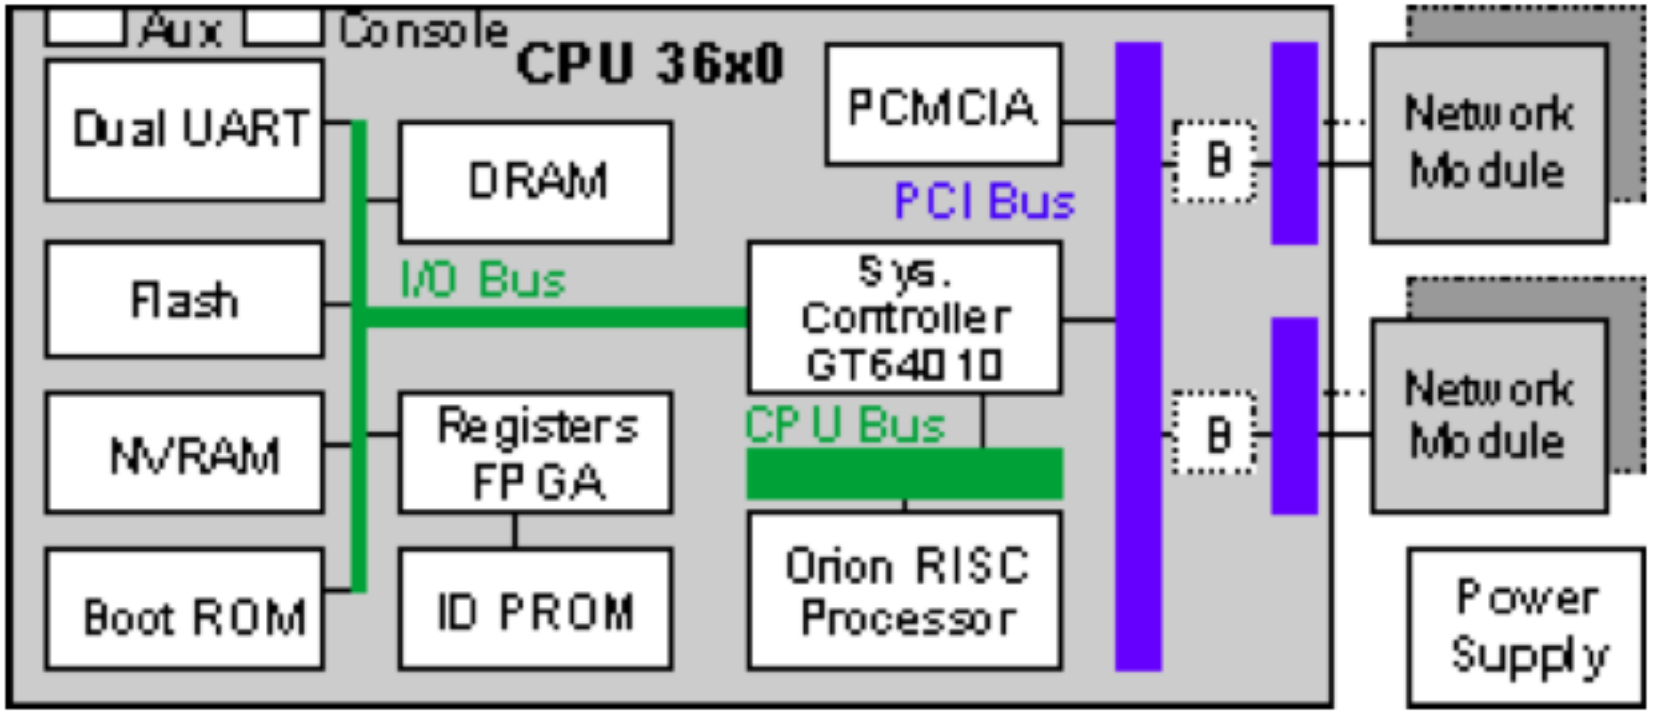

- Looks almost like an x86-based server to me
- Could I have this in a VM format, please?

Source: Cisco 3600 Series Router Architecture
http://www.cisco.com/c/en/us/support/docs/routers/3600-series-multiservice-platforms/7442-36xx-arch.html

Welcome to Network Function Virtualization

NFV 101

- Network services deployed in VM format
- Pilot in Deutsche Telekom (Terastream)
- Part of AT&T Domain 2.0 initiative
- Being standardized within an ETSI working group
- Another overhyped hot topic

Why?

- Flexibility and scalability
- Ease of deployment
- Faster time-to-market
- Hardware reuse

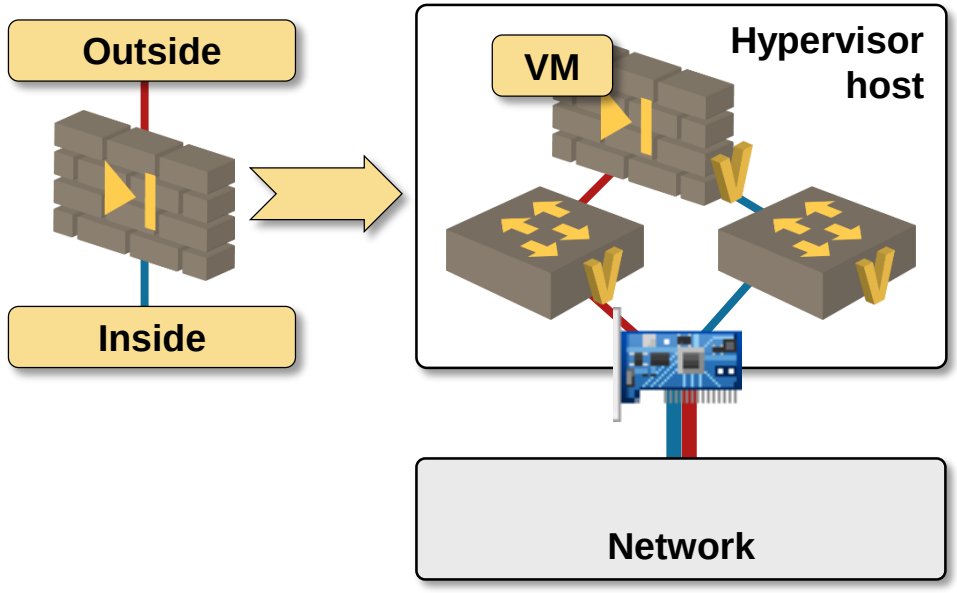

Why Are We Still Buying Boxes, Not Software?

