# Real-Life SDN – Seven Years Later

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

Focus

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

CISCO CCIE EMERITUS

vmware vEXPERT

# The Madness Started in March 2011

**SDN is the physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices**

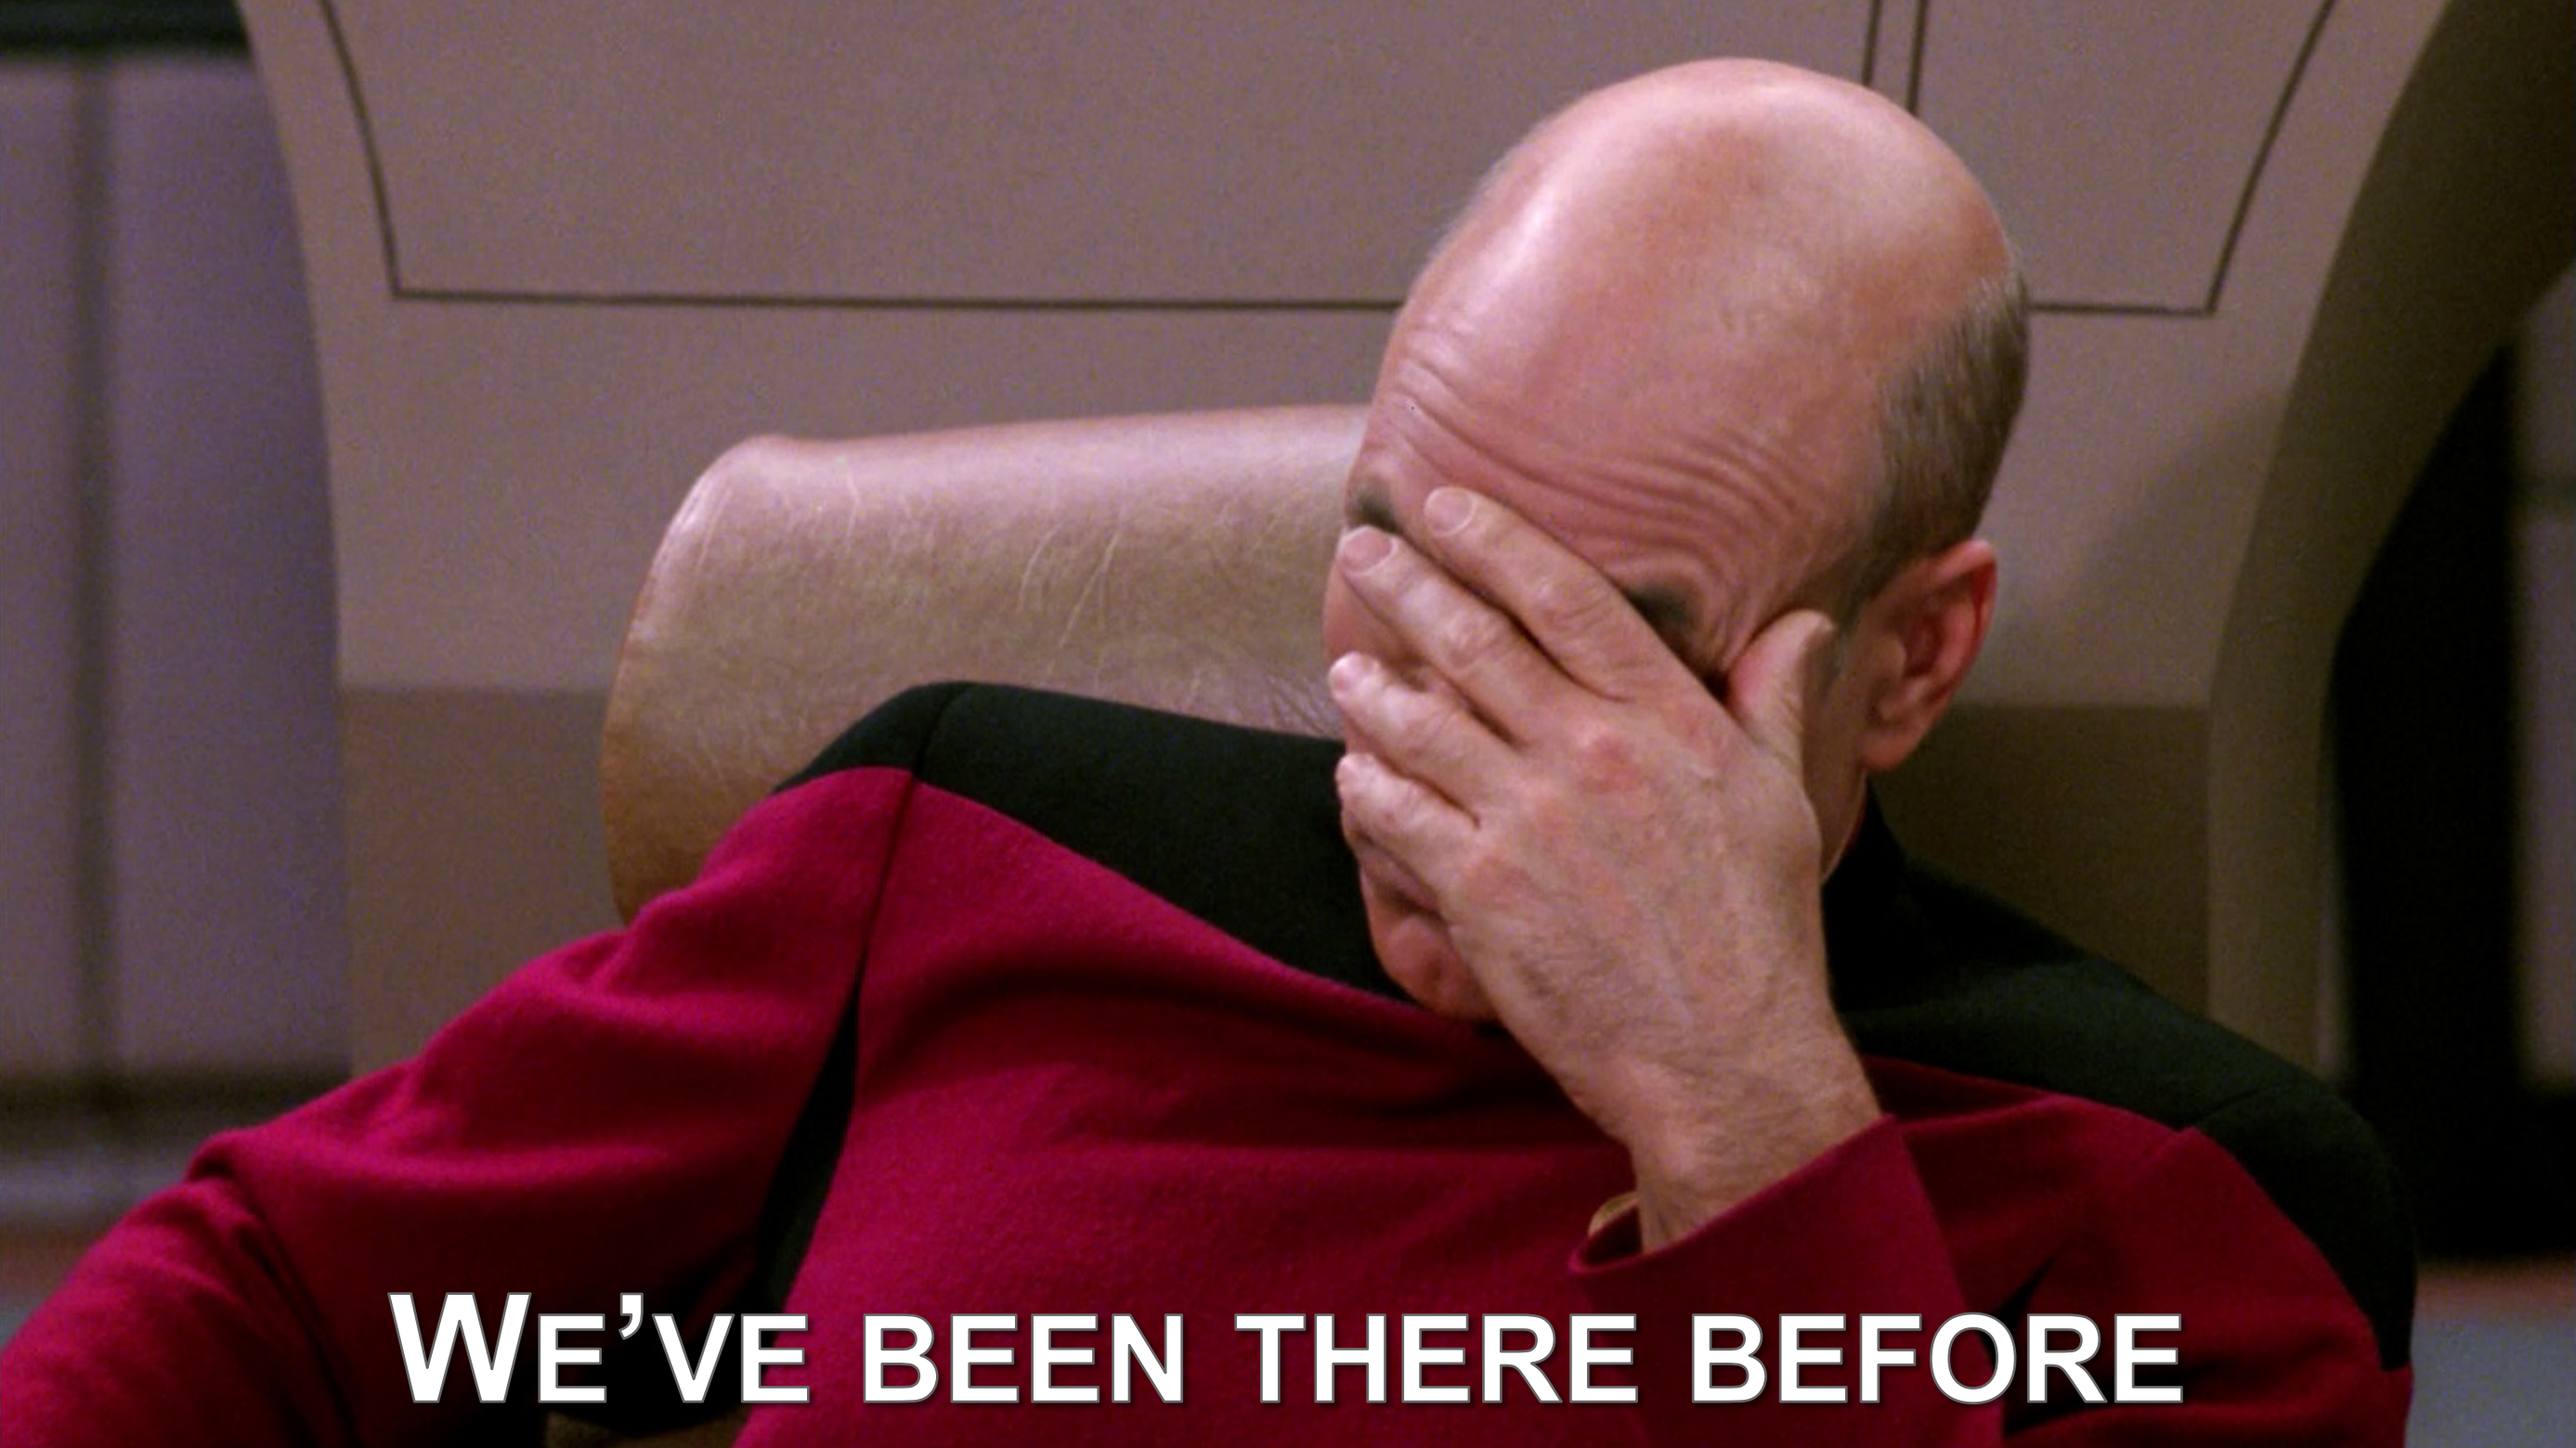
WE'VE BEEN THERE BEFORE

# Challenges of Centralized Control Plane

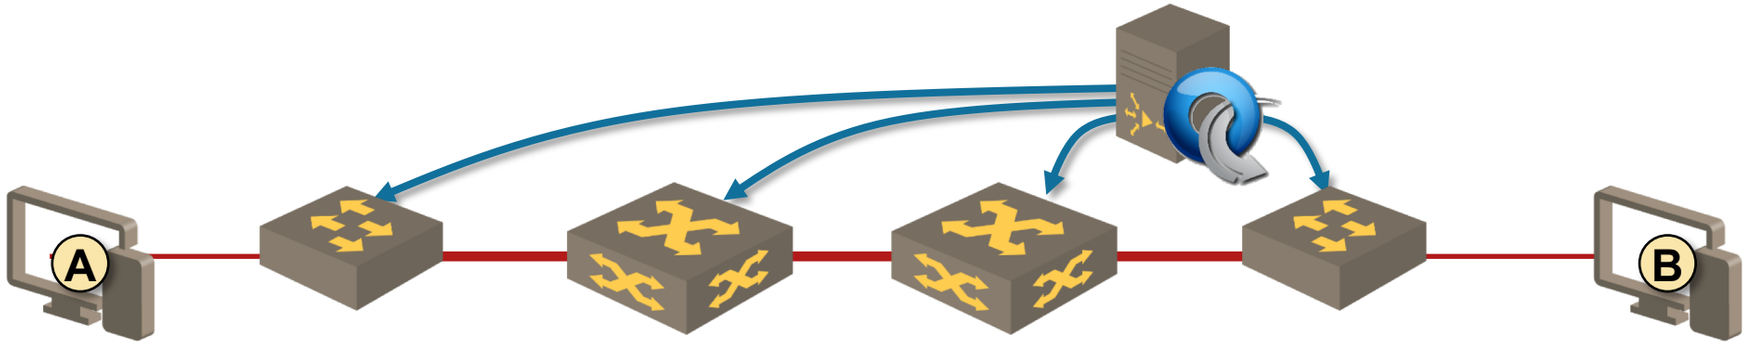


## Conceptual challenges

- Out-of-band control plane network
- No distributed intelligence ➔ no resilience to failures
- Controller is the central point of failure
- Total loss of control-plane protocols after a controller failure
- Lack of shared fate (requires end-to-end OAM)

## Real-life challenges

- Poor OpenFlow implementations (very limited multi-table support)
- Limited TCAM sizes (few thousands)
- Low TCAM update speed (less than thousand entries per second)
