The bridge to possible

Software-Defined Networking

Cristian Perissinotto
Technical Solution Architect
May 2022

Agenda

- SDN Introduction
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Agenda

- SDN Introduction
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Traditional Networking Paradigm

Control and Data Plane resides within Physical Device

Control Plane and Data Plane

Processing Plane	Where it runs	How fast these processes run	Type of processes performed
Control Plane	Switch CPU	In the order of thousands of packets per second	Routing protocols (i.e. OSPF, IS-IS, BGP), Spanning Tree, SYSLOG, AAA (Authentication Authorization Accounting), NDE (Netflow Data Export), CLI (Command Line interface), SNMP
Data Plane	Dedicated Hardware ASIC's	Millions or Billions of packets per second	Layer 2 switching, Layer 3 (IPv4 | IPv6) switching, MPLS forwarding, VRF Forwarding, QOS (Quality of Service) Marking, Classification, Policing, Netflow flow collection, Security Access Control Lists

What is Software-Defined Networking?

- SDN attempts to centralize network intelligence in one network component by disassociating the forwarding process of network packets (data plane) from the routing process (control plane)

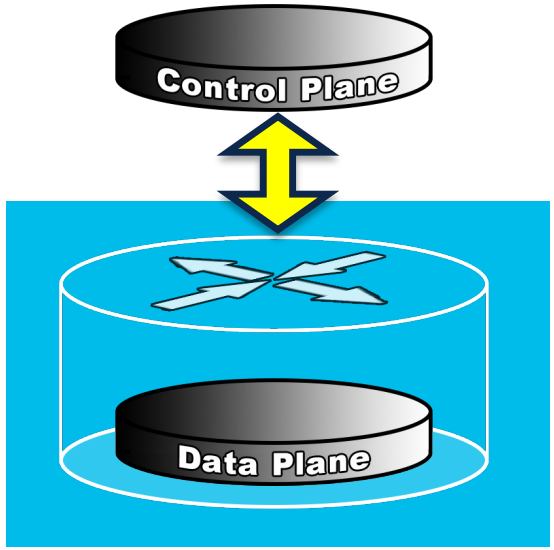

Control Plane runs external to the device in a central location, managing multiple devices

Agenda

- SDN Introduction
  - History
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

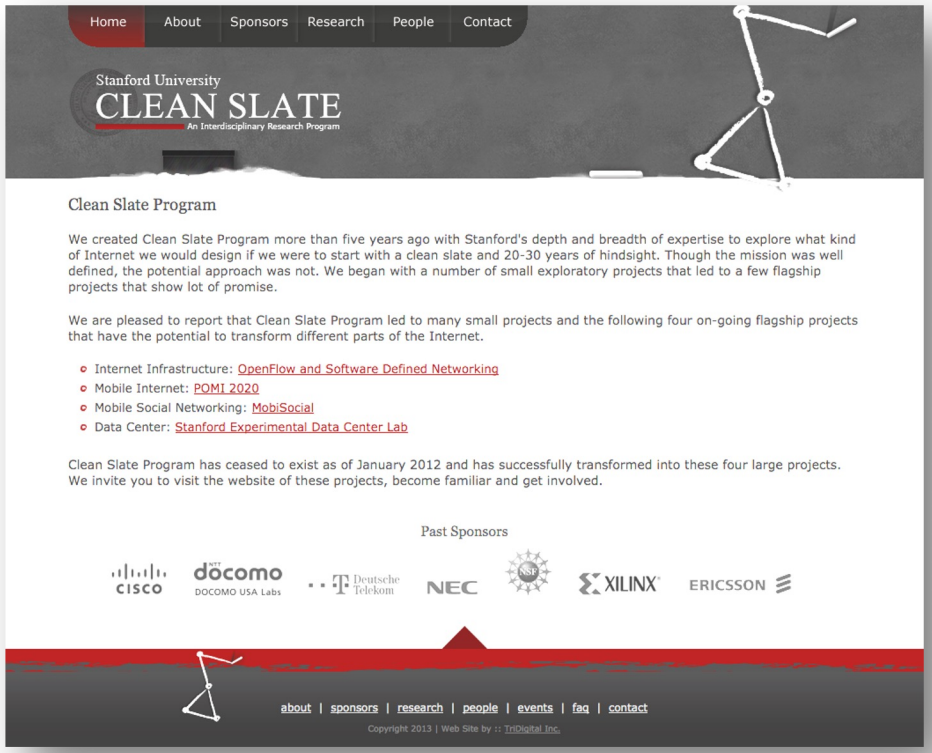

Stanford University – Clean Slate Project

"…explore what kind of Internet we would design if we were to start with a clean slate and 20-30 years of hindsight."

http://cleanslate.stanford.edu/

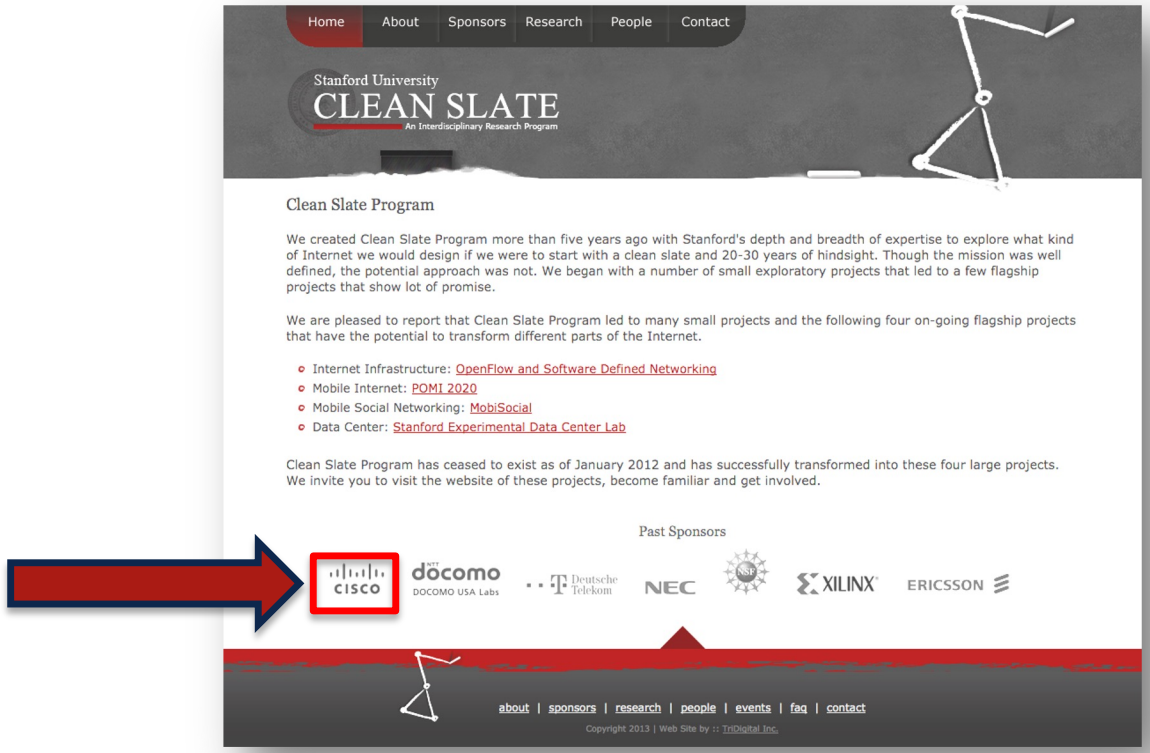

You might have noticed the Cisco Logo on the web page

Cisco R&D teams were engaged with Clean Slate since early days …

… Clean Slate led to the development of…

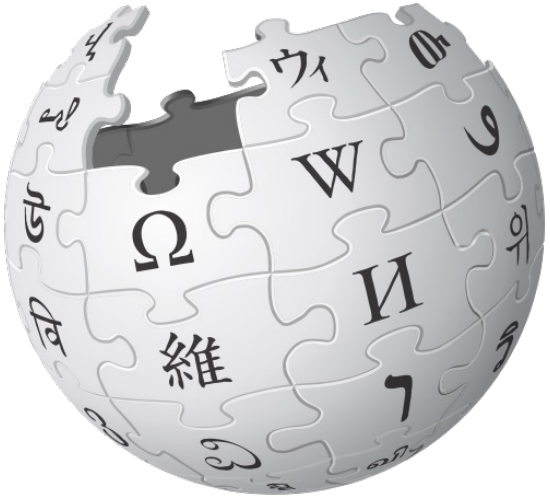

What is Openflow?

(per Wikipedia definition)

OpenFlow is a Layer 2 communications protocol that gives access to the forwarding plane of a network switch or router over the network

Open Flow Architecture

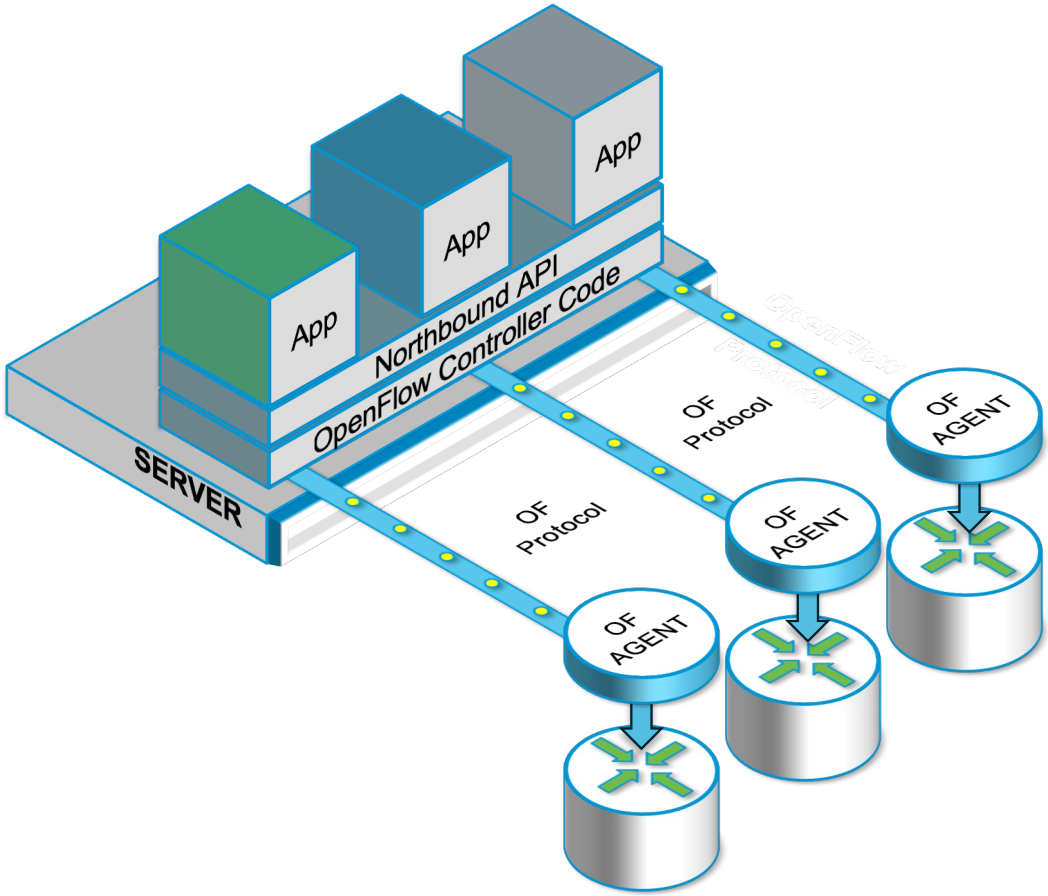

Open Flow Architecture includes four components

Open Flow Architecture

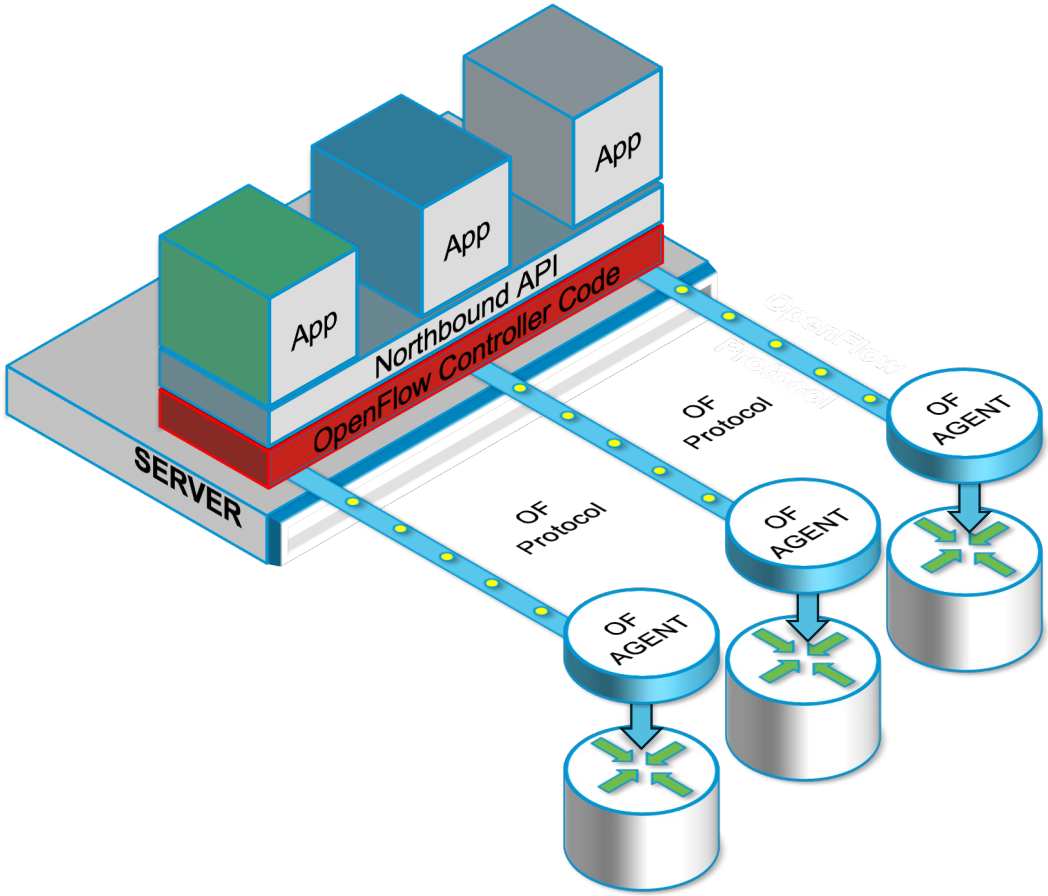

Open Flow Controller:

- Resides on a server
- Central administration and operations point for network elements
- Provides control plane functions for the network

Open Flow Architecture

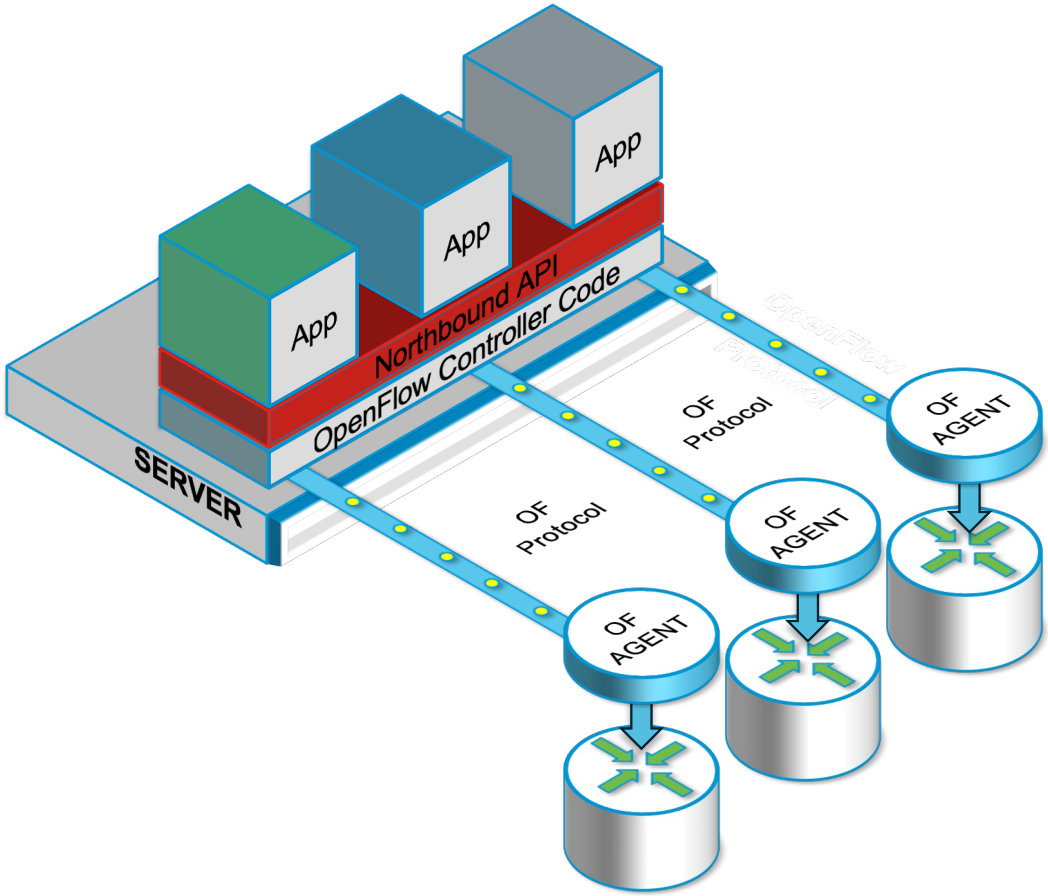

Northbound API:

- Integral part of the controller
- “Network enabled” applications can make use of Northbound API to request services from the network

Open Flow Architecture

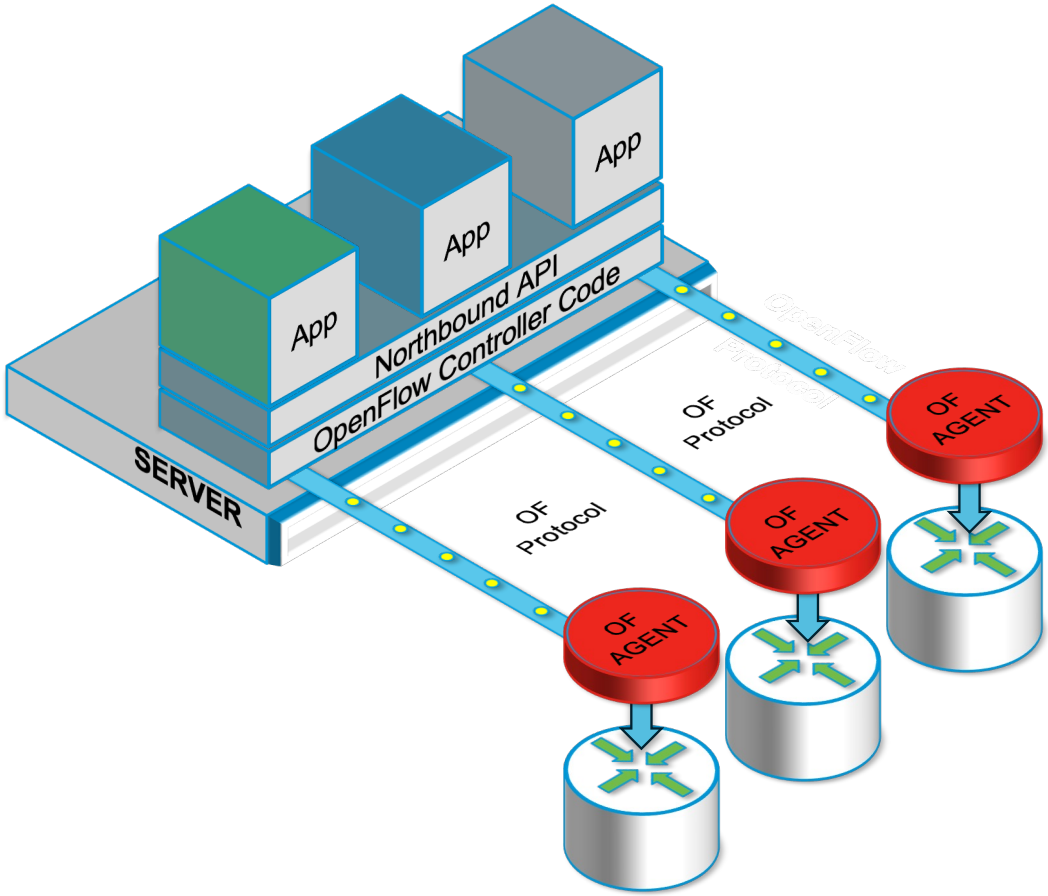

Openflow Device Agent:

- Run on the network device
- Receive instructions from Controller
- Program device tables

Open Flow Architecture

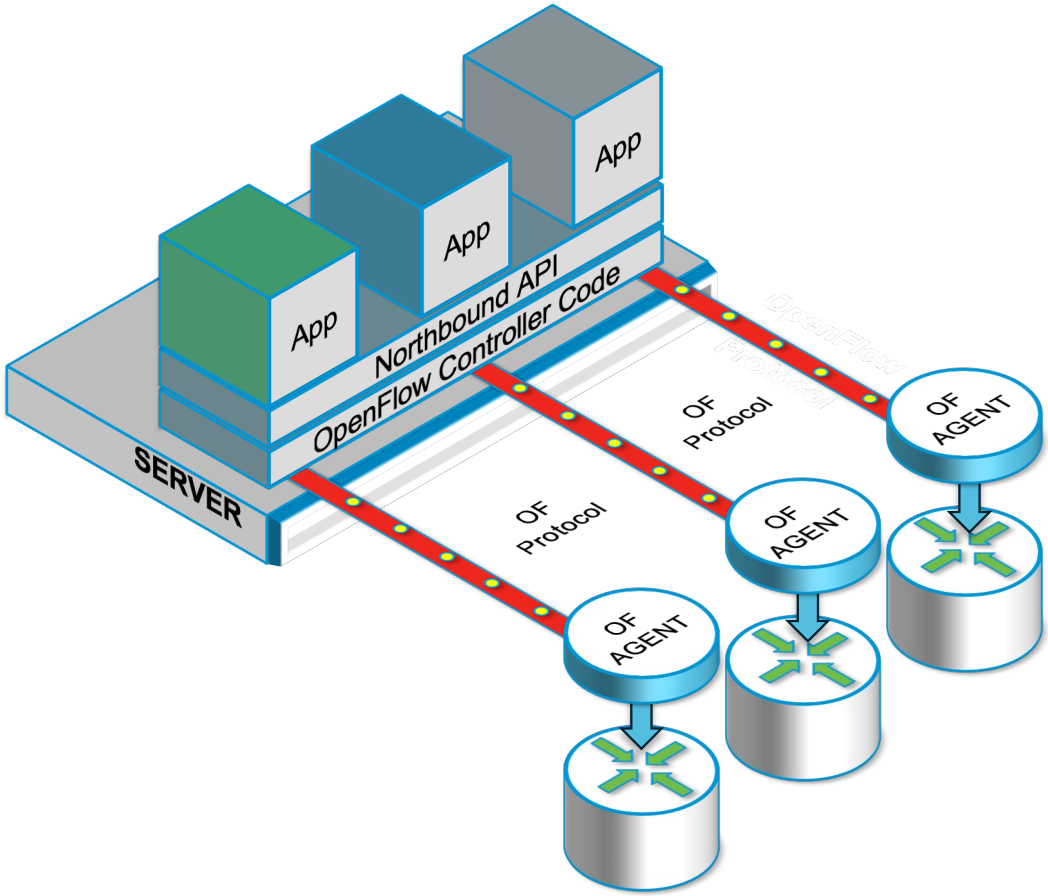

Openflow Protocol:

