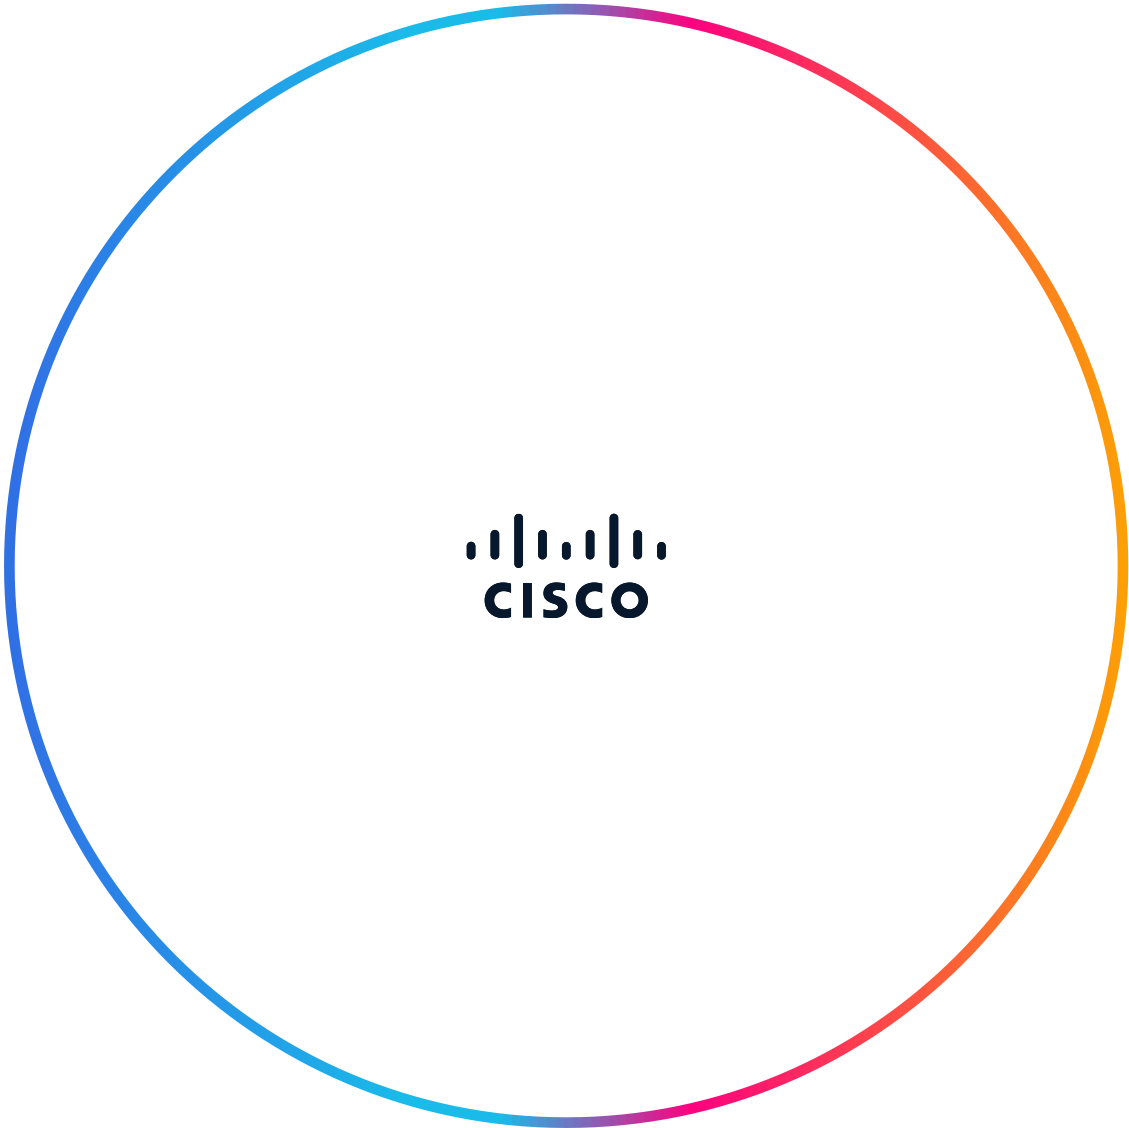


The Evolution of Network Engineering

Matt Leuschner
Solutions Engineer



Who am I?

Solutions Engineer

7 years at Cisco

24 years of experience in IT

Networking, virtualization, desktop services,
automation/coding

Focus on end-to-end infrastructure

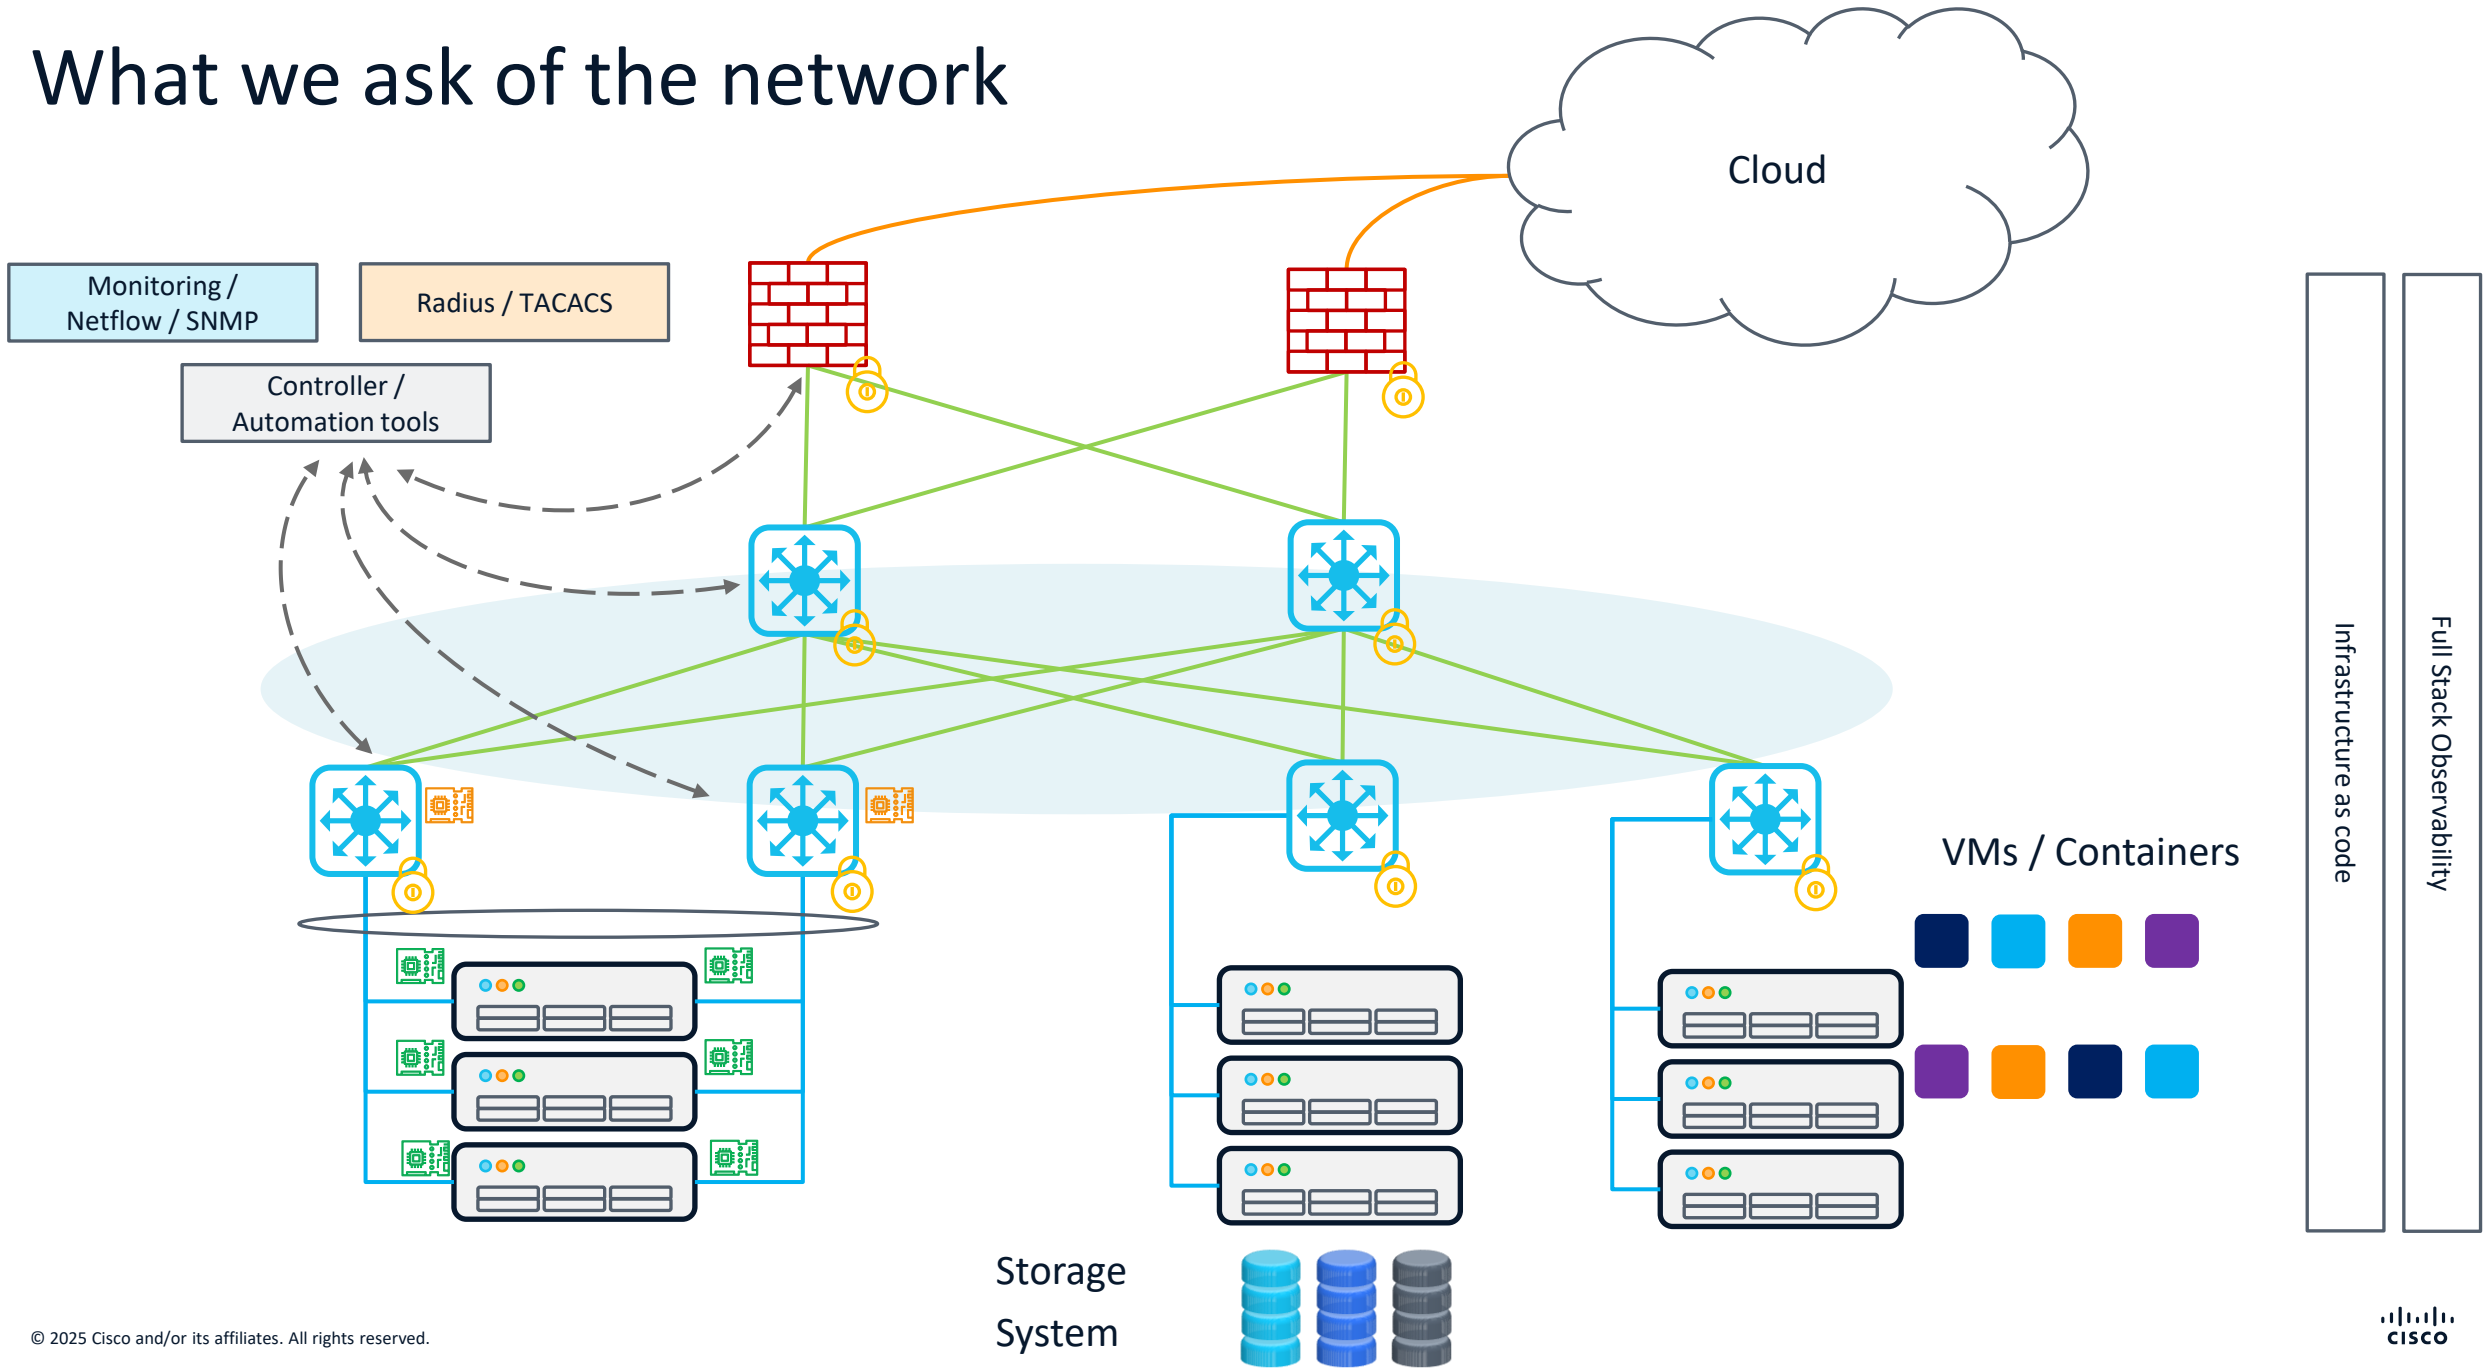
Probably the only Cisco presenter with an
unlicensed nuclear accelerator