- Efficiency
- Robustness
- Clear split of management responsibilities
- Inability to change existing sales models

Remember: when you buy a router, you pay (mostly) for the software, not hardware

Refresher: Management, Control and Data Plane

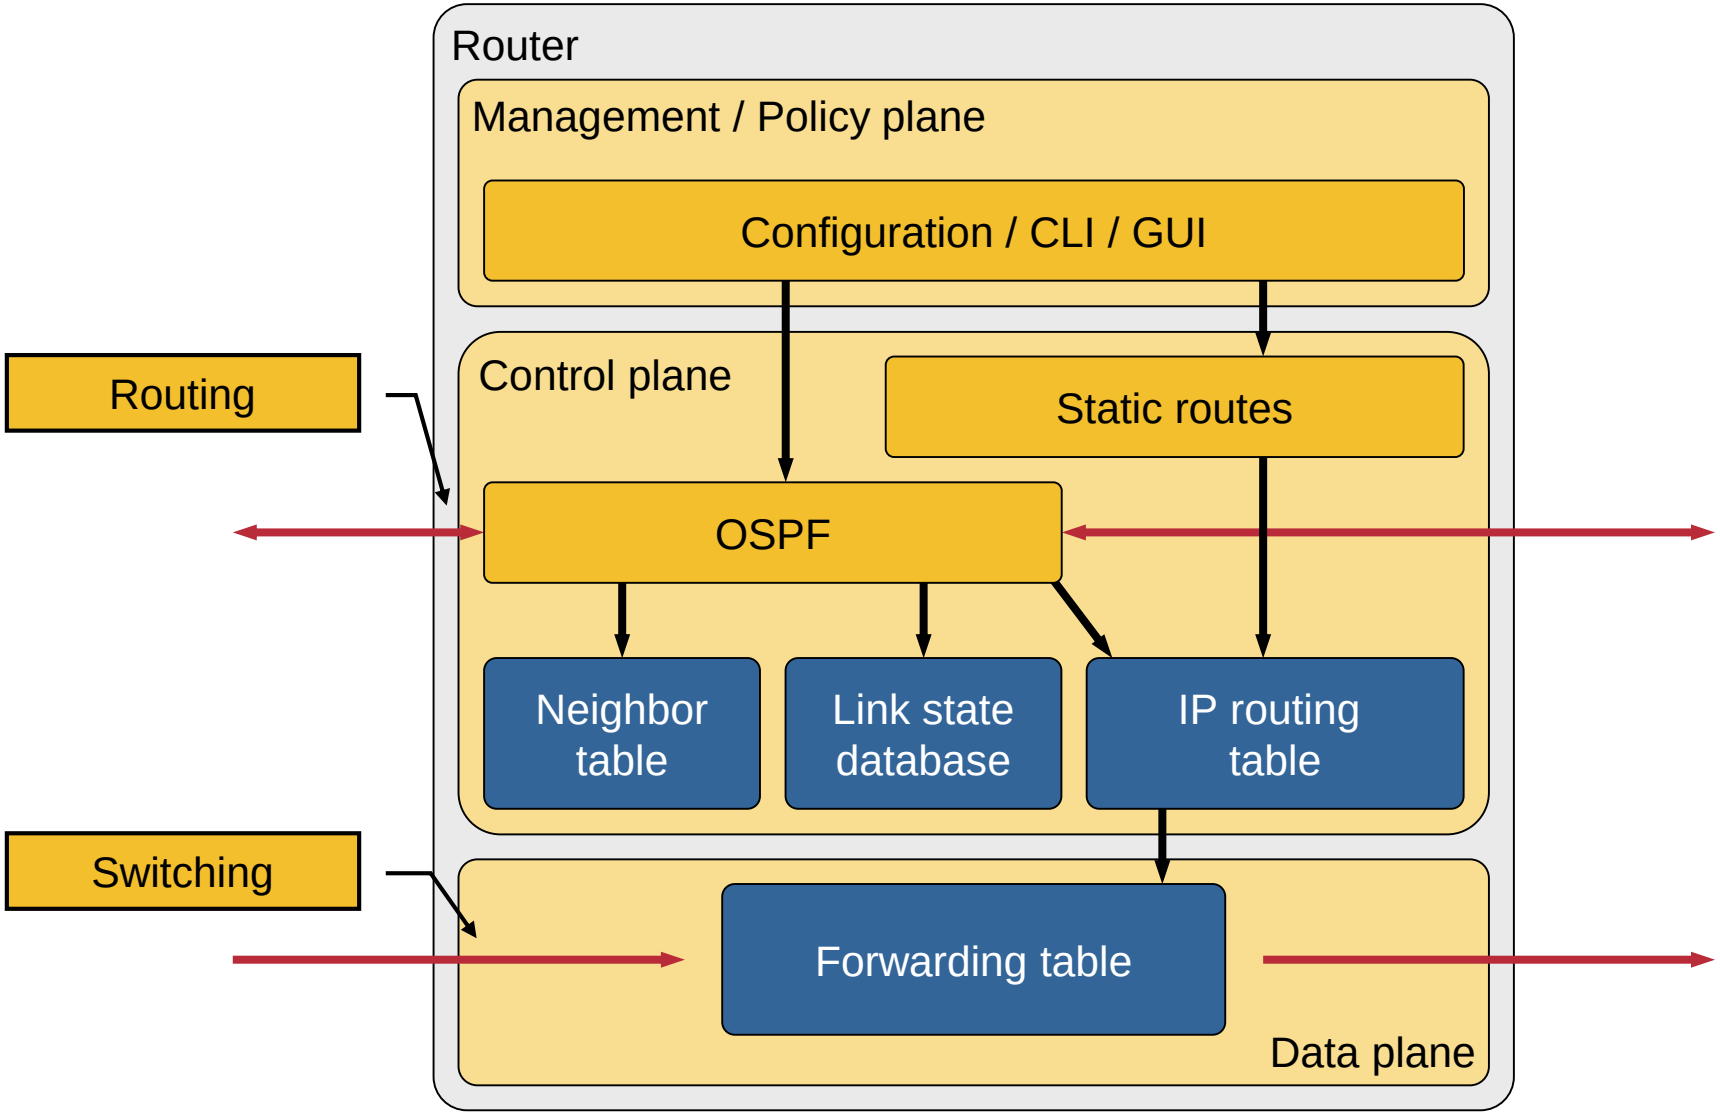

Performance Requirements

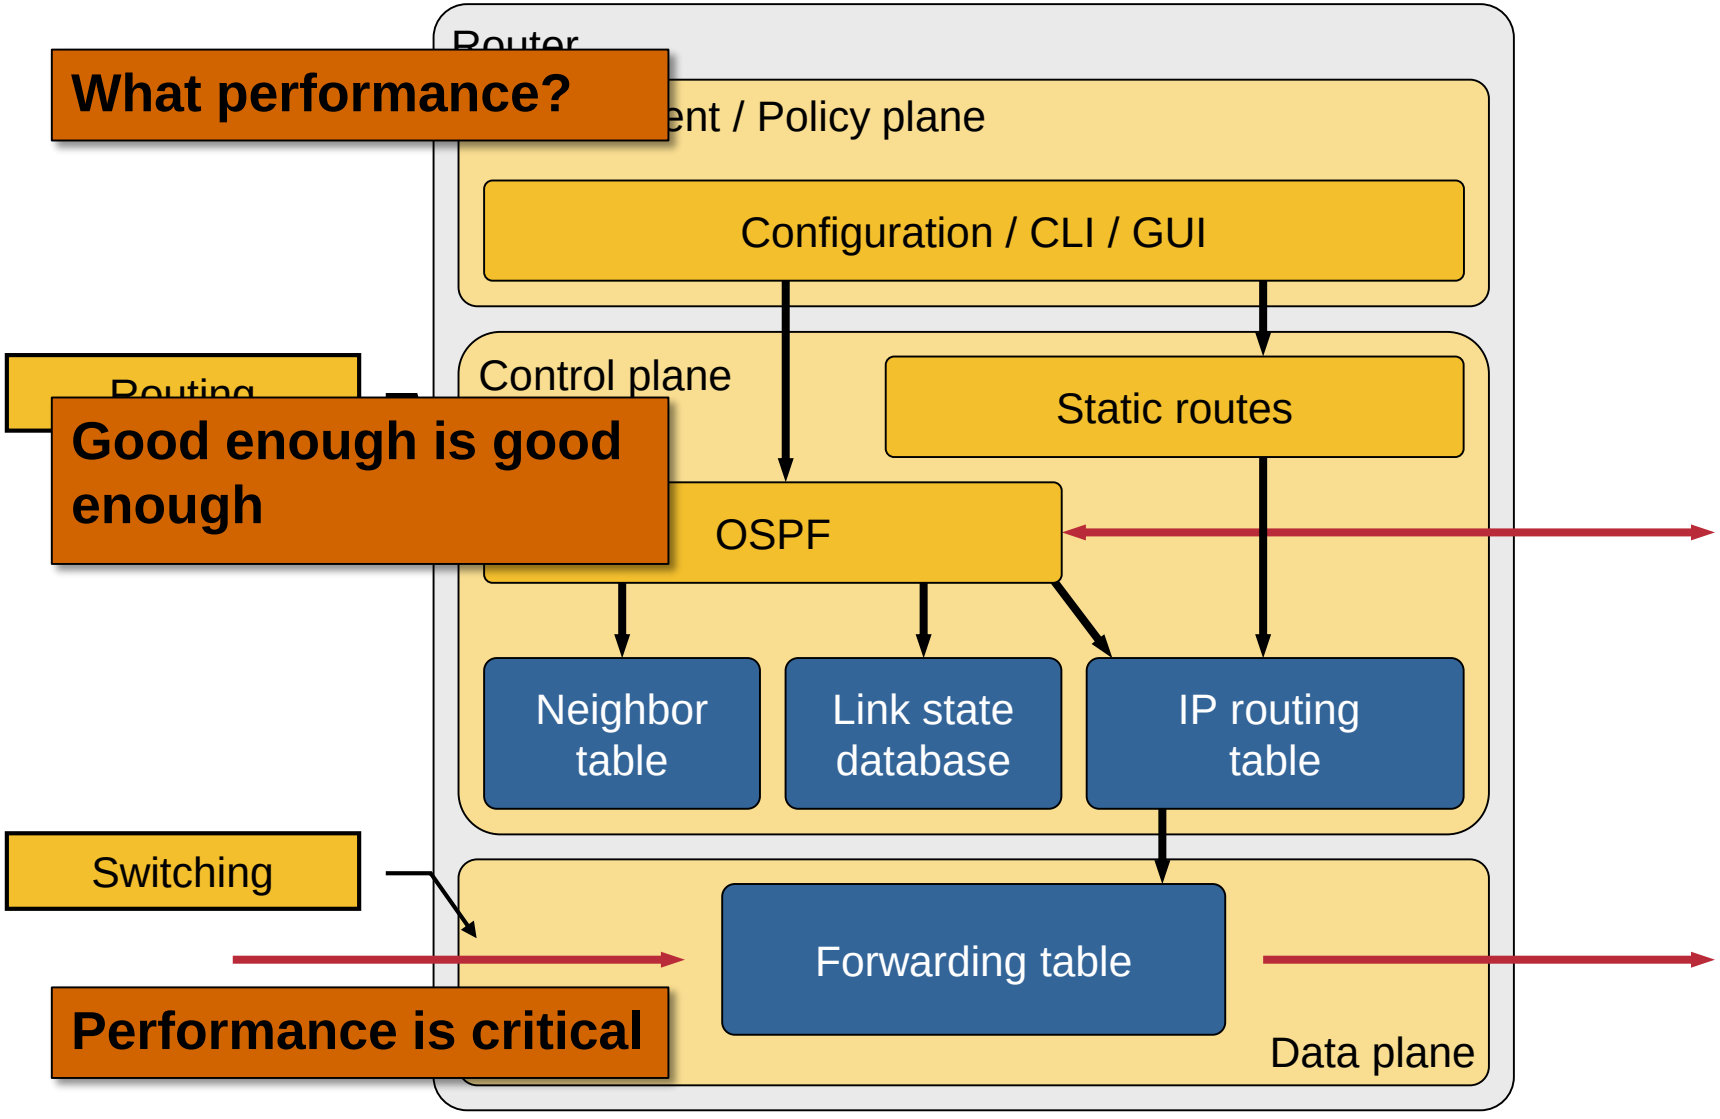

Performance Data Points

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 vShield Edge small instance (1 vCPU): 1-3 Gbps**
- Vmware NSX Edge Services Router: 10 Gbps firewall, 4 – 10 Gbps load balancer
- Palo Alto firewall: 1 Gbps @ 4 vCPU
- **Cisco CSR 1000V: 1 Gbps @ 4vCPU**
- **Vyatta 5600 series routers: 10 Gbps @ 1 core**