- The mechanism for the Openflow Controller to communicate with Openflow Agents

Openflow 1.0 Operation

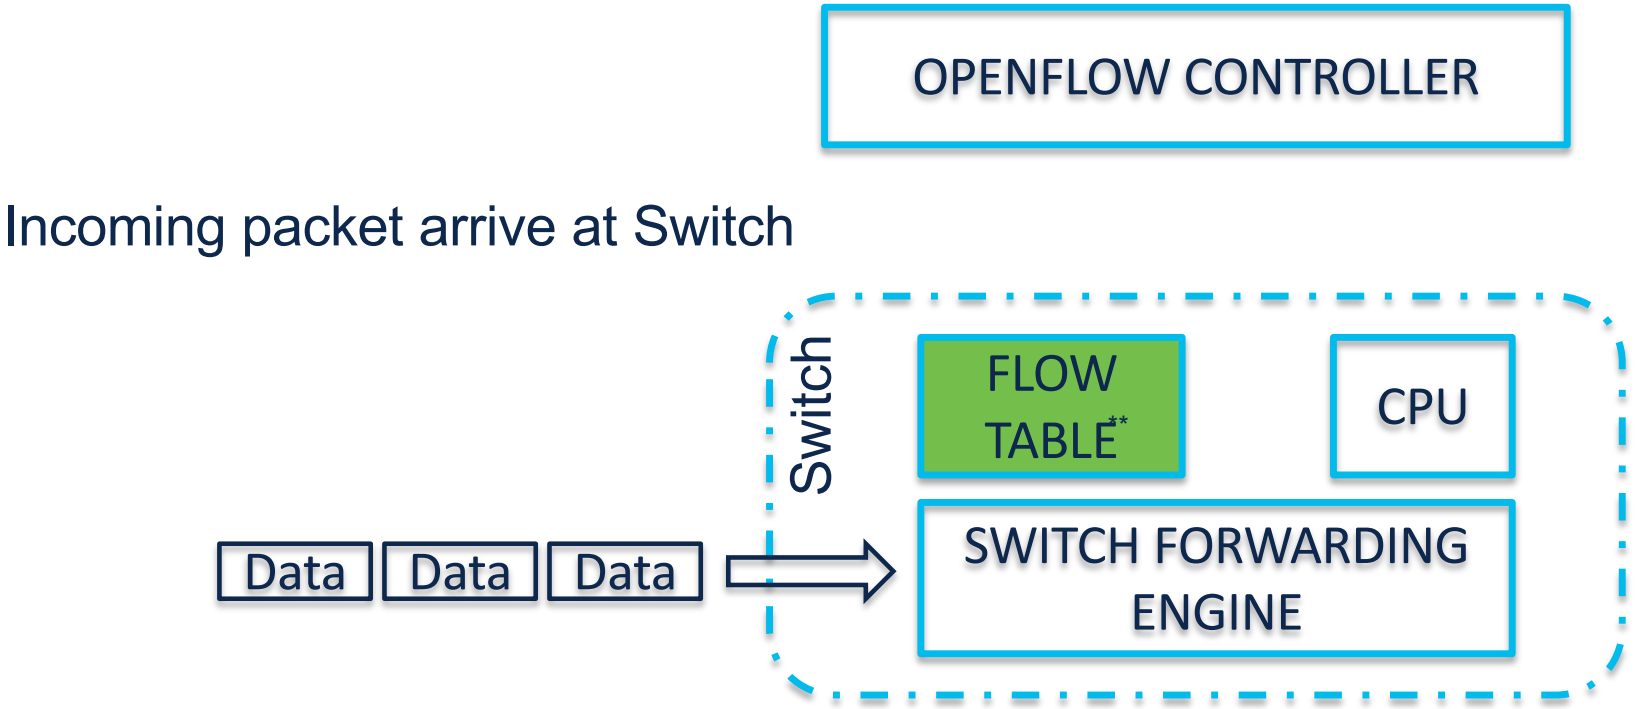

**Openflow 1.0 supports a lookup into a single flow table

Openflow 1.0 Operation

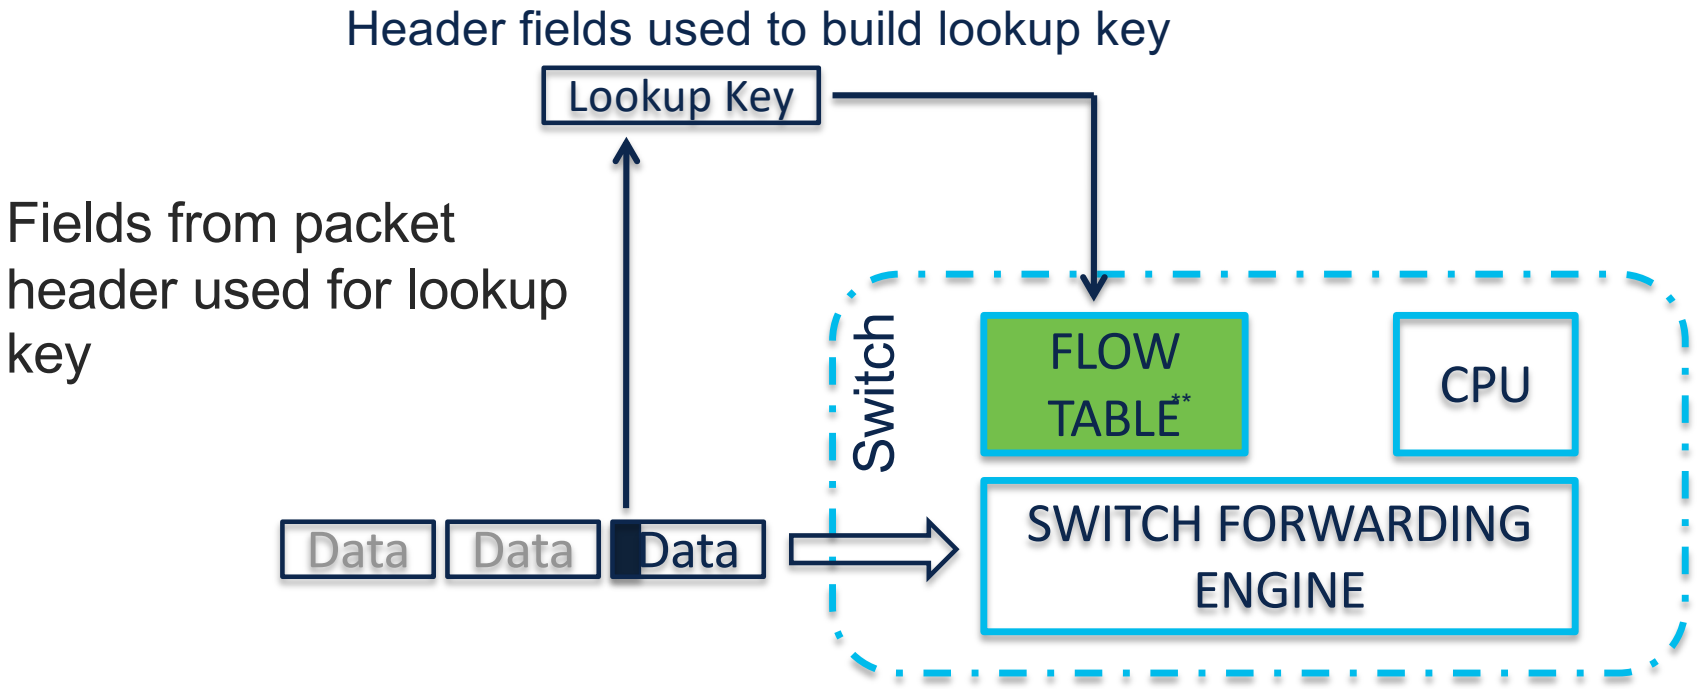

**Openflow 1.0 supports a lookup into a single flow table

Openflow 1.0 Operation

If no match, Controller programs switch flow table

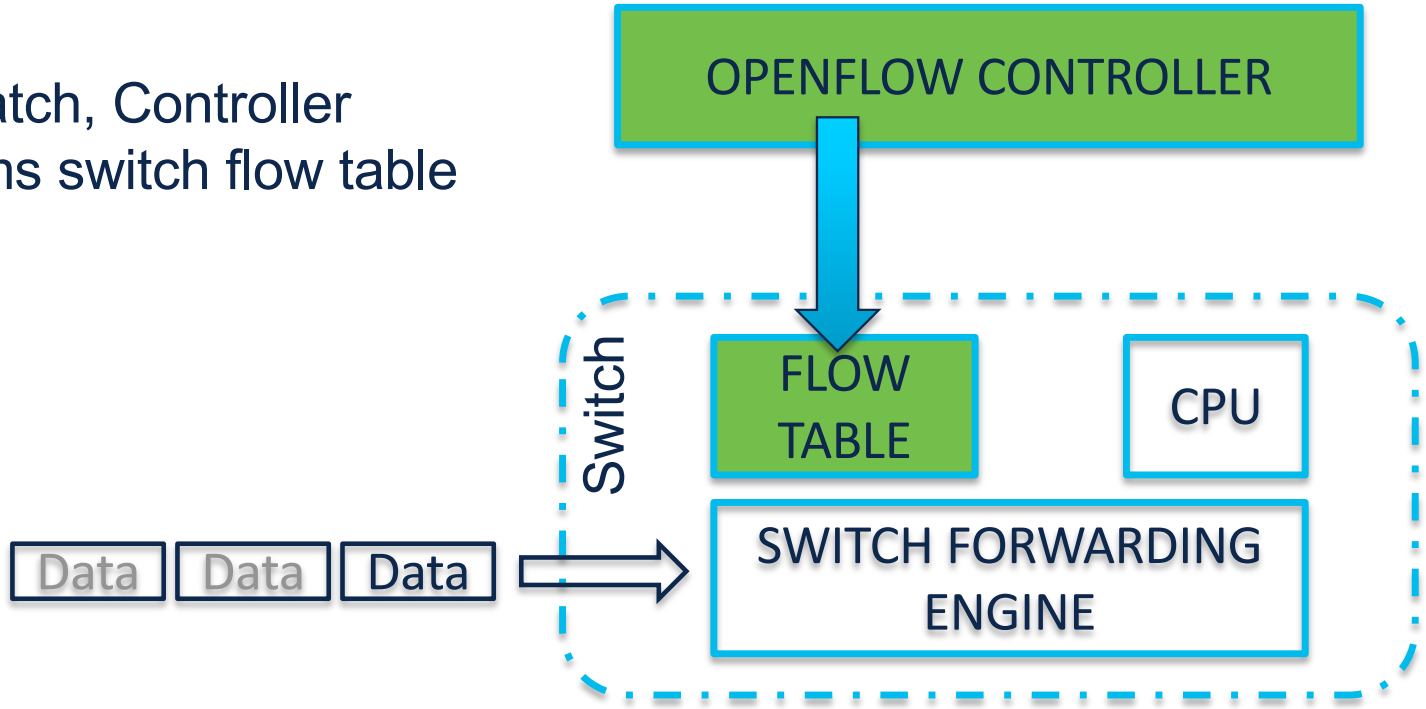

Openflow 1.0 Operation

Forwarding Engine forwards packets

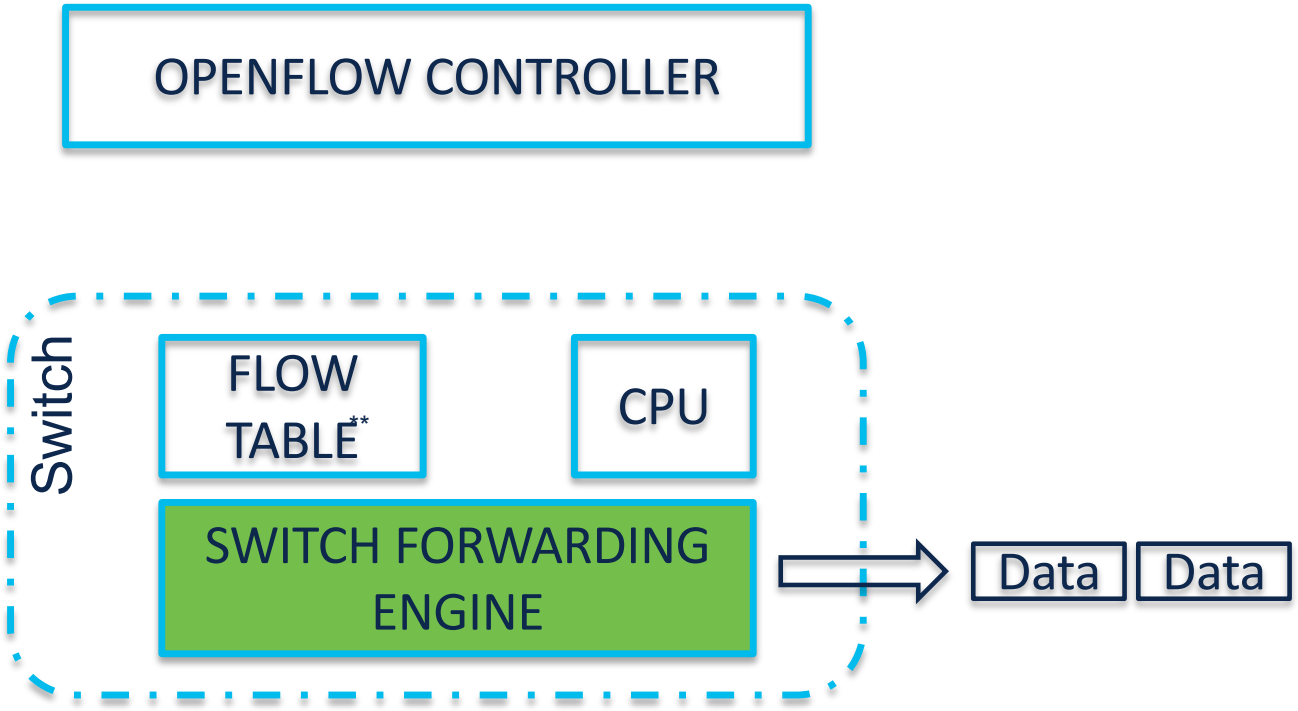

**Openflow 1.0 supports a lookup into a single flow table

Openflow 1.0 Operation

Flow Table in more detail…

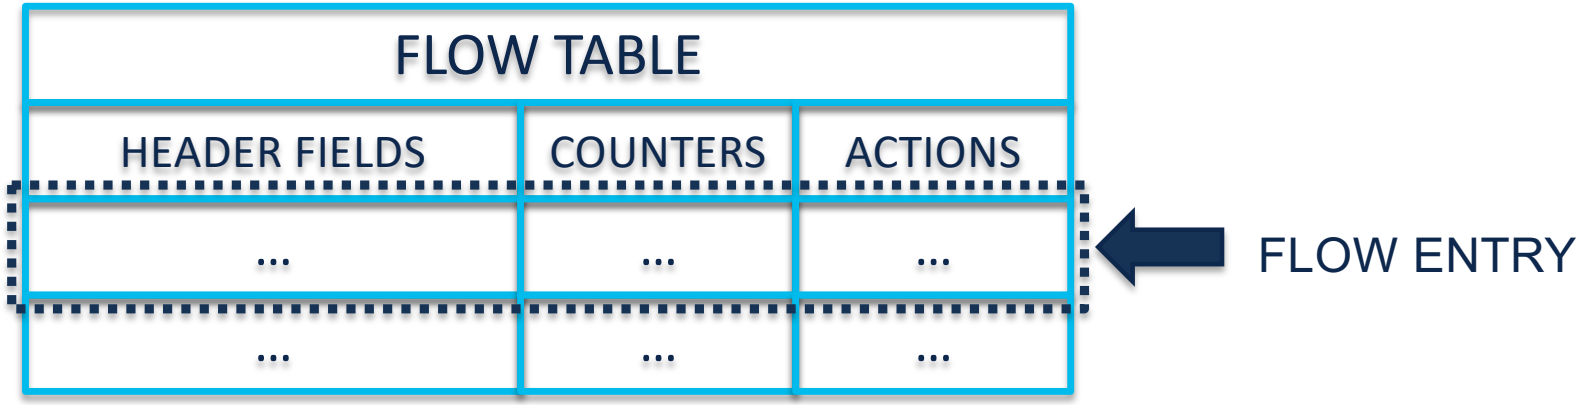

Flow “Entry” consists of one row in the Flow Table

Openflow 1.0 Operation

Flow Table in more detail…

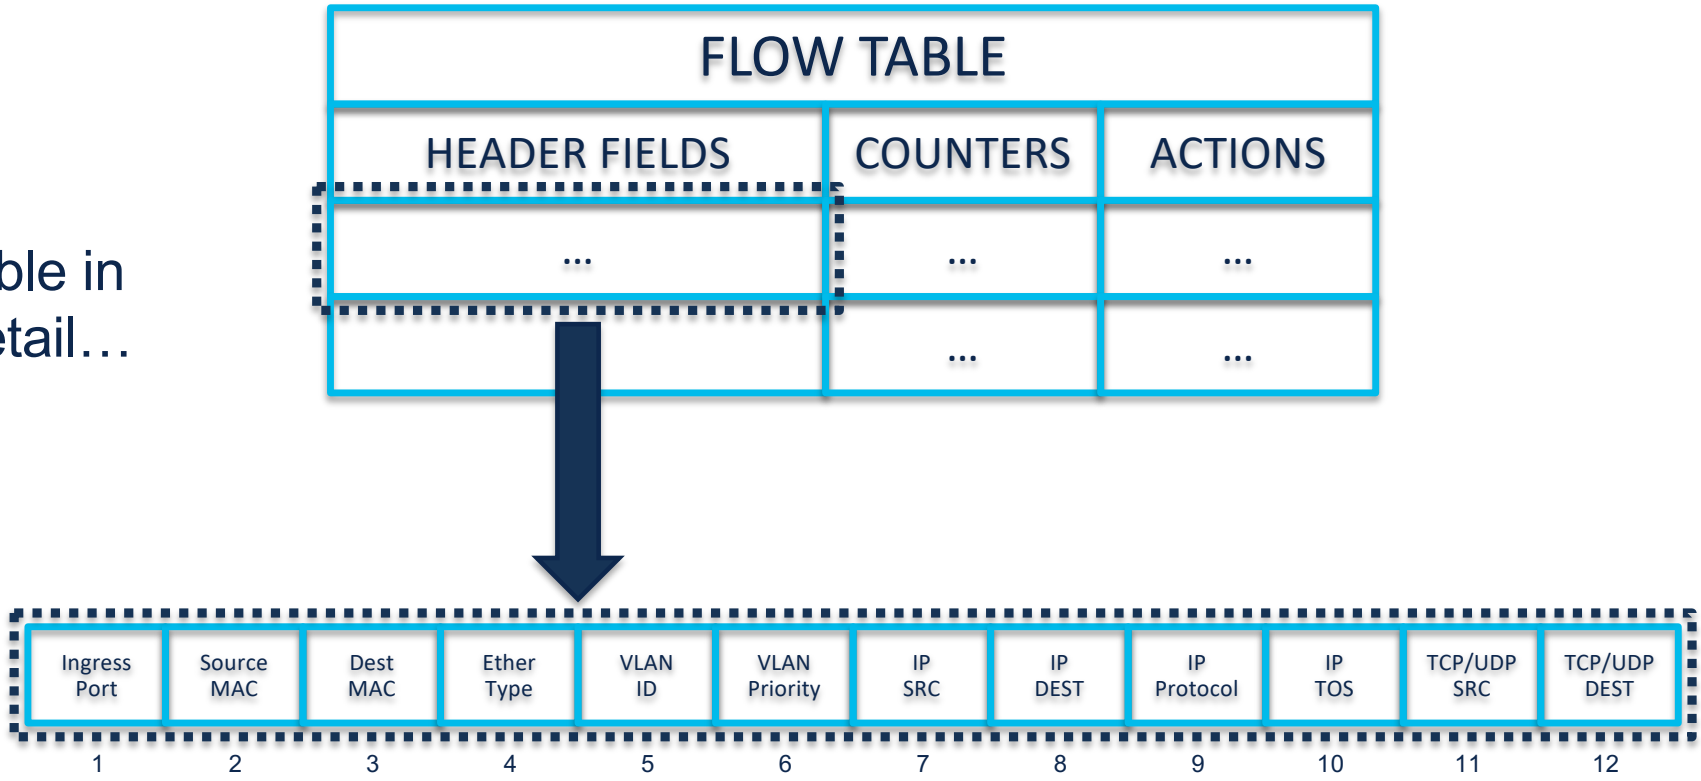

This is the “Famous” Openflow 12 Tuple

Openflow 1.0 Operation

Flow Table in more detail…

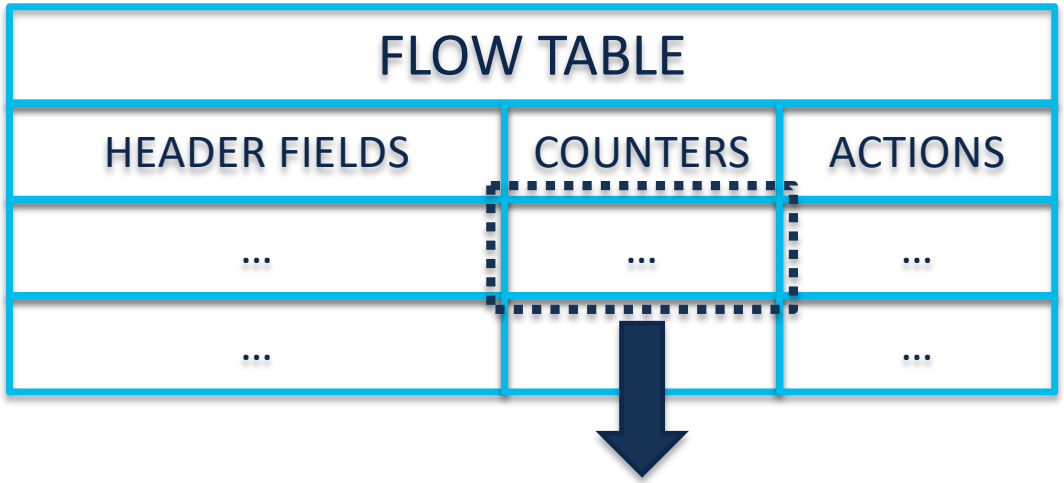

Per Table	
Active Entries	32 Bits
Packet Lookups	64 Bits
Packet Matches	64 Bits

Per Flow	
Received Packets	64 Bits
Received Bytes	64 Bits
Duration (seconds)	32 Bits
Duration (nanoseconds)	32 Bits

Per Queue	
Transmit Packets	64 Bits
Transmit Bytes	64 Bits
TX Overrun Errors	64 Bits

Per Port	
Received Packets	32 Bits
Transmit Packets	64 Bits
Received Bytes	64 Bits
Transmit Bytes	64 Bits
Received Drops	64 Bits
Transmit Drops	64 Bits
Received Errors	64 Bits
Transmit Errors	64 Bits
Received Frame Alignment Errors	64 Bits
RX Overrun Errors	64 Bits
RX CRC Errors	64 Bits
Collisions	64 Bits