The Pace of Network Change



What we ask of the network

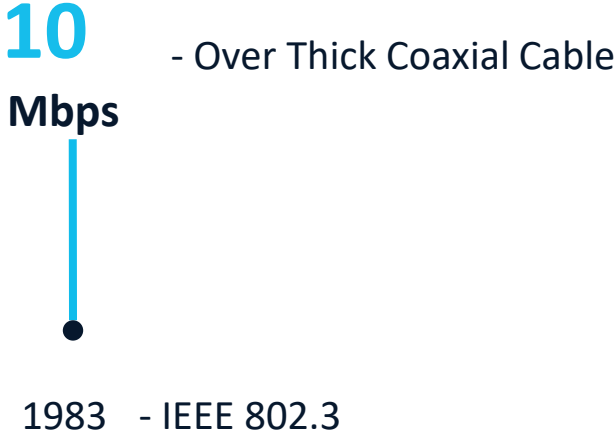


Evolution of the network

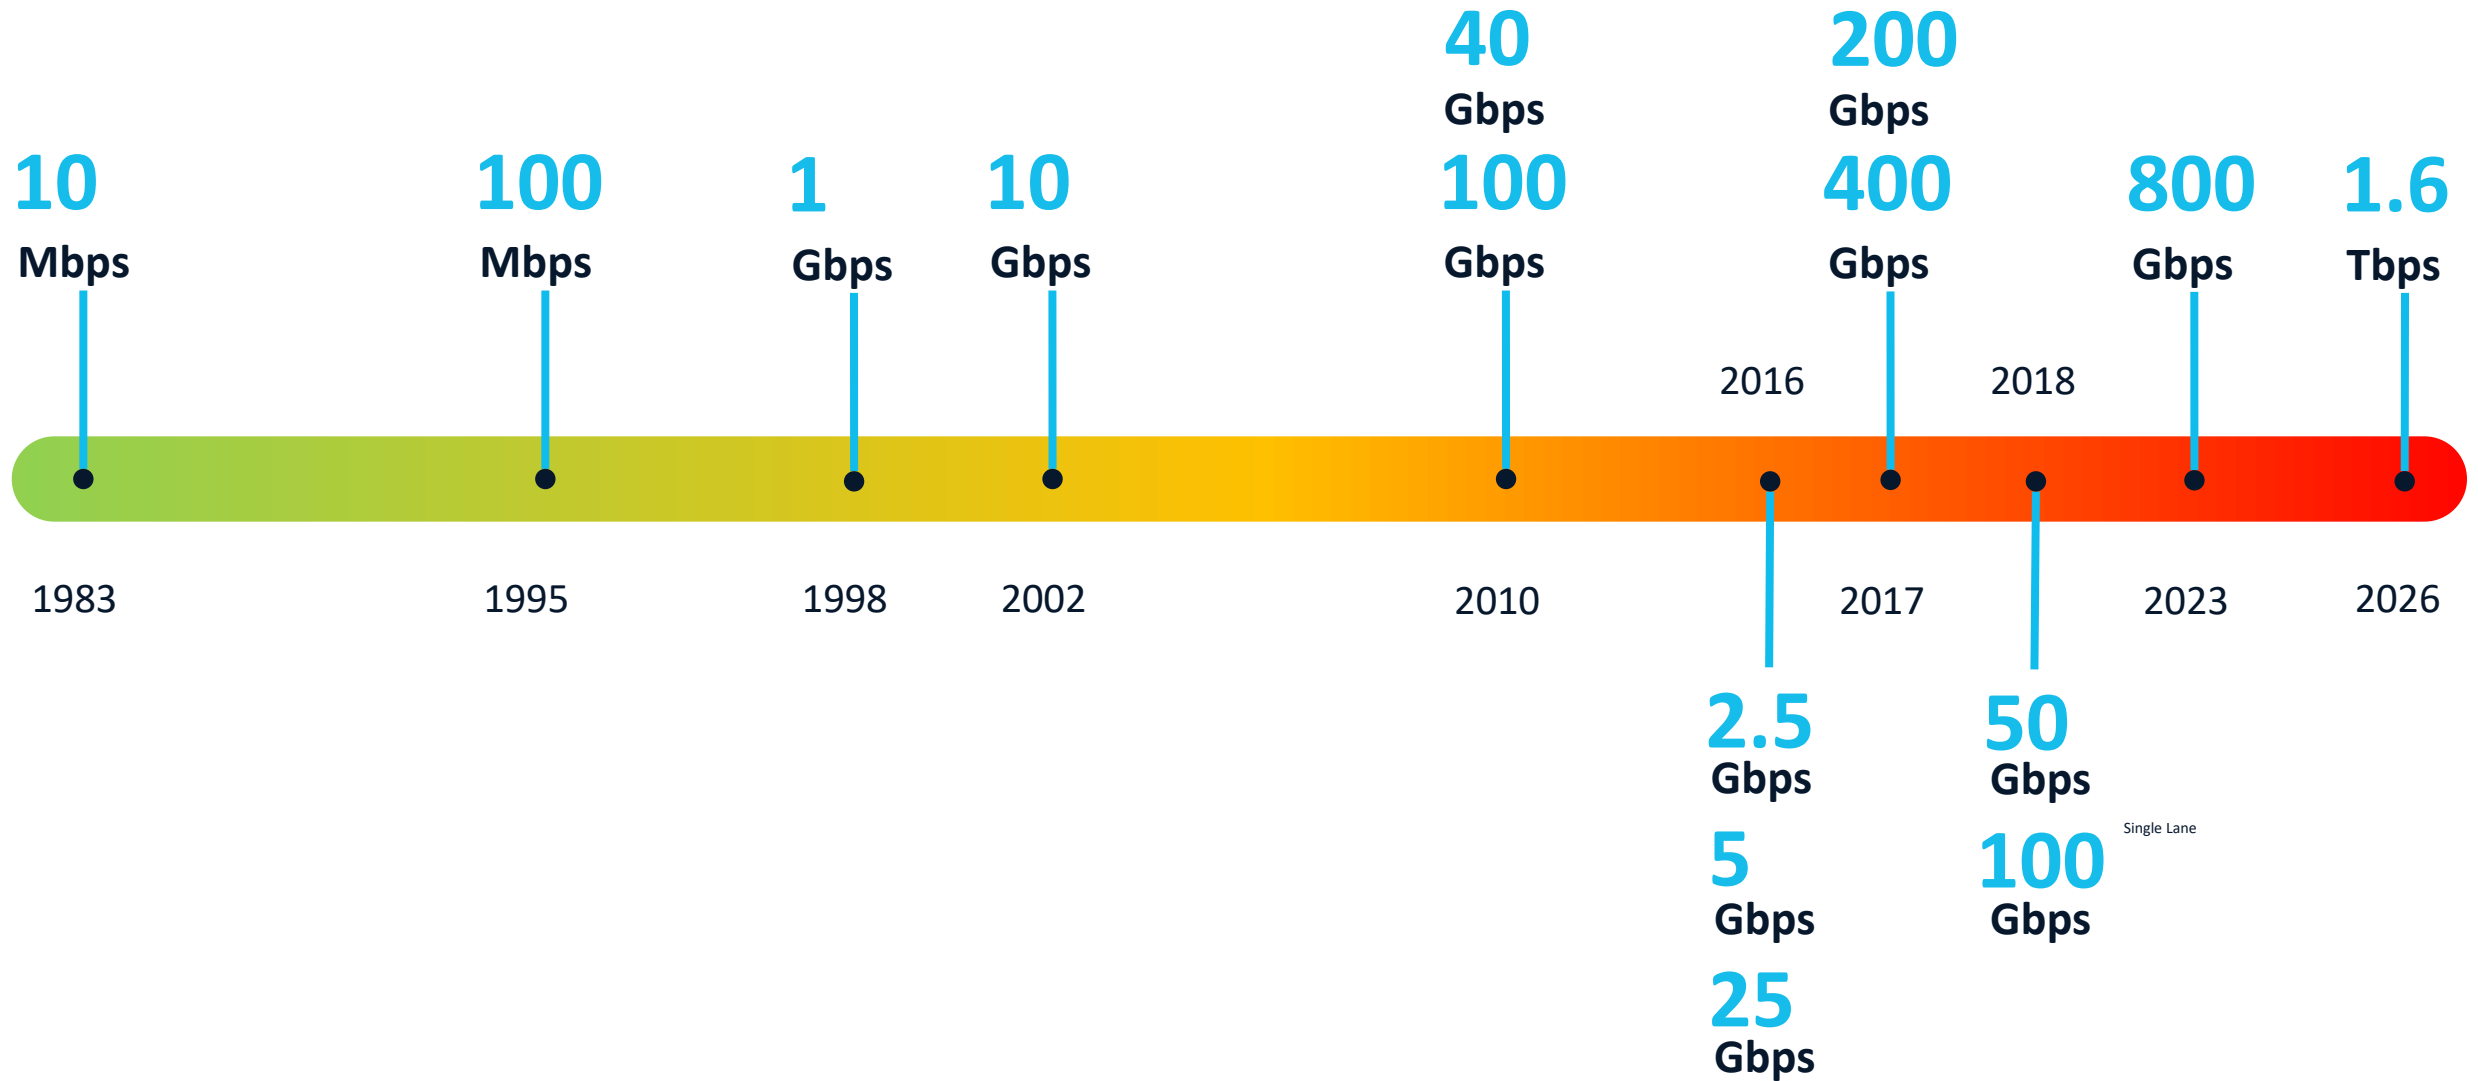
How has it evolved?