Hint: Don’t Use Linux TCP Stack

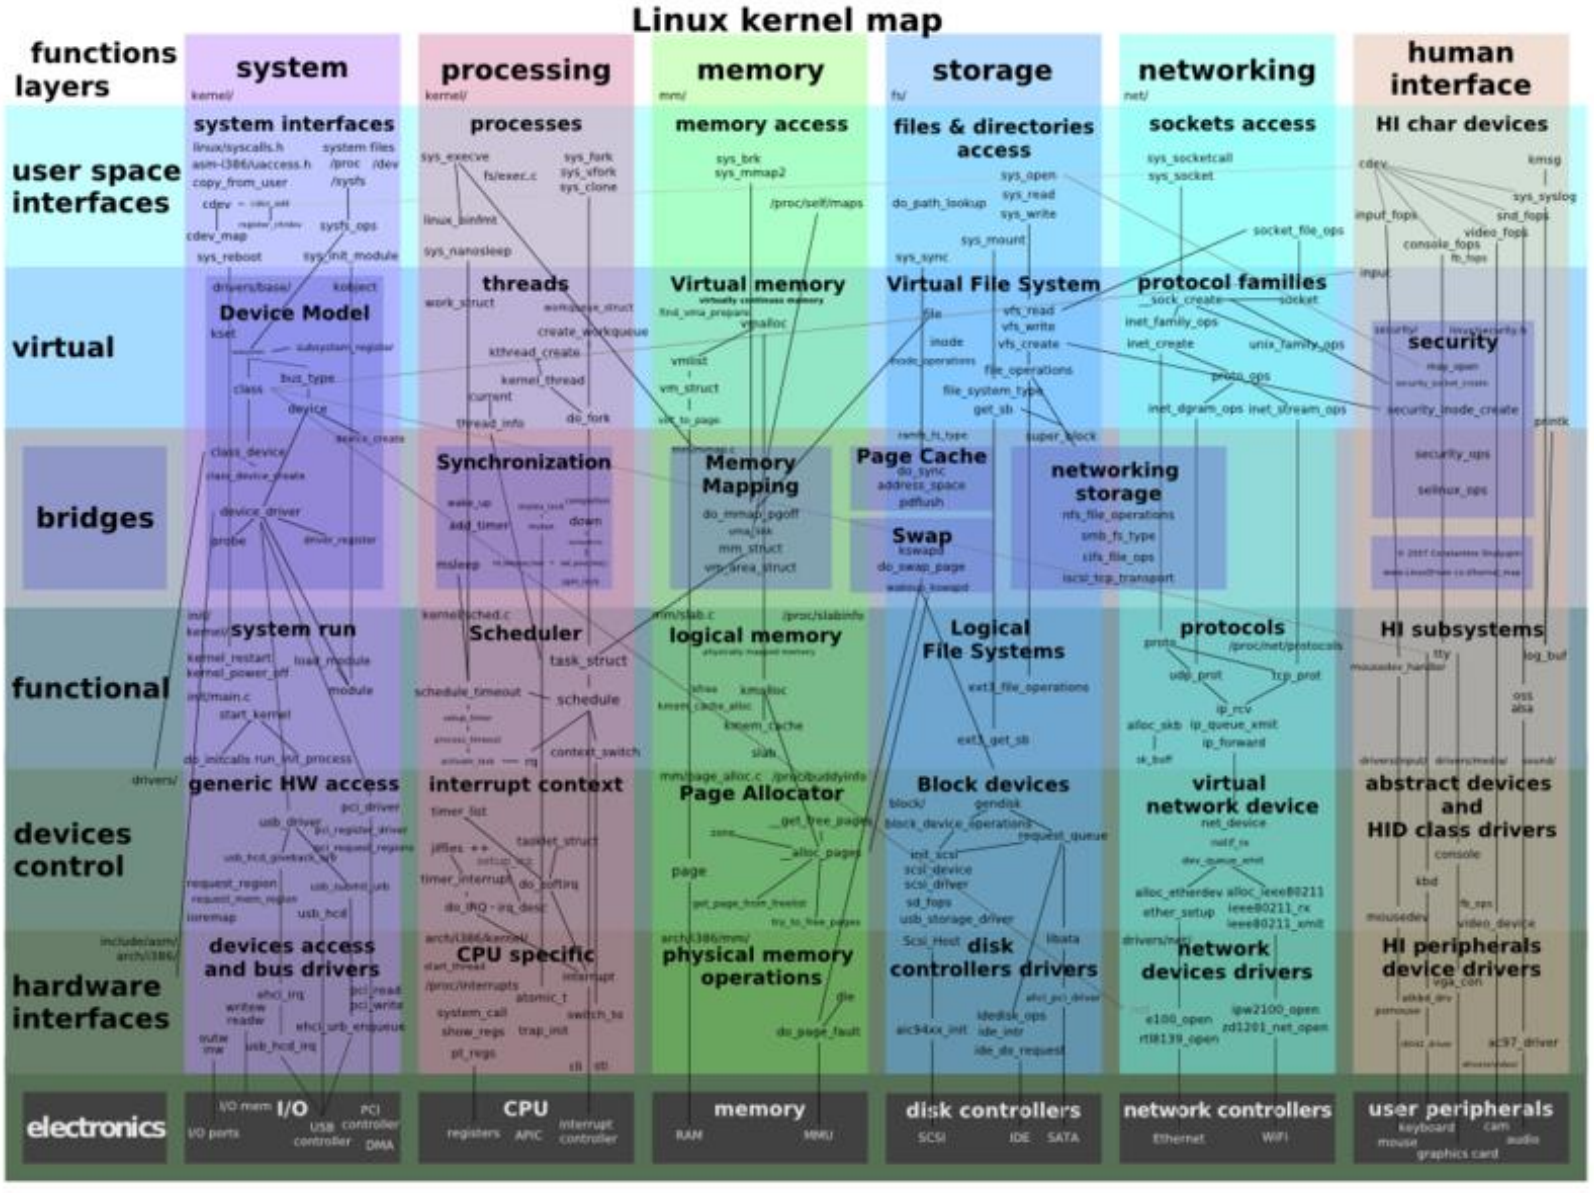

QoS Considerations

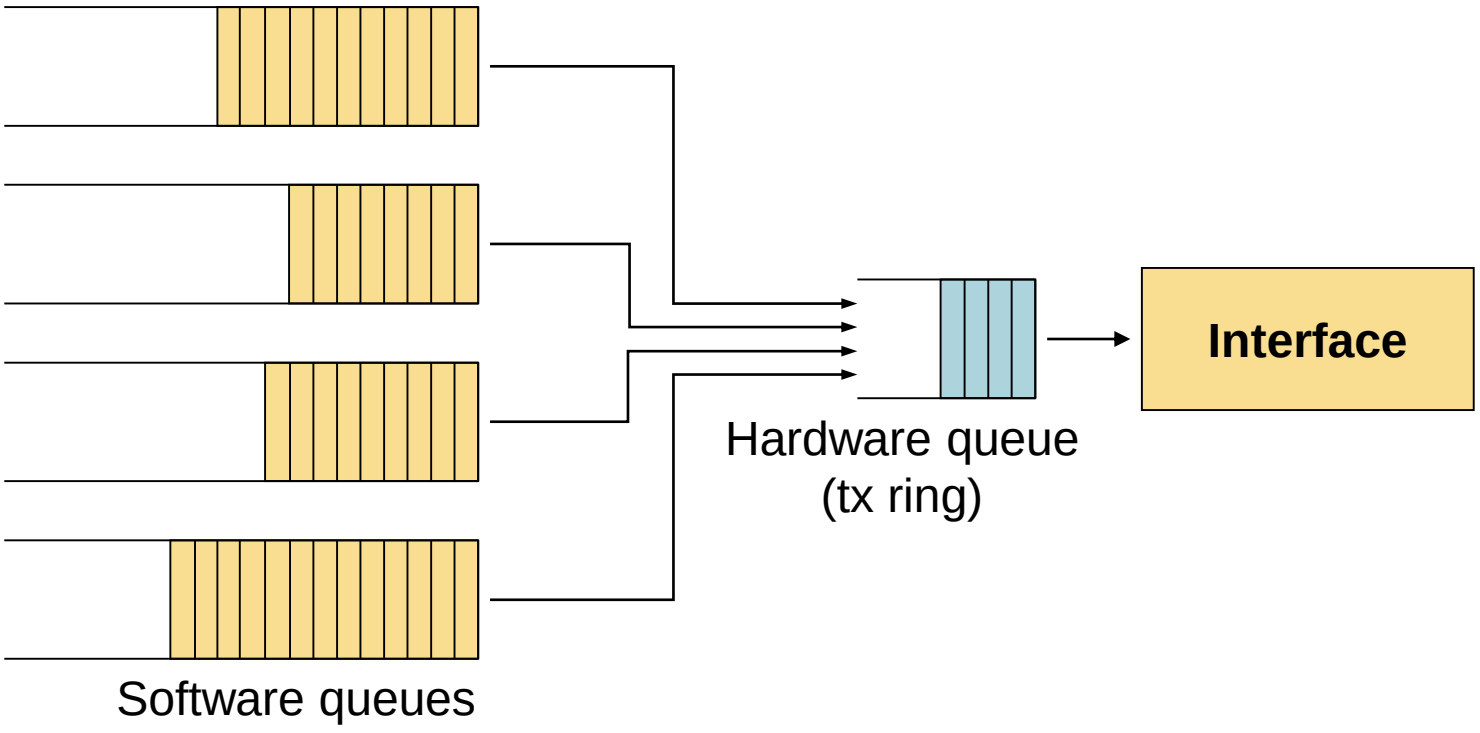

- Queuing needs hardware backpressure to work
- Virtual routers might not experience backpressure on VM NICs
- Solution: switch from queuing to shaping + queuing (similar to xDSL QoS)
- All other QoS mechanisms work as expected