Openflow 1.0 Operation

Flow Table in more detail…

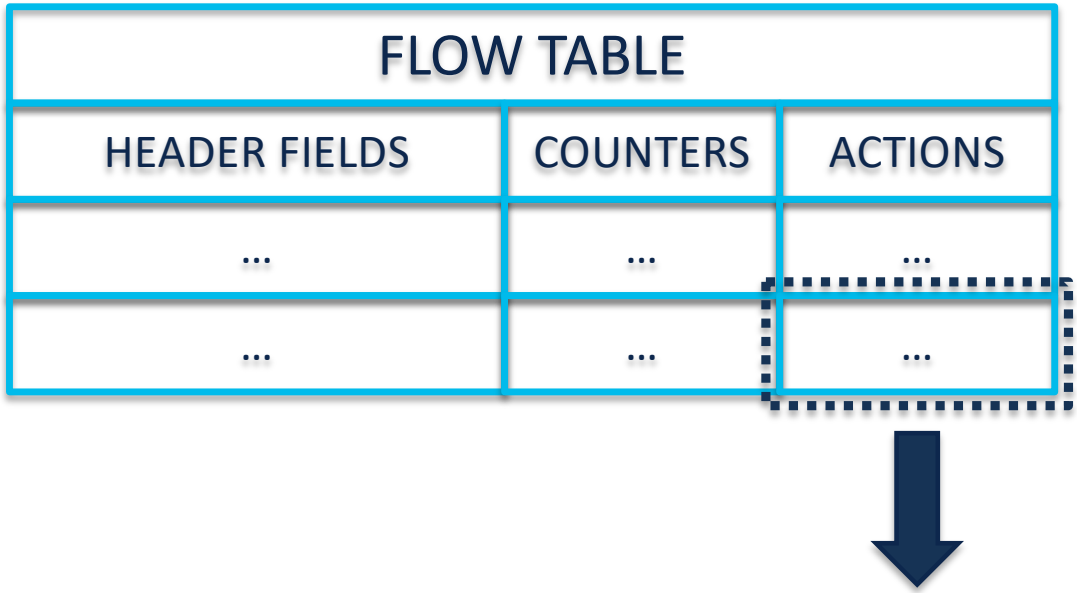

Multiple Actions available to be programmed
Let us explore those in more detail…

Openflow 1.0 Operation

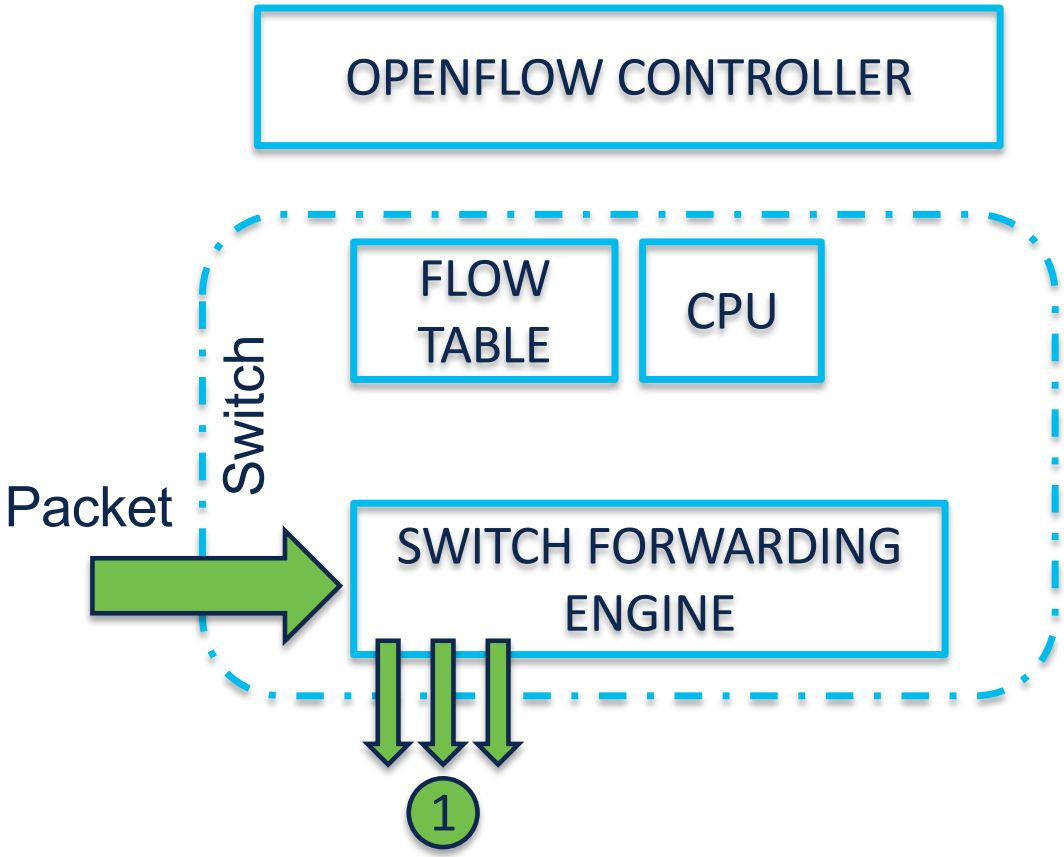

Required Action #1

Forward out all ports except input port

Openflow 1.0 Operation

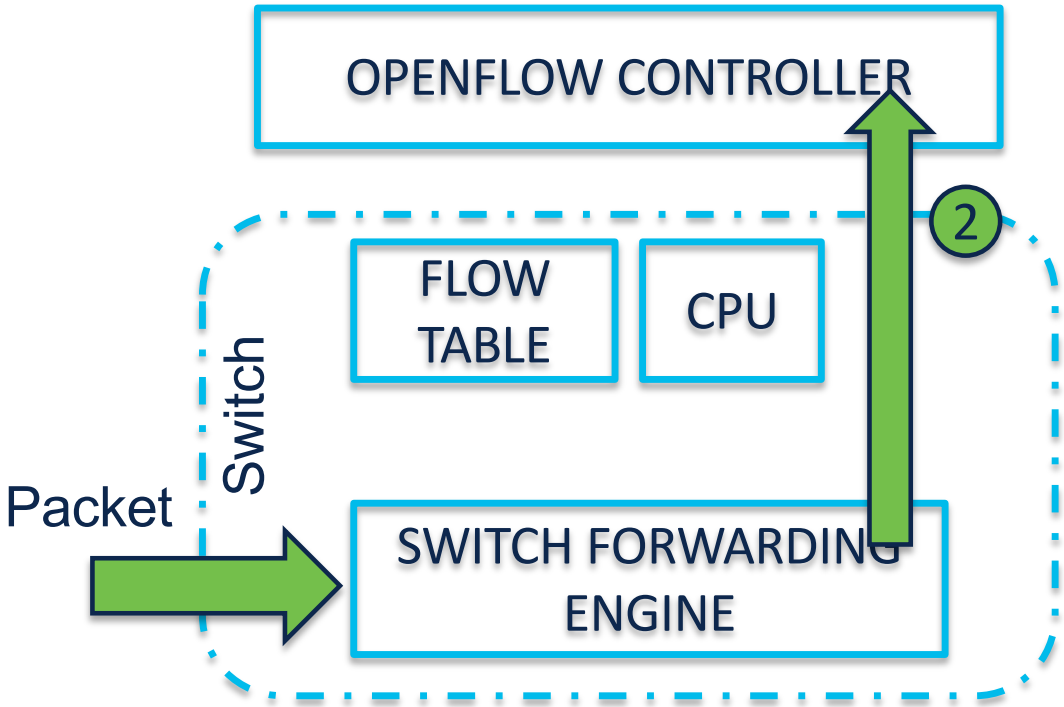

Required Action #2

Redirect to Openflow Controller

Openflow 1.0 Operation

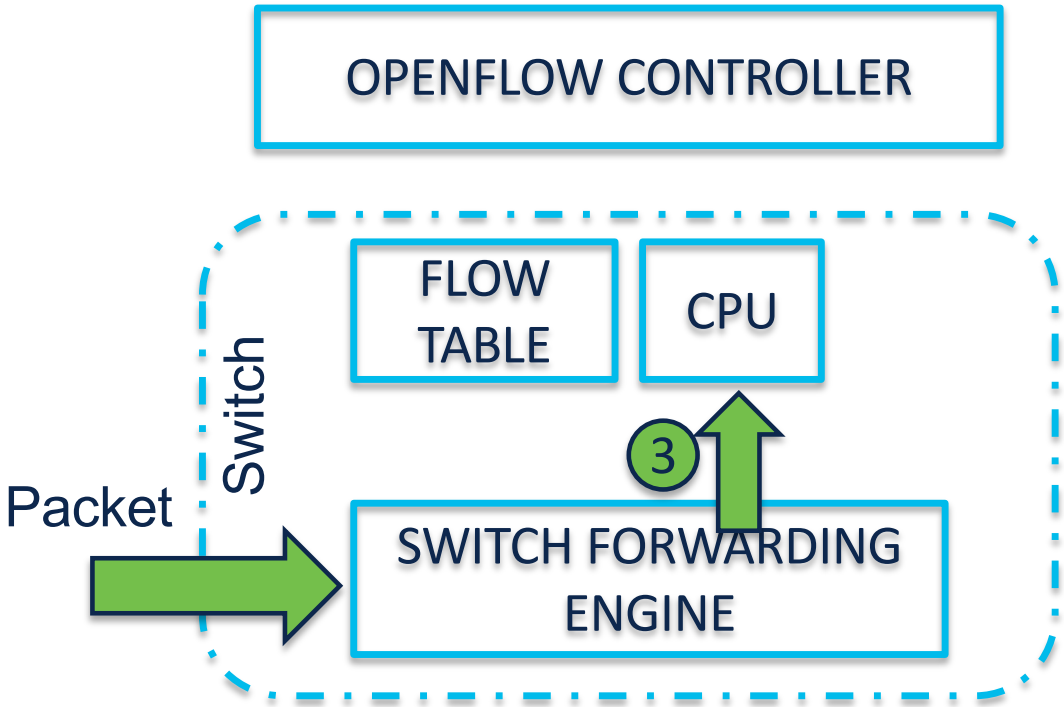

Required Action #3

Forward to local CPU

Openflow 1.0 Operation

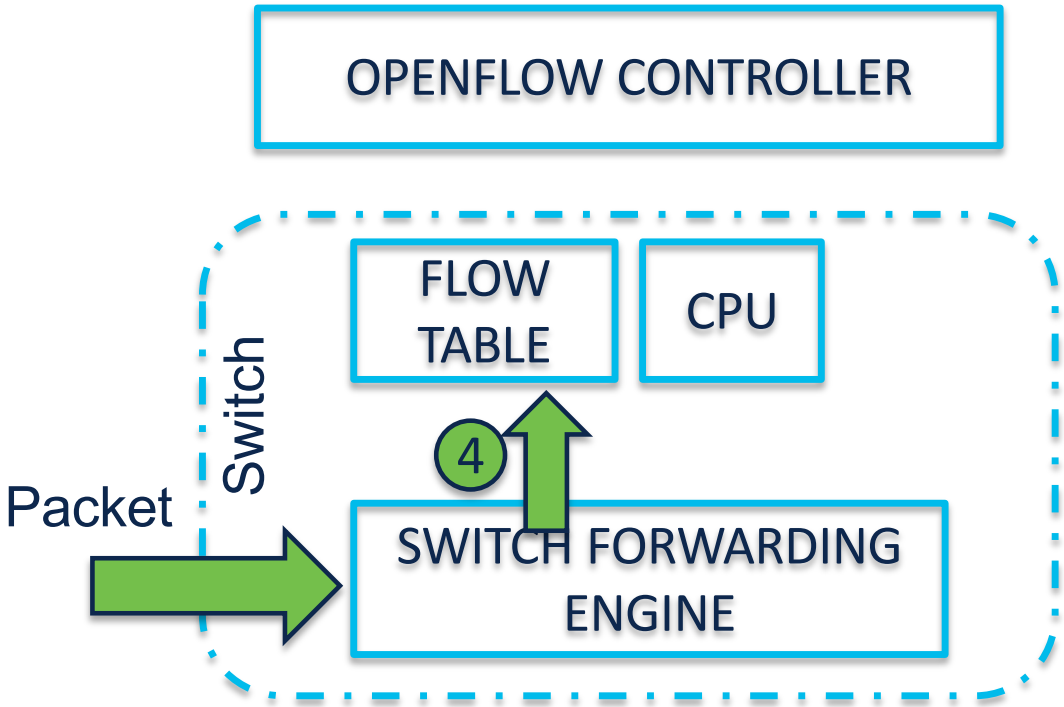

Required Action #4

Perform action in Flow Table

Openflow 1.0 Operation

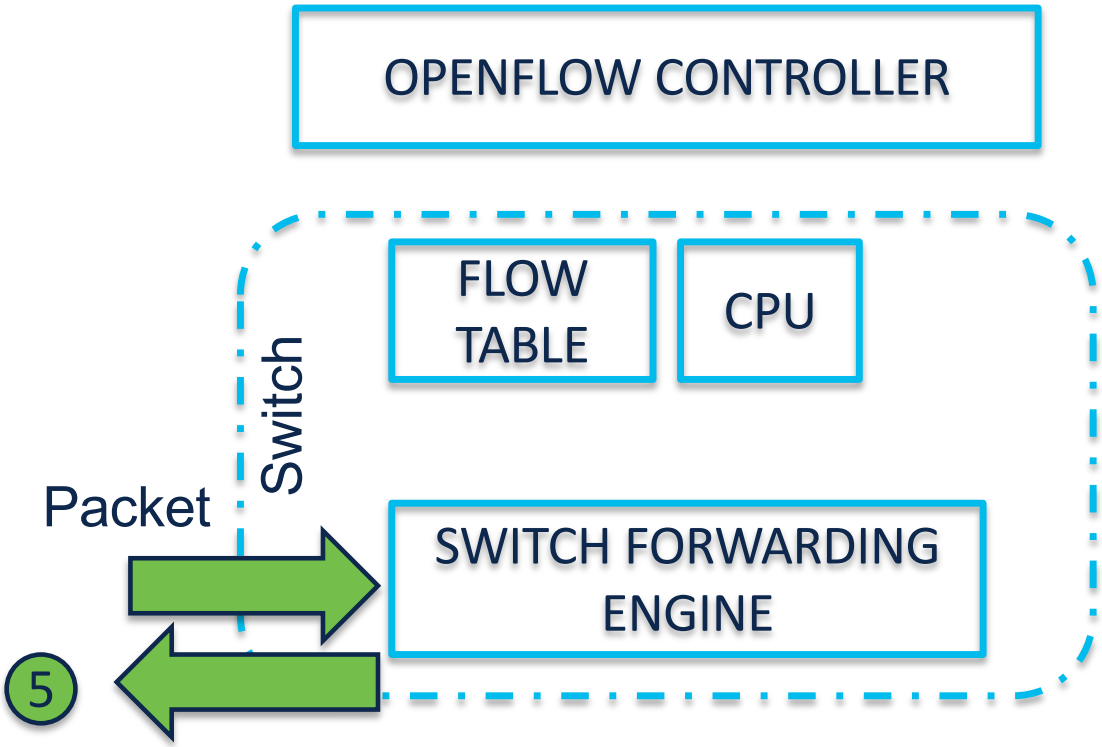

Required Action #5

Forward to Input Port

Openflow 1.0 Operation

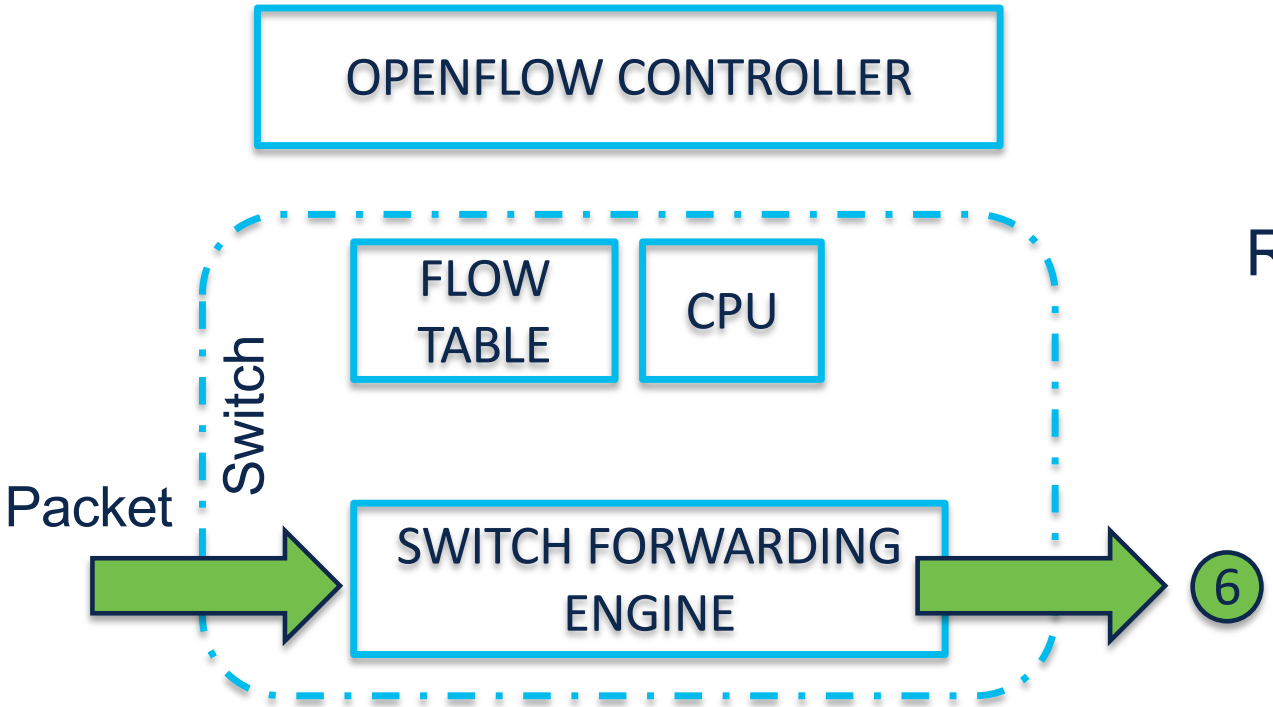

Required Action #6

Forward to Destination Port

Openflow 1.0 Operation

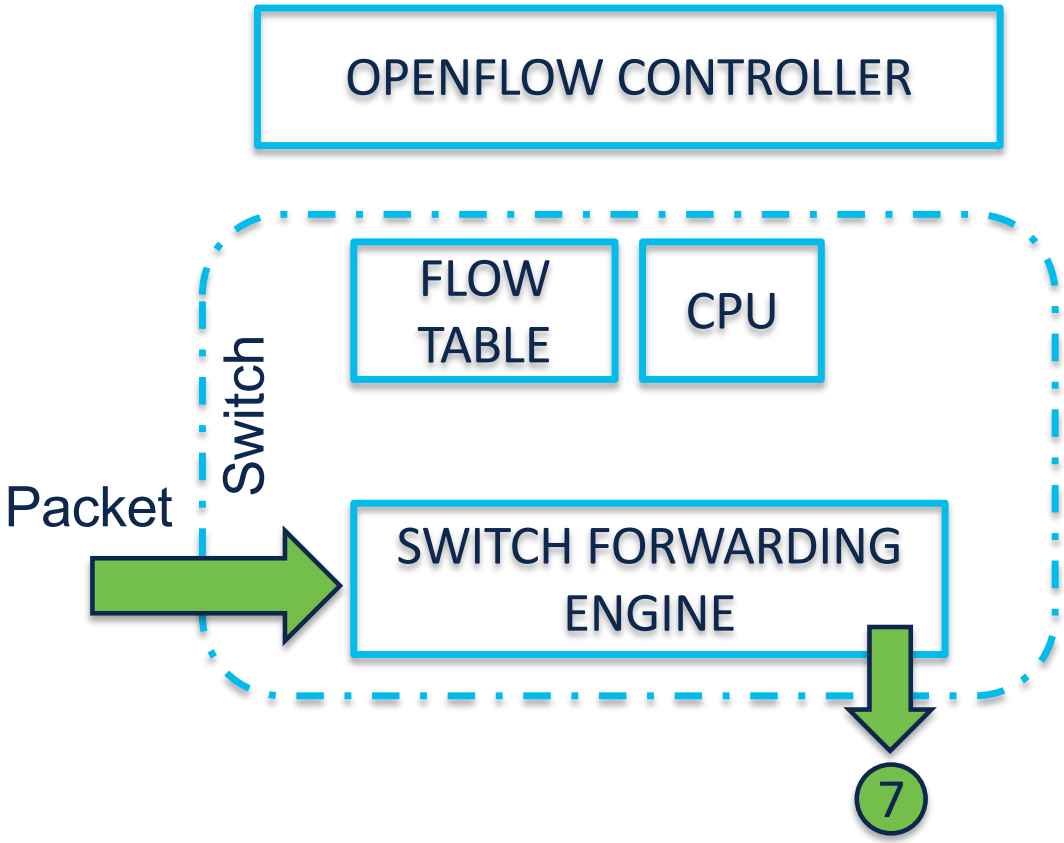

Required Action #7

Drop Packet

Openflow 1.0 Operation

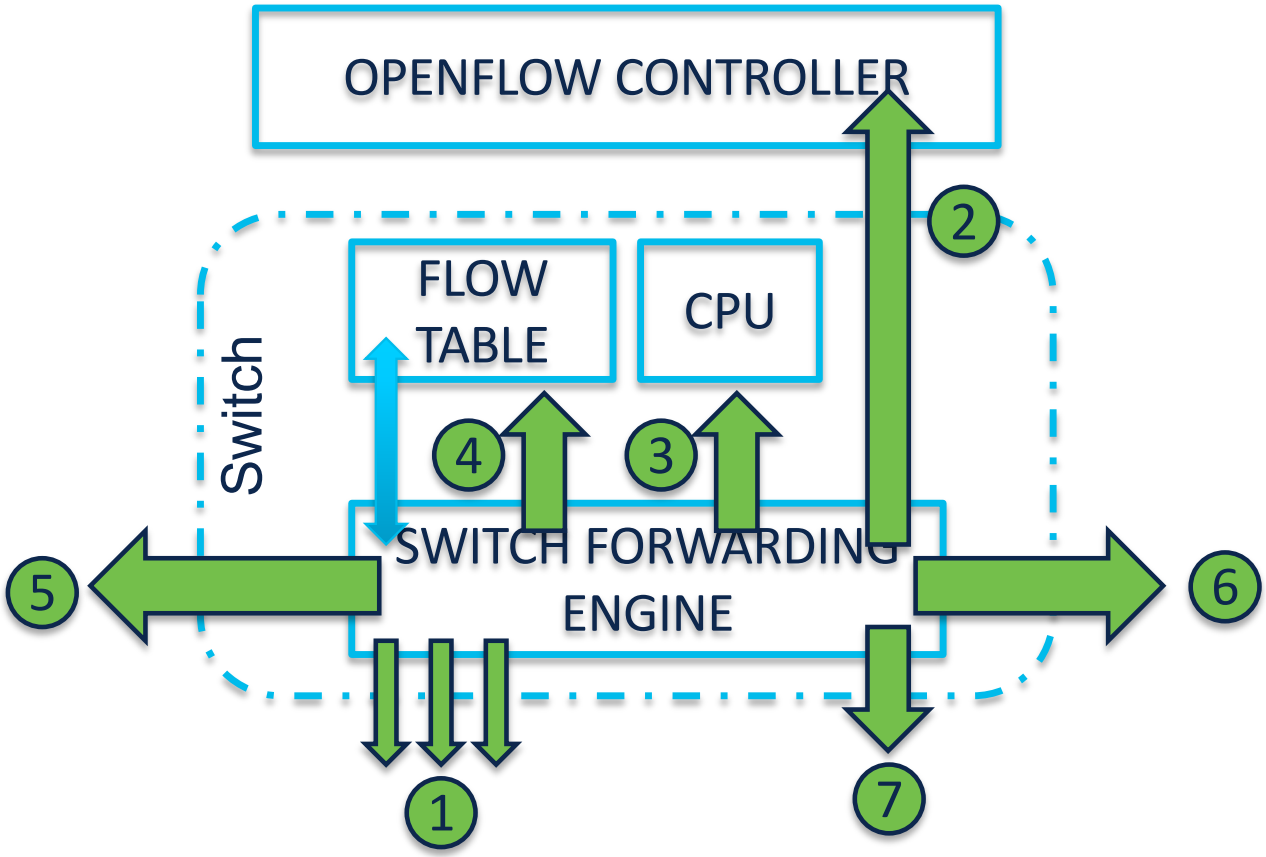

Required Actions	
1	Forward out all ports except input port
2	Redirect to Openflow Controller
3	Forward to local Forwarding Stack (CPU)
4	Perform action in flow table
5	Forward to input port
6	Forward to destination port
7	Drop Packet