The start of modern networking



Increase in Ethernet network speeds



Traditional 3 tier networking

STP – Spanning Tree Protocol

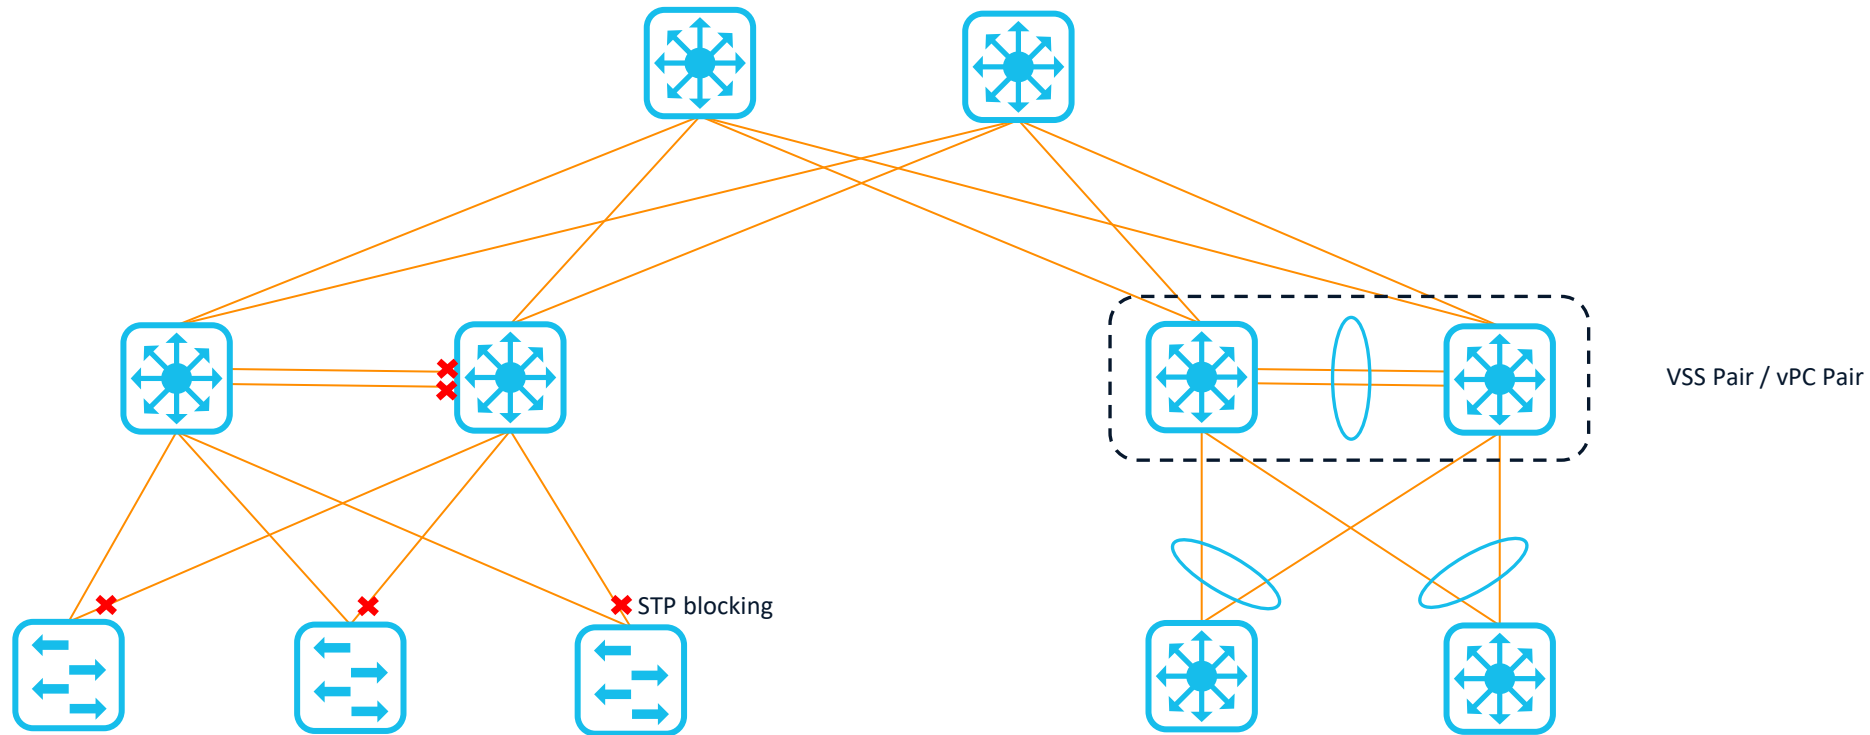
VSS – Virtual Switch Stack

vPC – Virtual Port Channel

Core

Distribution

Access



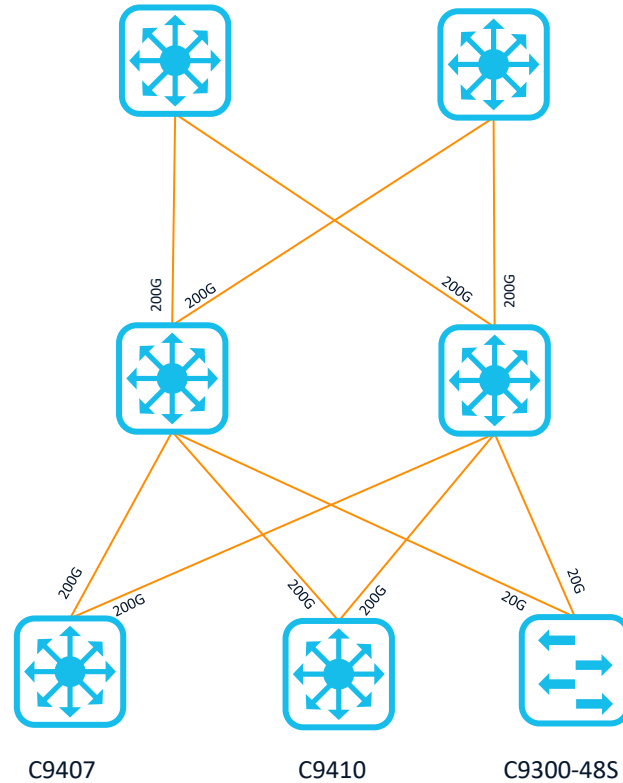
Spine and Leaf architecture

Border / Leaf

Spine

Leaf

2x100G per
supervisor (200G)



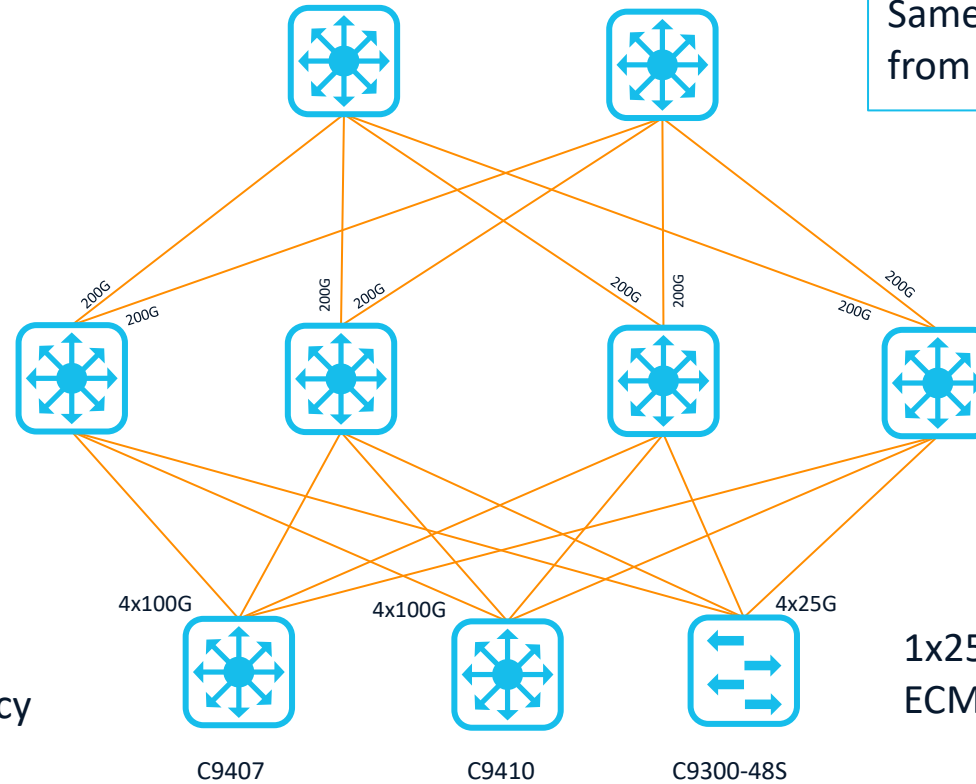
2x10G per
leaf uplink (40G)

Increased Fault Tolerance and Resiliency

Border / Leaf

Spine

Leaf

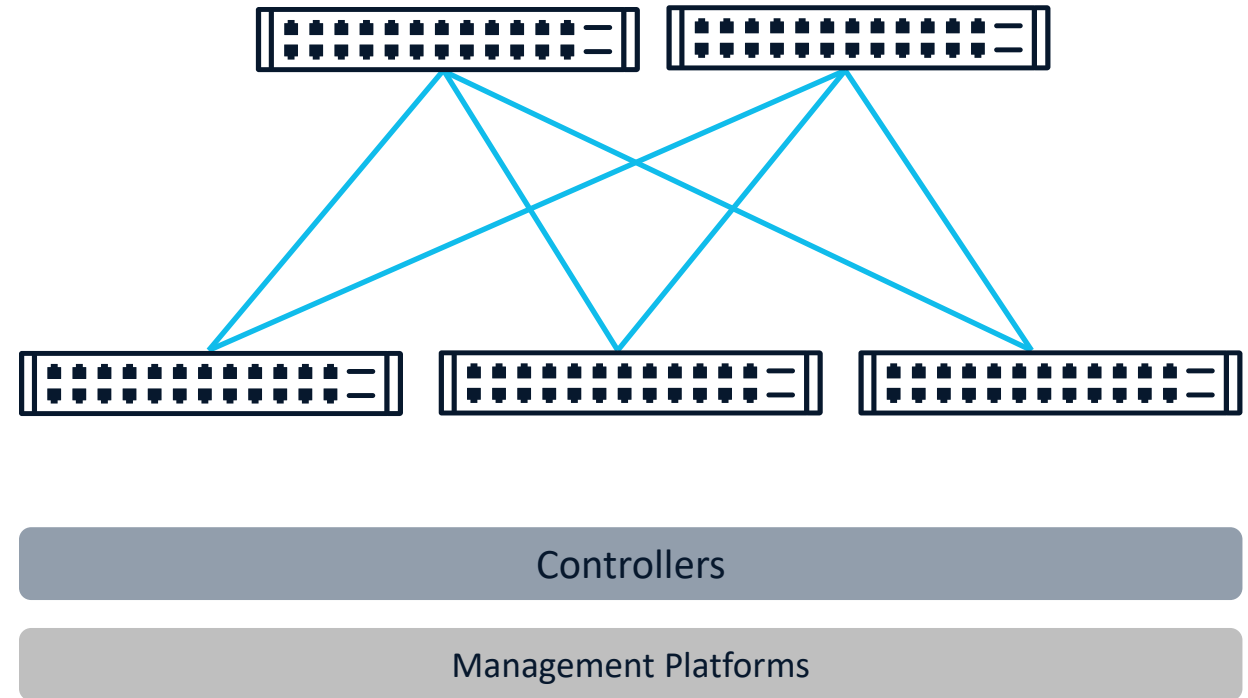
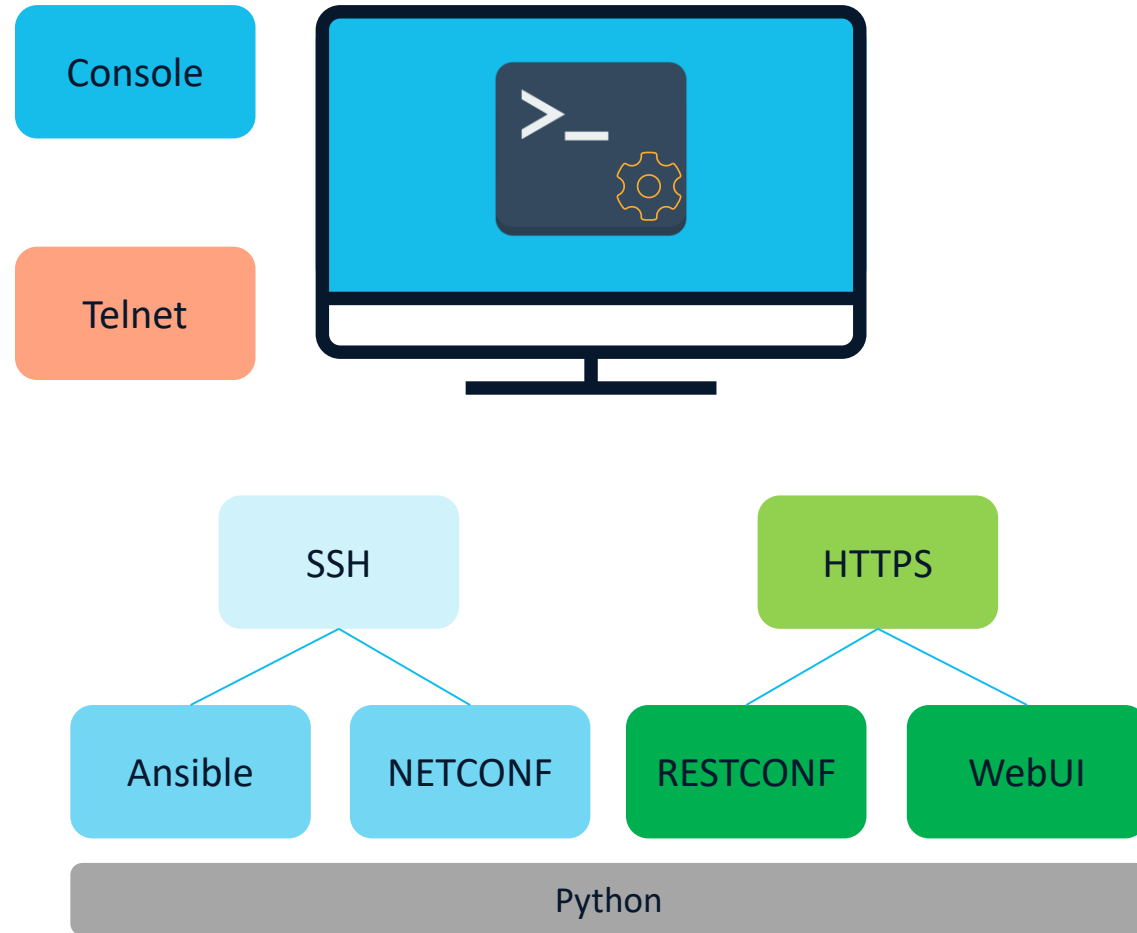


Same throughput, decreased impact from spine or link failures.