Use Cases

Customer Service Termination

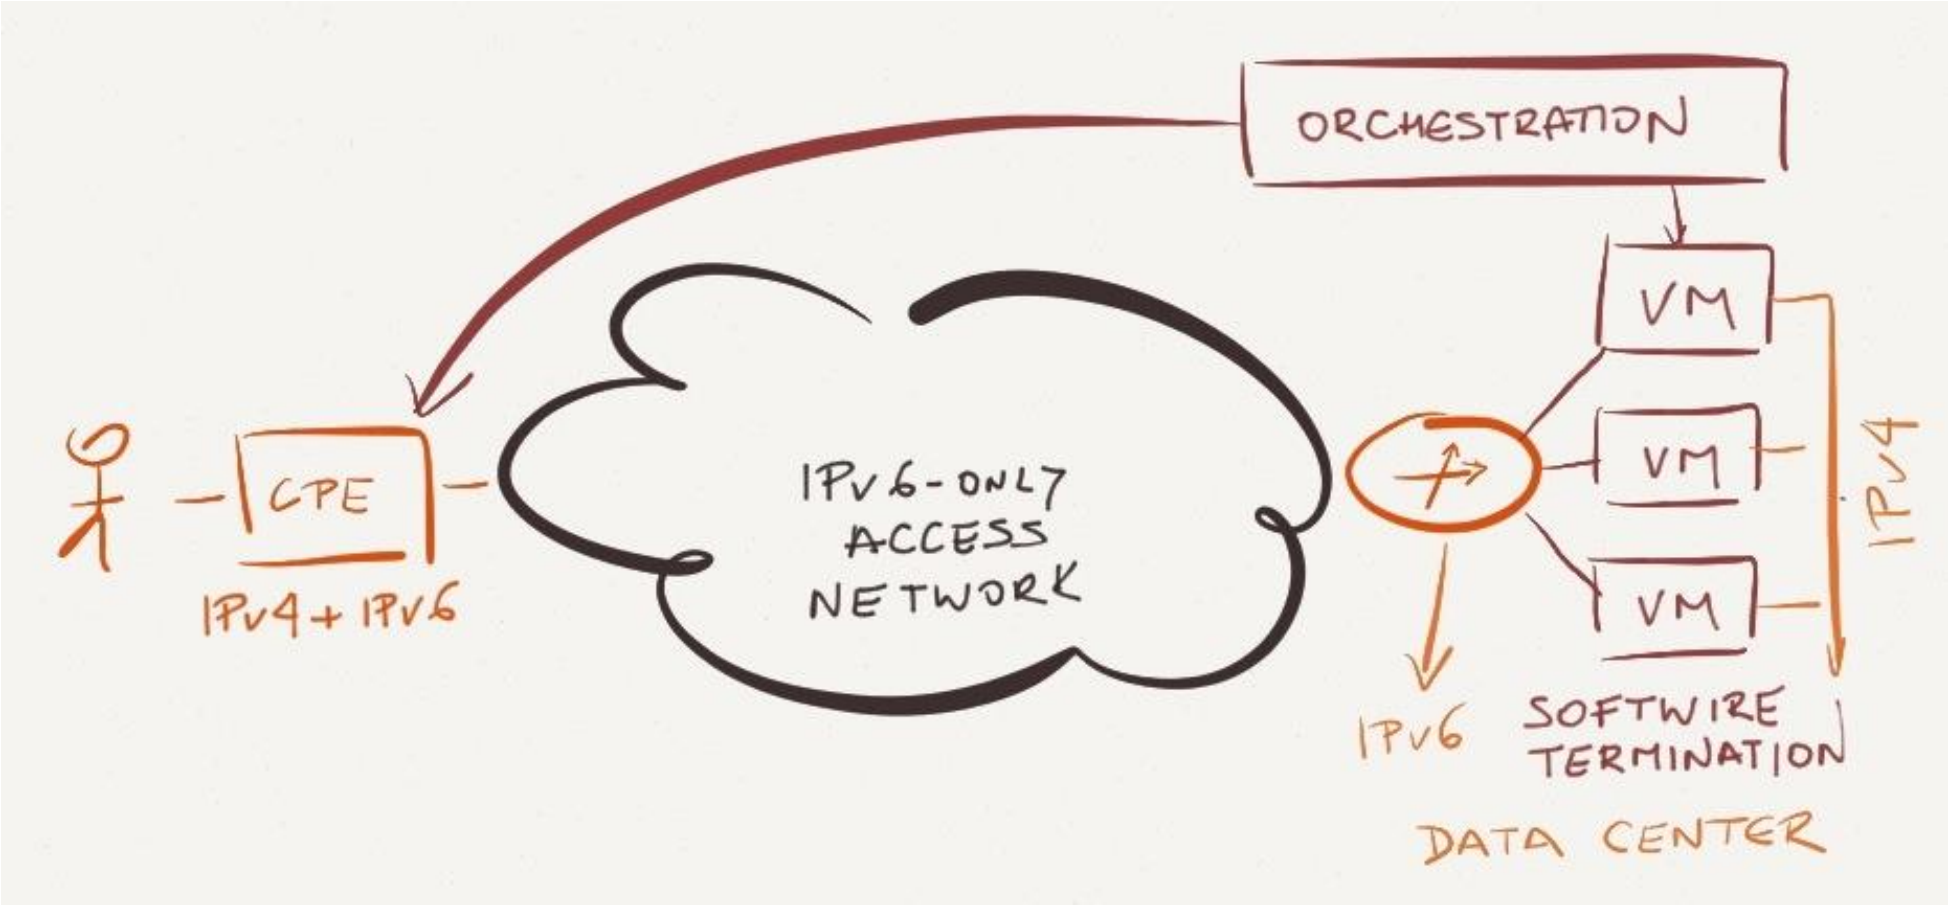

- PE-router / BRAS / GGSN in virtual machine format
- Deployed in Deutsche Telekom Terastream project

Hybrid Cloud Deployments

Principles

- Use a router with standard enterprise functionality
- Convert public cloud deployment into another site in enterprise WAN

Use cases

- VPN connectivity (cloud deployment = DMVPN hub)
- MPLS WAN endpoint
- Layer-2 or layer-3 extension for cloud migration (OTV or LISP)