Openflow 1.1 Operation

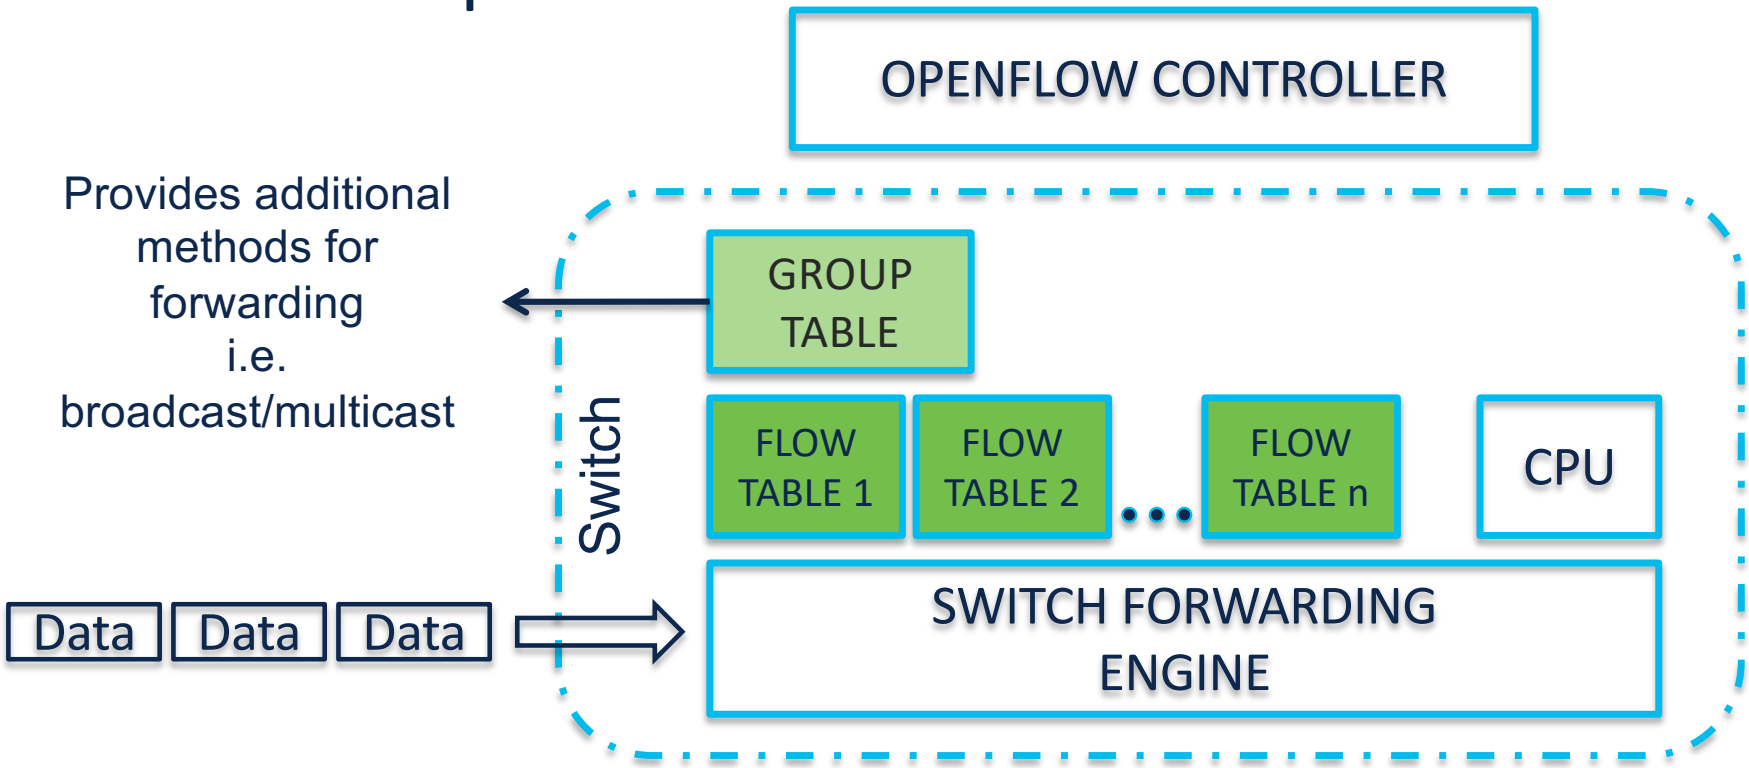

Openflow 1.1 Switch consists of one of more flow tables and a group table

Openflow 1.1 Operation

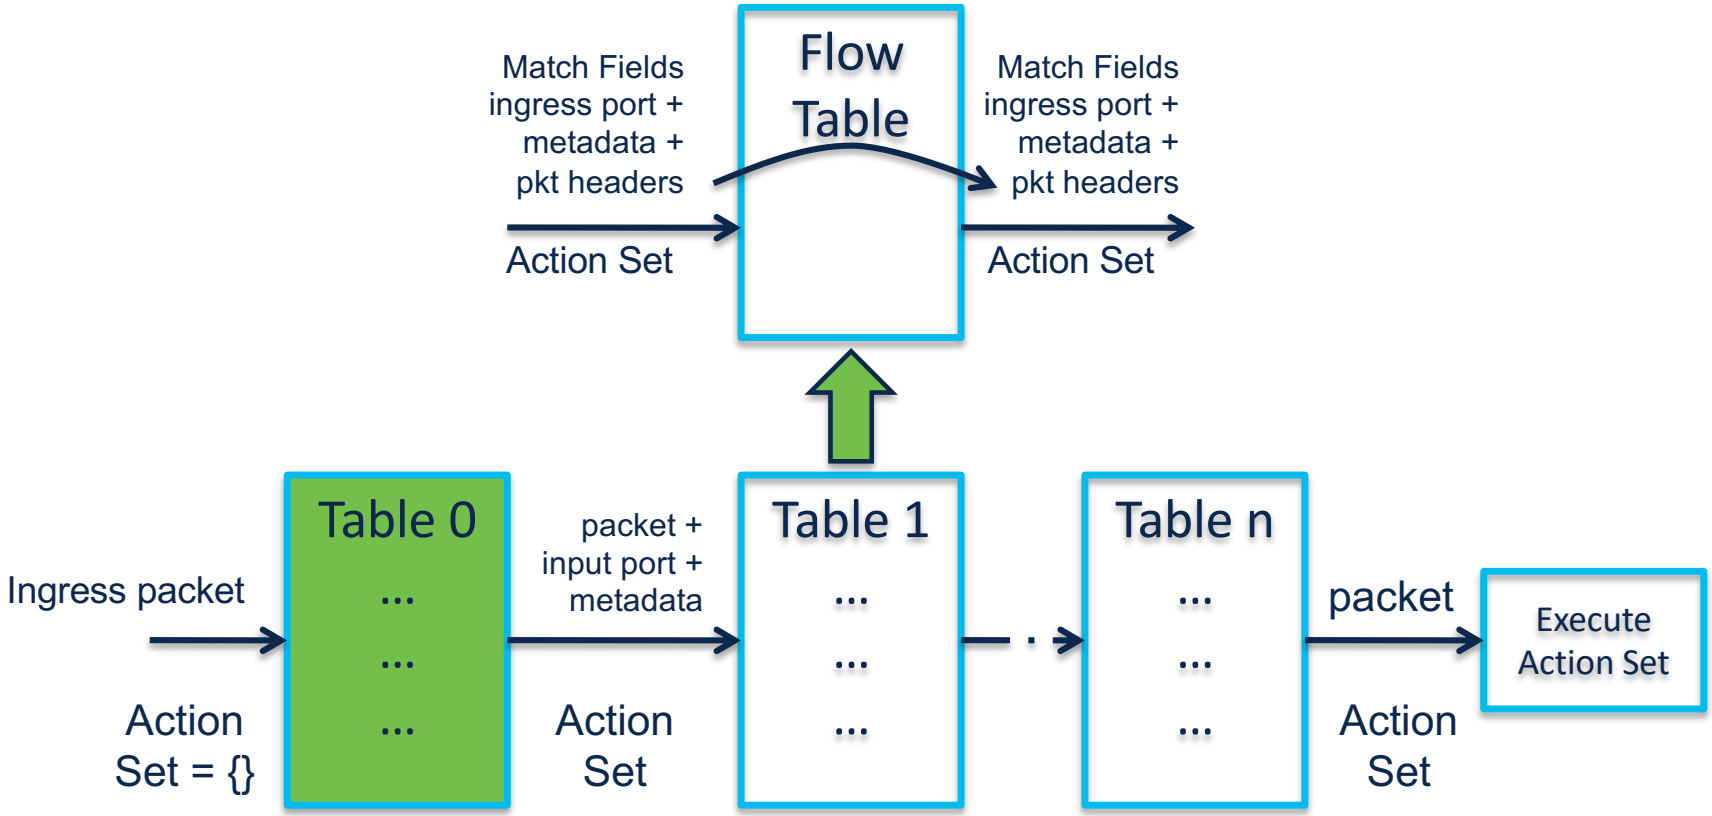

Matching starts at Table 1 and “may” continue to next table

Openflow 1.1 Operation

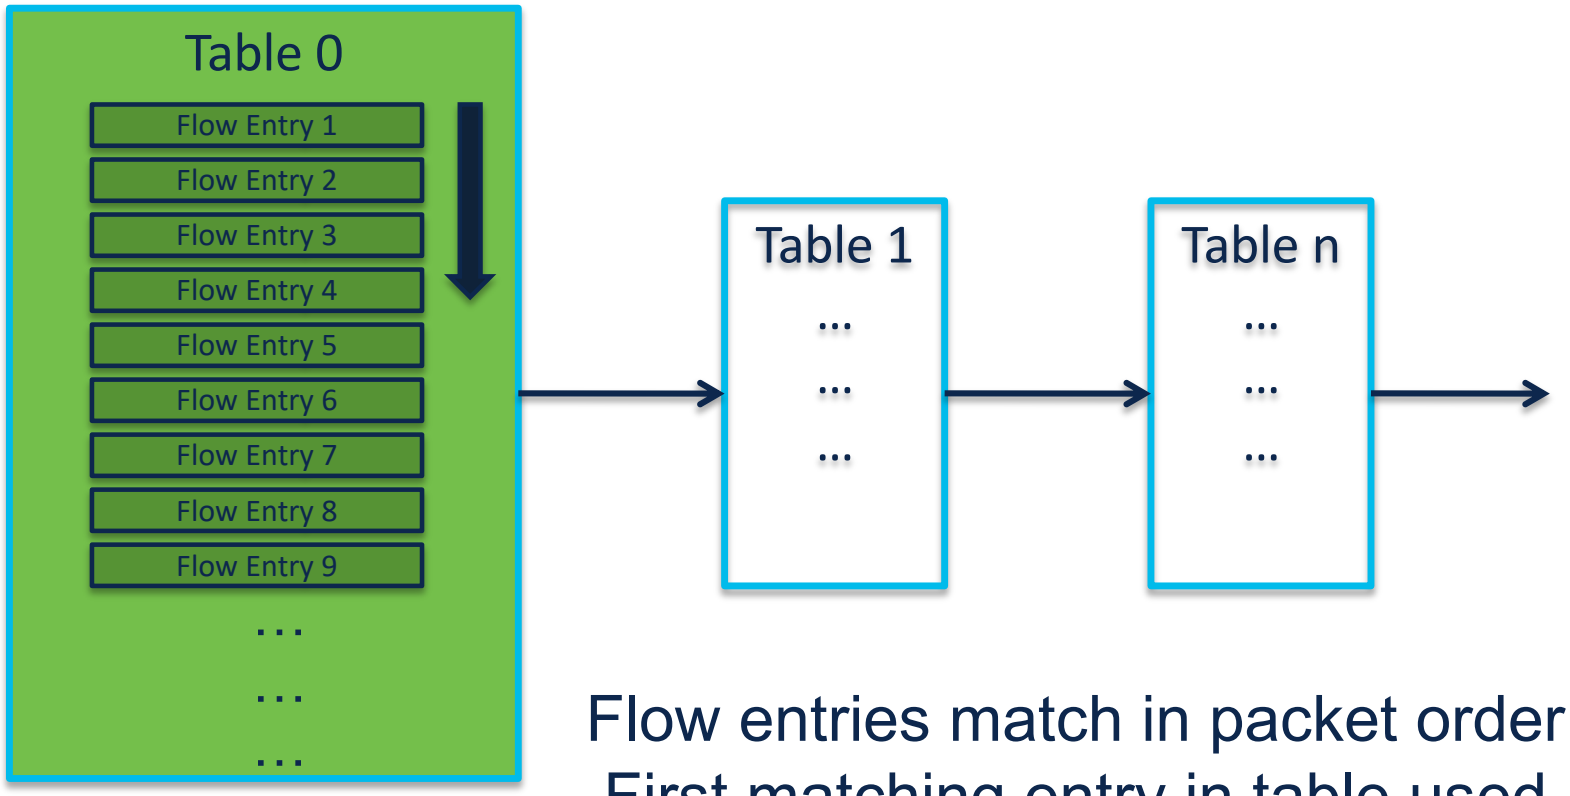

Flow entries match in packet order
First matching entry in table used

Openflow 1.1 Operation

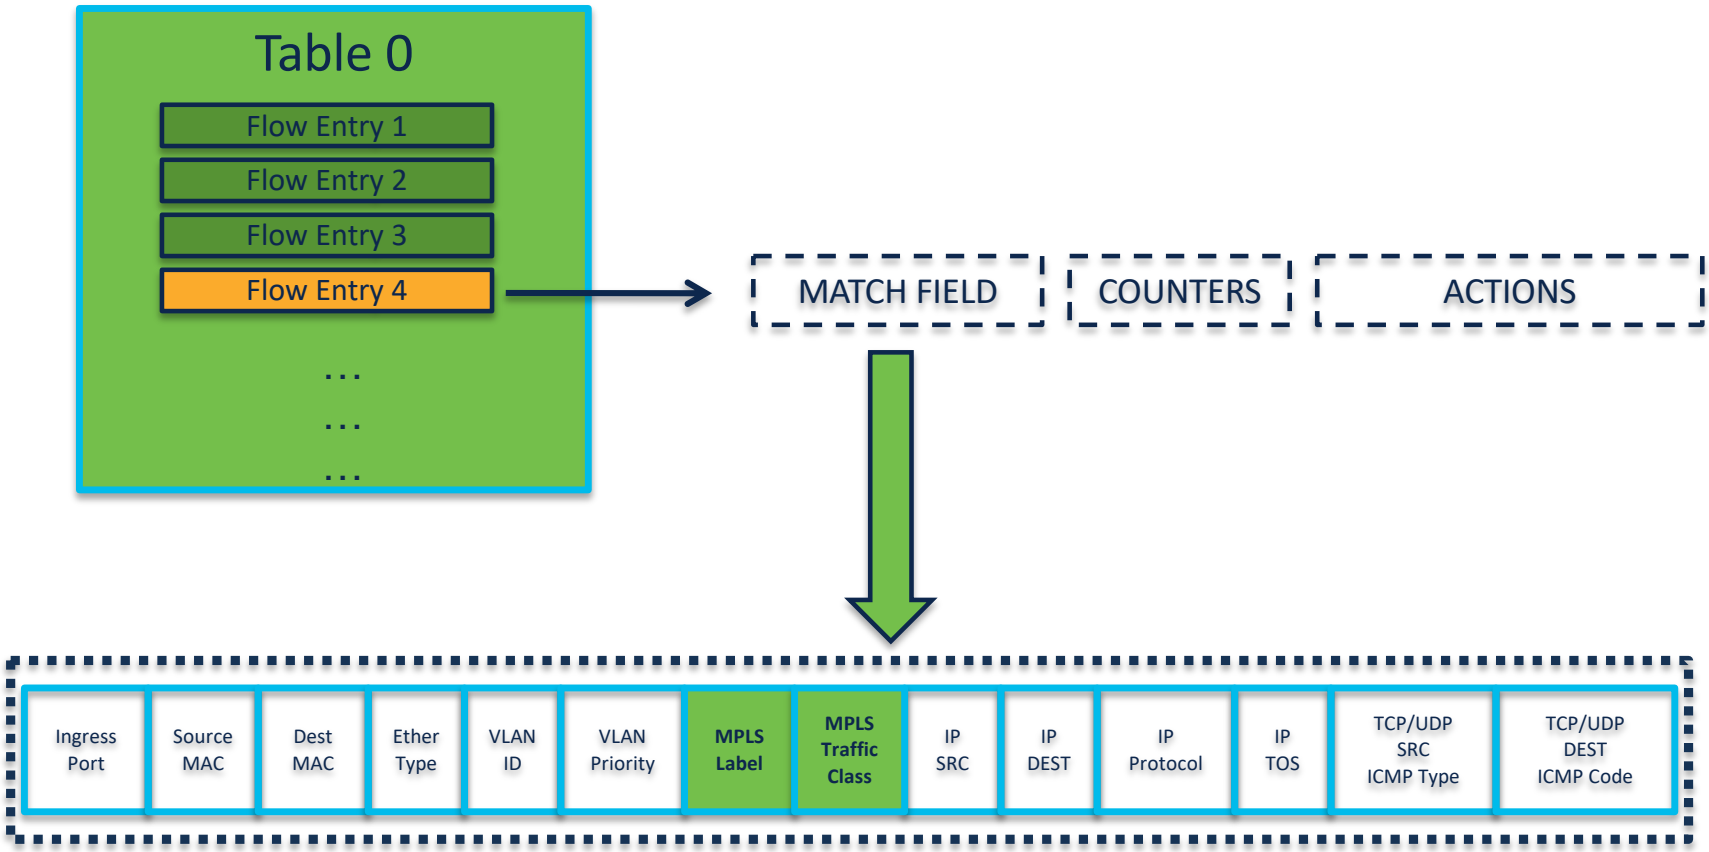

MPLS and VLAN Q-in-Q now supported in version 1.1

Openflow 1.1 Operation

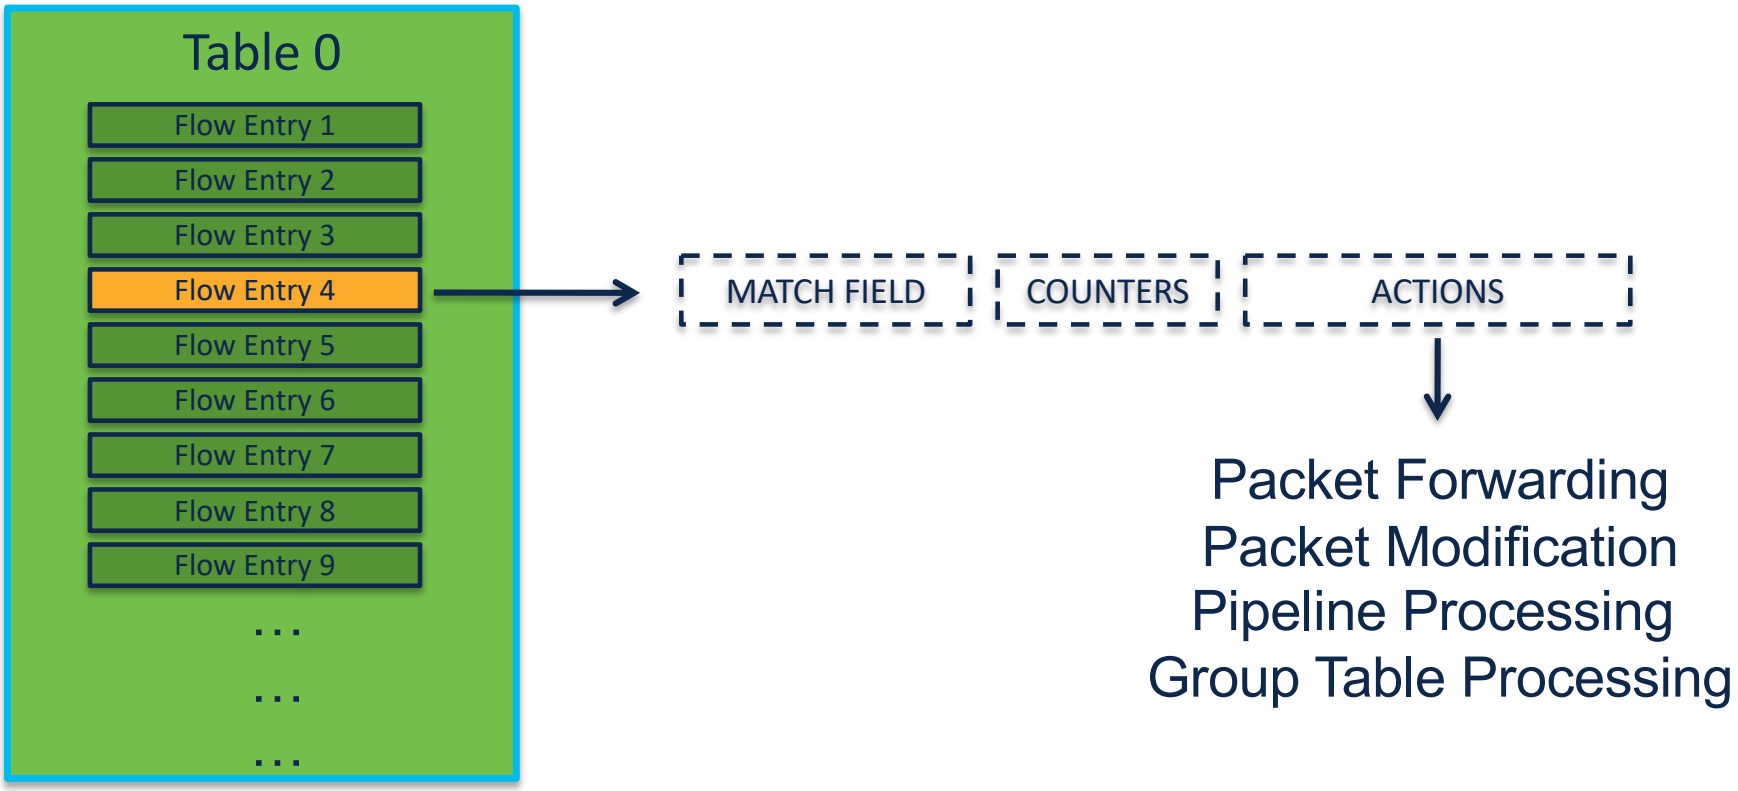

Actions in Flow Table define packet processing options

Openflow 1.1 Operation

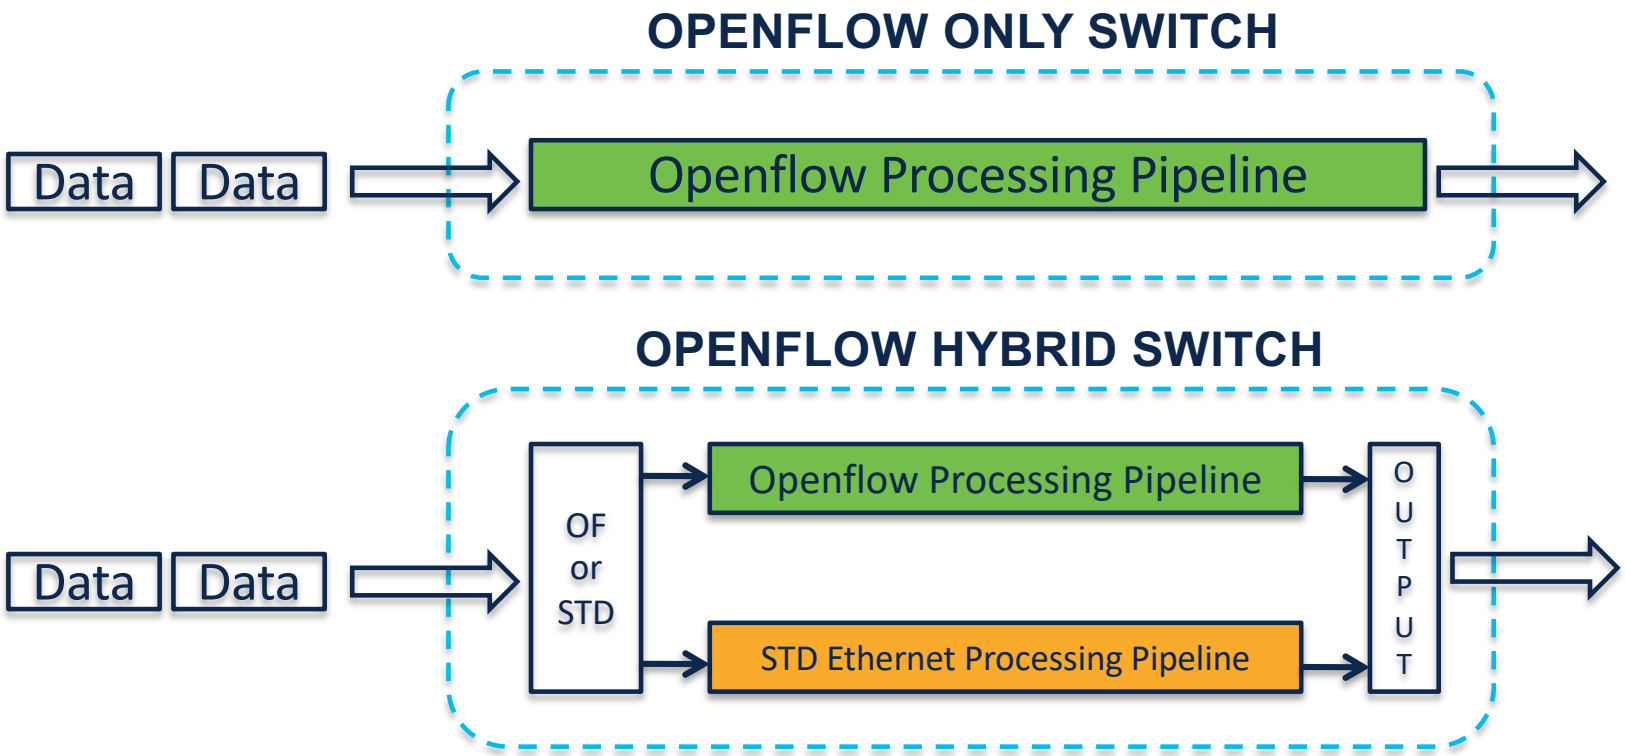

Openflow v1.1 defines two processing pipeline options
OPENFLOW ONLY and OPENFLOW HYBRID

Openflow 1.2 Operation

IPv6 now supported for lookup in flow table…

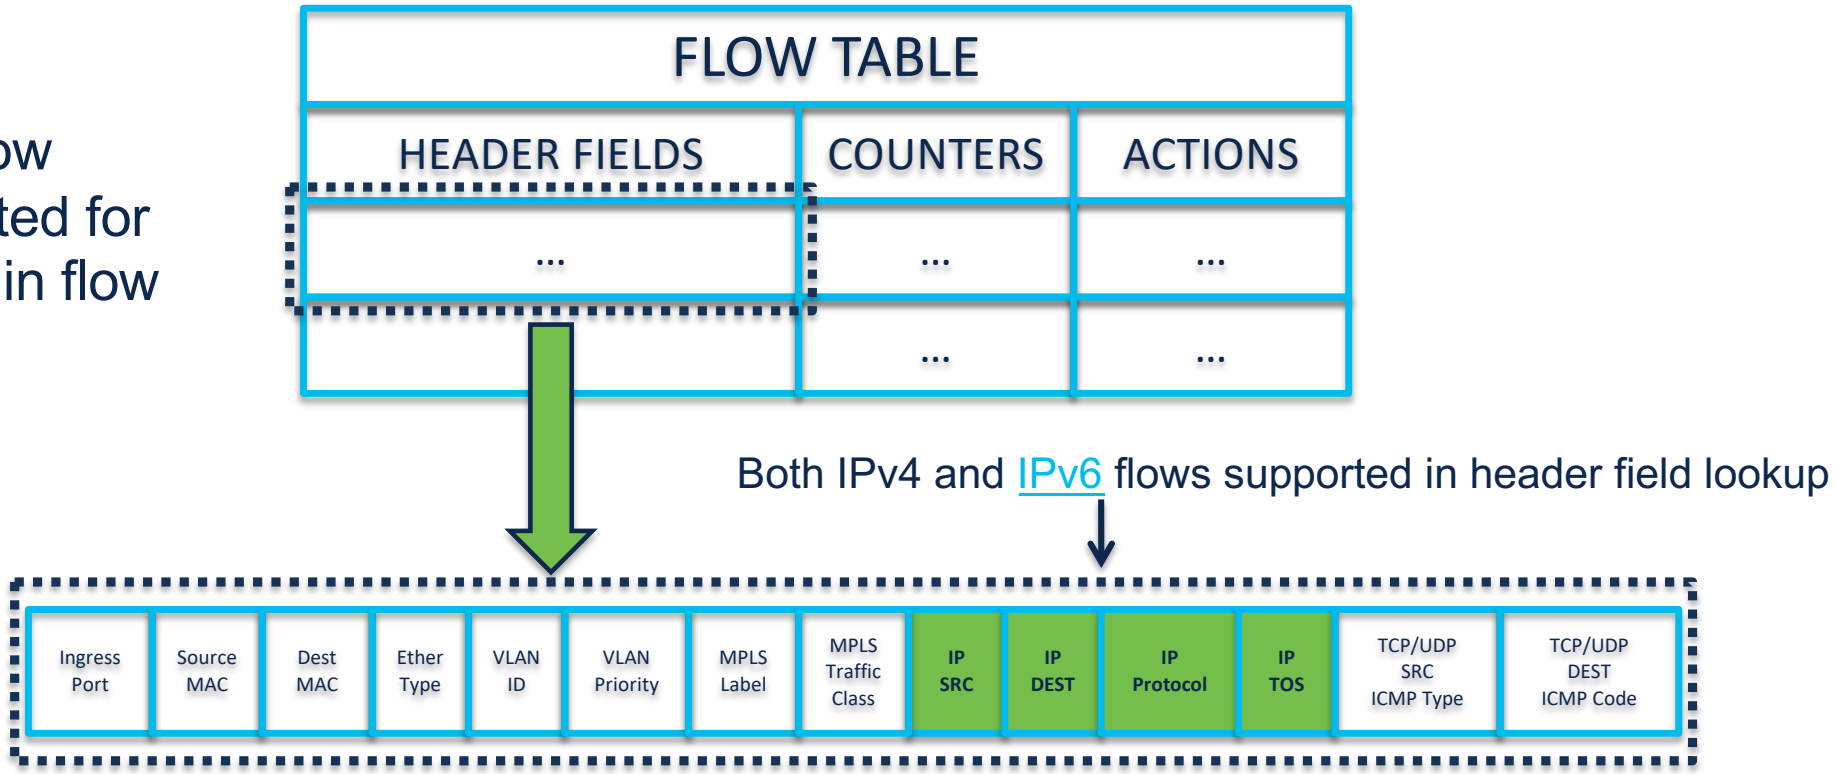

Openflow 1.3 Operation

Flow meter provides rate limiting (policing)

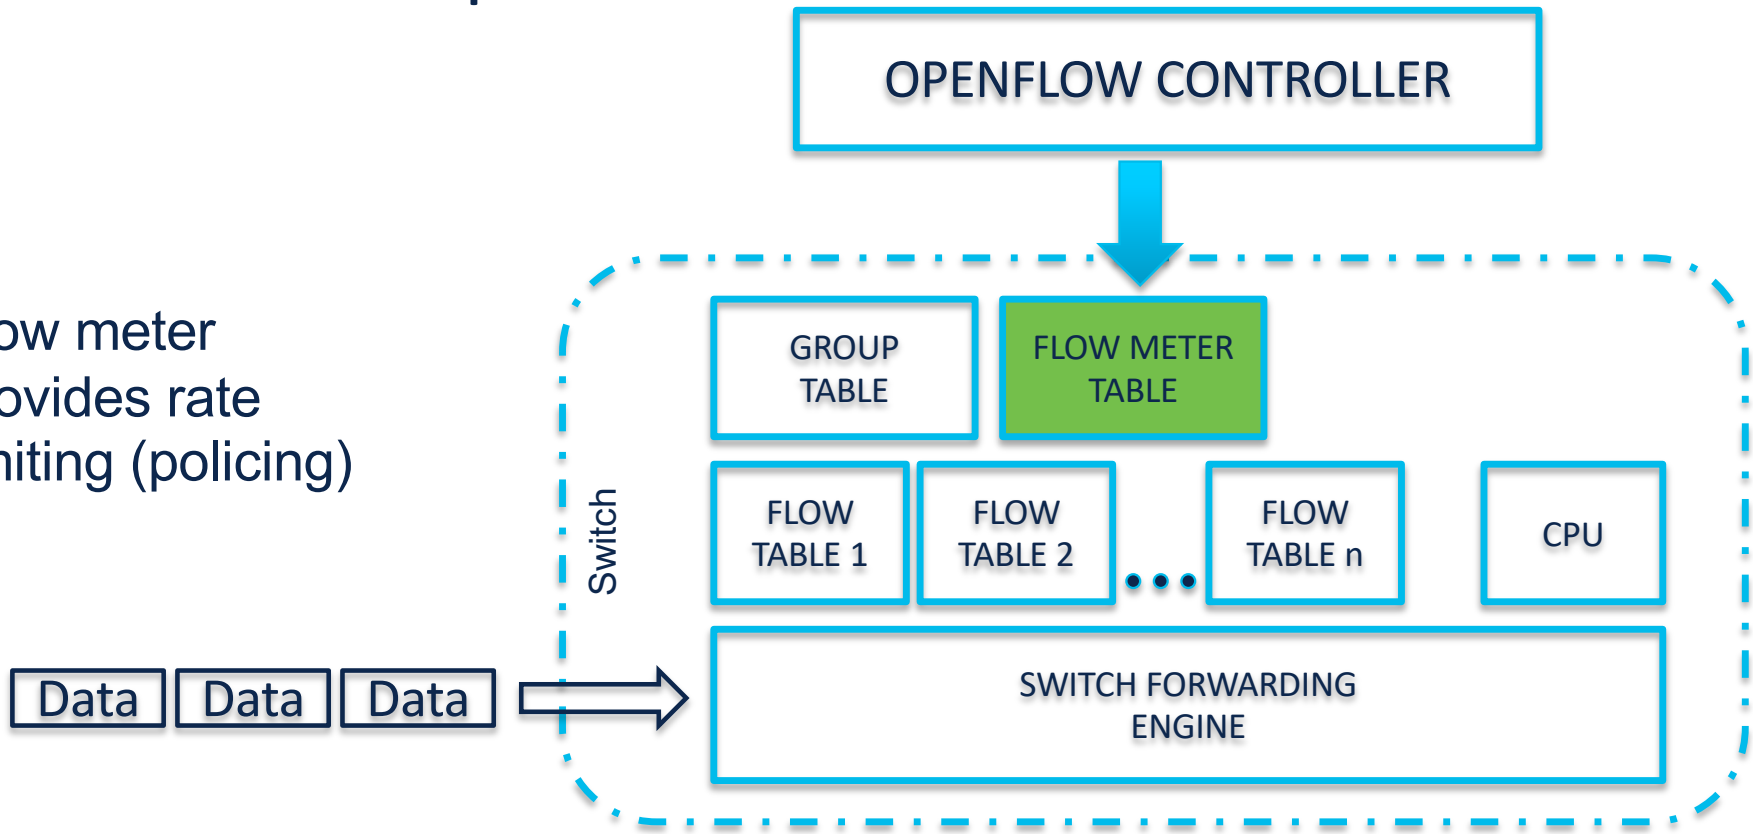

Openflow 1.3 Switch now adds a “flow meter” table

Openflow 1.3 Operation

Per Flow Meters supported in OF 1.3…

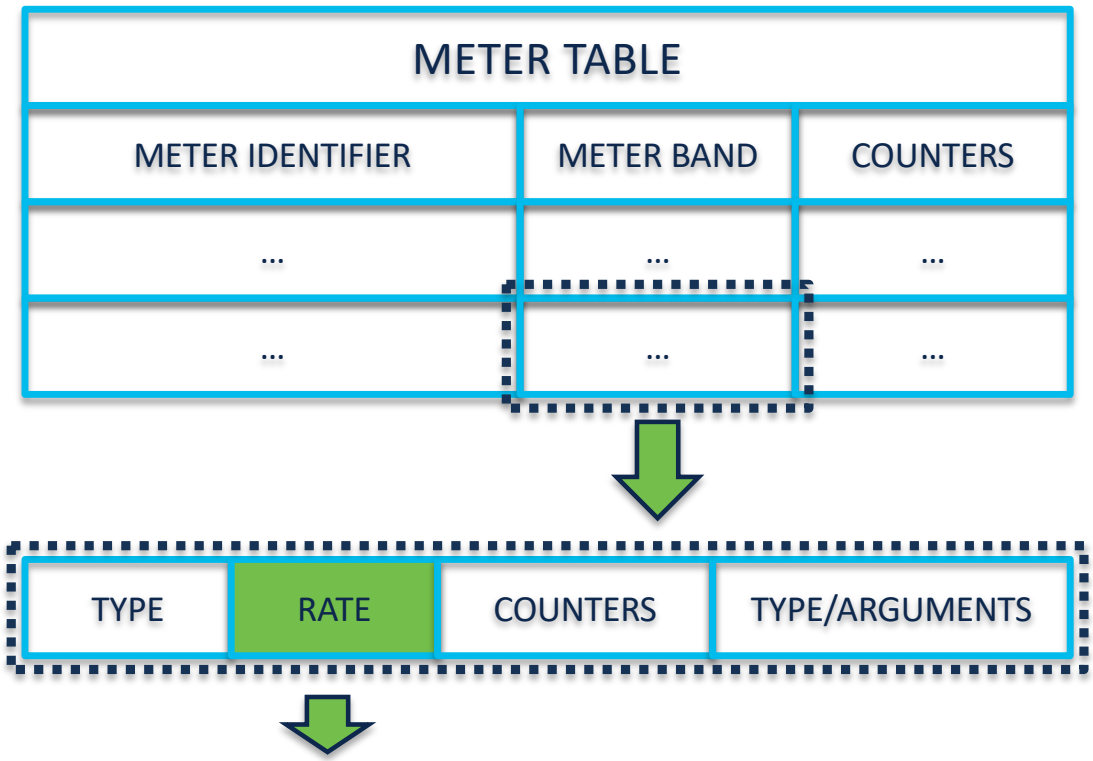

Controls the rate/flow of packets in a flow

Non Profit Consortium
Dedicated to "*the transformation of networks through SDN*"

Mission to "*commercialize and promote SDN…as a disruptive approach to networking…*"

Google's SDN Network

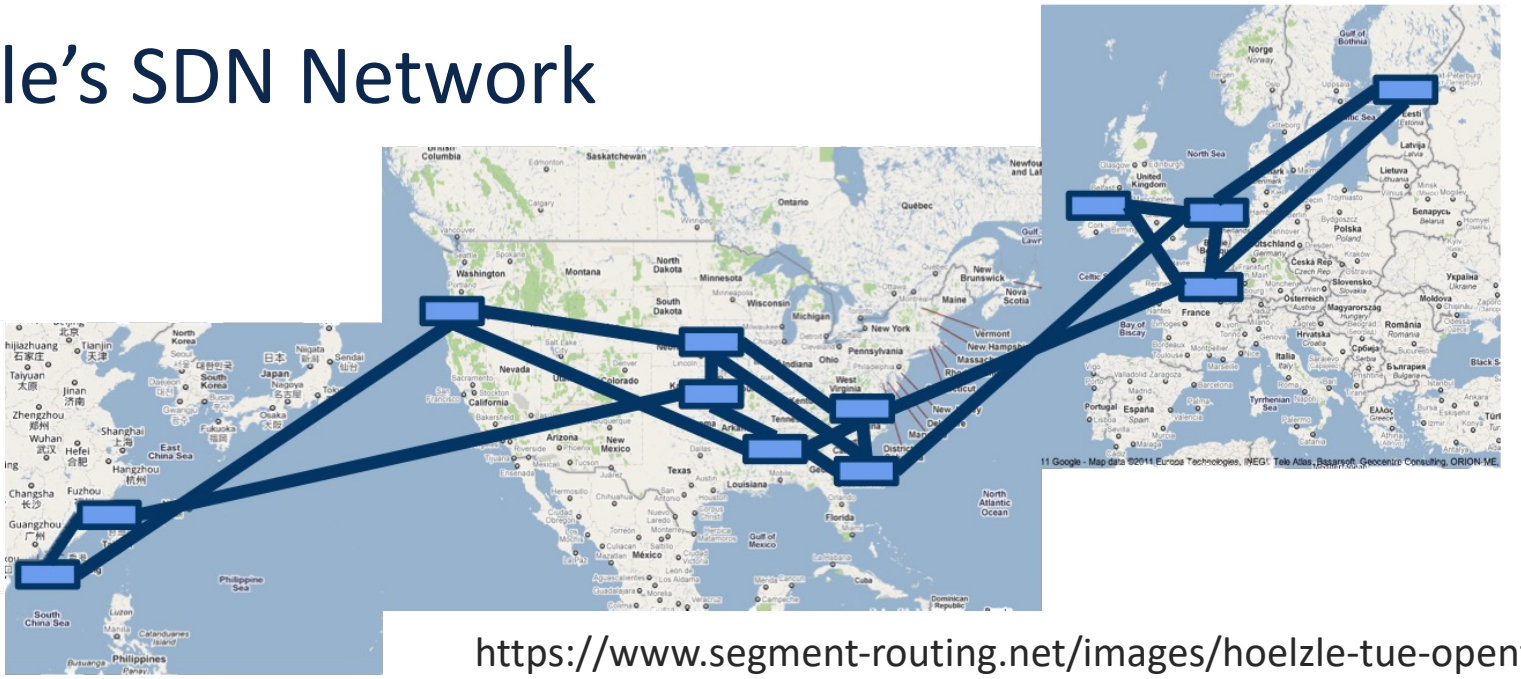

https://www.segment-routing.net/images/hoelzle-tue-openflow.pdf

Urs Holzle, Senior Vice President of Technology Infrastructure at Google
speaking in a keynote at the
second annual Open Networking Summit (April 2012)

G-Scale Example

- Google looks for a solution for its Data Center WAN Network called G-Scale
- Multiple sites located around the world
- Goal: manage the WAN as a fabric not as a collection of individual boxes:
  - Better network utilization
  - Faster convergence
  - Deterministic behavior
  - Controllers use modern server hardware
  - 50x better performance

G-Scale Example

- Build their own routers from merchant silicon
  - 100s of ports of nonblocking 10GE
  - OpenFlow support
- Open source routing stacks for BGP, ISIS
- Does not have all features
  - No support for AppleTalk…
- Multiple chassis per site
  - Fault tolerance
  - Scale to multiple Tbps

G-Scale Example

- Phase 1 (Spring 2010):
  - Introduce OpenFlow-controlled switches but make them look like regular routers
- No change from perspective of non-OpenFlow switches
- BGP/ISIS/OSPF now interfaces with OpenFlow controller to program switch state
- Phase 2 (until mid-2011): ramp-up
- Phase 3 (early 2012) full roll out of all G-Scale network
- Rolled out centralized TE
  - Optimized routing based on application-level priorities (currently 7)
- Globally optimized placement of flows