- Slow switch-to-controller channel due to underpowered switch CPUs

**More in *OpenFlow Deep Dive* webinar**

# Controller Plane Makes More Sense

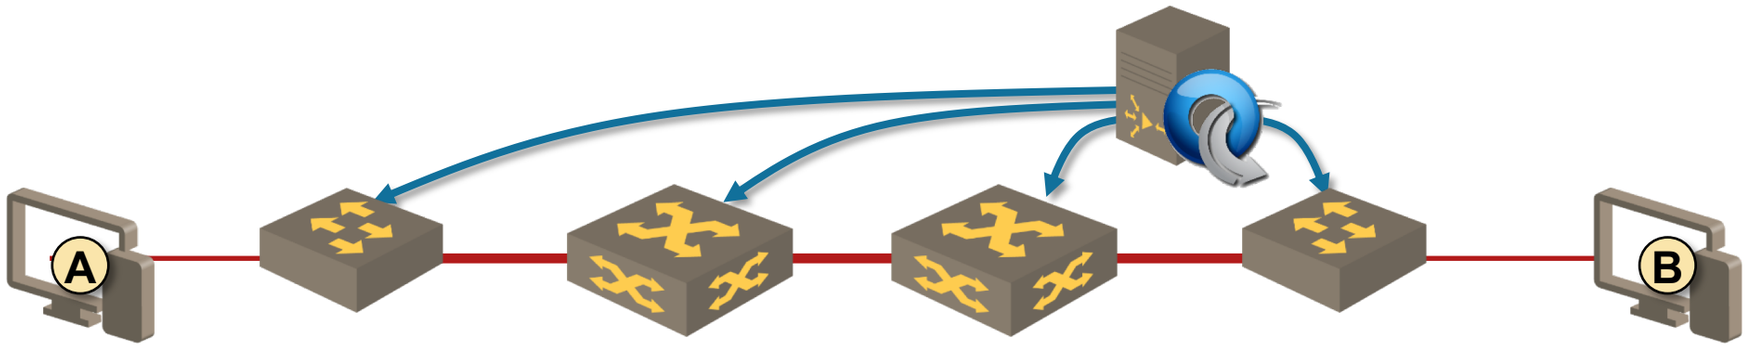


## SDN Architecture document

- … data plane may include the necessary minimum subset of control and management functions …
- … the controller plane may configure the data plane to respond autonomously to events such as network failures or to support functions delivered by, for example, LLDP, STP, BFD, or ICMP
- … number of functions with control aspects are widely considered as candidates to execute on network elements
- A more nuanced reading of the decoupling principle allows an SDN controller to delegate control functions to the data plane…

**More @ http://blog.ipspace.net/2015/08/sdn-onf-is-moving-to-logically.html**

SDN is the physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices

# Notable Exceptions

- Google WAN network (really?)
- Telstra traffic engineering
- Few Internet Exchange Points (IXP)
- Scale-out IDS infrastructure

# OpenFlow @ Google

Problem:

- Traffic engineering in inter-DC WAN backbone

Solution

- Custom data center edge switches
- Cluster of OpenFlow controllers in each data center
  ➔ data center edge switches behave like a single node (build-your-own-router)
- BGP and IS-IS between OpenFlow controllers
  ➔ classic routing between data centers as a mechanism-of-last-resort
- Centralized traffic engineering application
  ➔ path elements are downloaded into individual controllers

**Equivalent to IS-IS + BGP + PCEP**

# Speaking of Software-Defined

# Software-Defined? And What Were We Doing?

- Most networking devices have software
- Device behavior was always defined by its software
- Is it all hype … or just marketing gone bad?