1x100G per spine
ECMP + Redundancy

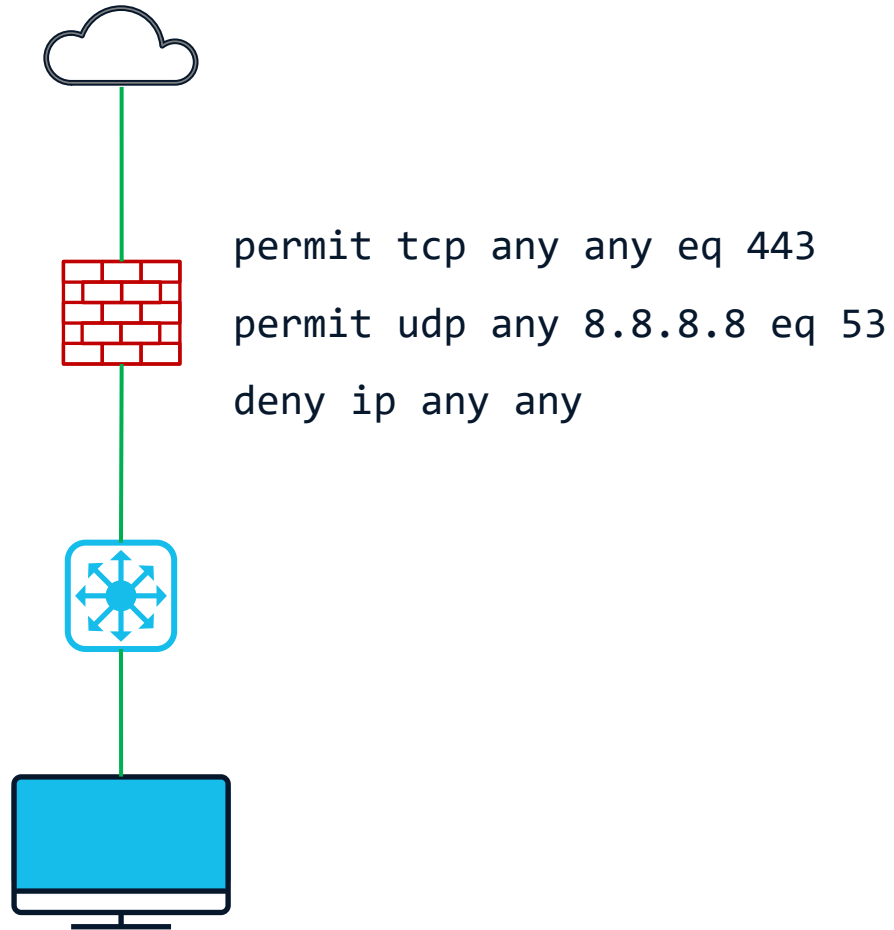
1x25G per spine
ECMP + Redundancy

Device Management Methods

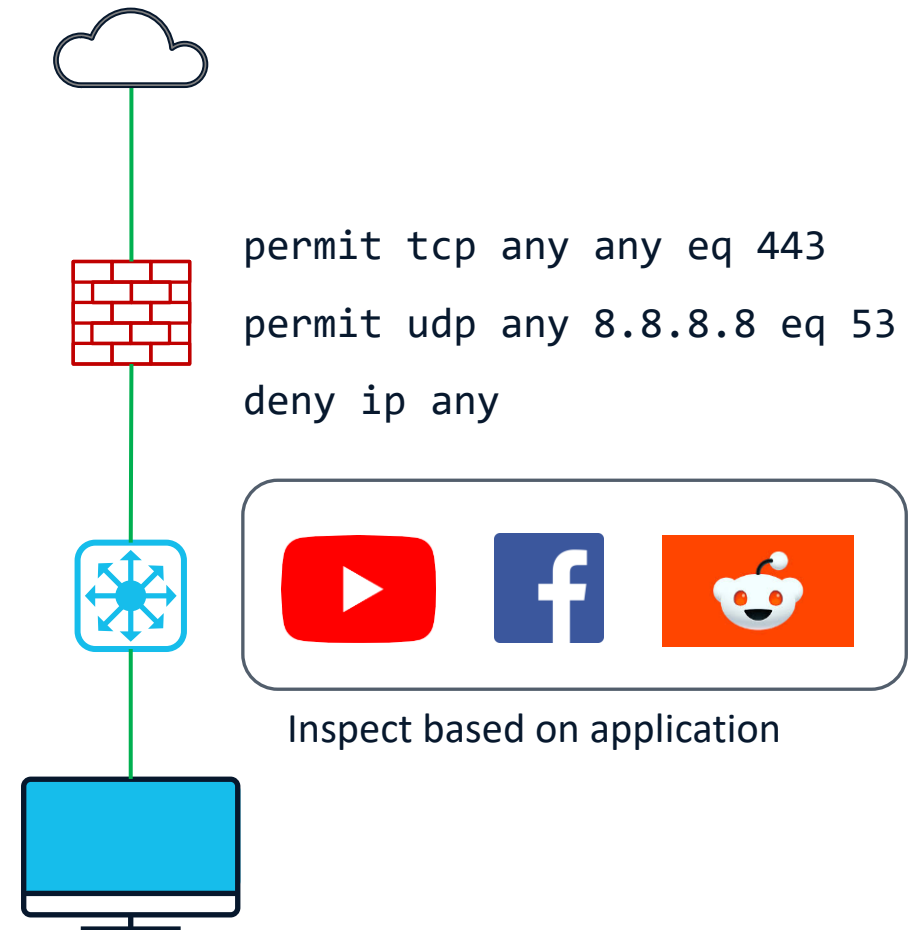


Security on the network

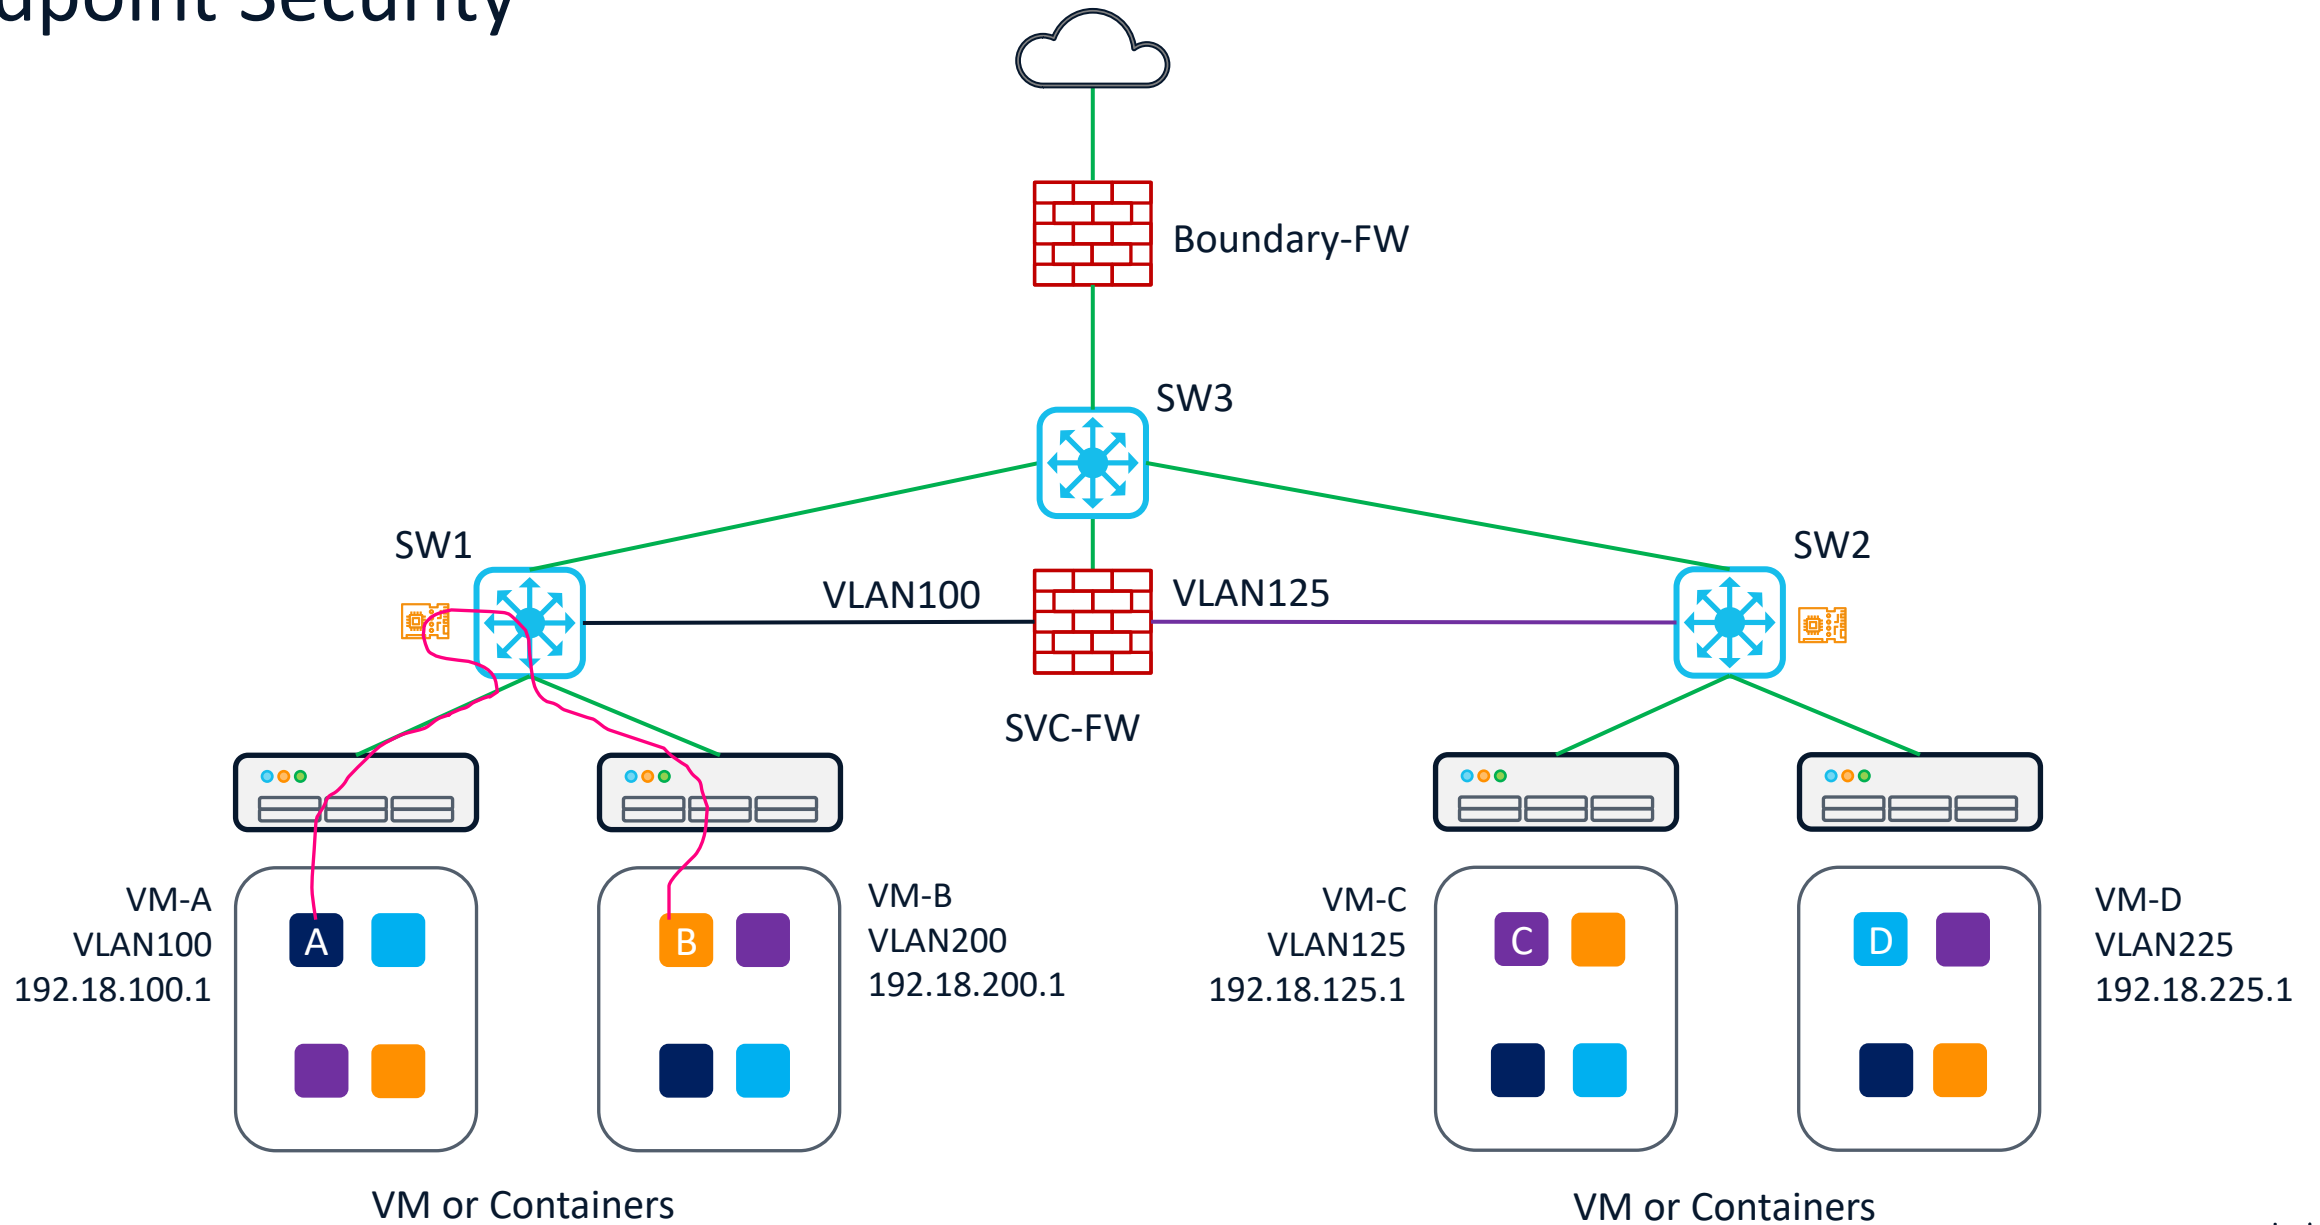
Routed Firewall



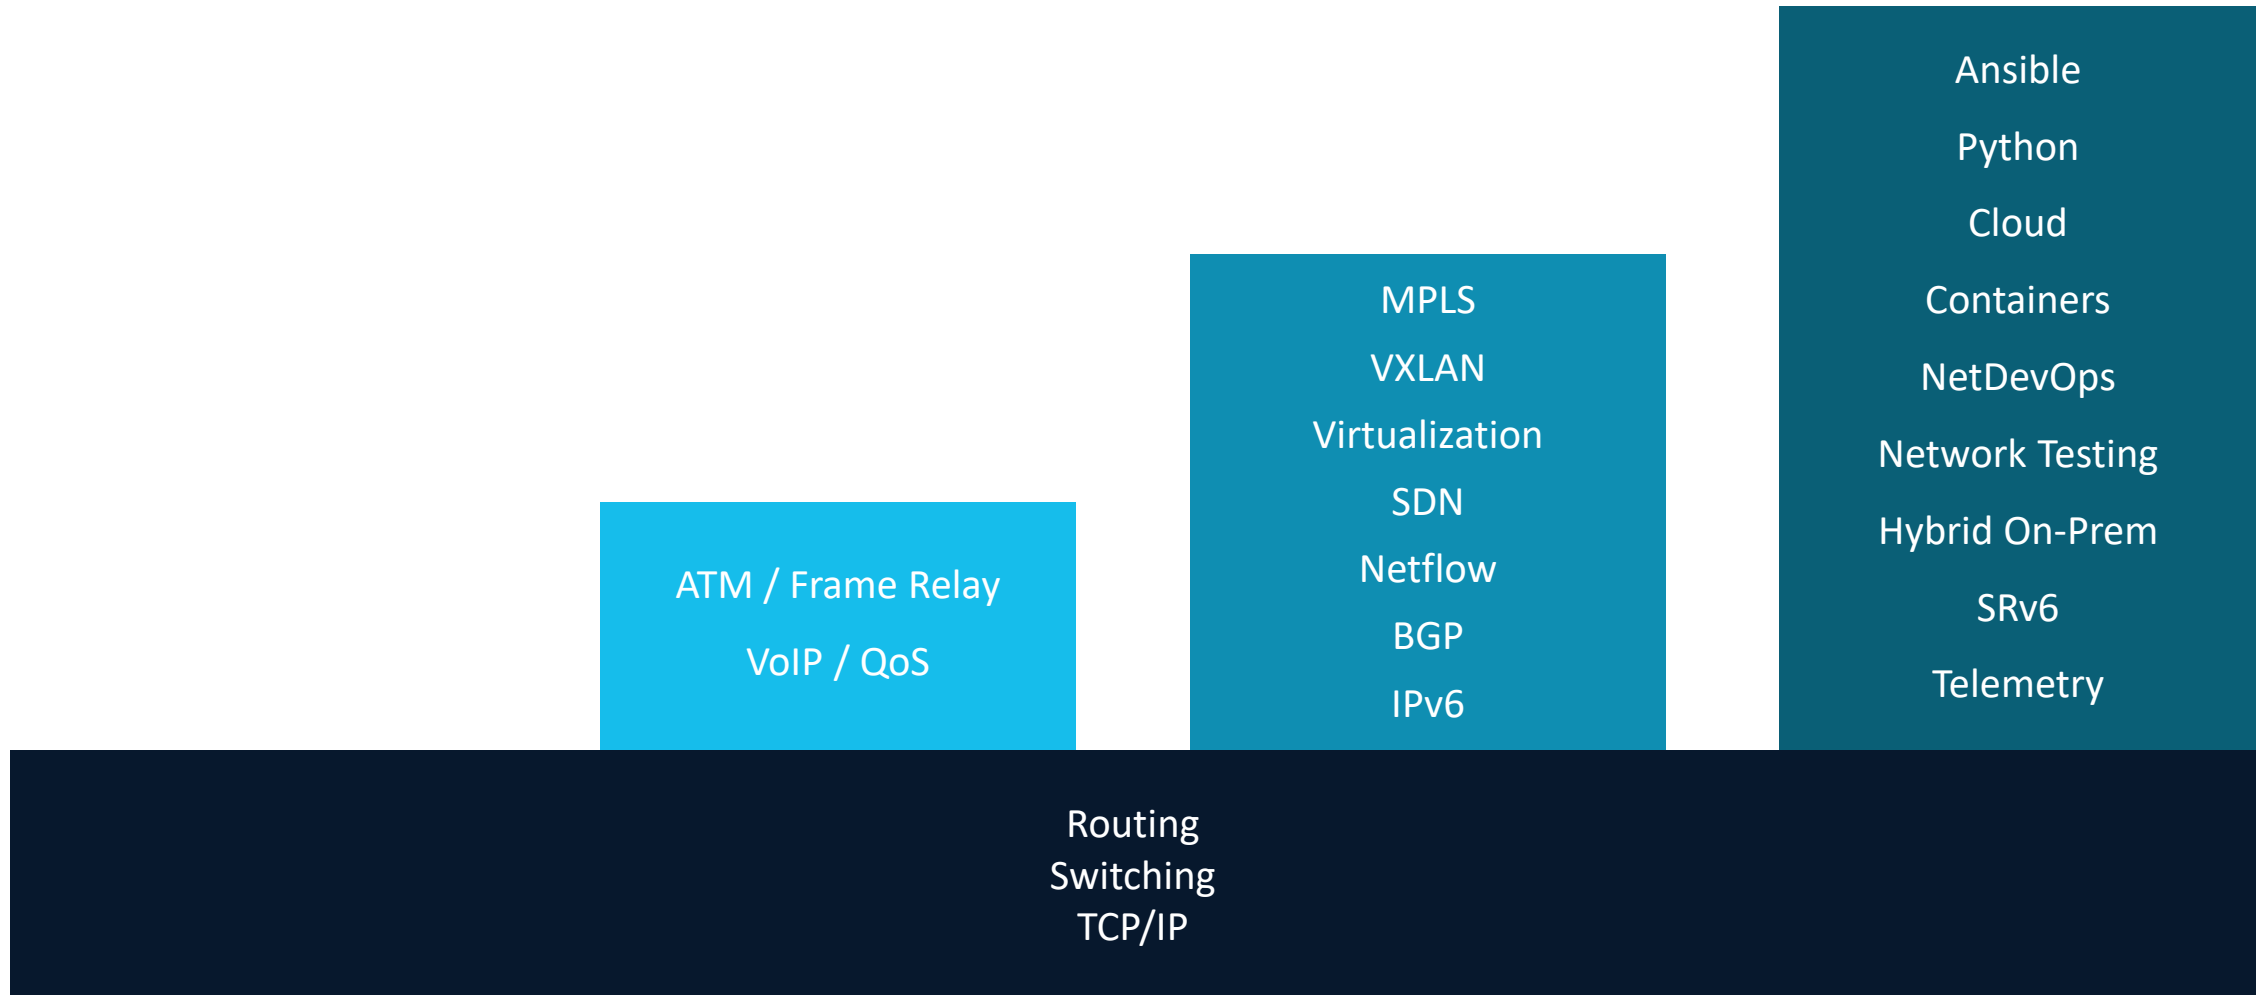
Next Generation Firewall



Endpoint Security



Expansion of network domain knowledge