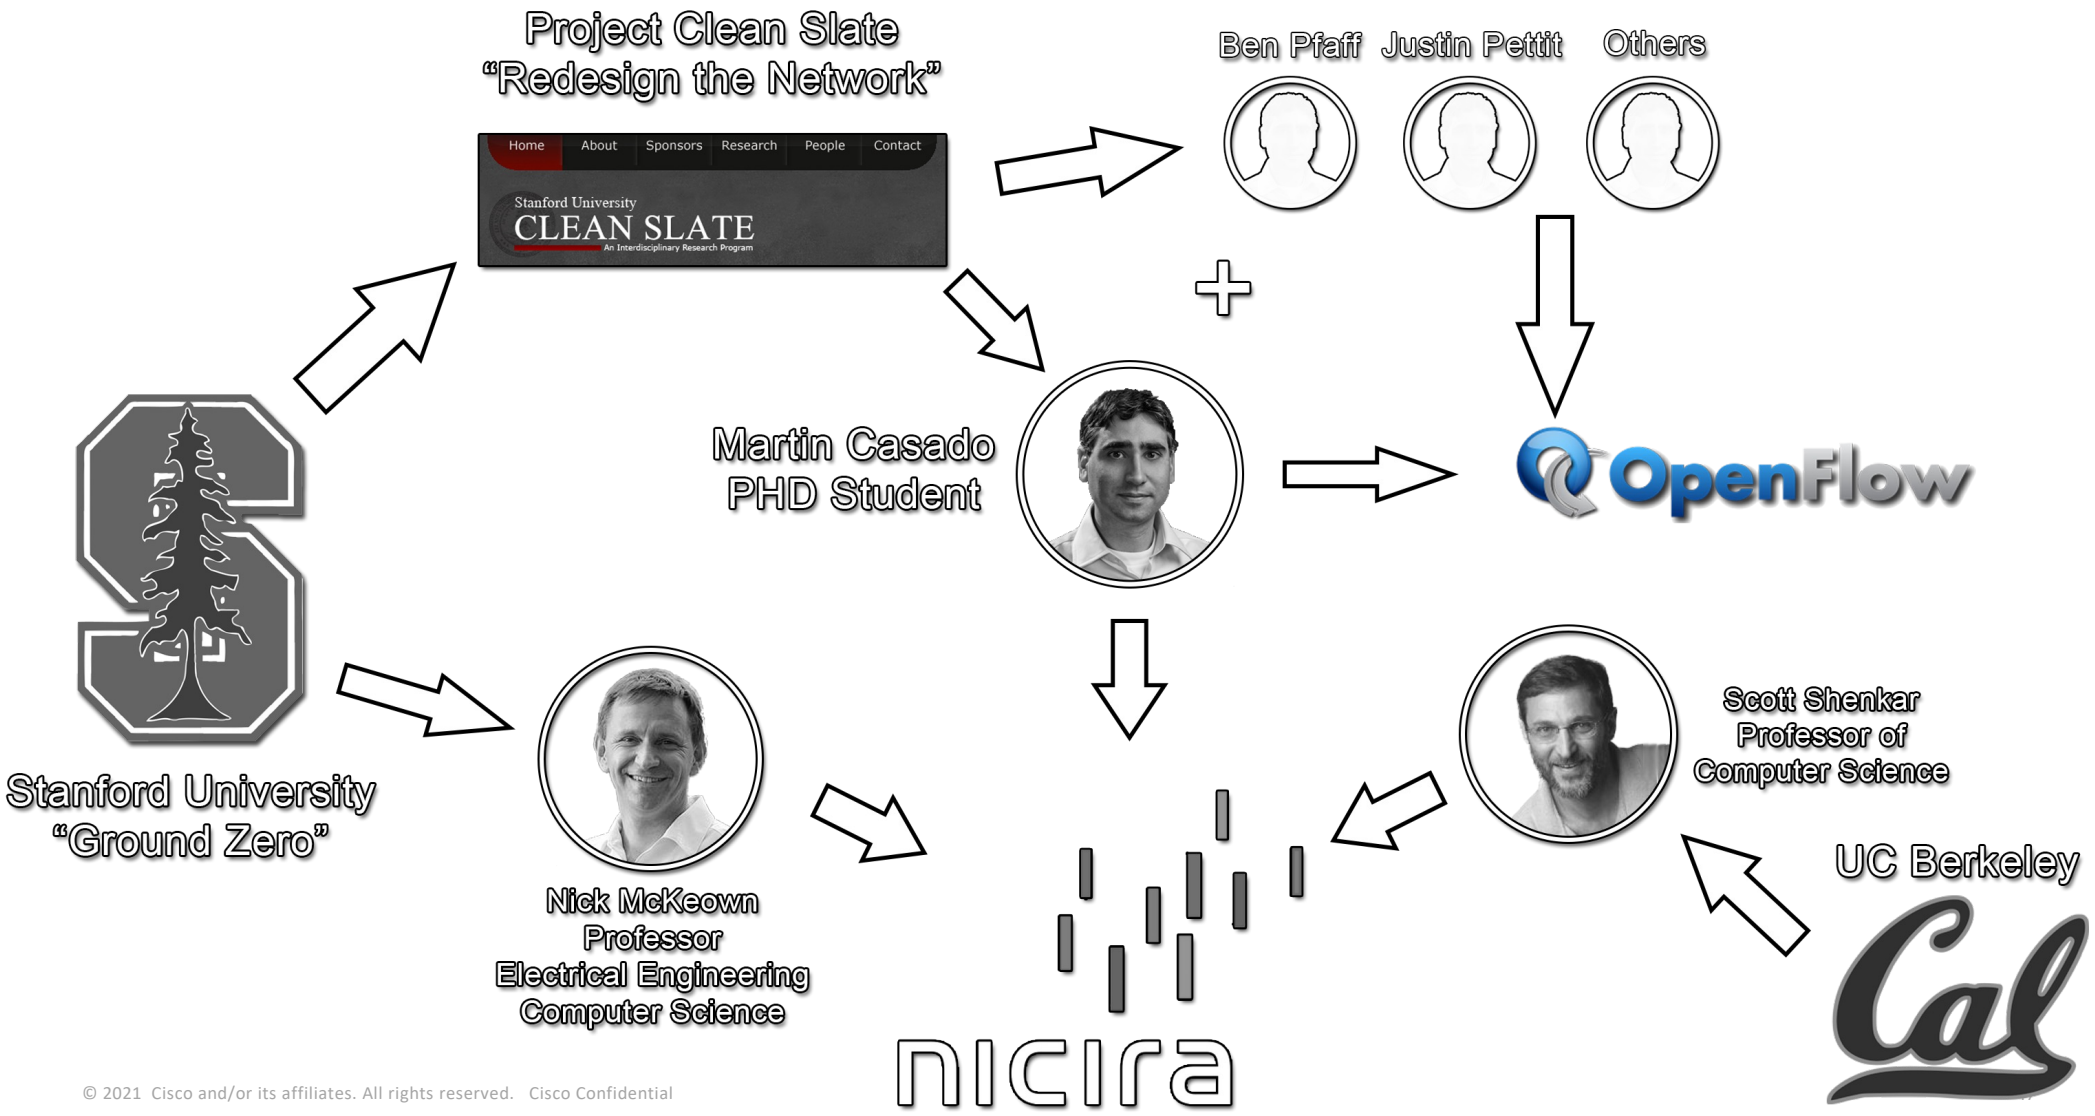

nicira

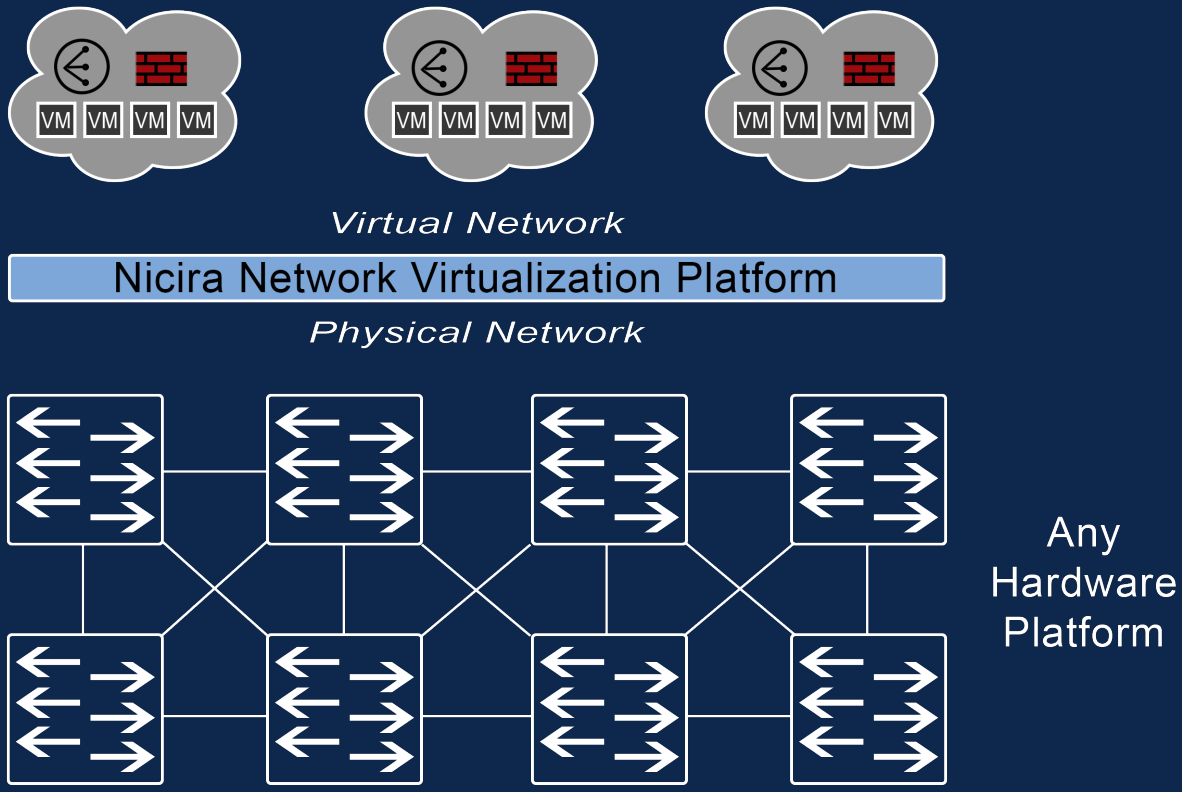

Nicira market a solution called NVP (Network Virtualization Platform) that provides their Overlay solution

You start with a Physical Switch Network

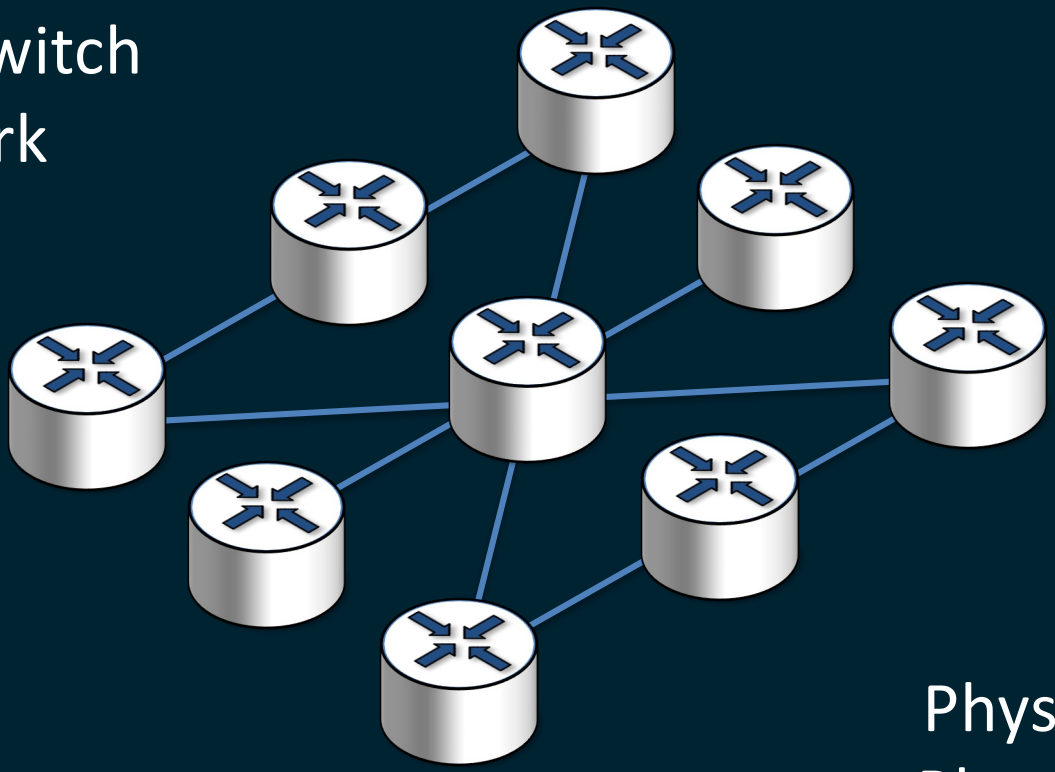

Physical Devices and Physical Connections

Then you add an overlay

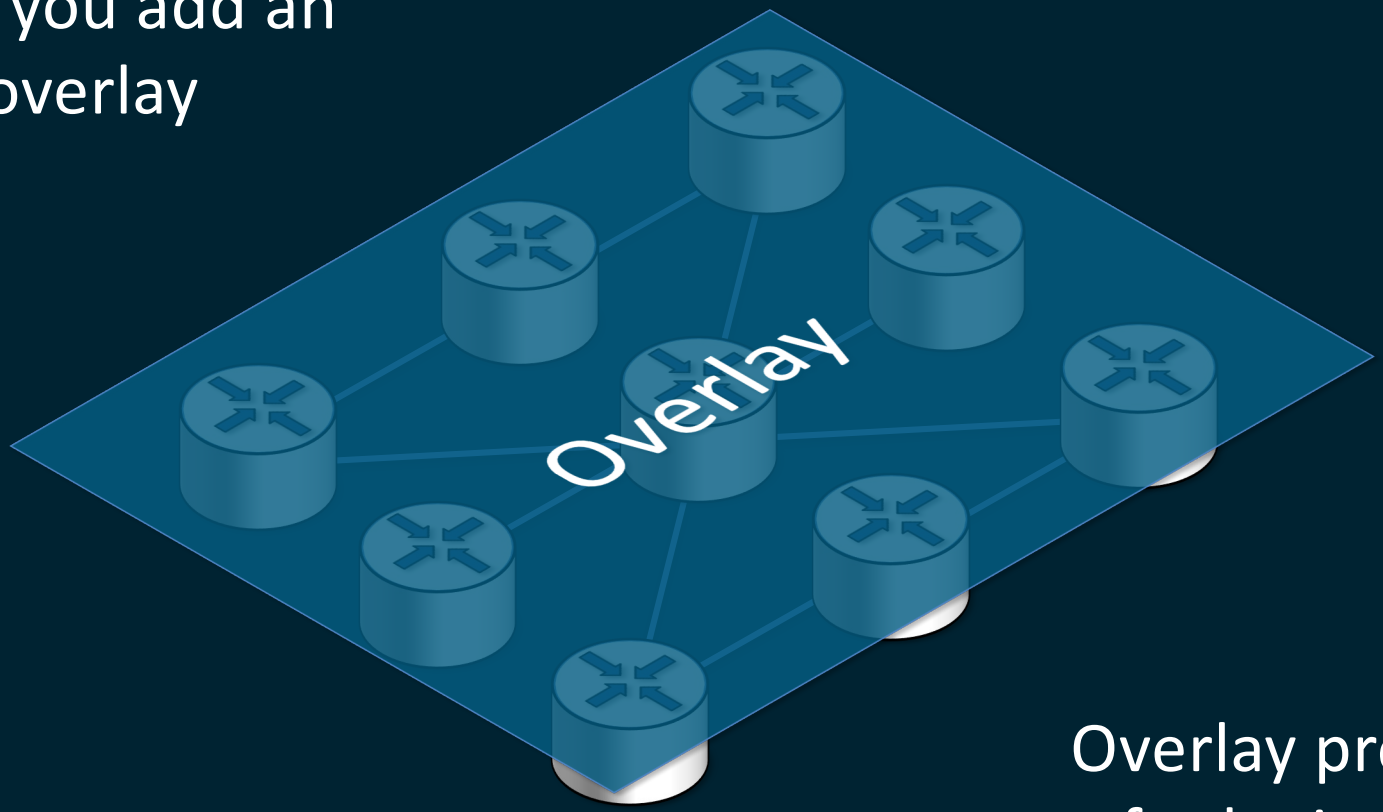

Overlay provides base for logical network

Logical “switch” devices overlay the physical network

They define their own topology

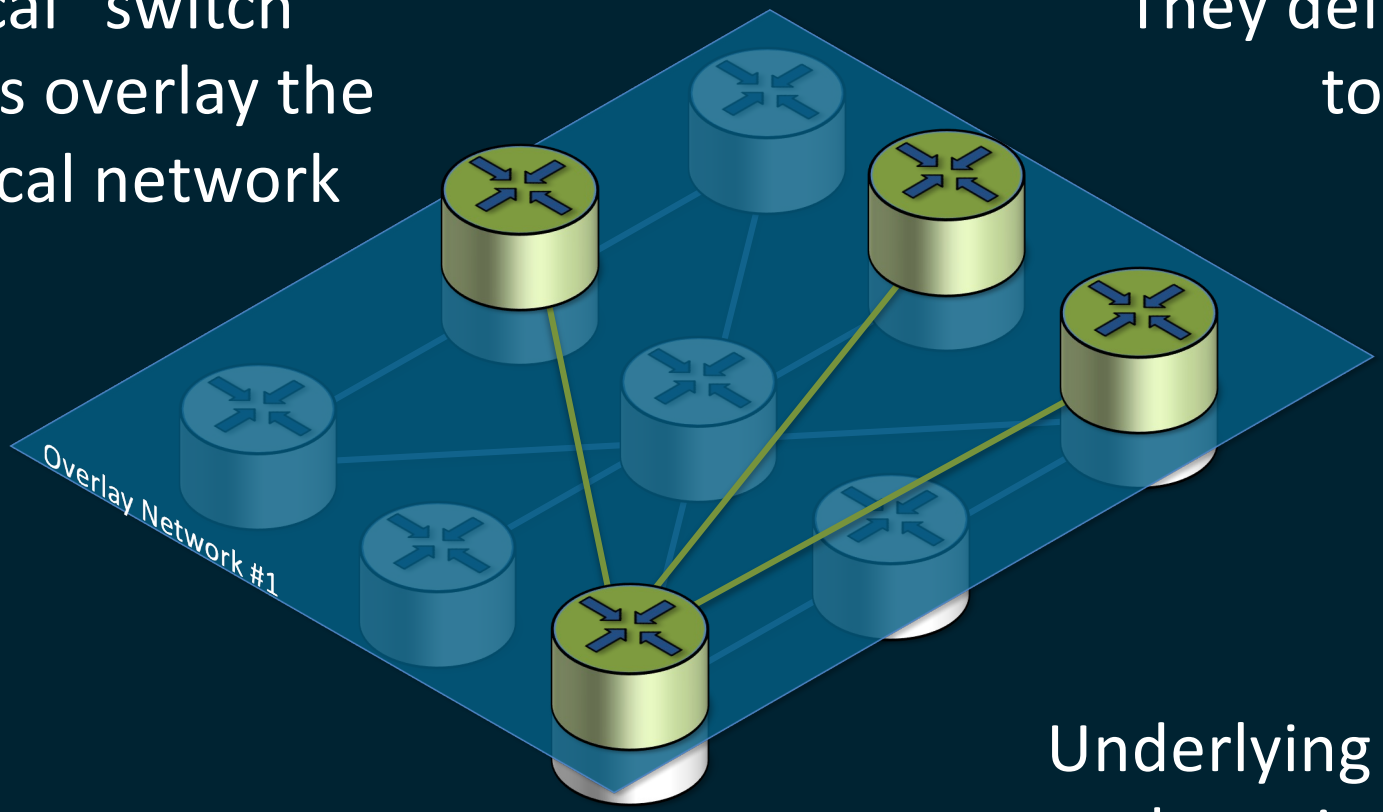

Underlying physical network carries data traffic for overlay network

Multiple “overlay” networks can co-exist at the same time

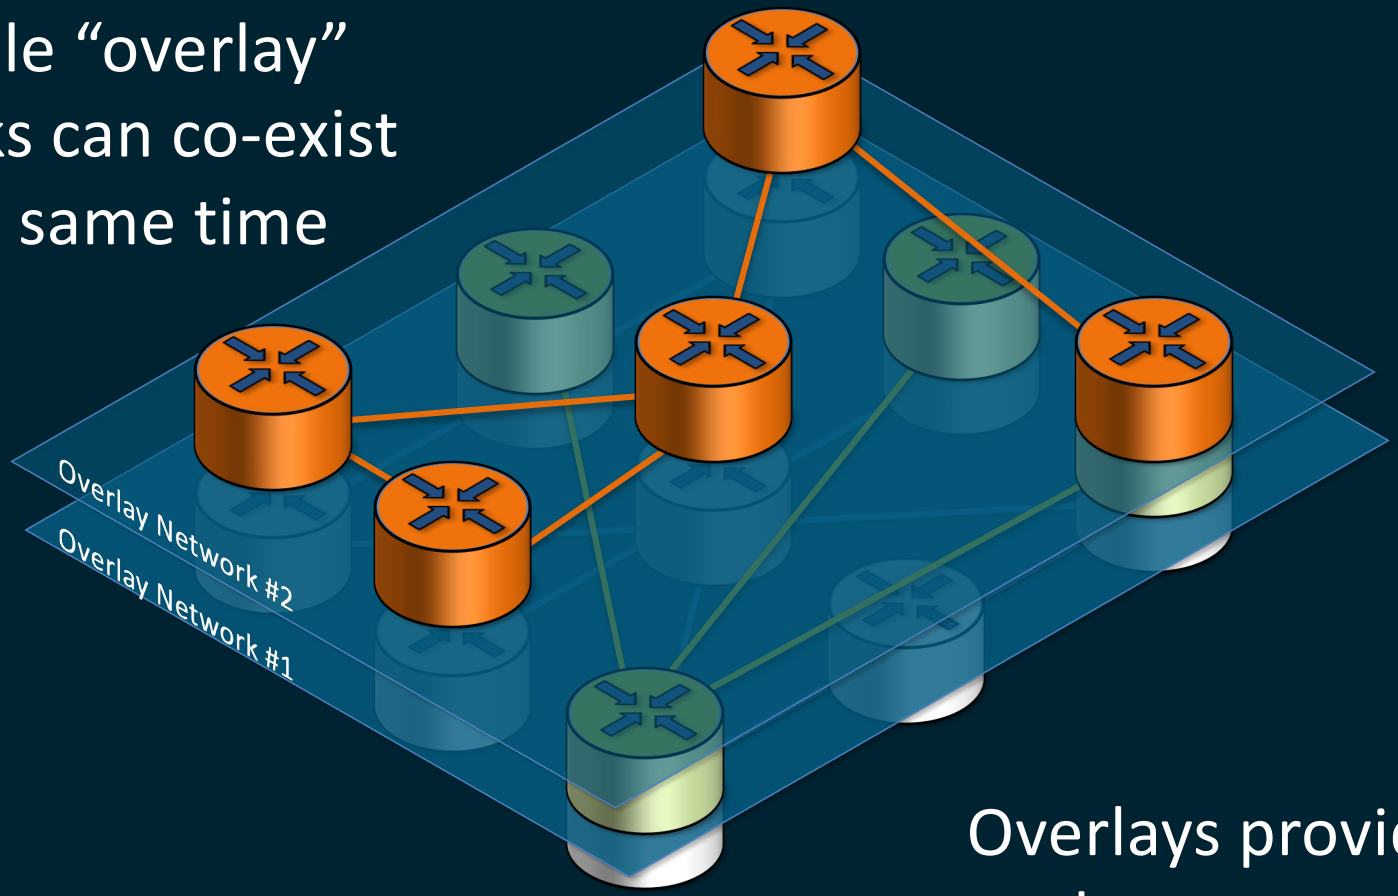

Overlays provides logical network constructs for different tenants (customers)

Main Benefit of Overlays?

Overlay Network can be created and torn down without changing underlying physical network

Nicira has a *Controller*

But its not an Openflow Controller

And they also have OVS

Open vSwitch – V for Virtual

OVS is a fully fledged Switch

Albeit it's a software based switch

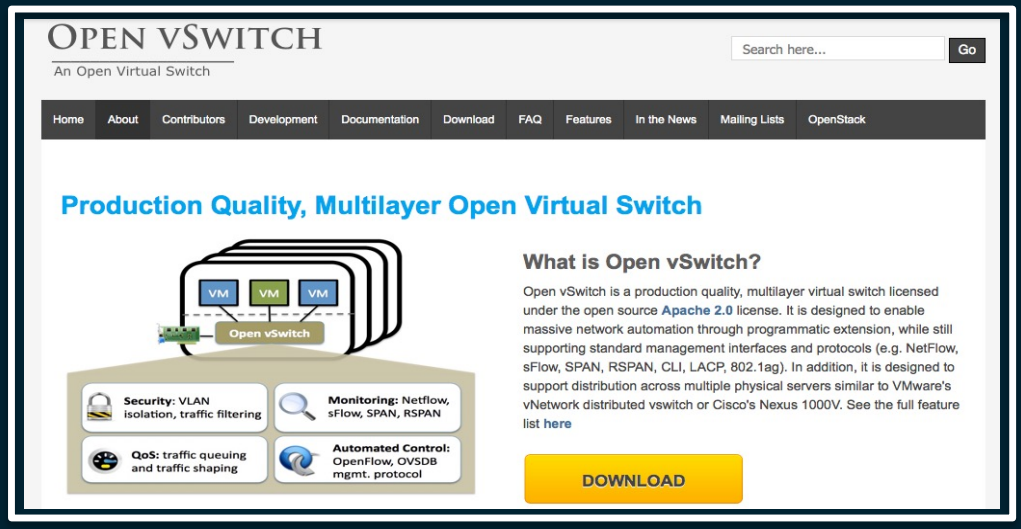

Nicira developed it and threw it out to Open Source

You can read more about it here - http://openvswitch.org/

OVS typically runs on a server**

It's integrated into the OS kernel and extends into the hypervisor

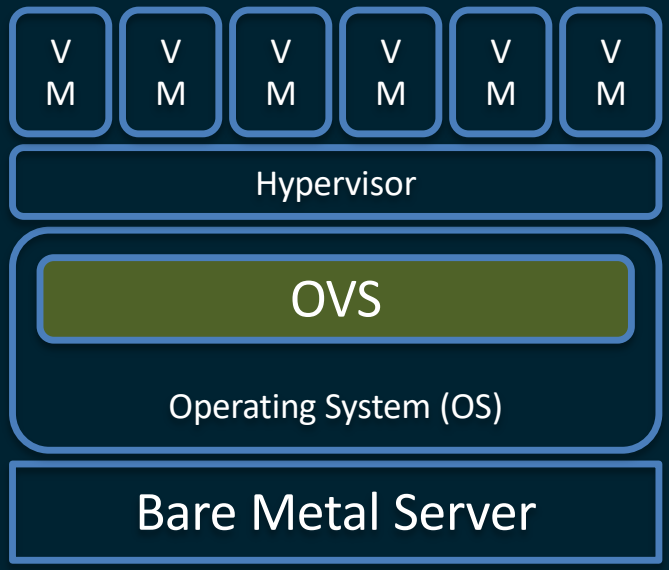

It provides *switching services* for the VM's on the same Bare Metal Server

** OVS can also run on a physical switch


$1.05 Billion Cash
+ $210 Million in stock
vmware® + nicira

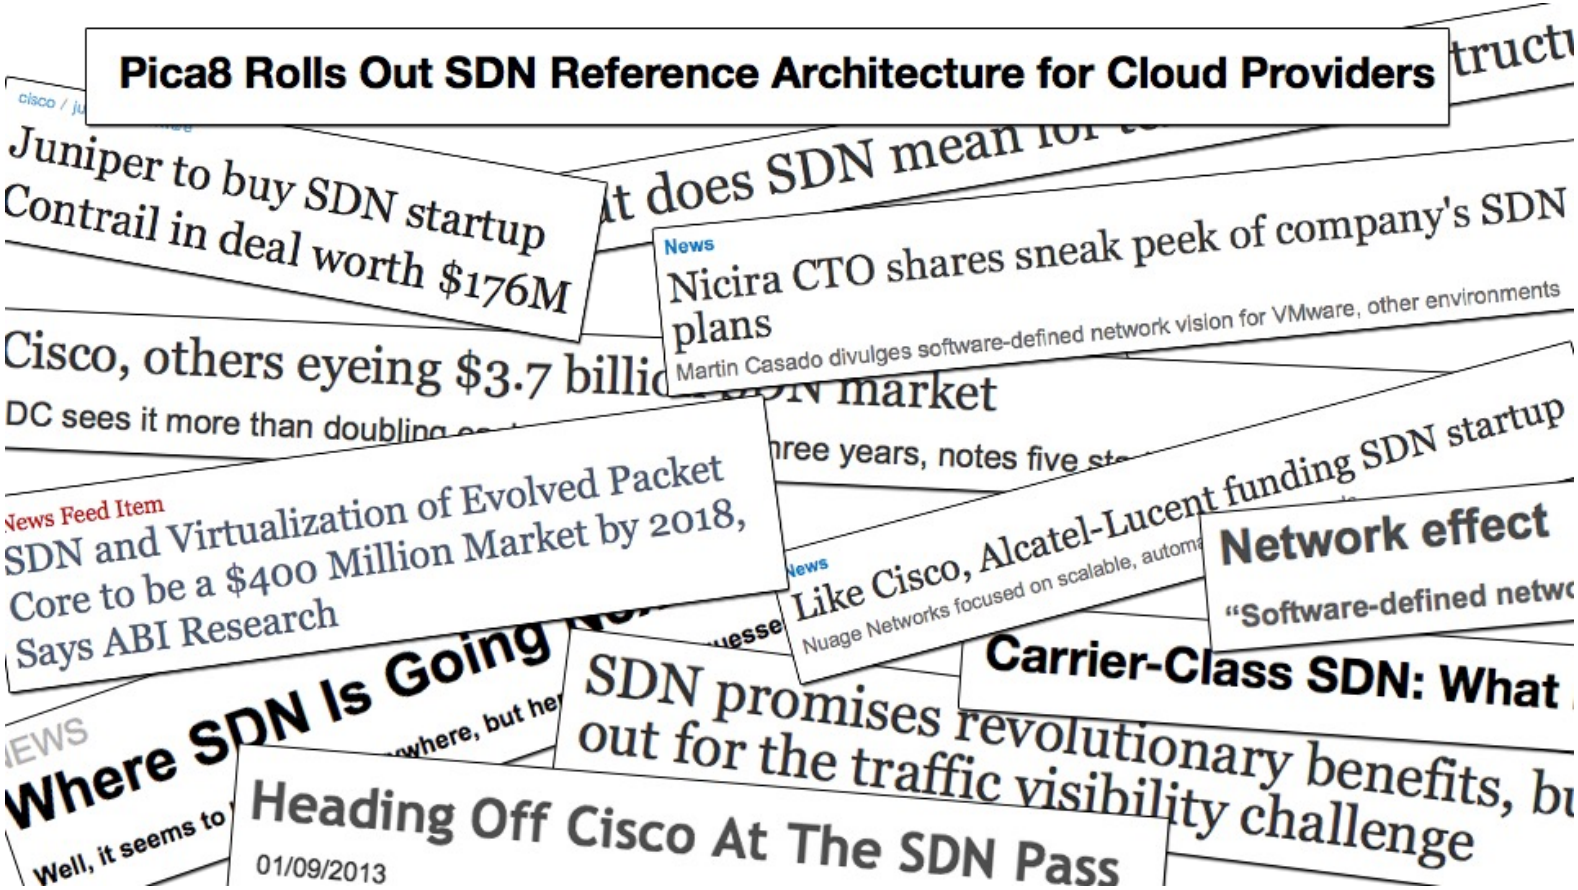

"SDN" gained massive industry mindshare

Cisco Commitment to Programmability

“Everything we build will be programmable”