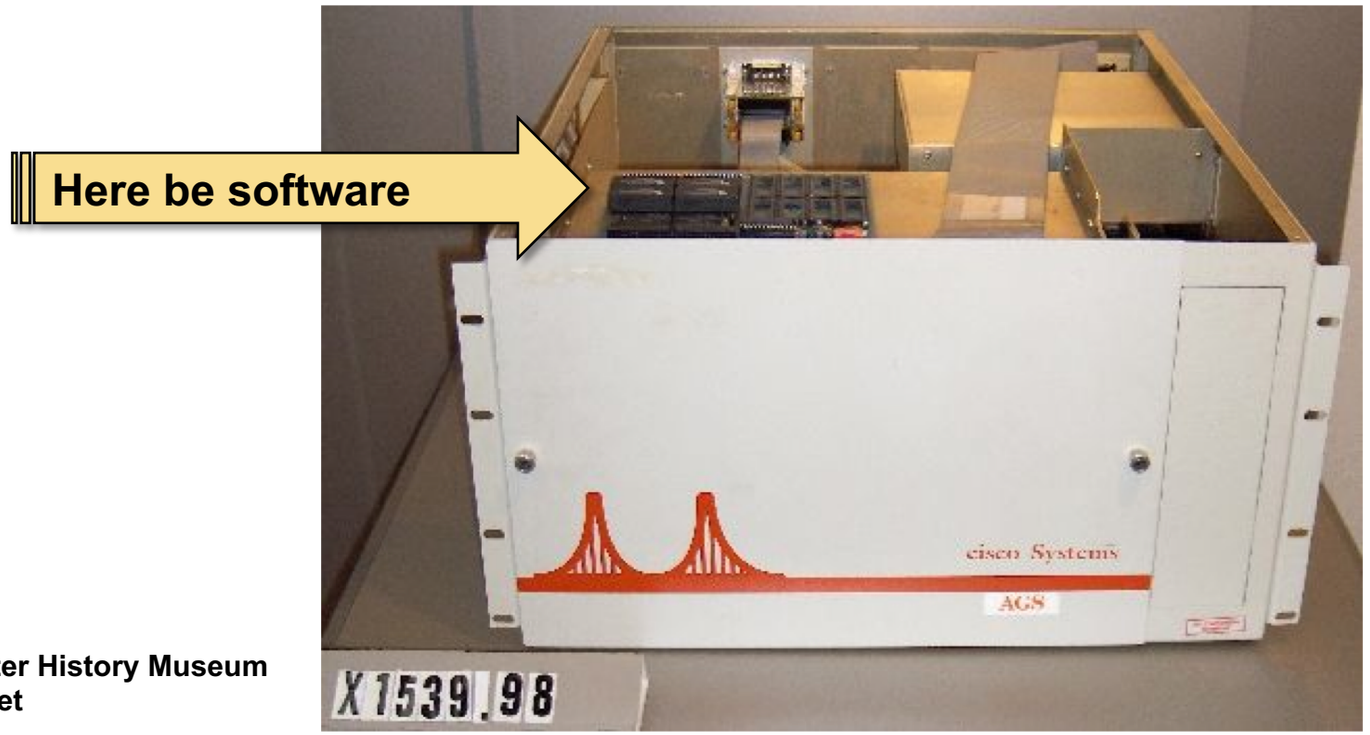


**Cisco AGS at Computer History Museum**
**Source: Evilrouters.Net**

# What Else Is Called SDN?

# SDN is packet forwarding done in software (on x86 platform)

# Huge Success (When Applied Correctly)

- 20 Gbps per core, 100+ Gbps per x86 server
- Innovative appliances (example: L2VPN over IPv6, 4-over-6 tunneling…)
- Major networking vendors offering virtualized devices with DPDK/6WIND or equivalent

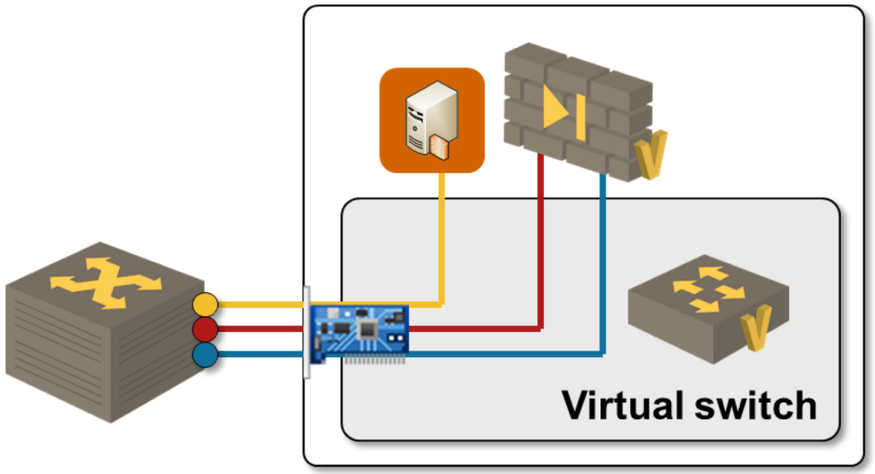


**More in *Network Function Virtualization* webinar**

# SDN is whitebox switching (running software on third-party cheap hardware)

# Software / Hardware Disaggregation

- Hardware costs are 30-40% of the product costs (gross margin of networking vendors is above 60%)
- Software and support are the really expensive parts (and yet we're all buying boxes)
- Why can't we buy hardware and software as separate items?