- Network services control point (DHCP, NAT, BGP RR)

Unified Resources at Remote Sites

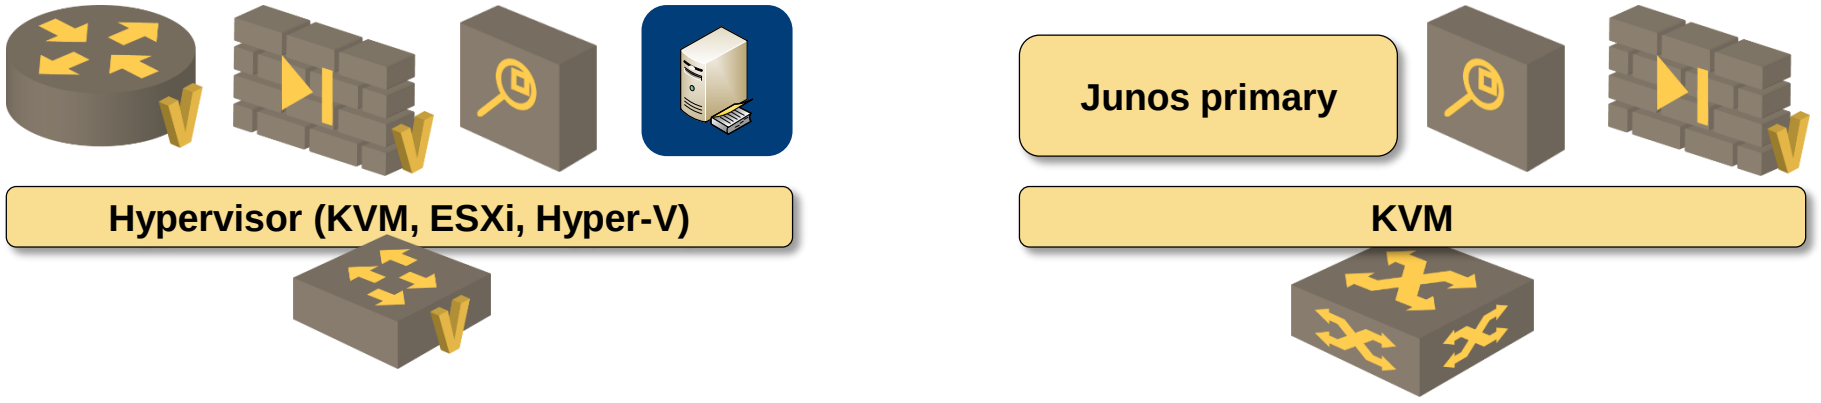

- Numerous virtual functions running on the same physical hardware
- Server, AD/DNS/DHCP servers, router, firewall, WAN accelerator, call manager

Implementation options

- Traditional x86 server
- Modern data center ToR switch (example: Arista, Juniper QFX5100)
- Router blade (example: Cisco ISR)