Chuck Robbins
CEO, Cisco
October 2015

to be Programmable

Application Programming Interface (API)

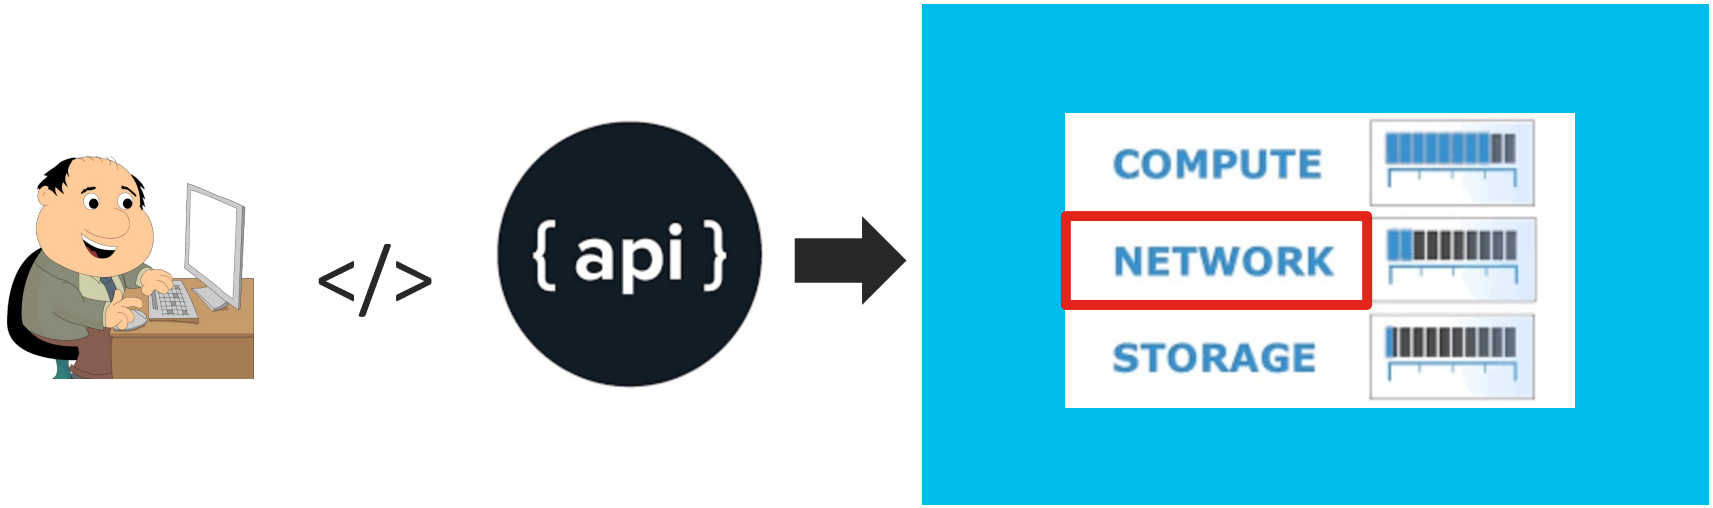

"It's a way for two pieces of software to talk to each other"

What is an API?

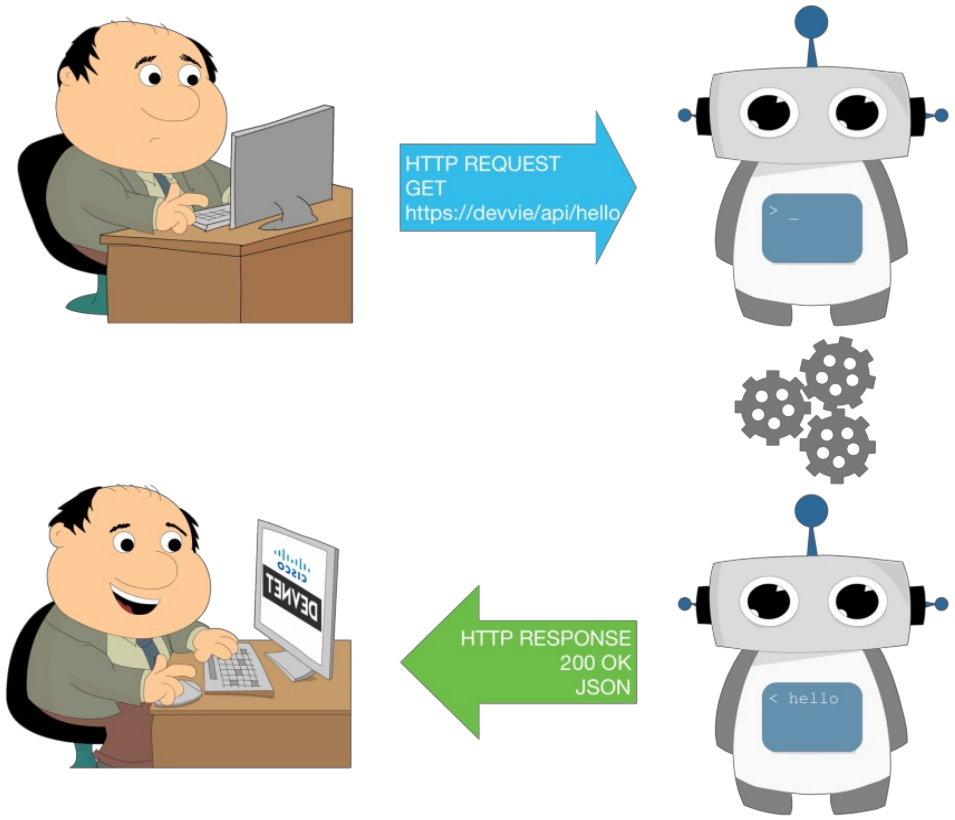

API for Network Configuration

- A Network API for network configuration requires three components:
  - Data Models
  - Data Encoding
  - Transport Protocols

What is a Data Model?

A data model is simply a well understood and agreed upon method to describe "something". As an example, consider this simple "data model" for a person.

- ***Person***
  - **Gender** - male, female, other
  - **Height** - Feet/Inches or Meters
  - **Weight** - Pounds or Kilos
  - **Hair Color** - Brown, Blond, Black, Red, other
  - **Eye Color** - Brown, Blue, Green, Hazel, other

YANG Data Models

- YANG data model: network-centric data modeling language defined in RFC 6020 specifically built for used to model configuration and state data manipulated by the NETCONF protocol, NETCONF operations, and NETCONF notifications
- Used by both Netconf and Restconf
- Human readability is highest priority
- Example YANG vlan definition:

```
list vlan-list {
    key "id";
    leaf id {
     description
       "a single VLAN id (allowed value range 1-4094) \
        or Comma-separated VLAN id range. \
        e.g. 99 or 1-30 or  1-20,30,40-50";
     type union {
       type uint16 {
         range "1..4094";
       }
       type ios-types:range-string;
     }
    }
```

```
leaf name {
      description
        "Ascii name of the VLAN";
      type string {
        length "1..100";
      }
      must "/ios:native/ios:vtp/ios-vtp:version = 3 or string-length(.) <= 32";
    }
    leaf state {
      description
        "Operational state of the VLAN";
      type enumeration {
        enum "active";
        enum "suspend";
      }
    }
```

Encoding Formats

"lightweight, text-based, language-independent ***data interchange formats***"

XML vs JSON

lightweight, text-based, language-independent data interchange formats

<tag>value</tag>

```
<interfaces xmlns:=“[…]yang:ietf-interfaces”>
   <interface>

      <name>eth0</name>
      <type>ethernetCsmacd</type>
      <location>0</location>
      <enabled>true</enabled>
      <if-index>2</if-index>

   </interface>
</interfaces>
```

”key”: ”value”

```
{
   "ietf-interfaces:interfaces": {
      "interface": [
         {
            "name": "eth0”,
            "type": "ethernetCsmacd”,
            "location": "0”,
            "enabled": true,
            "if-index": 2
         }
      ]
   }
}
```

Transport Protocols

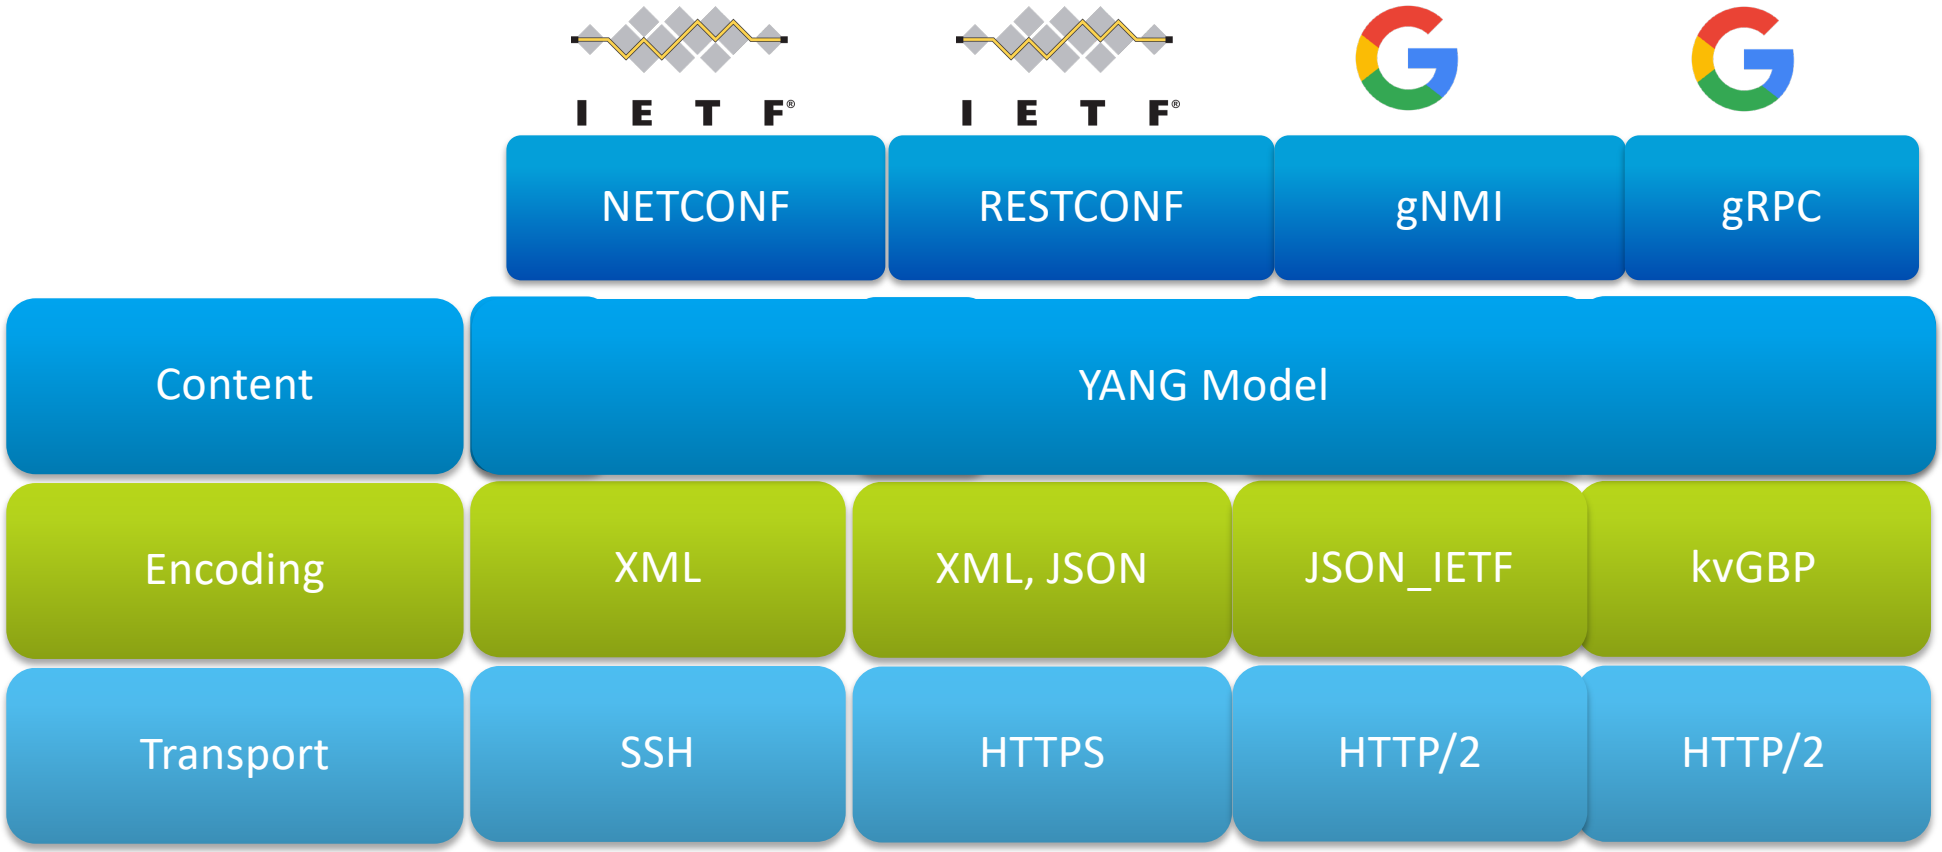

Agenda

- SDN Introduction
  - Network Functions Virtualization
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Traditional Physical Networks

- Shared resource supported many “service instances”
- Upfront procurement /purchase
- High Opex - difficult to automate or apply consistent policy
- Rigorous processes to maintain 99.999% uptime (change & release management)
- Easier to add a service than remove one (orphaned FW ACL rule)
- “Peak load” capacity planning model required
- Slow and expensive to create new **service-design** or add new **service-features**
- Operational “domaining” necessary to manage complexity – entrenched “silo’d” approach

What is NFV (Network Functions Virtualization)?

- "Virtualize" some network function. That's it.
- These "Virtualized Network Functions" are called VNFs.
- ETSI defines VNF as a "Service"…

NFV: This whole process of virtualizing things

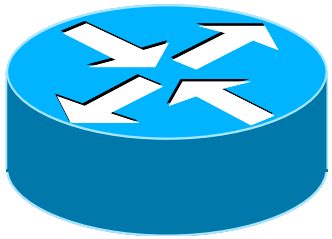

Simple Example

Cisco Router

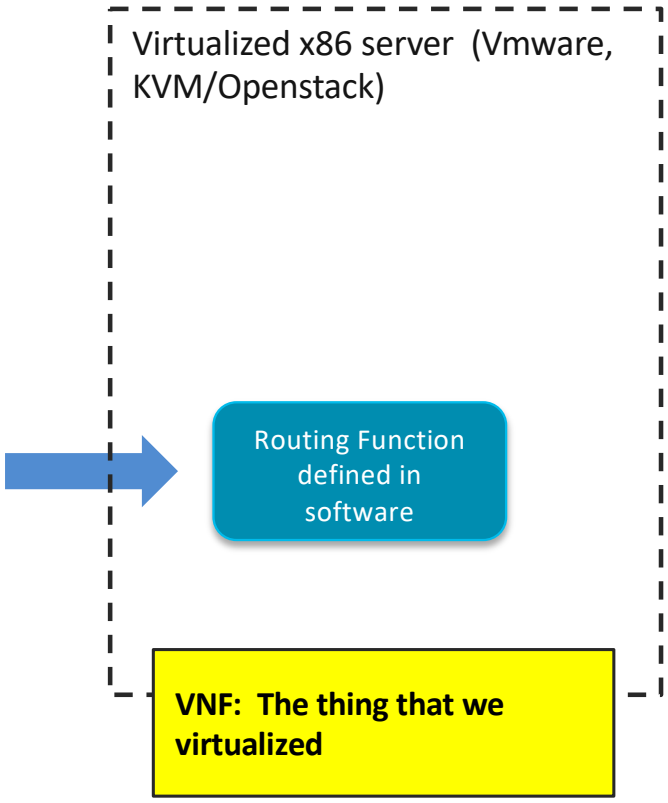

Network Functions Virtualisation (NFV) – Initial Goals

NFV = Transition of network infrastructure services to run on virtualised compute platforms – typically x86

Enablers

- Hypervisor and cloud computing technology
- Improving x86 h/w performance
- Optimised packet processing and coding techniques
- Network industry standardising on Ethernet
- SDN based orchestration

Value Proposition

- Shorter innovation cycle
- Improved service agility
- Reduction in CAPEX and OPEX

Applications

- Potentially all network functions

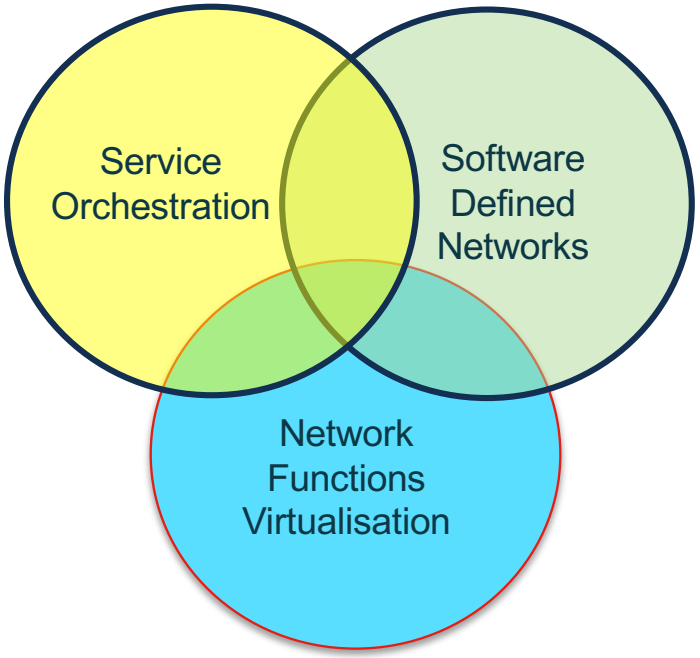

Extract from "Network Functions Virtualisation – Introductory White Paper"

SDN and NFV

- Perform different functions
- Can and DO co-exist

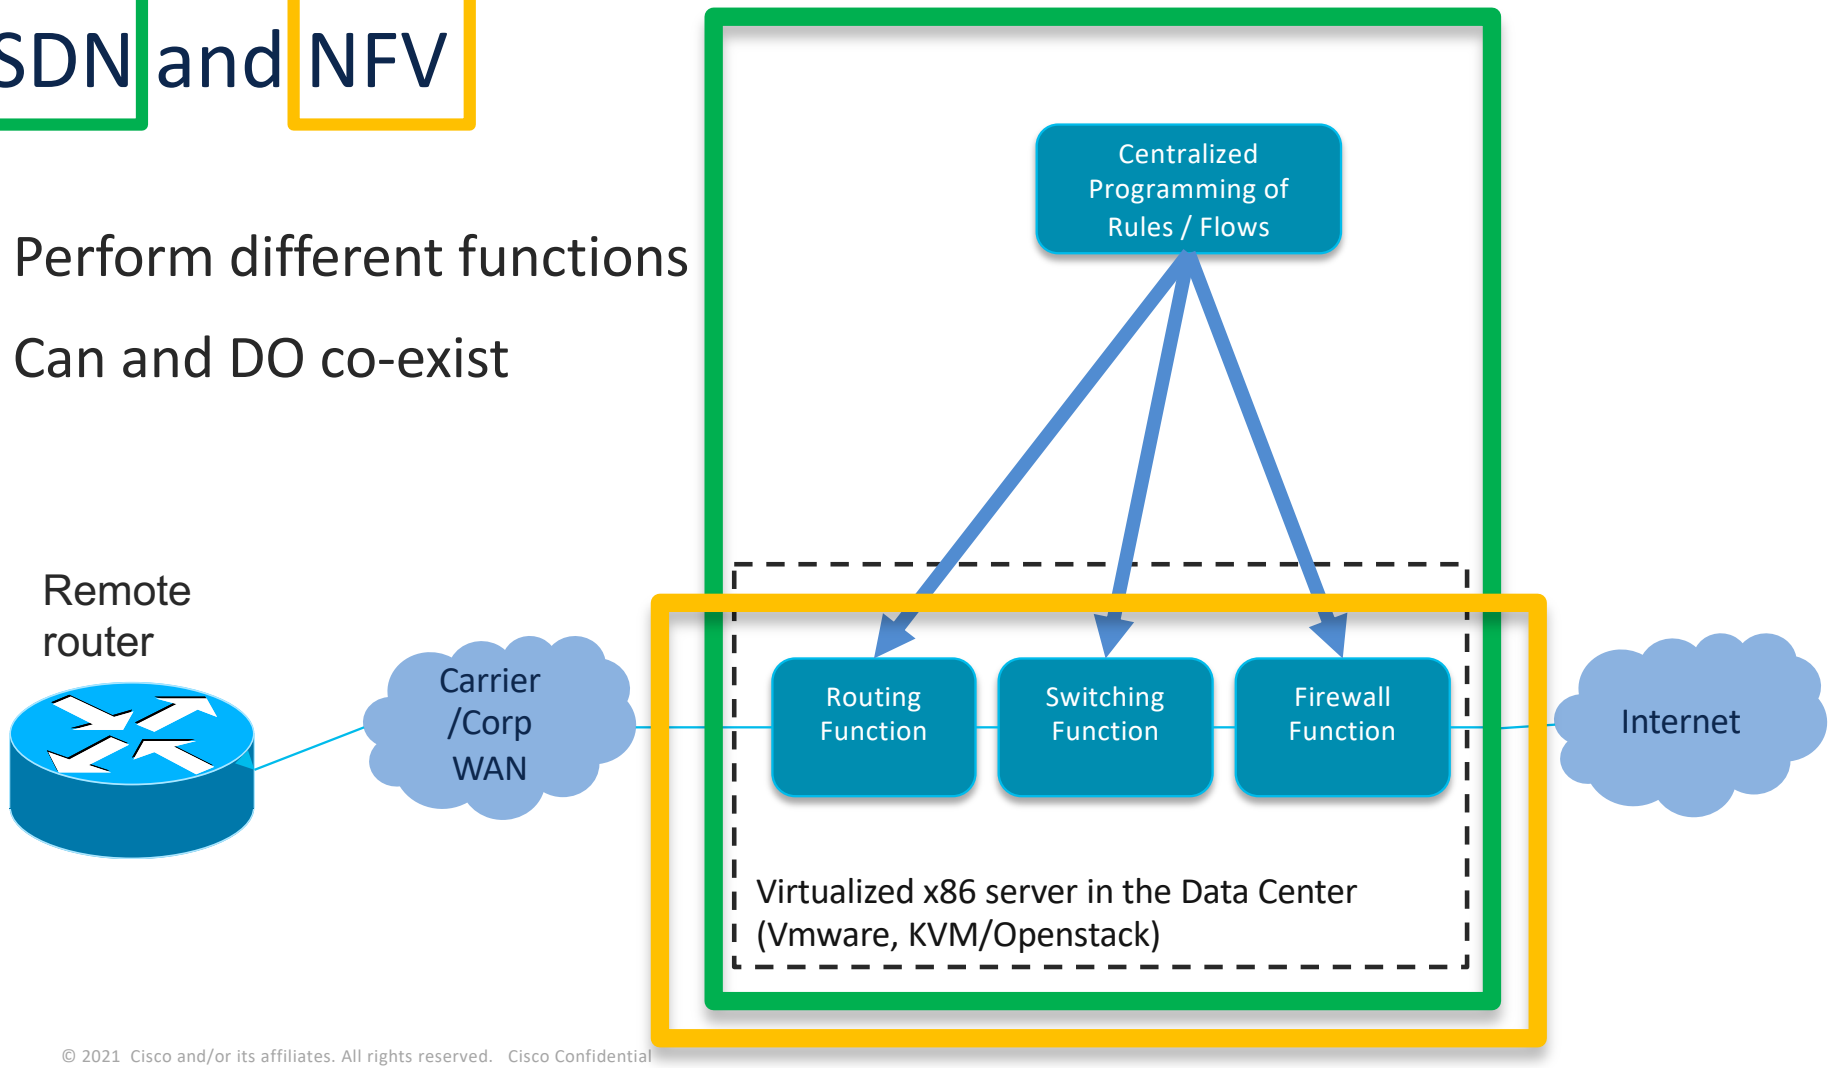

NFV and SDN

Category	SDN	NFV
Why/what	Separation of control and data, centralisation of control and dynamic programmability of Network	Relocation of network functions from dedicated appliances to generic servers
Primary Verticals focus	Campus, data centre, cloud	Service Provider network
Hardware focus	Proprietary or commodity servers and switches	Commodity servers and switches
Initial Applications	Cloud orchestration and networking	Routers, firewalls, gateways, CDN, WAN accelerators, SLA assurance etc.
New Protocols	OpenFlow	Multiple programmability options
Standards Body	Open Networking Forum (ONF)	ETSI NFV Working Group