Benefits:

- Increased flexibility (reuse the same hardware)
- Simplified sparing

SDN is whitebox switching (running software on third-party cheap hardware)

# Real Benefits

## Install your own software on networking devices

- Control-plane daemons
- Customized telemetry
- Push agents

## Linux everywhere

- Unified management of servers and network devices
- Common tooling
- Common control-plane functionality (including shared bugs)

**More in *Open Networking* and *Cumulus Linux* webinars, explore also *Software Gone Wild* podcast**

# Success Story: Extensible SDN Internet Router

**Concept:**

- Use a ToR switch as Internet edge router (cheap 100 GE)
- Hardware forwarding table is an order of magnitude too small

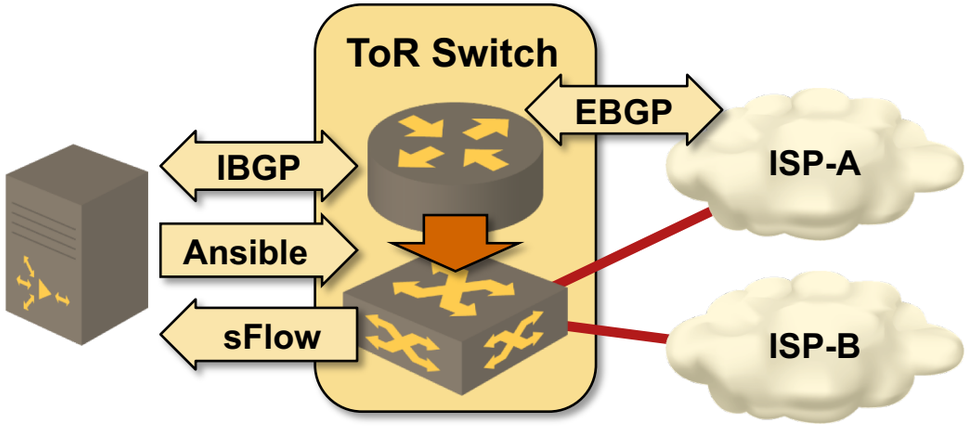


**Solution:**

- Use default routing + specific routes for high-volume prefixes

**Implementation:**

- SDN controller correlates sFlow data and BGP information
- SDN controller finds top-N prefixes and generates list of interesting prefixes
- BGP RIB-to-FIB filter is pushed to the ToR switch
- ToR switch copies selected prefixes from BGP table into the hardware forwarding table

**More in *SDN Use Cases* webinar and *Software Gone Wild* Episodes 19 & 42**

# SDN is programmable access to network devices

SDN is programmable access to network devices

# … But at Least the Hype Moved Us Forward


Yesterday

- Telnet or SSH
- Inconsistent CLI with no status codes or unified error handling
- Printouts changing in every software release (and breaking automation scripts)

What we got (in small doses)

- NETCONF
- REST APIs using JSON or XML data format

**More in *Network Automation* webinars and online course**

SDN is an approach to computer networking that allows network administrators to manage network services through abstraction of lower level functionality

# Isn’t That Just a Glorified Orchestration System?

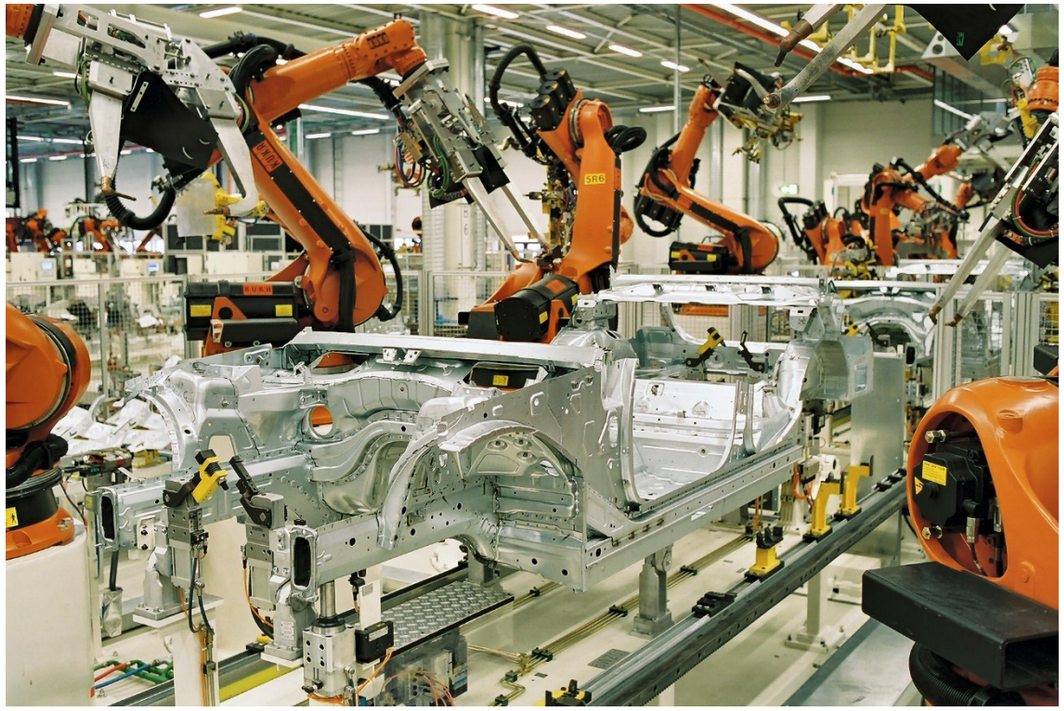

# SDN Principles Revisited

What we would love to have

- Automated and consistent network services deployment
- Consistent policies
- End-to-end visibility
- Decisions made on centralized view of end-to-end visibility
- Automatic programming or configuration of network devices
- Automated response to events or changes in traffic or topology