Control Plane Functionality

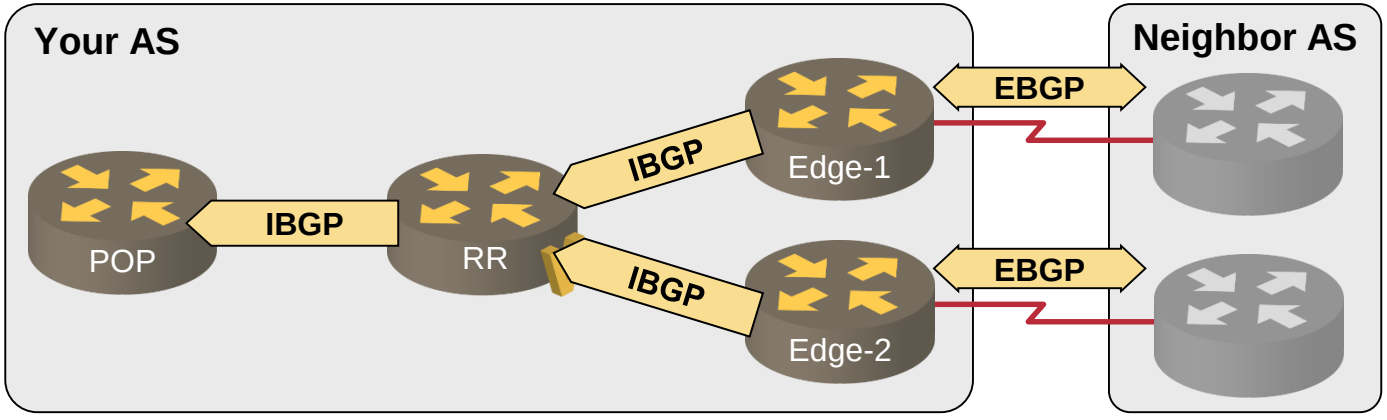

BGP Route Reflector is an ideal use case for virtual routers

- Very limited forwarding performance (control plane traffic only)
- Good enough control plane performance
- Deployable on any x86 hardware

Conclusions: Do Virtual Routers Make Sense?

TL;DL: Yes

- Low-speed environment (up to a few Gbps)
- Features that don’t require hardware acceleration (RSA key exchange)
- Control plane functionality

Not suitable for

- Bandwidth requirements higher than ~10 Gbps (40-50 Gbps in future)
- High number of ports

SDN, OpenFlow and NFV Resources on ipSpace.net

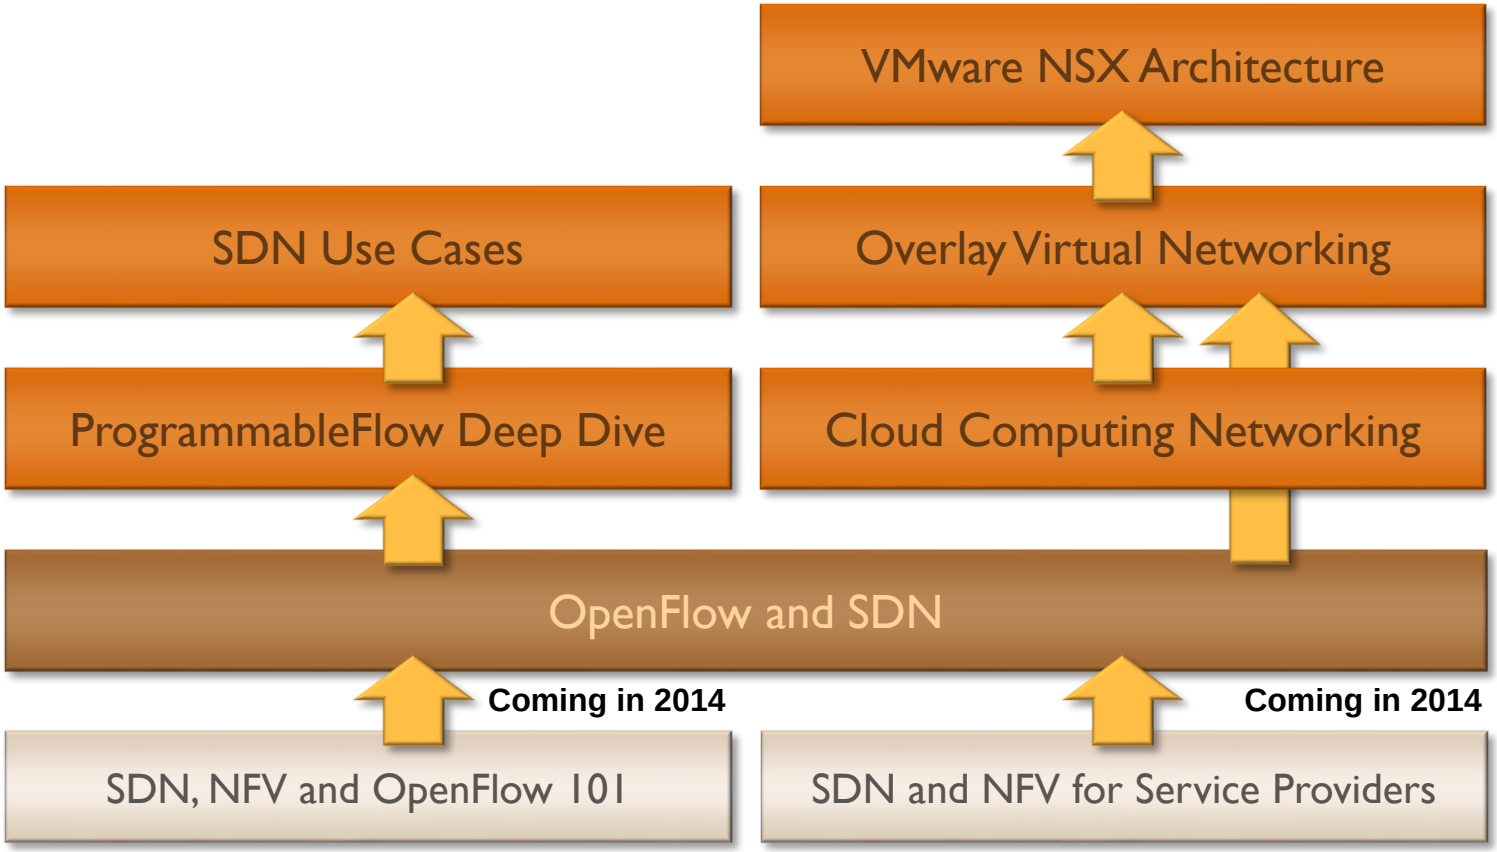

Trainings

- Live sessions
- On-Site workshops
- Recordings

Other resources

- Consulting
- Books and case studies
- Subscriptions

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

Questions?

Send them to ip@ipSpace.net or @ioshints