How do we build NFV? Use ETSI Model

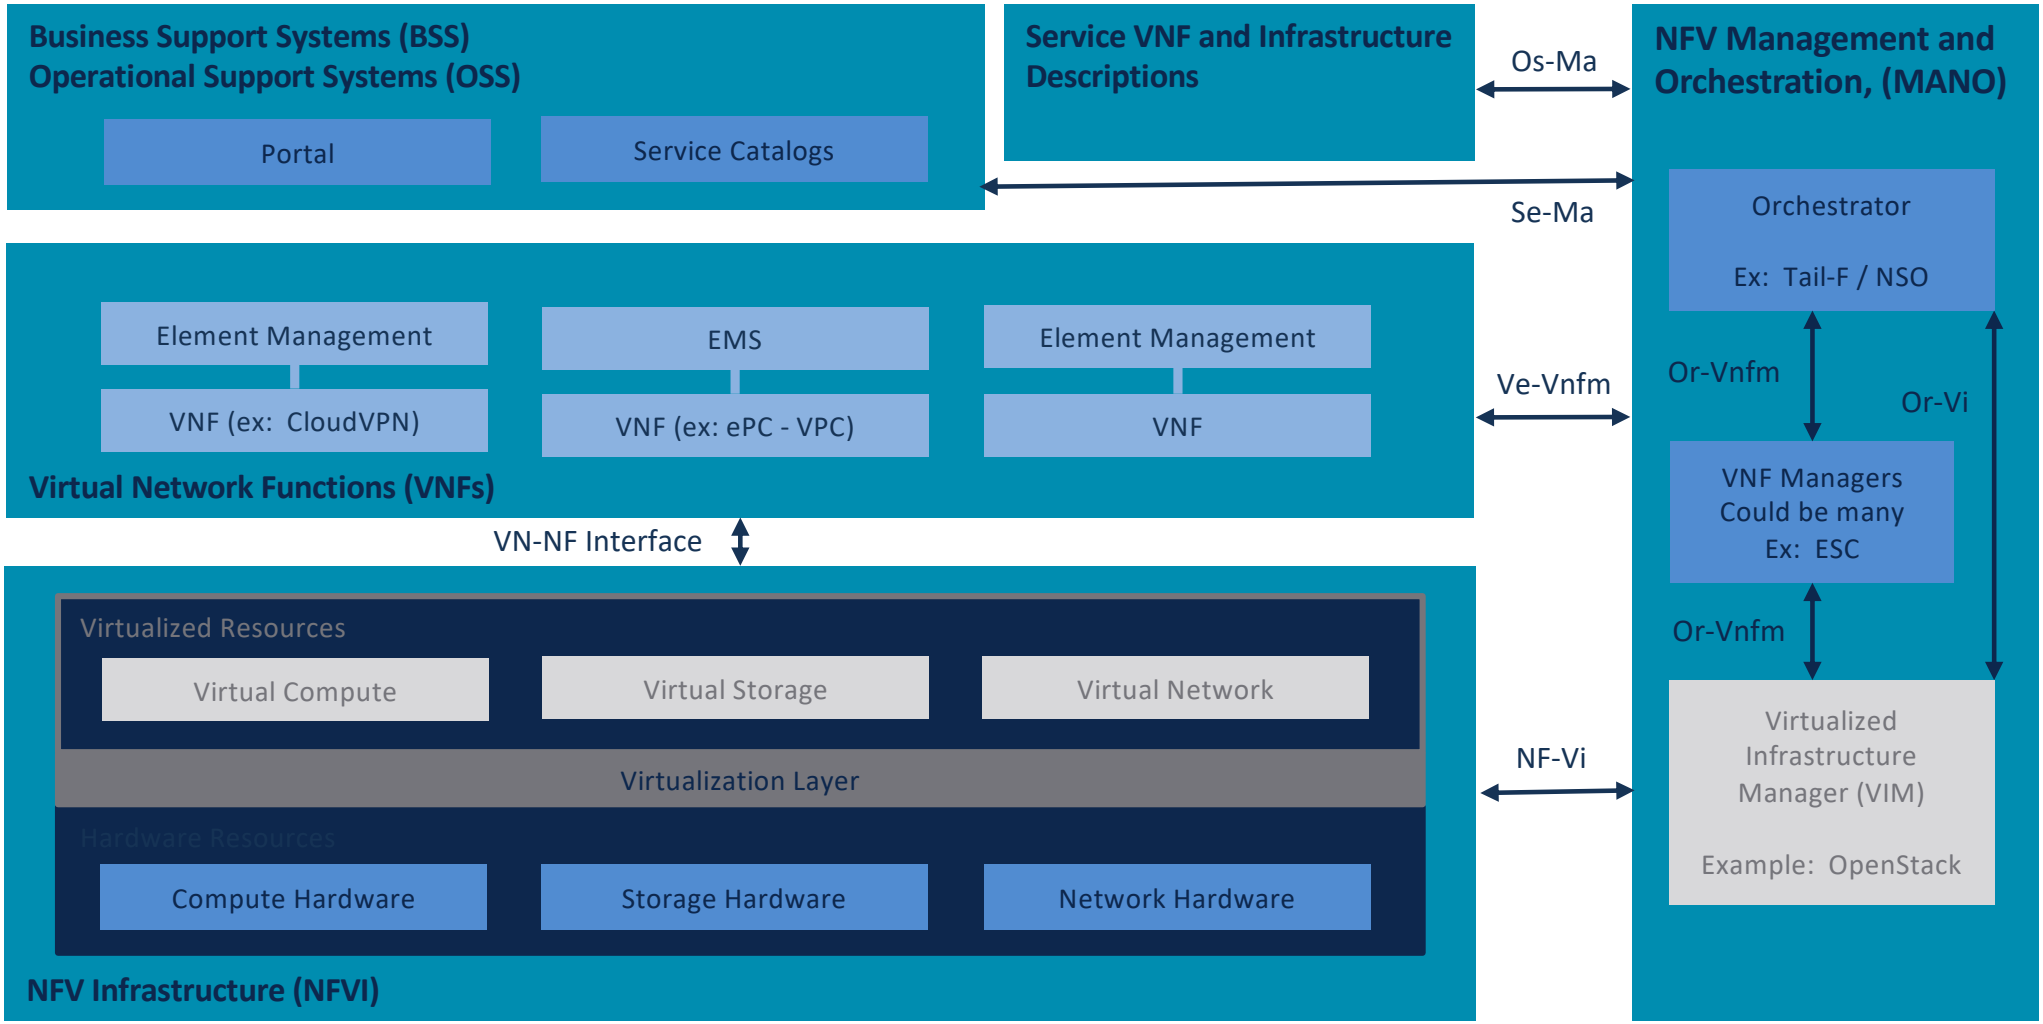

Agenda

- SDN Introduction
  - Openstack
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Openstack is an IAAS (Infrastructure As A Service) cloud computing project

It is also referred to as a Cloud Operating System

"…provides a means to control (administer) compute, storage, network and virtualization technologies…"

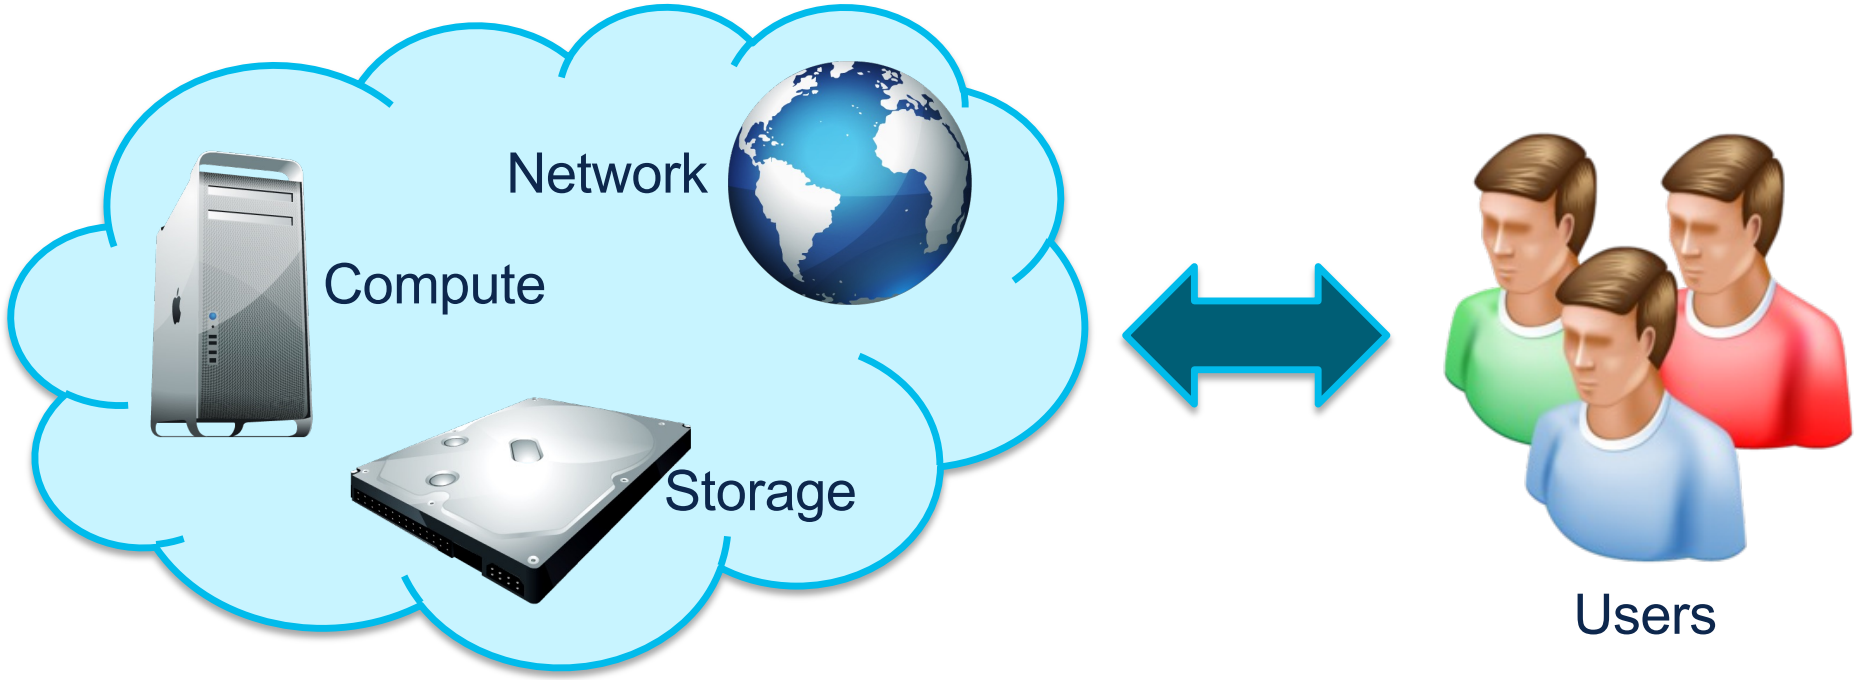

Cloud Computing provides a set of resources and services through the internet

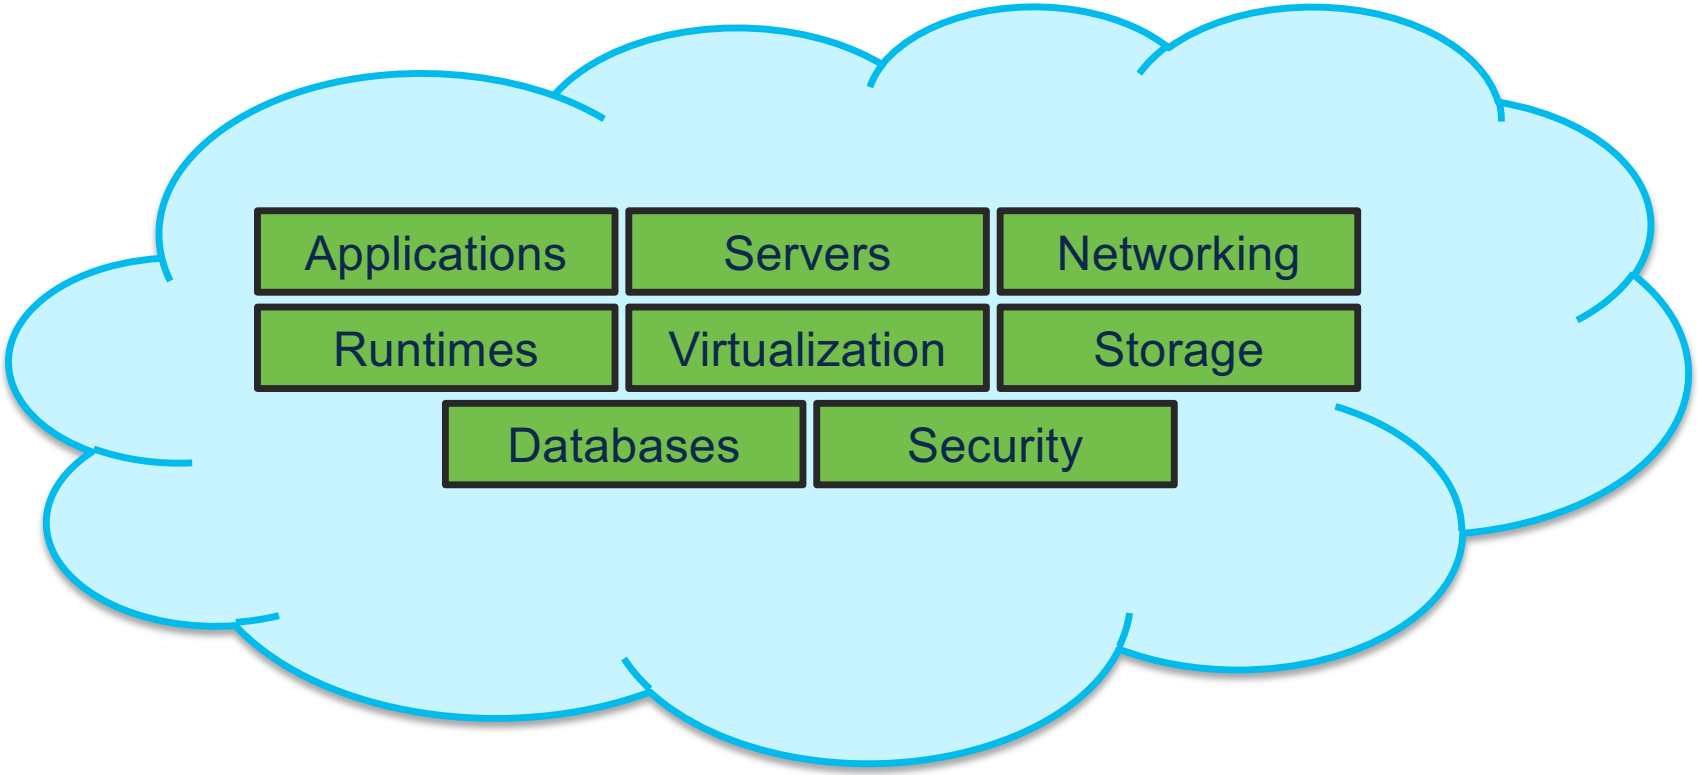

At a more detailed level, there are many resources inside the cloud

What resources you manage inside the cloud defines the following…

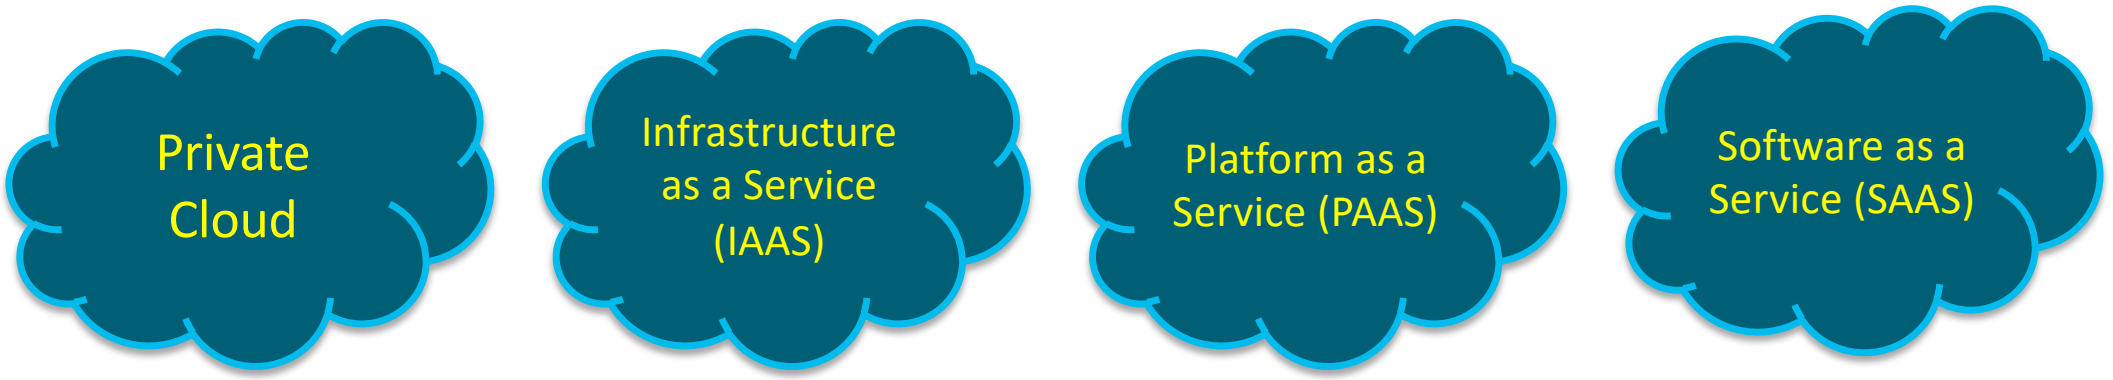

How do these differ from one another?

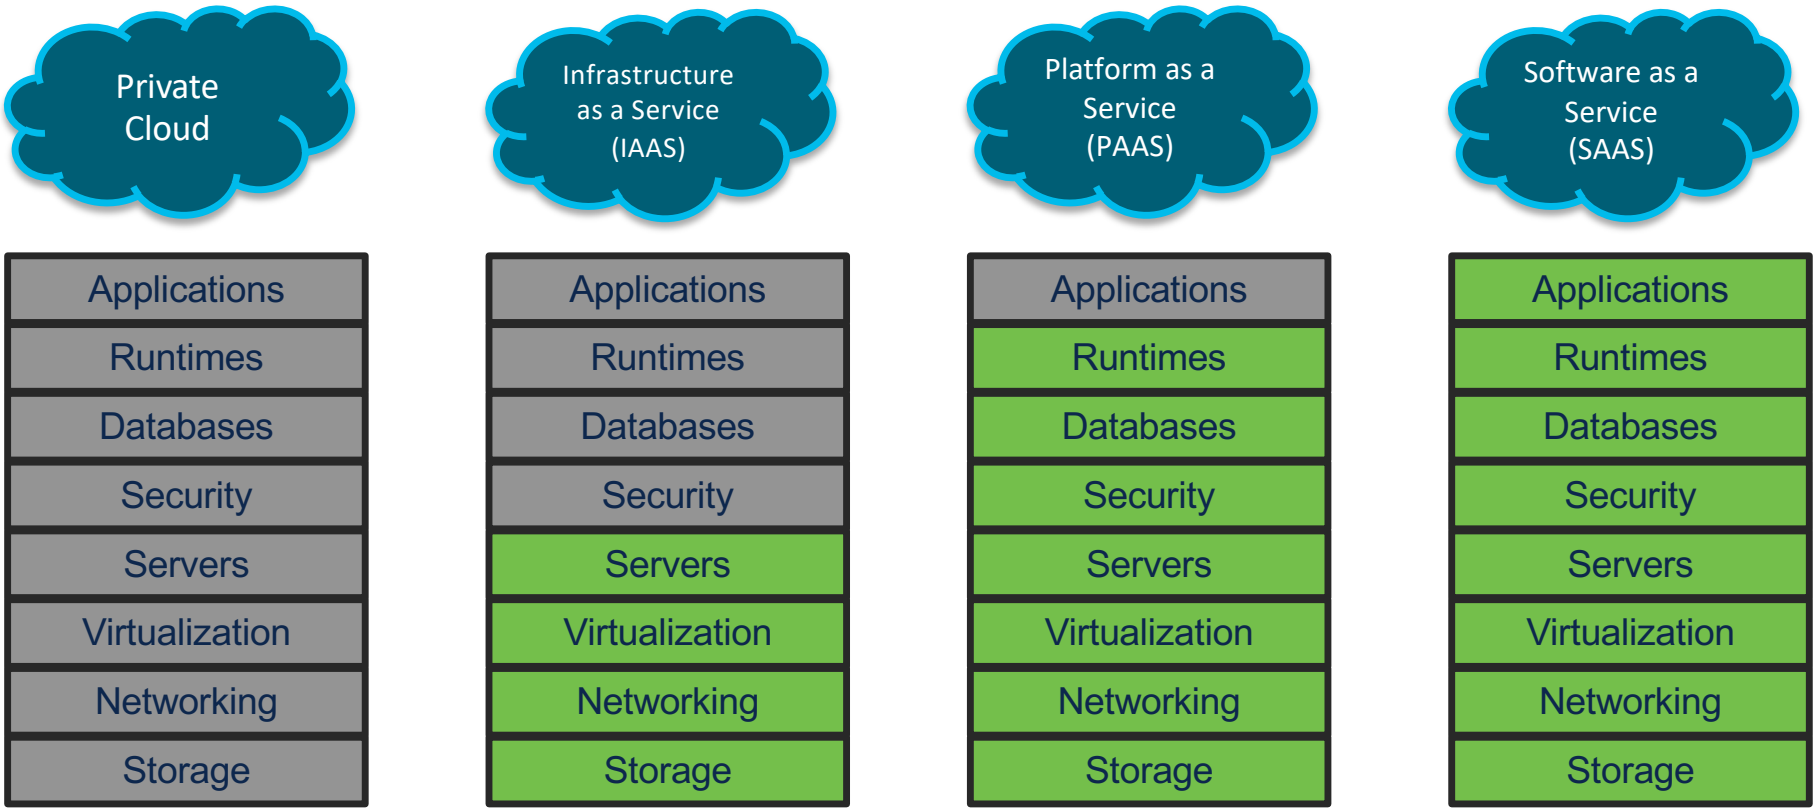

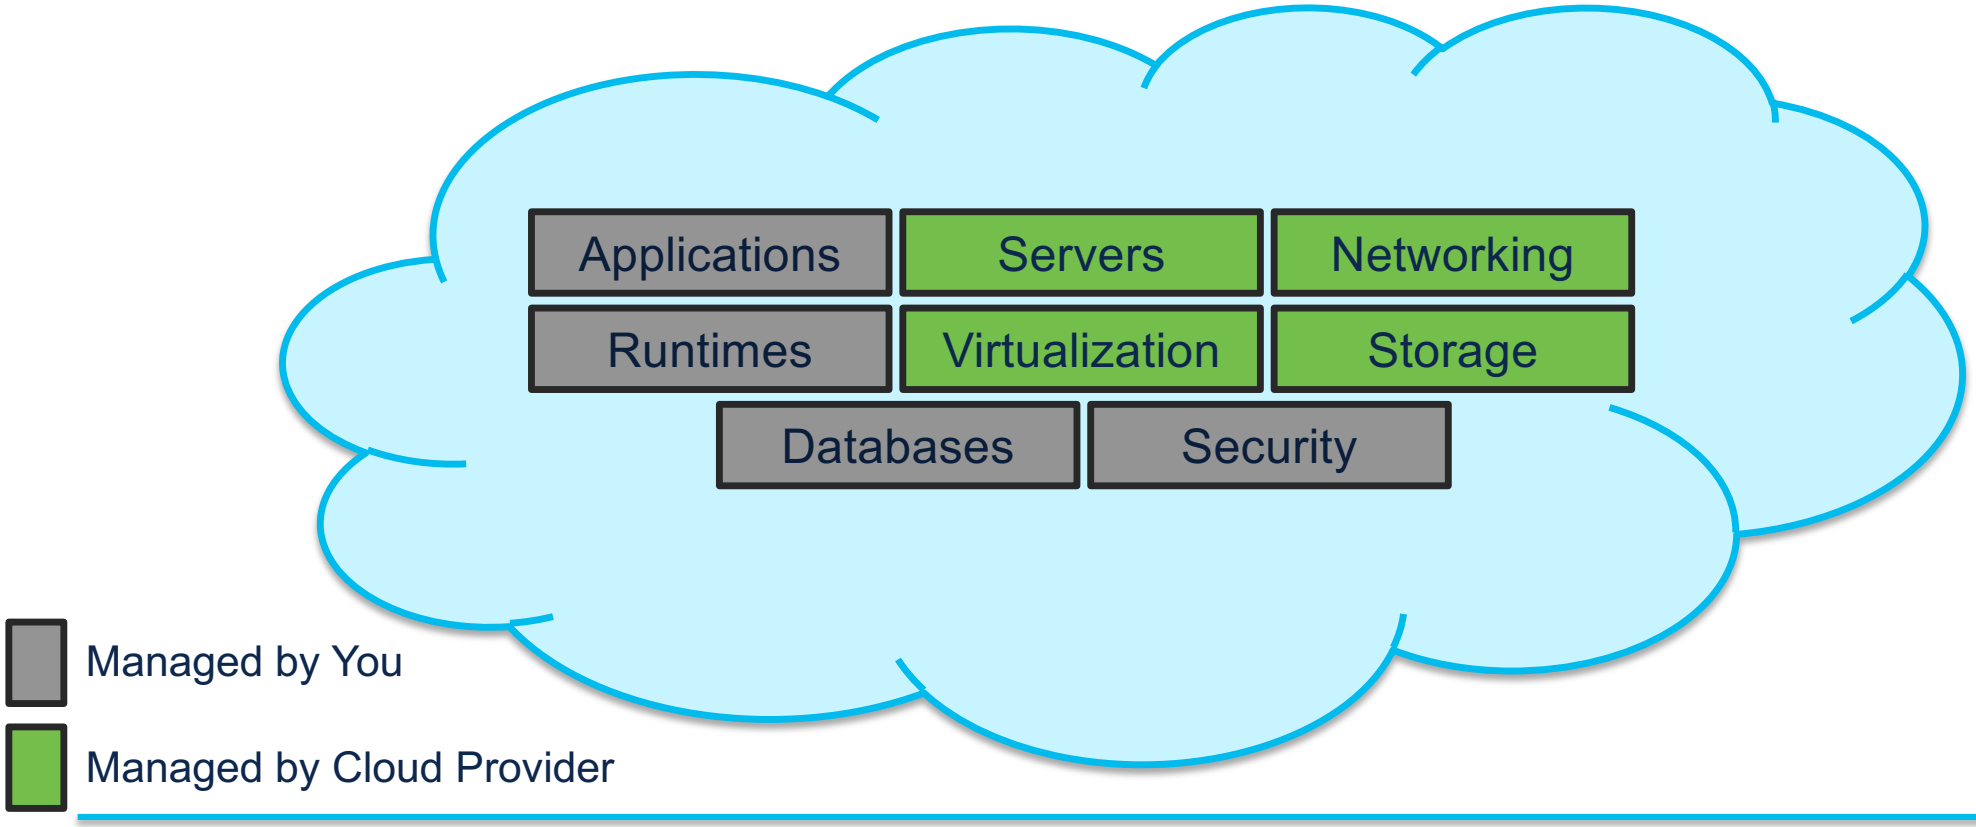

With IAAS, compute, storage, networking and virtualization resources are managed by the Cloud Provider (this defines them as an IAAS provider)