Essential Skills

Skill sets for network engineering

Core

- IPv4 / IPv6, TCP, UDP
- Routing Protocols, OSPF / ISIS
- Switching protocols, VLANs, VXLAN
- Multi-Protocol BGP

Infrastructure

- Linux networking
- Cloud Networking (AWS / Azure / GCP)
- Containers
- Virtualization, RoCEv2

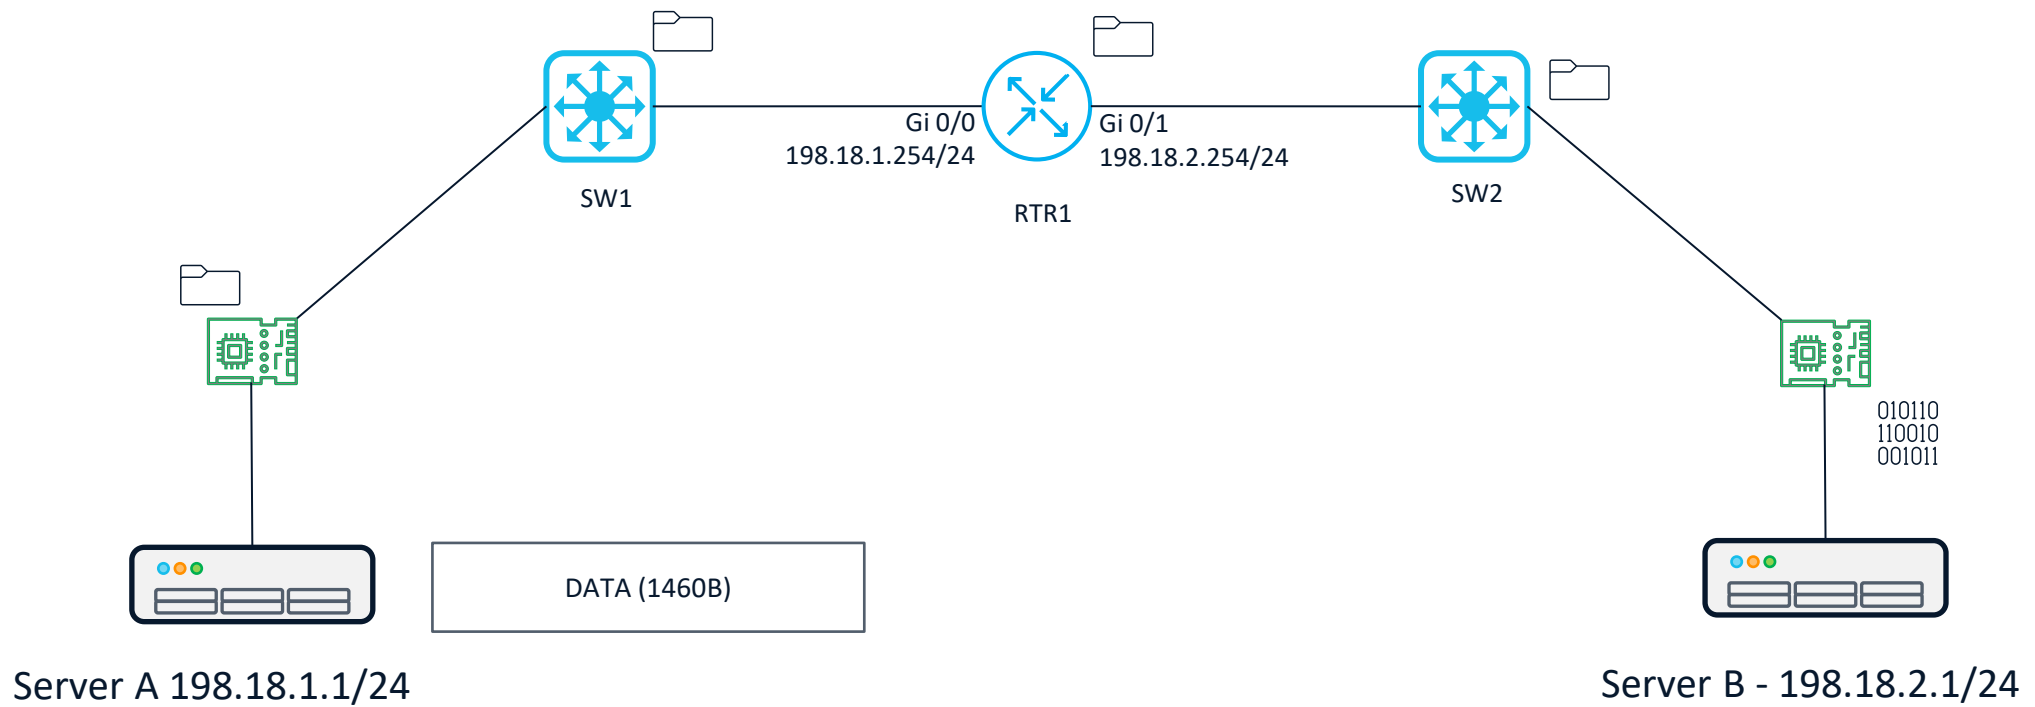
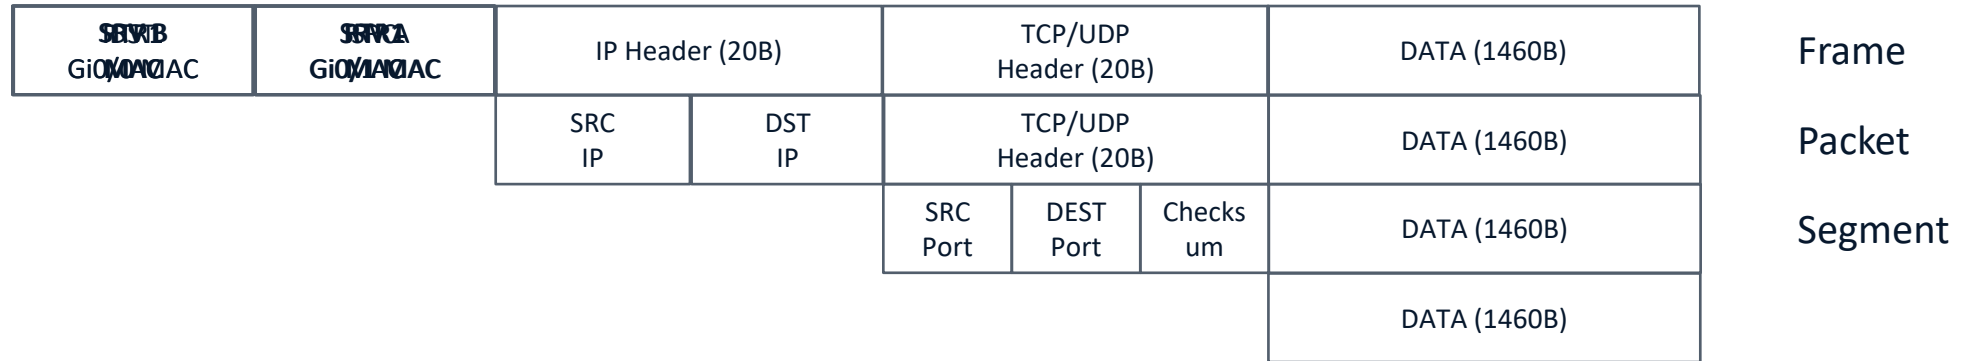
Troubleshooting