# Build or Buy?

# Others Made It Work…

# … When Will You?

# More Information

# Learn At Your Own Pace

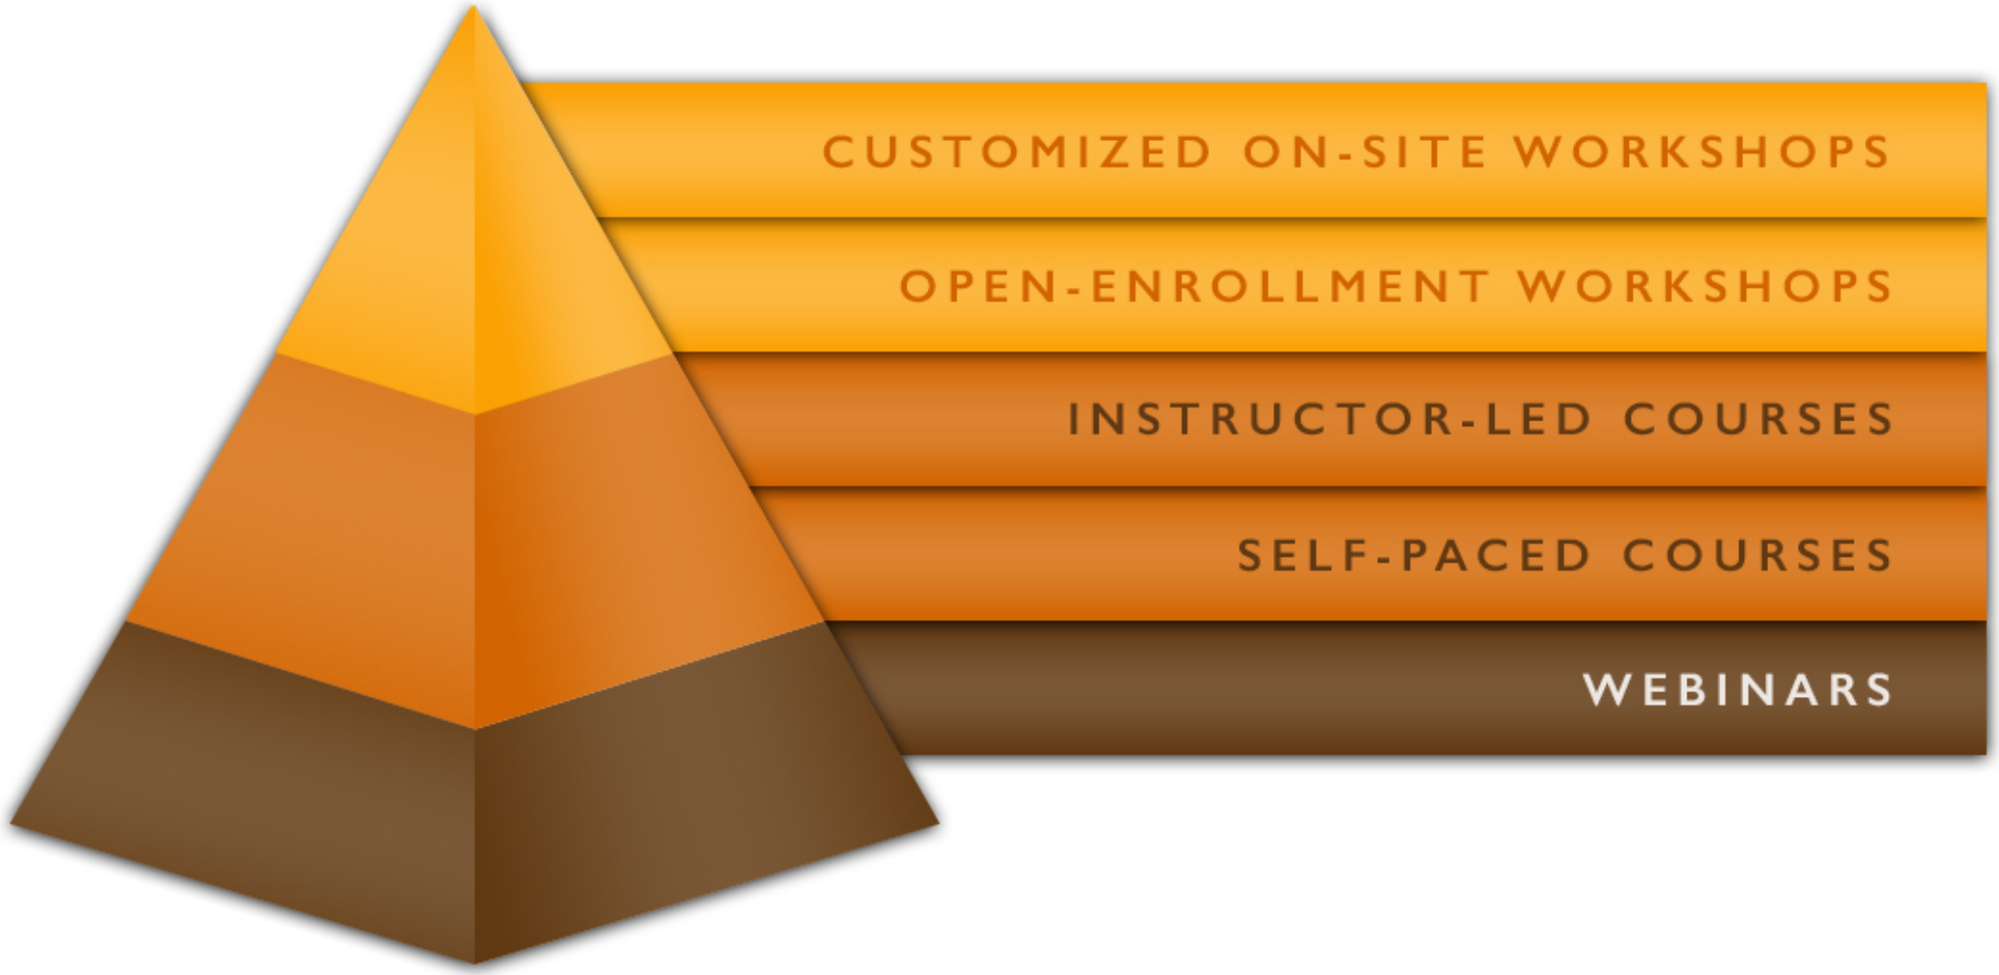