- TCPdump/Wireshark
- Traffic Flow
- Network testing
- Debug commands

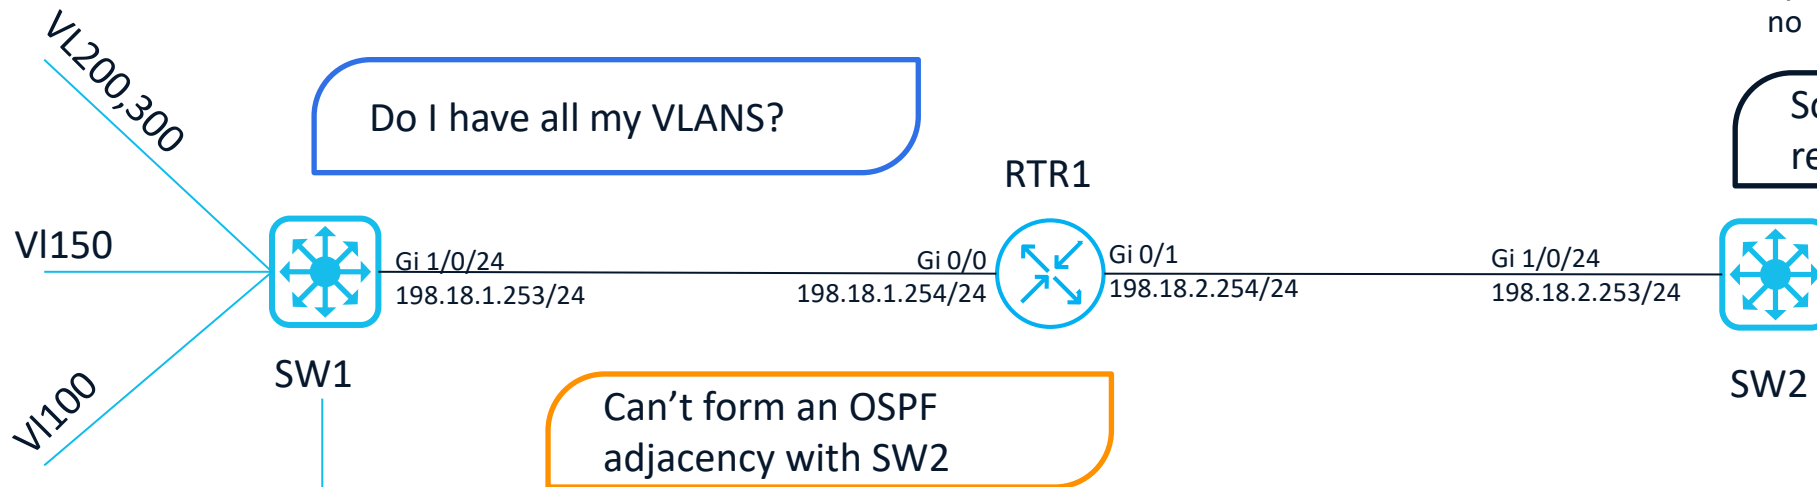
Management / Automation

- Programmability / Scripts / Tools / APIs
- Data Structures (XML, JSON, YAML)
- Telemetry
- Testing / Pipelines (CI/CD)

Understand the life a packet



Is the network configured correctly?



```
SW2# show run int vlan 101
!
interface vlan 101
  ip address 192.18.101.254/24
  no shut
SW2# show run interface g1/0/24
!
interface GigabitEthernet1/0/24
  ip address 192.18.2.126/25
  ip ospf 1 area 10
  no shutdown
```

```
SW1# show vlan
vlan name
-----
100 research
150 data
300 compute
```

```
RTR1# show ip route
C          192.18.1.254
C          192.18.2.254
O          192.18.100.0/24 [110/11] via 192.168.18.1.254
```

```
RTR1# show run interface g0/1
!
interface GigabitEthernet0/1
  ip address 192.18.2.254/24
  ip ospf 10 area 1
  shutdown
```

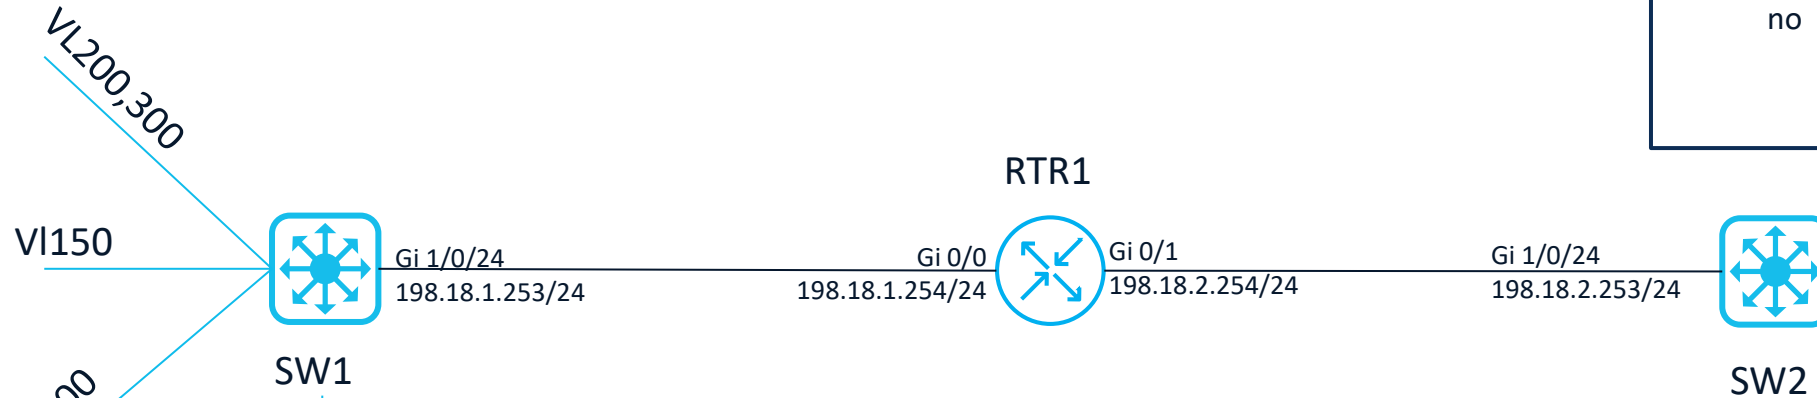
Where are the problems?

Do I have all my VLANS?

Can't form an OSPF adjacency with SW2

Some things wrong, I can't reach the router

4.) Wrong IP or subnet



```
SW2# show run int vlan 101
!
interface vlan 101
 ip address 192.18.101.254/24
 no shut
SW2# show run interface g1/0/24
!
interface GigabitEthernet1/0/24
 ip address 192.18.2.126/25
 ip ospf 1 area 10
 no shutdown
```

SW1# show vlan

```
vlan name
-----
100 research
150 data
300 compute
```

1.) Missing vlan 200

RTR1# show ip route

```
C      192.18.1.254
C      192.18.2.254
O      192.18.100.0/24 [110/11] via 192.168.18.1.254
```

RTR1# show run interface g0/1

```
!
interface GigabitEthernet0/1
 ip address 192.18.2.254/24
 ip ospf 10 area 1
 shutdown
```

2.) Interface in shutdown state

3.) Interface in the wrong area

Find the missing route

5

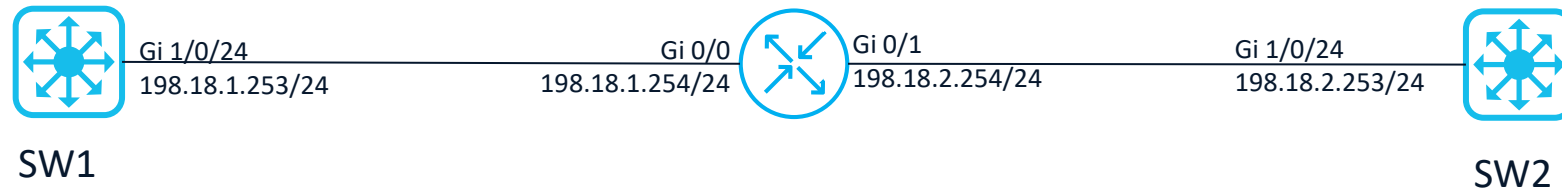
4

3

2

1

Learn network state and compare



Learn details about how the network operates at a point in time

- VLANs
- Interface configuration and state
- Route Tables
- Routing protocols and process details

Compare changes in state from one period to the next

- Do all VLANs exist?
- Did an interface change?
- Did I lose a route?
- Did something in the process or protocol change?

Networking testing gets you faster resolution

```
read-only_playbooks > pyats > state_check > diff_interface_iosxe_browsers-keep_
1 --- ./state/interface_02_28_2025/interface_iosxe_browsers-keep_ops.
2 +++ ./state/interface_03_01_2025/interface_iosxe_browsers-keep_ops.
3 info:
4   Loopback3:
5   - enabled: True
6   + enabled: False
7   ipv6:
8     2001:DB8:CAFE:3033:1000::1/128:
9     - status: valid
10    + status: tentative
11    FE80::21E:E6FF:FE8B:ED00:
12    - status: valid
13    + status: tentative
14    - oper_status: up
15    + oper_status: down
16    - accounting:
17    - ip:
18    -   chars_in: 25608
19    -   chars_out: 25608
20    -   pkts_in: 264
21    -   pkts_out: 264
22    - ipv6:
23    -   chars_in: 0
24    -   chars_out: 1976
25    -   pkts_in: 0
26    -   pkts_out: 26

read-only_playbooks > pyats > state_check > diff_routing_iosxe_browsers-kee
1 --- ./state/routing_02_28_2025/routing_iosxe_browsers-keep_ops.
2 +++ ./state/routing_03_01_2025/routing_iosxe_browsers-keep_ops.
3 info:
4   vrf:
5   default:
6     address_family:
7     ipv4:
8     routes:
9     - 10.10.11.1/32:
10    -   active: True
11    -   metric: 2
12    -   next_hop:
13    -     next_hop_list:
14    -     1:
15    -     index: 1
16    -     next_hop: 10.1.10.0
17    -     outgoing_interface: GigabitEthernet1
18    -     updated: 06:55:53
19    -   route: 10.10.11.1/32
20    -   route_preference: 110
21    -   source_protocol: ospf
22    -   source_protocol_codes: 0
23    - 10.30.33.1/32:
24    -   active: True
25    -   next_hop:
26    -     outgoing_interface:
27    -     Loopback3:
28    -     outgoing_interface: Loopback3
29    -   route: 10.30.33.1/32
30    -   source_protocol: connected
31    -   source_protocol_codes: C

read-only_playbooks > pyats > state_check > diff_ospf_iosxe-ke
1 --- ./state/ospf_02_28_2025/ospf_iosxe_browsers-ke
2 +++ ./state/ospf_03_01_2025/ospf_iosxe_browsers-ke
3 info:
4   vrf:
5   default:
6     address_family:
7     ipv4:
8     instance:
9     100:
10    areas:
11    0.0.0.0:
12    interfaces:
13    Loopback3:
14    - enable: True
15    + enable: False
```