# SDN Webinars

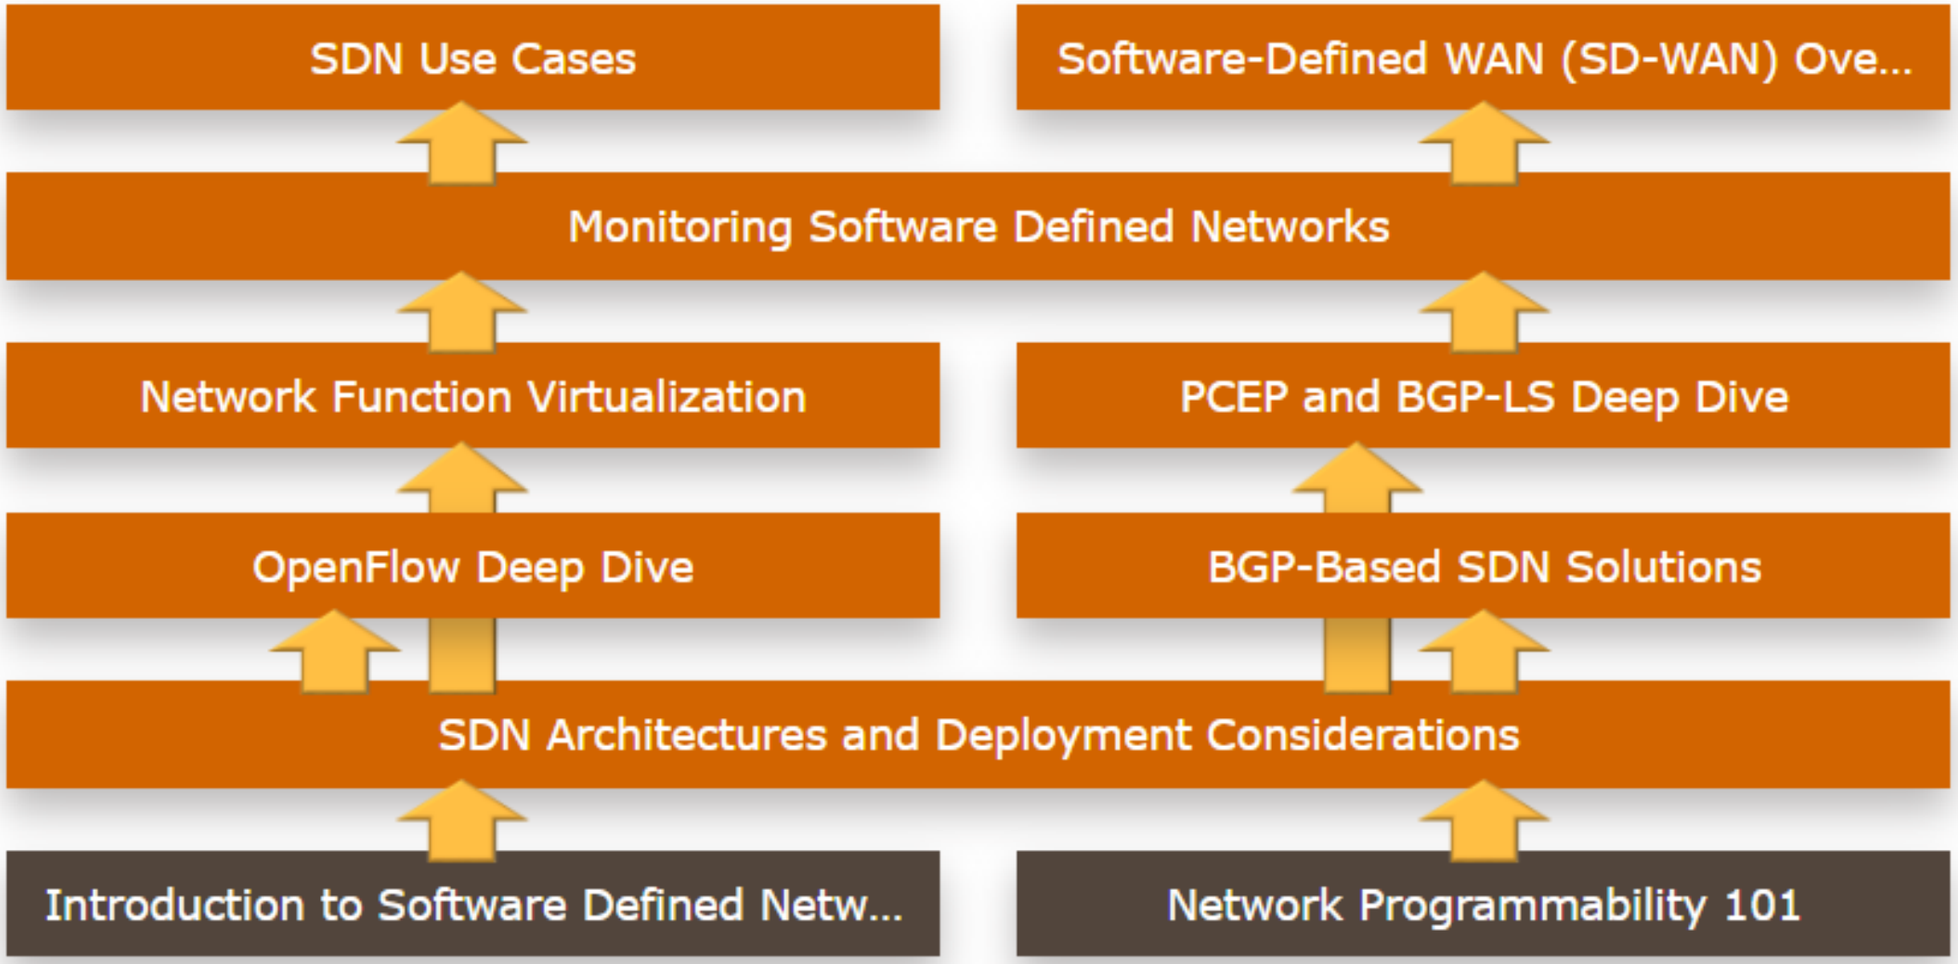

# Network Automation Webinars

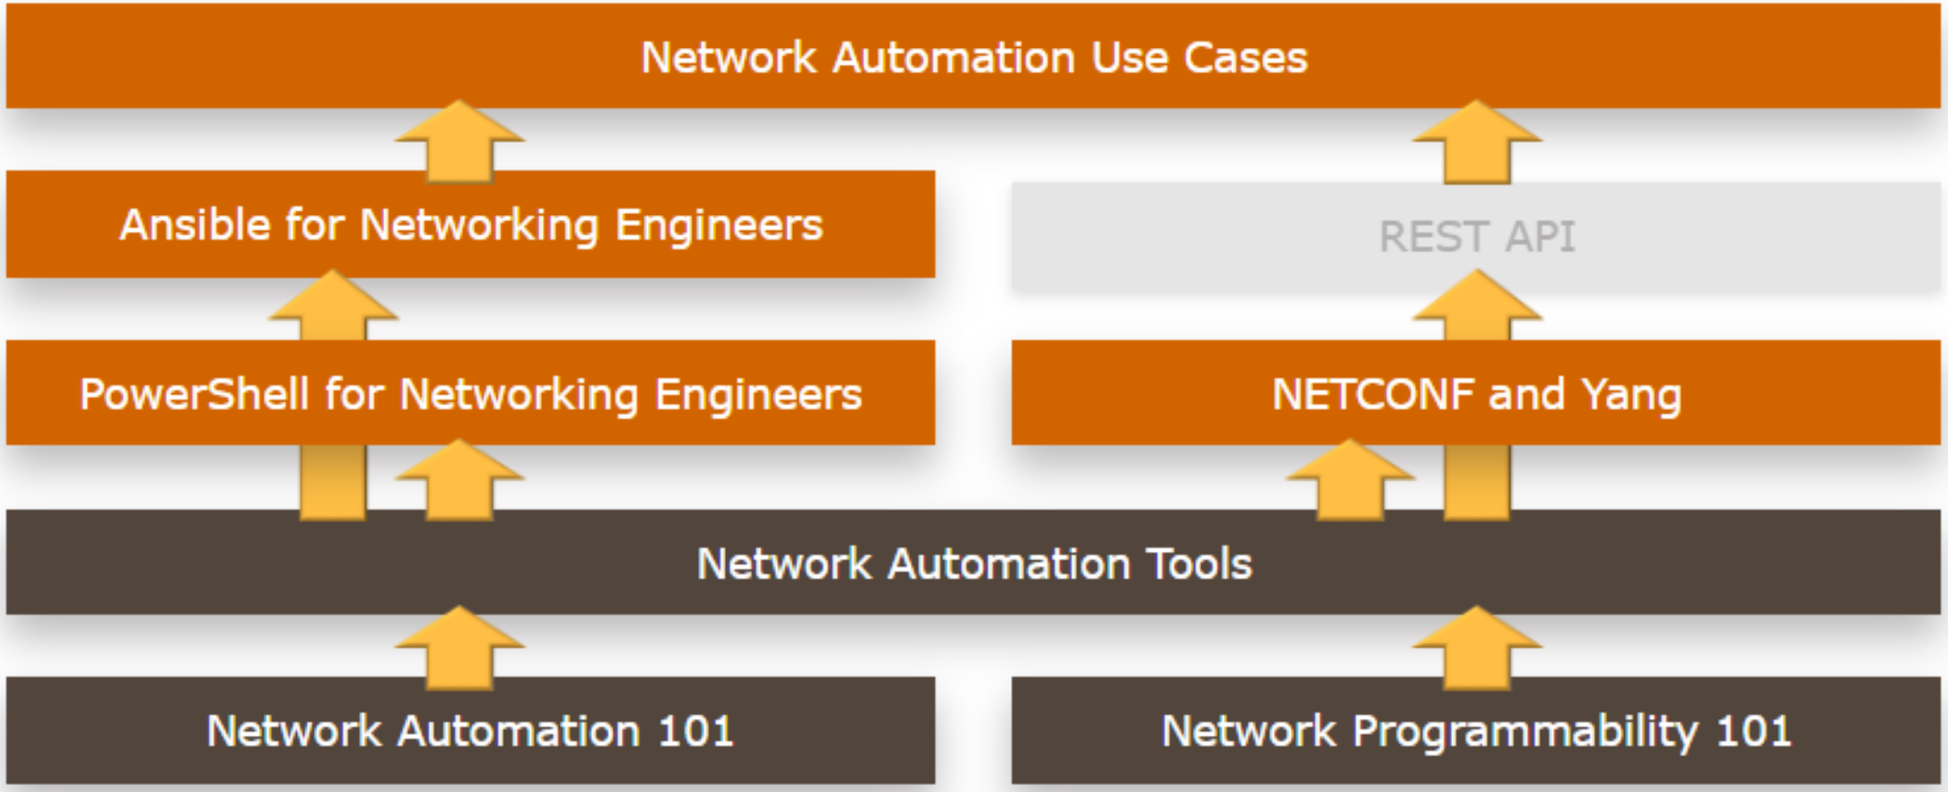

# Questions?

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

| | |
|---|---|
| Automation: | ipSpace.net/NetAutSol |
| Data center: | ipSpace.net/NextGenDC |
| Webinars: | ipSpace.net/Webinars |
| Consulting: | ipSpace.net/Consulting |