Caused by a disabled interface

Route Table Changes

- enabled: True
+ enabled: False

- oper_status: up
+ oper_status: down

- enable: True
+ enable: False

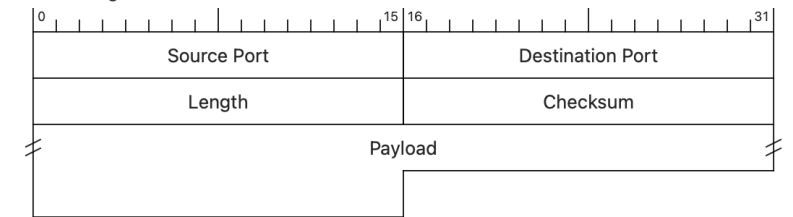
Packet Capture

```

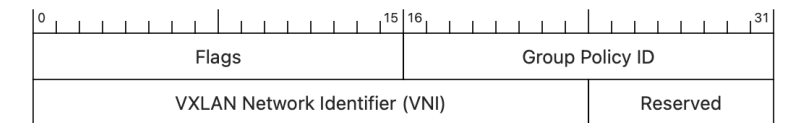
> Frame 333: Packet, 140 bytes on wire (1120 bits), 140 bytes captured (1120 bits)
√ Ethernet II, Src: aa:bb:cc:00:96:20 (aa:bb:cc:00:96:20), Dst: aa:bb:cc:00:99:01 (aa:bb:cc:00:99:01)
  > Destination: aa:bb:cc:00:99:01 (aa:bb:cc:00:99:01)
  > Source: aa:bb:cc:00:96:20 (aa:bb:cc:00:96:20)
  Type: IPv4 (0x0800)
  [Stream index: 0]
√ Internet Protocol Version 4, Src: 10.168.100.101, Dst: 10.168.100.103
  0100 .... = Version: 4
  .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
  Total Length: 126
  Identification: 0x7dd4 (32212)
  > 000. .... = Flags: 0x0
  ...0 0000 0000 0000 = Fragment Offset: 0
  Time to Live: 254
  Protocol: UDP (17)
  Header Checksum: 0x607e [validation disabled]
  [Header checksum status: Unverified]
  Source Address: 10.168.100.101
  Destination Address: 10.168.100.103
  [Stream index: 0]
> User Datagram Protocol, Src Port: 4789, Dst Port: 4789
> Virtual eXtensible Local Area Network
√ Ethernet II, Src: aa:bb:cc:80:9a:00 (aa:bb:cc:80:9a:00), Dst: aa:bb:cc:80:99:00 (aa:bb:cc:80:99:00)
  > Destination: aa:bb:cc:80:99:00 (aa:bb:cc:80:99:00)
  > Source: aa:bb:cc:80:9a:00 (aa:bb:cc:80:9a:00)
  Type: IPv4 (0x0800)
  [Stream index: 1]
√ Internet Protocol Version 4, Src: 172.16.1.210, Dst: 172.17.1.10
  0100 .... = Version: 4
  .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
  Total Length: 76
  Identification: 0x4372 (17266)
  > 010. .... = Flags: 0x2, Don't fragment
  ...0 0000 0000 0000 = Fragment Offset: 0
  Time to Live: 63
  Protocol: UDP (17)
  Header Checksum: 0x9d31 [validation disabled]
  [Header checksum status: Unverified]
  Source Address: 172.16.1.210
  Destination Address: 172.17.1.10
  [Stream index: 1]
> User Datagram Protocol, Src Port: 123, Dst Port: 60442
> Network Time Protocol (NTP Version 4, server)

```

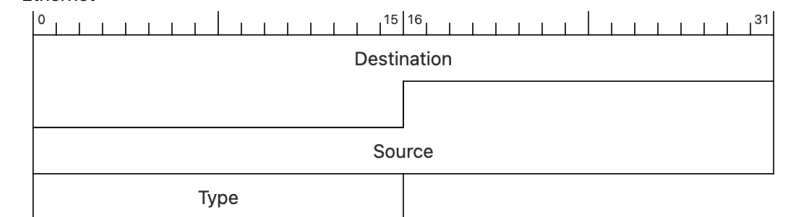
User Datagram Protocol



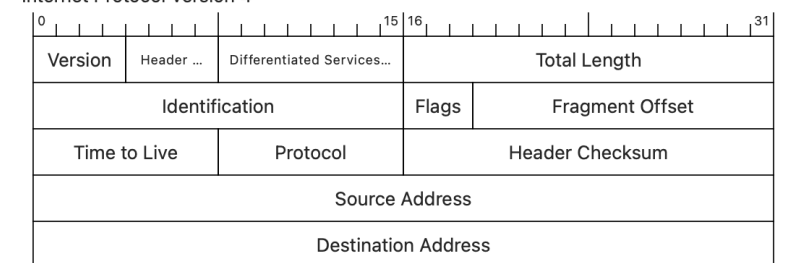
Virtual eXtensible Local Area Network



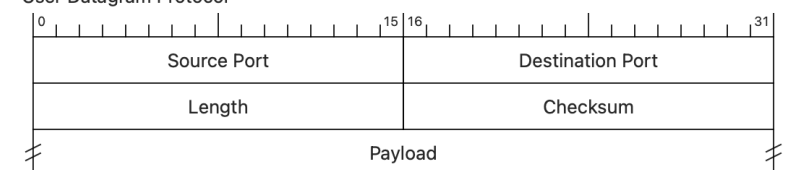
Ethernet



Internet Protocol Version 4



User Datagram Protocol



Data types

Unstructured CLI data

```
interface GigabitEthernet8
vrf forwarding management
ip address 172.16.1.101 255.255.255.0
negotiation auto
no mop enabled
no mop sysid
end
```

Structured JSON data

```
{
  "Cisco-IOS-XE-native:GigabitEthernet": [
    {
      "name": "8",
      "vrf": {
        "forwarding": "management"
      },
      "ip": {
        "address": {
          "primary": {
            "address": "172.16.1.101",
            "mask": "255.255.255.0"
          }
        }
      },
      "mop": {
        "enabled": false,
        "sysid": false
      },
      "Cisco-IOS-XE-ethernet:negotiation": {
        "auto": true
      }
    }
  ]
}
```

Structured XML data

```
<GigabitEthernet>
  <name>8</name>
  <vrf>
    <forwarding>management</forwarding>
  </vrf>
  <ip>
    <address>
      <primary>
        <address>172.16.1.101</address>
        <mask>255.255.255.0</mask>
      </primary>
    </address>
  </ip>
  <mop>
    <enabled>false</enabled>
    <sysid>false</sysid>
  </mop>
  <negotiation>
    <auto>true</auto>
  </negotiation>
</GigabitEthernet>
```

Structured data is more efficient to use...

CLI

```
c8kv-internet#sho run int g 8 | inc address  
ip address 172.16.1.101 255.255.255.0
```

As good as it gets without
additional text manipulation

Both XML and JSON outputs can be further parsed for specific values

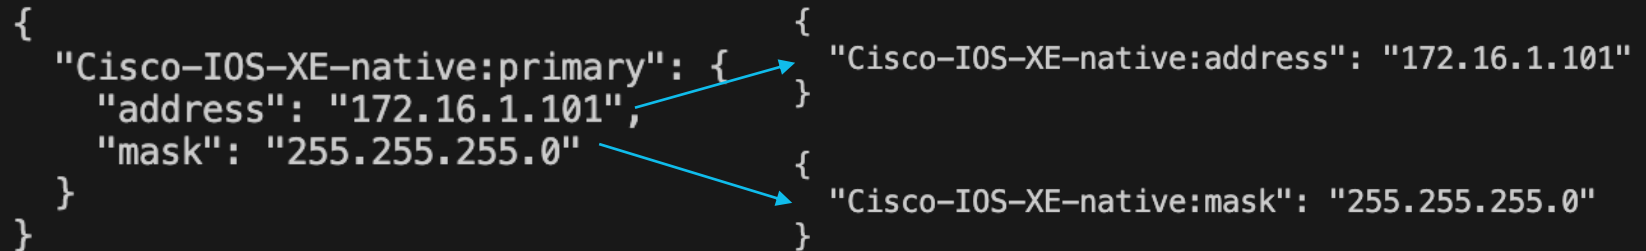
XML

```
<primary>  
  <address>172.16.1.101</address>  
  <mask>255.255.255.0</mask>  
</primary>
```



JSON

```
{  
  "Cisco-IOS-XE-native:primary": {  
    "address": "172.16.1.101",  
    "mask": "255.255.255.0"  
  }  
}
```



and we can do more with the data

Device Details

Device Name: c8kv-internet

MGMT IP Address: 172.16.1.101
Operating System: IOS-XE 17.13
Domain Name: automation.lab
Model: C8000V

Interfaces:

Interface	Status	IP Address	Dot1Q VLAN	Description
GigabitEthernet 1	Up	DHCP configured		WAN/NAT interfaces (DHCP)
GigabitEthernet 2	Up	192.168.100.254		Link to n9kv-sw1
GigabitEthernet 3	Up			ROAS interface
GigabitEthernet 3.100	Up	10.100.1.254	100	
GigabitEthernet 3.200	Up	10.200.1.254	200	
GigabitEthernet 4	Down			
GigabitEthernet 5	Down			
GigabitEthernet 6	Down			
GigabitEthernet 7	Down			
GigabitEthernet 8	Up	172.16.1.101		

AI in networking



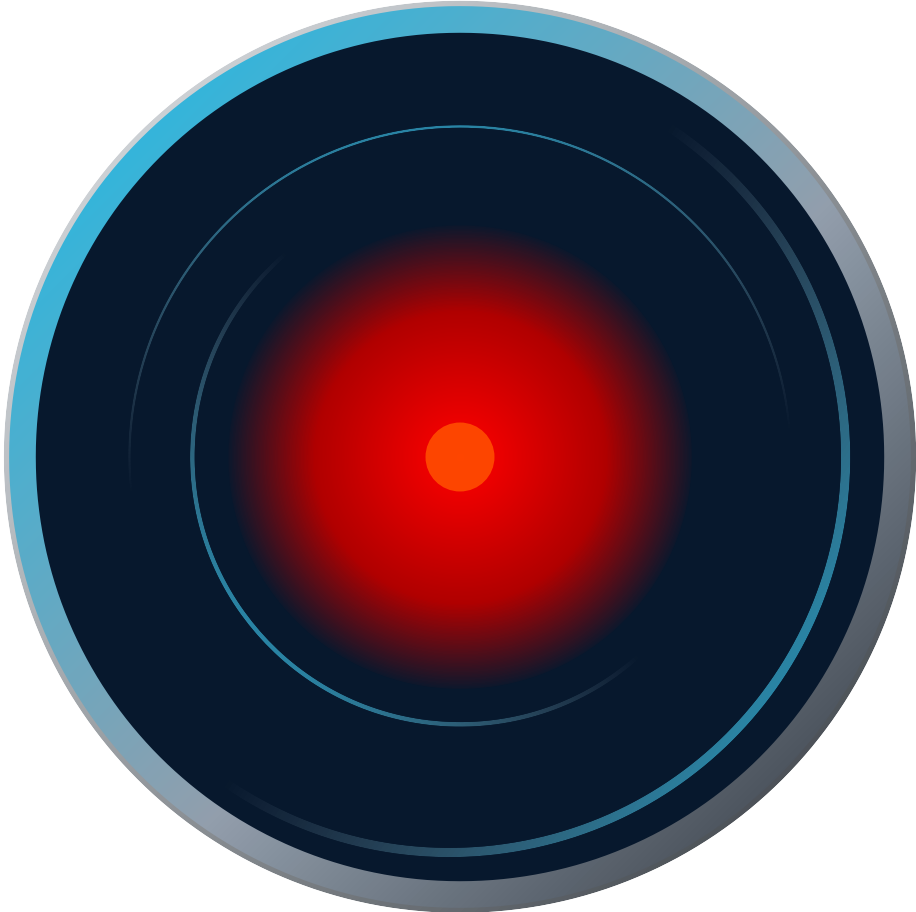
“AI is not a tool. AI is a worker, using a tool”

Jensen Huang @ NVIDIA GTC DC 2025 Keynote

Some of the AI workers available commercially



ChatGPT



 Claude

 Gemini



From assistants to agents

Agentic AI

- Has some level of autonomy to act on tasks, within specific guardrails, using provided or available tools.
- For IT, networking, security can help manage increasing alert fatigue.

Prompting still matters; agent scope is critical.

- Defining things like role, process, or how it should output data.
- Sets boundaries on the agent. Limits risk, set guardrails, protects against failure/surprising activity.

Model awareness

- Know what a model can and can't do reliably from a domain capabilities standpoint.

Agentic Automation

AI needs the right tools to do the job

- Netmiko, NAPALM, pyATS, Ansible allow agents to interact with network devices, learn state, and make changes.
- MCP (Model Context Protocol) is becoming the connective tissue between models and our tools, data, and devices.

Agentic workflows bring new challenges

- Once an agent can call APIs and push config, its tool access becomes an attack surface.
- Scope agents with least privilege, gate changes behind approval, and log every tool call.

Where AI helps in networking today

Use it to learn
networking
(or anything IT
related)

Use it as a
troubleshooting
aid

Use it to help you
understand
configurations

Allowing agents to
operate devices

(within specified
guardrails)

Use it to learn networking



You 8:32 PM

Help me learn computer networking



You 8:40 PM

Why can't a 100Gbps QSFP be used as 10x10G?



You 8:37 PM

How does a switch ASIC work?



You 8:41 PM

Can we get rid of networks entirely?

Learn concepts, and identify areas to deep dive.



CIRCUIT 8:33 PM Edited

Computer networking is the practice of connecting two or more computing devices to enable them to communicate and share resources. This interconnected system allows for the transmission and exchange of information, forming the backbone of modern digital communication, including the internet.

To understand computer networking, it's helpful to grasp several fundamental concepts:

Basic Concepts

- **What is a Network?** A network is a collection of autonomous computers and other devices (nodes) linked together to exchange data. These nodes can include computers, printers, or any device capable of sending or receiving data.
- **Types of Networks:** Networks are often categorized by the geographical area they cover:
 - **Personal Area Network (PAN):** A short-range network connecting personal devices like smartphones, tablets, and computers, typically covering less than 10 meters (about 33 feet) using technologies like Bluetooth or Wi-Fi.
 - **Local Area Network (LAN):** Connects computers and devices within a small, confined area such as a home, office, school, or hospital, often using Wi-Fi or Ethernet cables. LANs allow for resource sharing like printers and files.
 - **Metropolitan Area Network (MAN):** Spans a larger area than a LAN, typically covering a city or a large campus, ranging from 5–50 km.
 - **Wide Area Network (WAN):** Covers large geographical areas, connecting multiple LANs across cities, states, or countries, often using telephone lines, radio waves, or satellites. The internet is the largest WAN.
 - **Wireless Local Area Network (WLAN):** A type of LAN that uses wireless communication instead of cables.
 - **Virtual Private Network (VPN):** Extends a private network across a public network, enabling users to send and receive data securely as if their devices were directly connected to the private network. VPN protocols create secure, encrypted connections over the public internet.
- **Network Devices:** Essential components that facilitate communication within a network include:
 - **Routers:** Route data from one network to another.
 - **Switches:** Connect devices within a network.
 - **Network Interface Card (NIC):** Creates a physical connection to the network for each device (node).
 - **IP Address:** A unique number assigned to every network device in an Internet Protocol (IP) network, identifying its host network and location.
 - **MAC Address:** A unique identifier assigned to a NIC by its manufacturer.
- **Network Topologies:** Refer to the arrangement of hosts and hardware within a network architecture.



CIRCUIT 8:40 PM Edited

A 100Gbps QSFP (Quad Small Form-factor Pluggable) transceiver generally cannot be used as 10x10G because of its internal architecture and the way its optical or electrical lanes are designed.

Here's a breakdown of why this is the case:

1. Lane Configuration of 100Gbps QSFP (QSFP28):

- Most commonly, 100Gbps QSFP modules, specifically QSFP28, are designed with **four electrical lanes**, each operating at **25Gbps** ($4 \times 25\text{Gbps} = 100\text{Gbps}$).
- These four lanes can be transmitted over four separate optical fibers (e.g., using MPO/MTP connectors) or aggregated onto a single pair of fibers using wavelength division multiplexing (WDM) technologies.
- Some 100Gbps modules might also use two lanes at 50Gbps ($2 \times 50\text{G}$).

2. Physical and Electrical Limitations:

- The internal circuitry and optical components within a QSFP28 module are engineered to handle four 25Gbps streams. To break out into 10 individual 10Gbps links, the module would need to have 10 distinct electrical and optical channels, which it typically does not possess.
- The connector interface on the QSFP28 module itself is designed for 4 lanes, not 10.

3. Breakout Capabilities:

- A common breakout scenario for a 100Gbps QSFP28 module is to split it into **4x25Gbps SFP28** connections using a QSFP28-to-4xSFP28 breakout cable. This is possible because the module inherently has four 25Gbps lanes.
- For **40Gbps QSFP+** modules, which typically operate as 4x10Gbps, it is possible to break them out into **4x10Gbps SFP+** connections. This is a frequent use case for connecting a 40G switch port to four 10G devices.

4. Achieving 10x10G:

- To get 10x10G connections, you would typically need to aggregate multiple 10Gbps SFP+ transceivers (ten of them) into a higher-density switch or use a different type of module or switch architecture that supports such a configuration.
- There are no standard QSFP form factors that natively support a 10x10G breakout from a single module.

As a troubleshooting aid



You 8:47 PM

Tell me about this error %BGP_SESSION-5-ADJCHANGE: neighbor [192.168.100.201](#) L2VPN Evpn topology base removed from session



CIRCUIT 8:47 PM Edited

The log message `%BGP_SESSION-5-ADJCHANGE: neighbor 192.168.100.201 L2VPN Evpn topology base removed from session` indicates a change in the BGP (Border Gateway Protocol) session with a specific neighbor, specifically concerning the exchange of EVPN (Ethernet VPN) routing information.

Let's break down each part of this Cisco IOS/IOS-XE log message:

- **%BGP_SESSION-5-ADJCHANGE:**
 - **%BGP_SESSION:** This is the facility code, indicating that the message is related to the BGP session process.
 - **5:** This is the severity level, which is "Notification." This means it's an informational message about a significant but often normal event, not necessarily an error that requires immediate intervention unless it's flapping frequently.
 - **ADJCHANGE:** This is the mnemonic, standing for "Adjacency Change," which means there has been a change in the state or capabilities of the BGP session with a neighbor.
- **neighbor 192.168.100.201:** This specifies the IP address of the BGP peer with which the adjacency change occurred.
- **L2VPN Evpn topology base removed from session:** This is the core of the message and tells you precisely what changed:
 - **L2VPN Evpn:** This refers to the Layer 2 VPN Ethernet VPN address family. EVPN is a technology that allows you to extend Layer 2 connectivity over an IP or MPLS network, often used in data centers and service provider environments. BGP is used to exchange the routing information (MAC addresses, IP addresses, VNI mappings, etc.) for EVPN.
 - **topology base removed from session:** This indicates that the capability to exchange EVPN routes (specifically the "topology base," which refers to the fundamental EVPN routing information) has been withdrawn or is no longer active for this BGP session with neighbor `192.168.100.201`.

Understanding Configurations



You 10:26 AM

Can you briefly summarize what this configuration is doing?

```
interface Ethernet1/1
no switchport
ip unnumbered Loopback0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 100 area 0
!
```



CIRCUIT 10:26 AM Edited

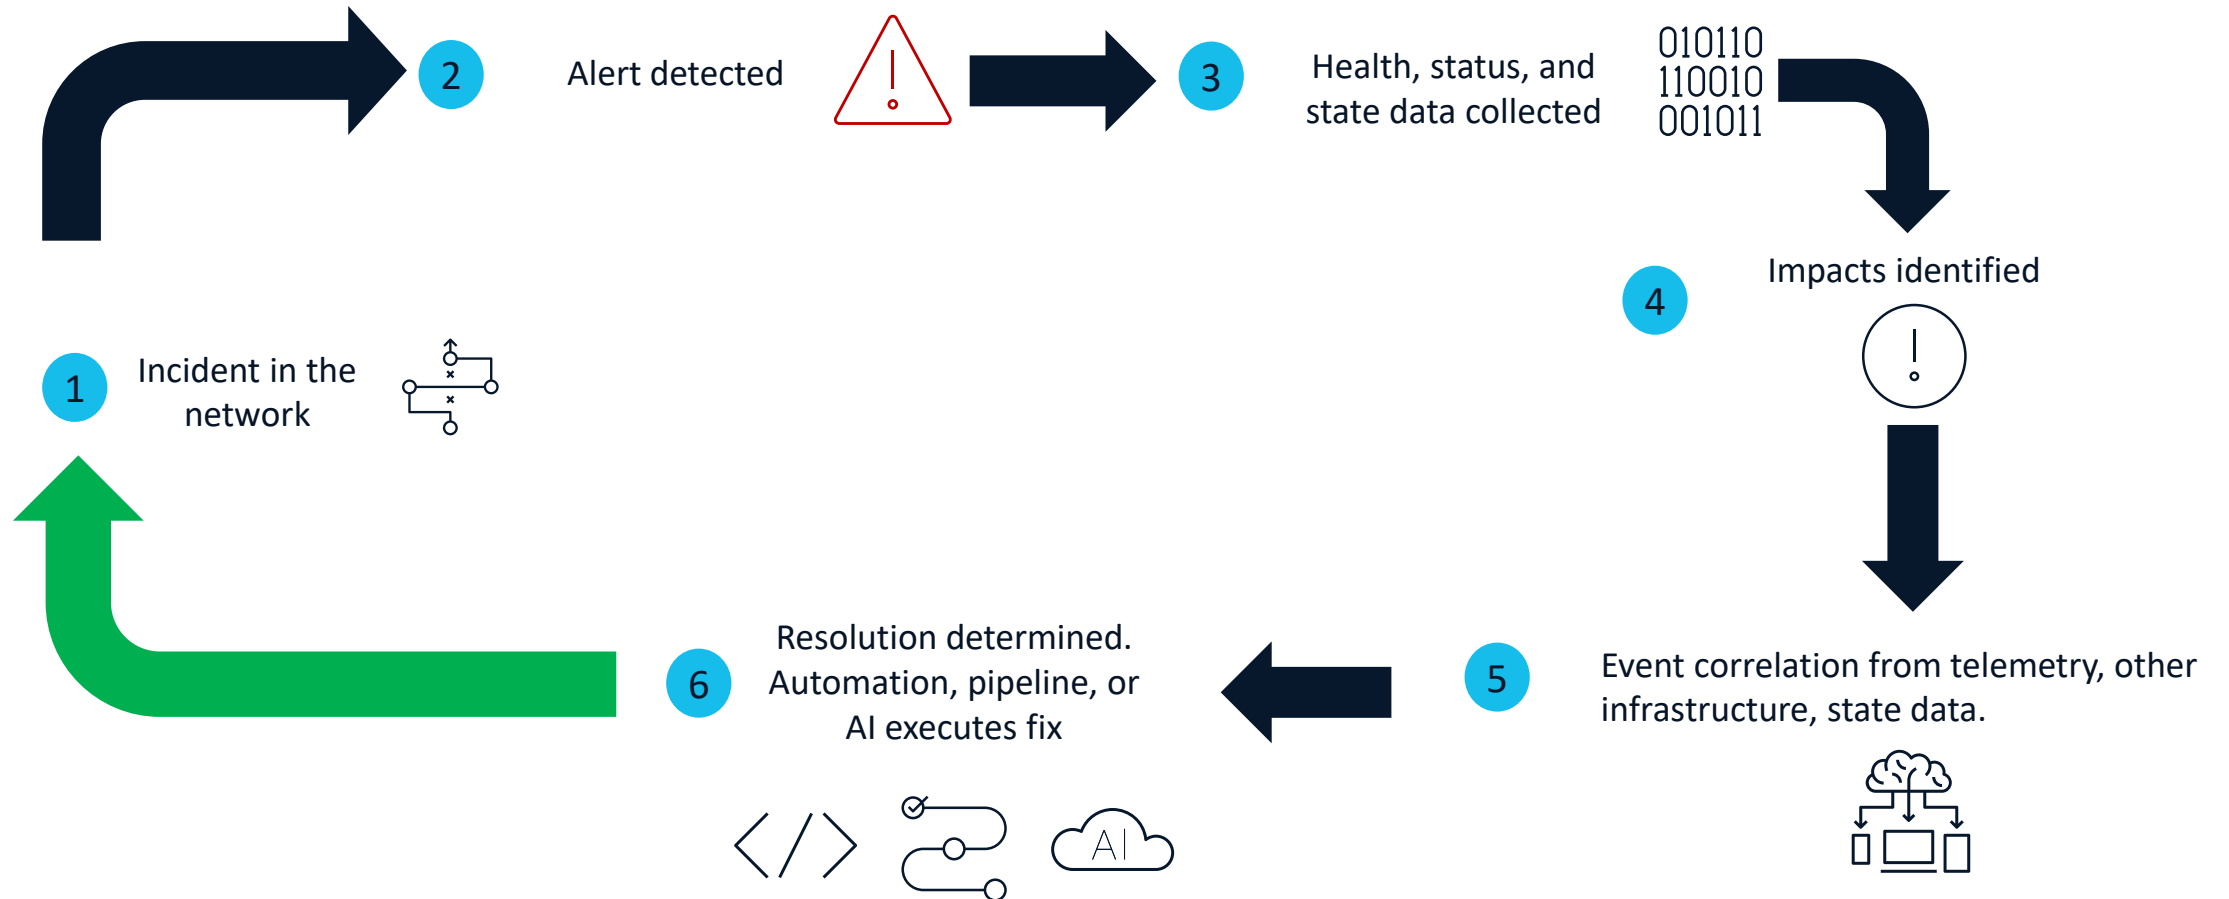
This configuration snippet sets up the `Ethernet1/1` interface as a **Layer 3 routed port** with the following key functions:

- It **borrows its IP address** from the `Loopback0` interface, a common practice to conserve IP addresses and ensure interface stability.
- It enables **PIM (Protocol Independent Multicast) in sparse mode** on this interface, allowing it to participate in multicast routing.
- It configures **OSPF (Open Shortest Path First) routing** on the interface, specifically treating it as a **point-to-point network type** and assigning it to **OSPF process 100 in Area 0 (the OSPF backbone)**.

How is AI going to change things?

Where is this going?

Closed loop automation - Root Cause Analysis and closed loop fix



Securing Infrastructure

The capability is real, not hypothetical

- Anthropic's Mythos is changing the cybersecurity landscape
 - Chained together vulnerabilities to exploit a browser.
https://www.anthropic.com/research/mythos-preview?utm_source=tdrai
- Researcher used Claude Opus 4.6 and found a 23-year-old bug in Linux Kernel
 - <https://www.redhat.com/en/blog/navigating-mythos-haunted-world-platform-security>

What it means for us

- Attackers get the same capability. AI reduces the barrier to entry to creating attacks on IT infrastructure.
- Network and server infrastructure will have more frequent patching.

Final thoughts

Useful resources

- Cisco Modeling Labs Free Tier!
 - <https://software.cisco.com/download/home/286193282/type/286326381/release/CML-Free>
- Cisco Live – On Demand Library
 - <https://www.ciscolive.com/on-demand/on-demand-library.html>
- Cisco U. – Occasional free RevUp to ReCert content
 - <https://u.cisco.com>
 - Current RevUp to ReCert material: **Designing Cisco Security Infrastructure (41 CEs), thru July 13th 2026!**
- Cisco Developer (formerly DevNet)
 - <https://developer.cisco.com>

Keep in touch

LinkedIn

- <https://www.linkedin.com/in/matt-leuschner/>

Twitter/X

- https://x.com/rogue_router

YouTube

- https://www.youtube.com/@stayfresh_networks

Email

- mleuschn@cisco.com