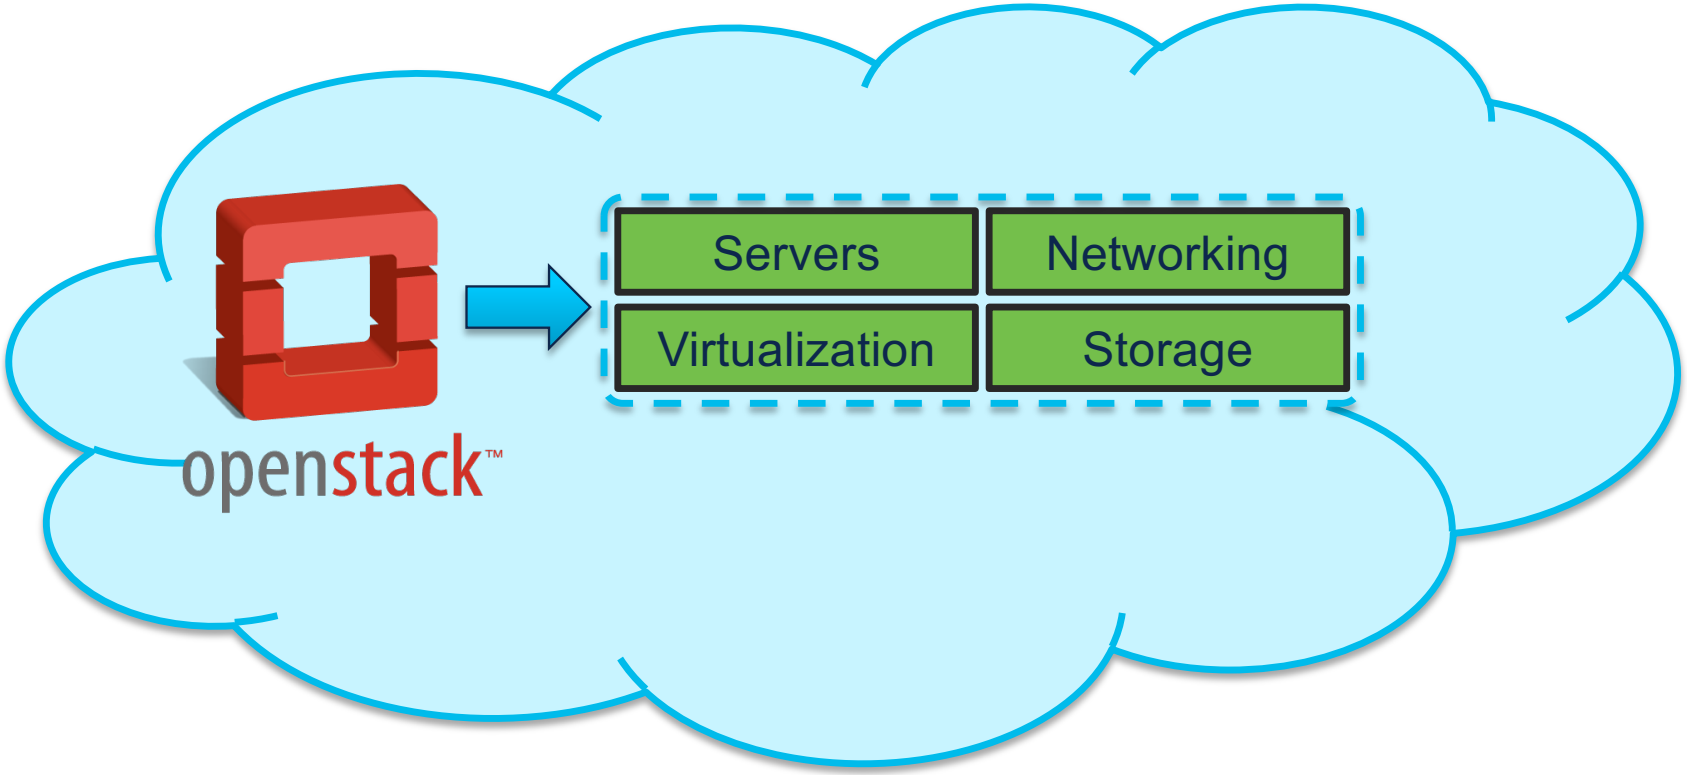

Openstack lets the provider manage these resources

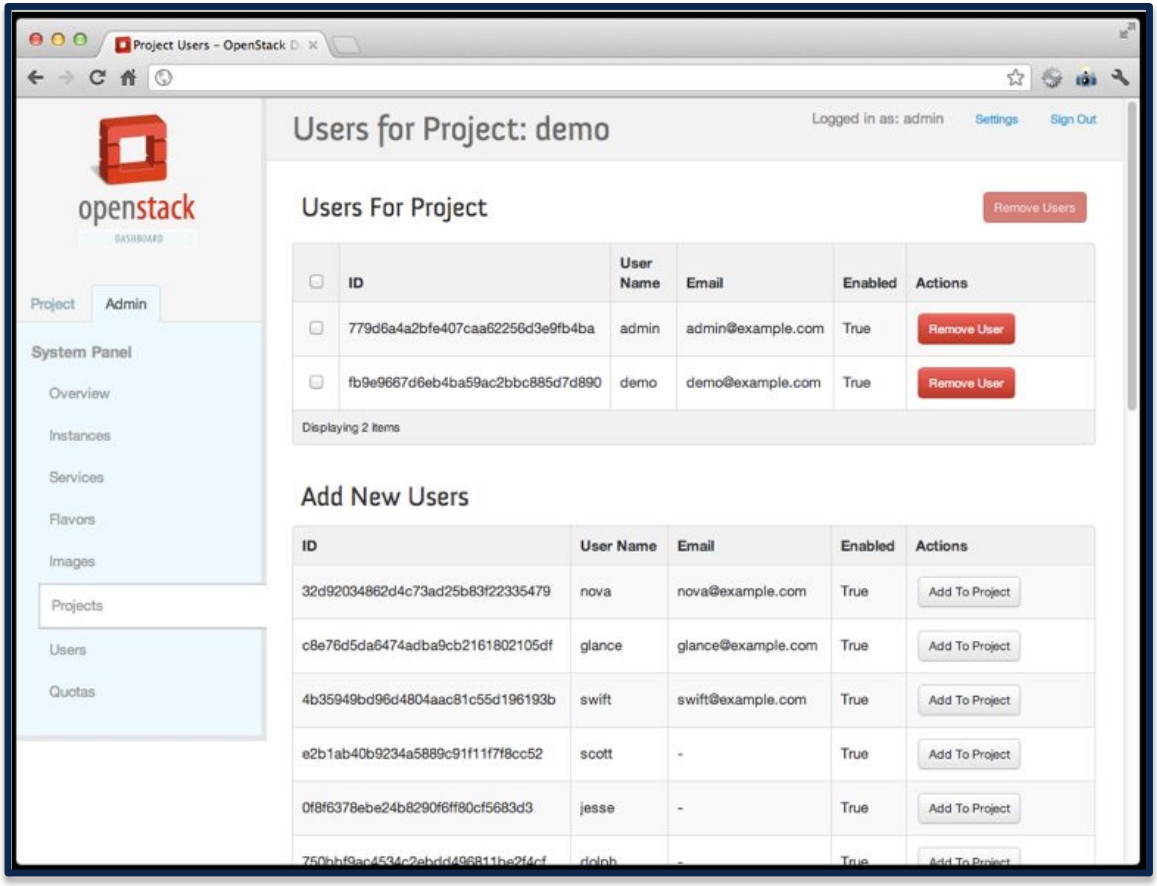

Openstack provides a nice web based front end to manage those cloud resources…

Openstack consists of a number of components

Openstack Compute (NOVA)

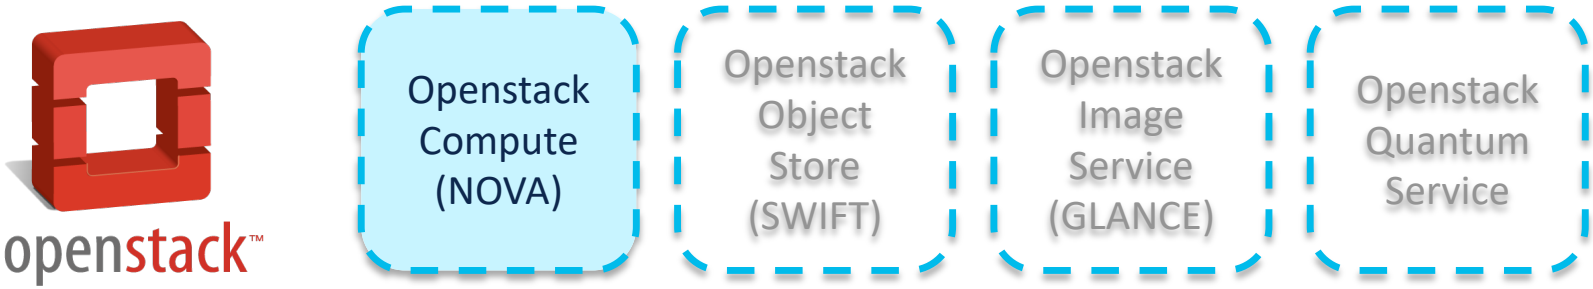

Allows the administrator to create and manage Virtual Machines (VM's) using various (stored) machine images

Object Store (SWIFT)

Provides the ability to store objects – basically it is the component that is responsible for managing storage and reading/writing objects to that storage

An object could be a video file, a document, a picture, a database… basically anything that you would normally store on your computer

Image Store (GLANCE)

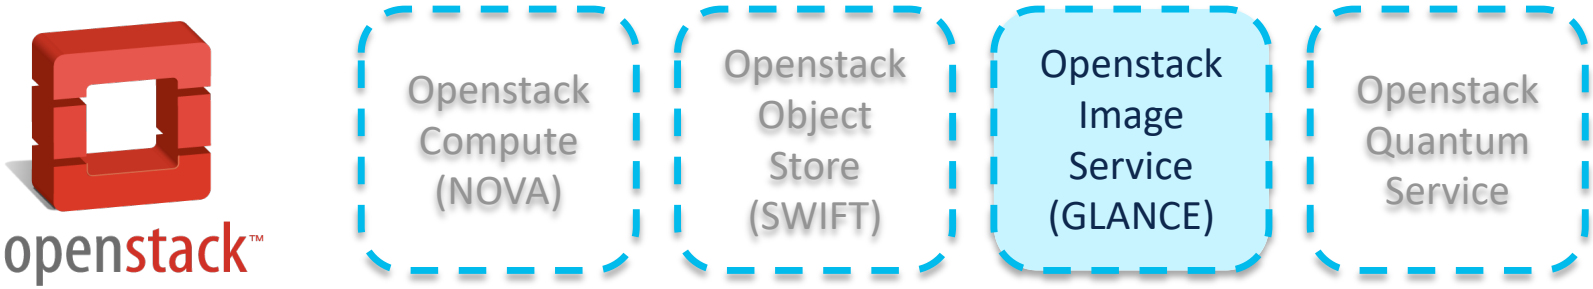

This is the component responsible for managing the different operating system images (Windows, Linux, etc) that NOVA uses to create virtual machine's

Network Service (QUANTUM)

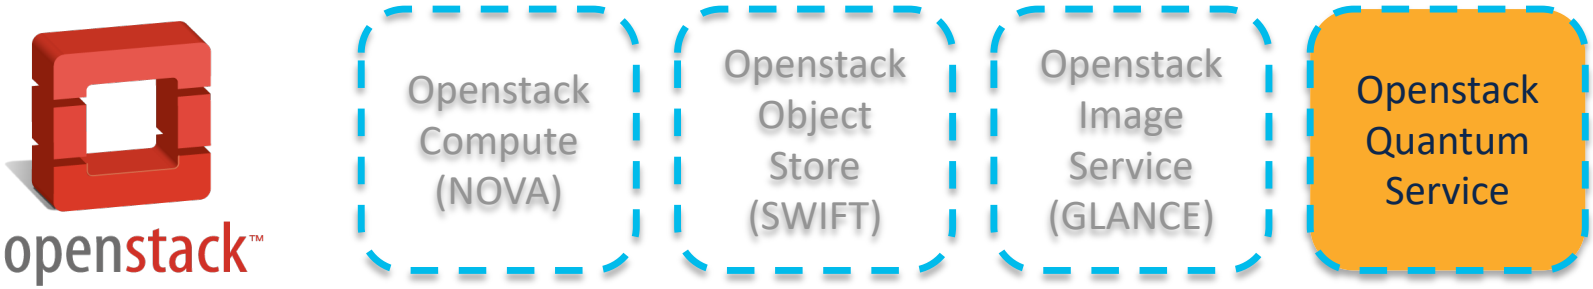

Allows the administrator to create and manage virtual networks

This is the piece that has relevance to our SDN story

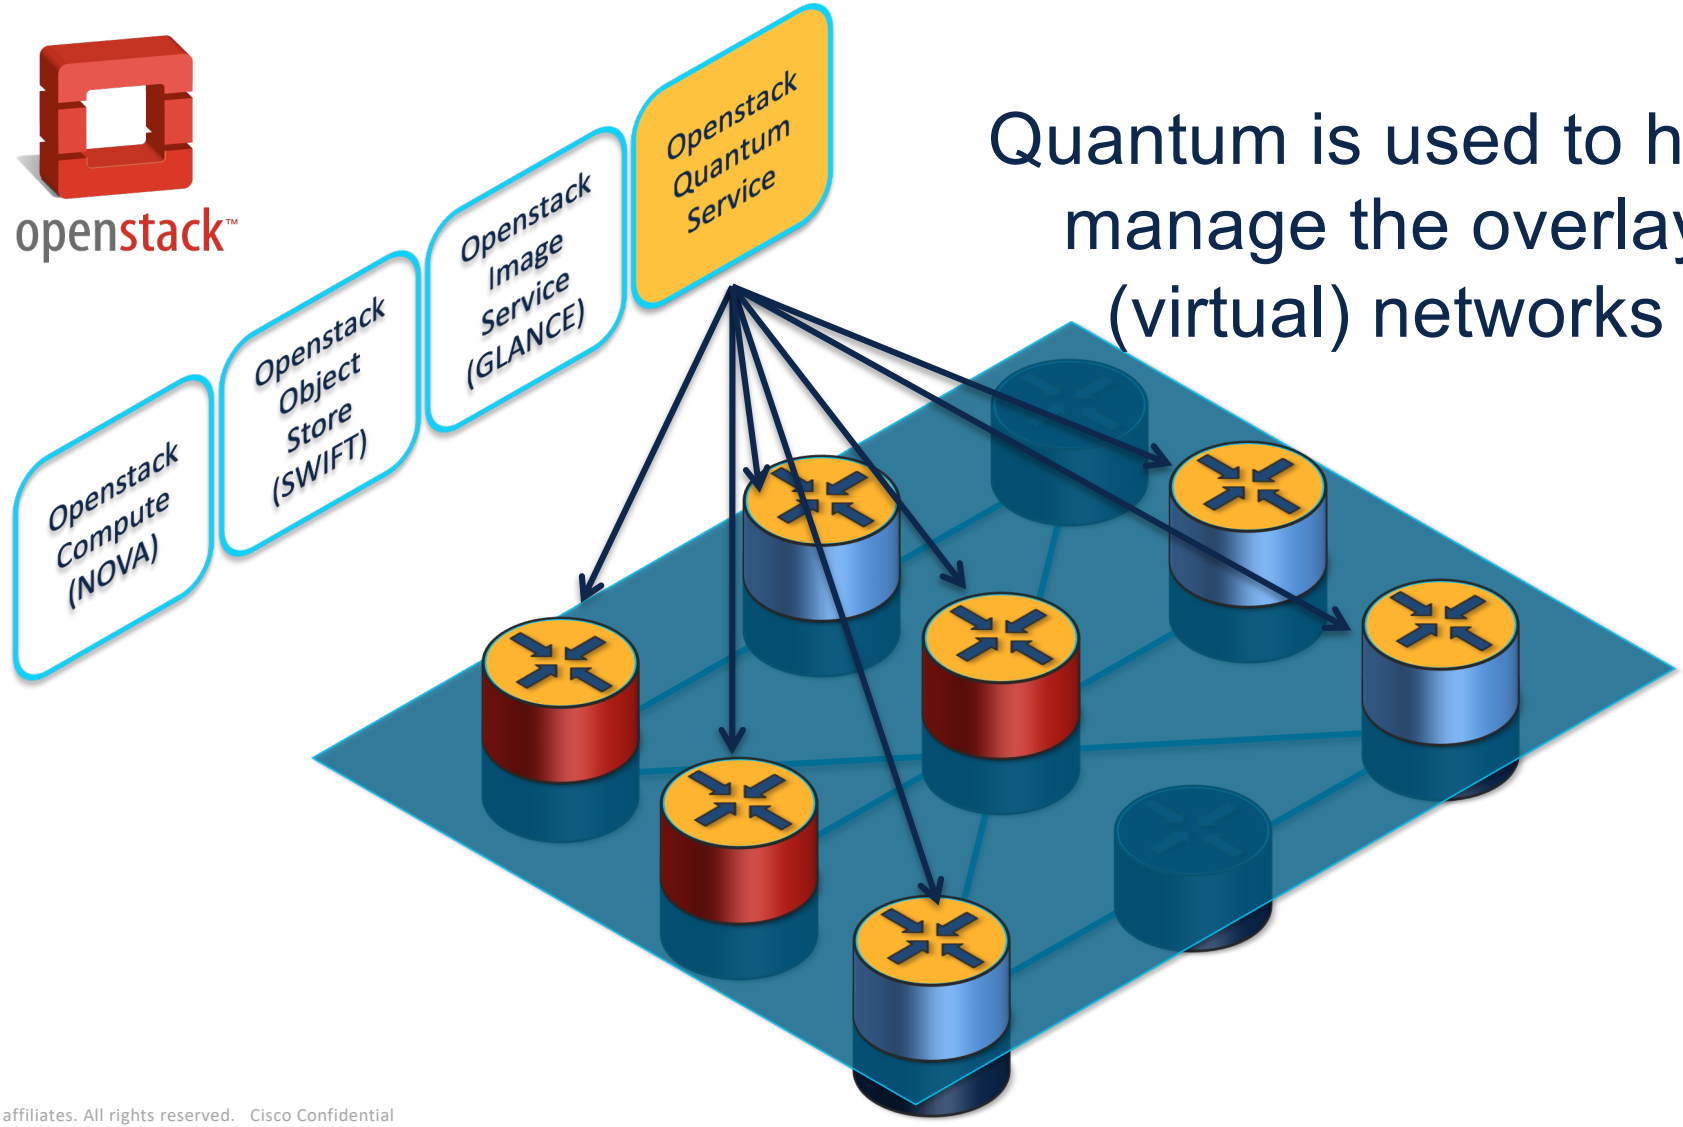

Quantum is used to help manage the overlay (virtual) networks

Agenda

- SDN Introduction
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Cisco’s Enterprise SDN Strategy

Policy and Intent to Unlock the Power of your Distributed System

Unlock the Power that Exists in the Network through **Abstraction, Automation, and Policy Enforcement**

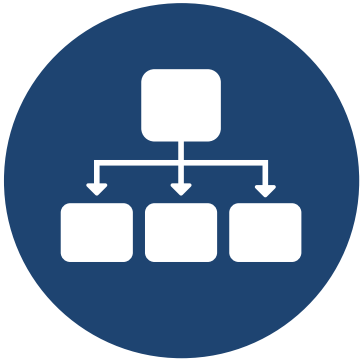

Leverage the Power of Existing **Distributed Systems**

Enable Network Wide Fidelity to an Expressed Intent **(Policy)**

The Cisco Approach to SDN Implementation

Control Plane — Where/How to Send packet

Data plane — Forwarding Packets

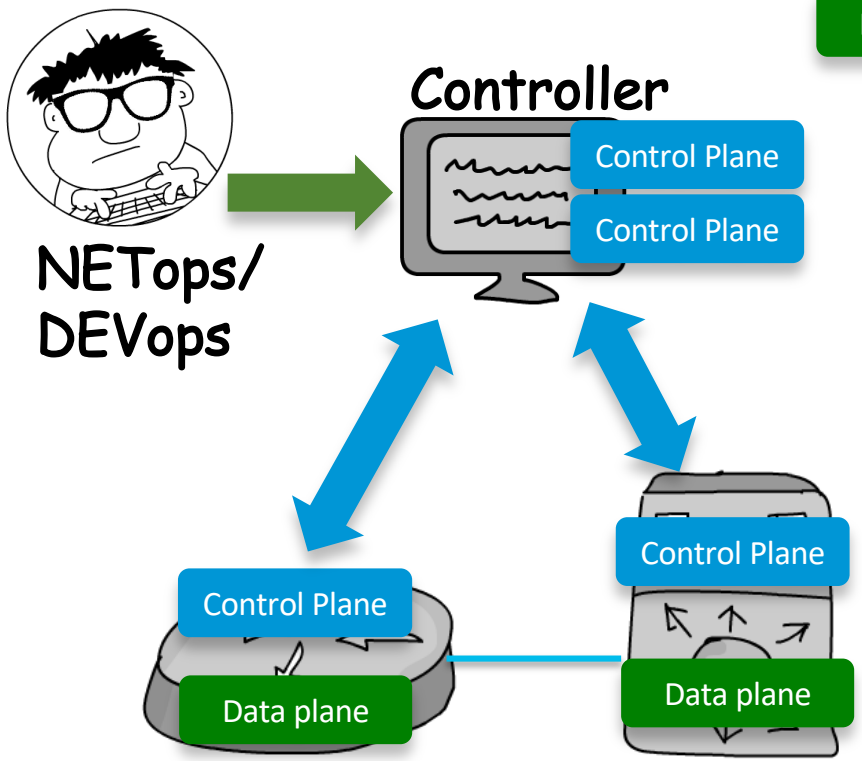

Evolution **NOT REVOLUTION**

EVOLVE FOR EMERGING REQUIREMENTS	• Operational Simplicity • Programmability • Application Aware
PRESERVE WHATS WORKING	• Resiliency • Scale & Security • Rich Feature Set

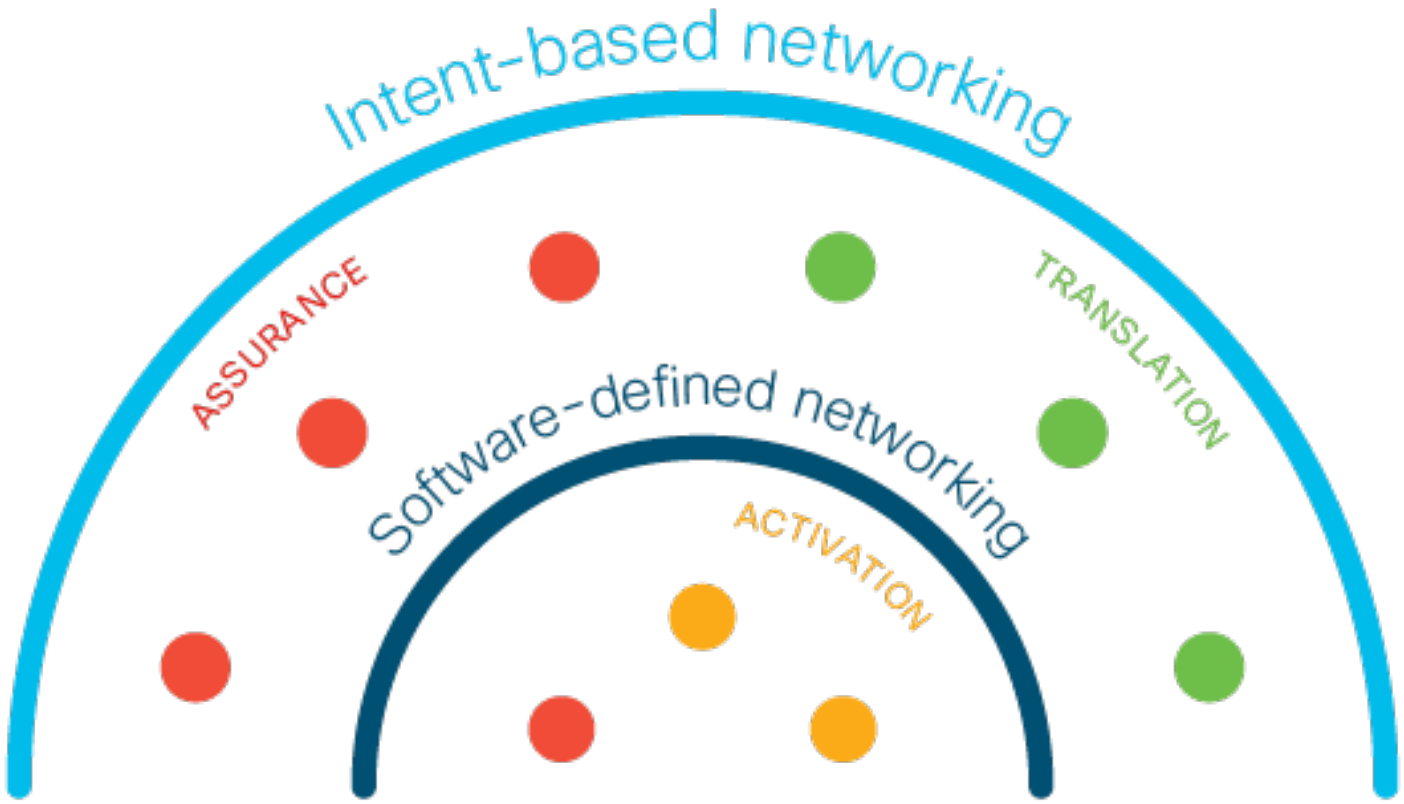

https://blogs.cisco.com/analytics-automation/why-is-intent-based-networking-good-news-for-software-defined-networking

The Intent-based Networking Model

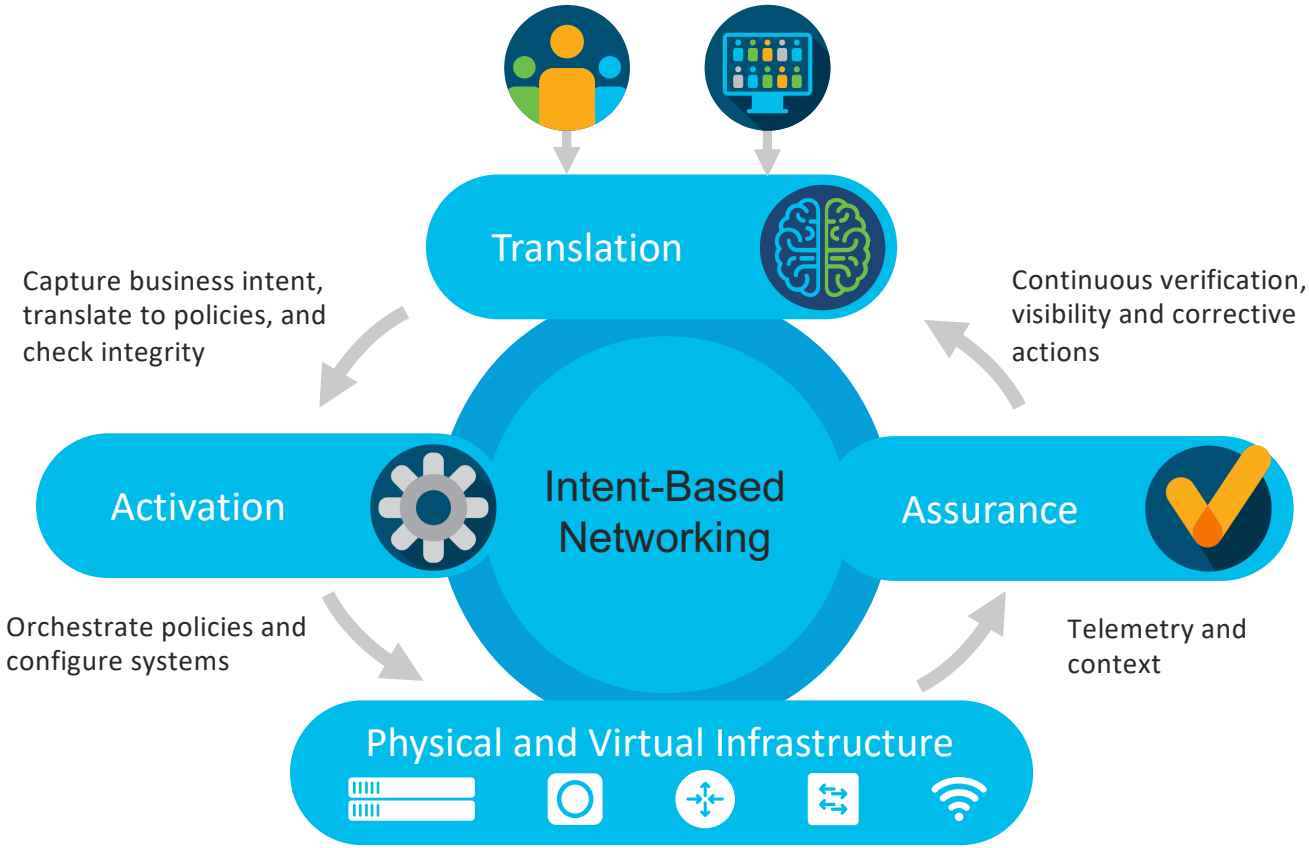

Cisco Digital Network Architecture

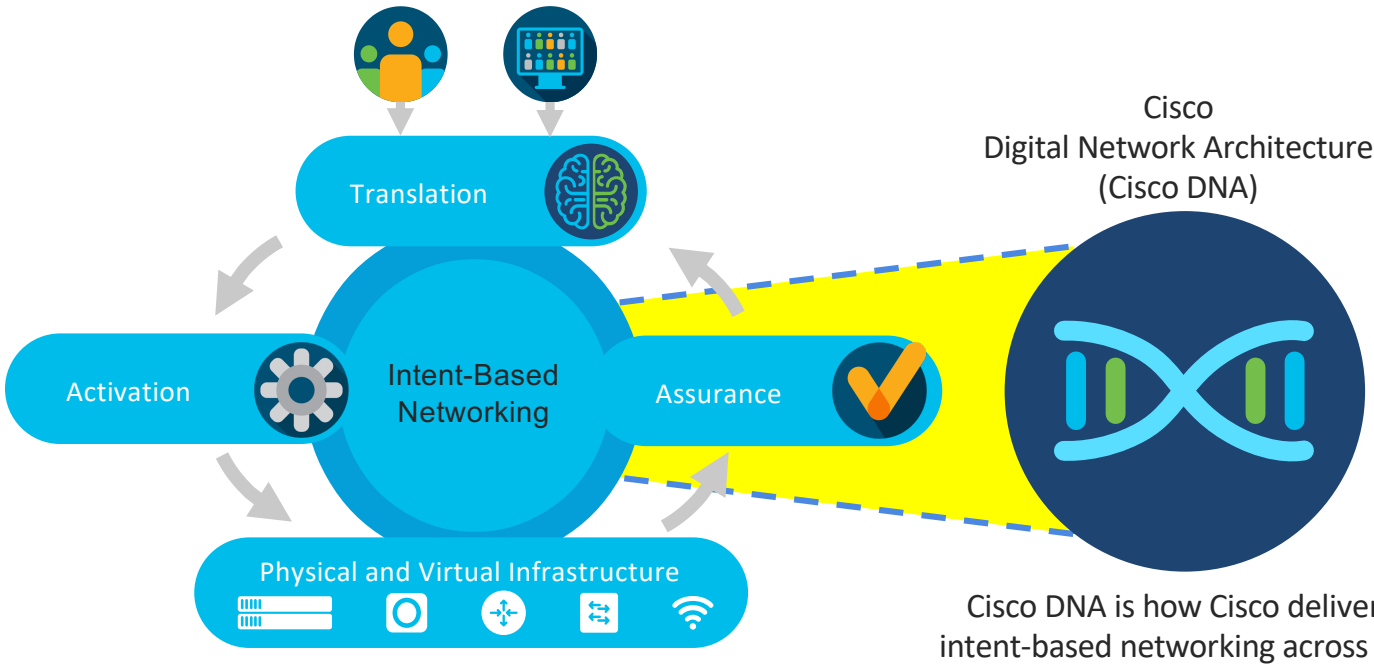

Cisco
Digital Network Architecture
(Cisco DNA)

Cisco DNA is how Cisco delivers intent-based networking across the campus, branch, WAN, and extended enterprise

Intent-Based Networking Multi-Domain Integration

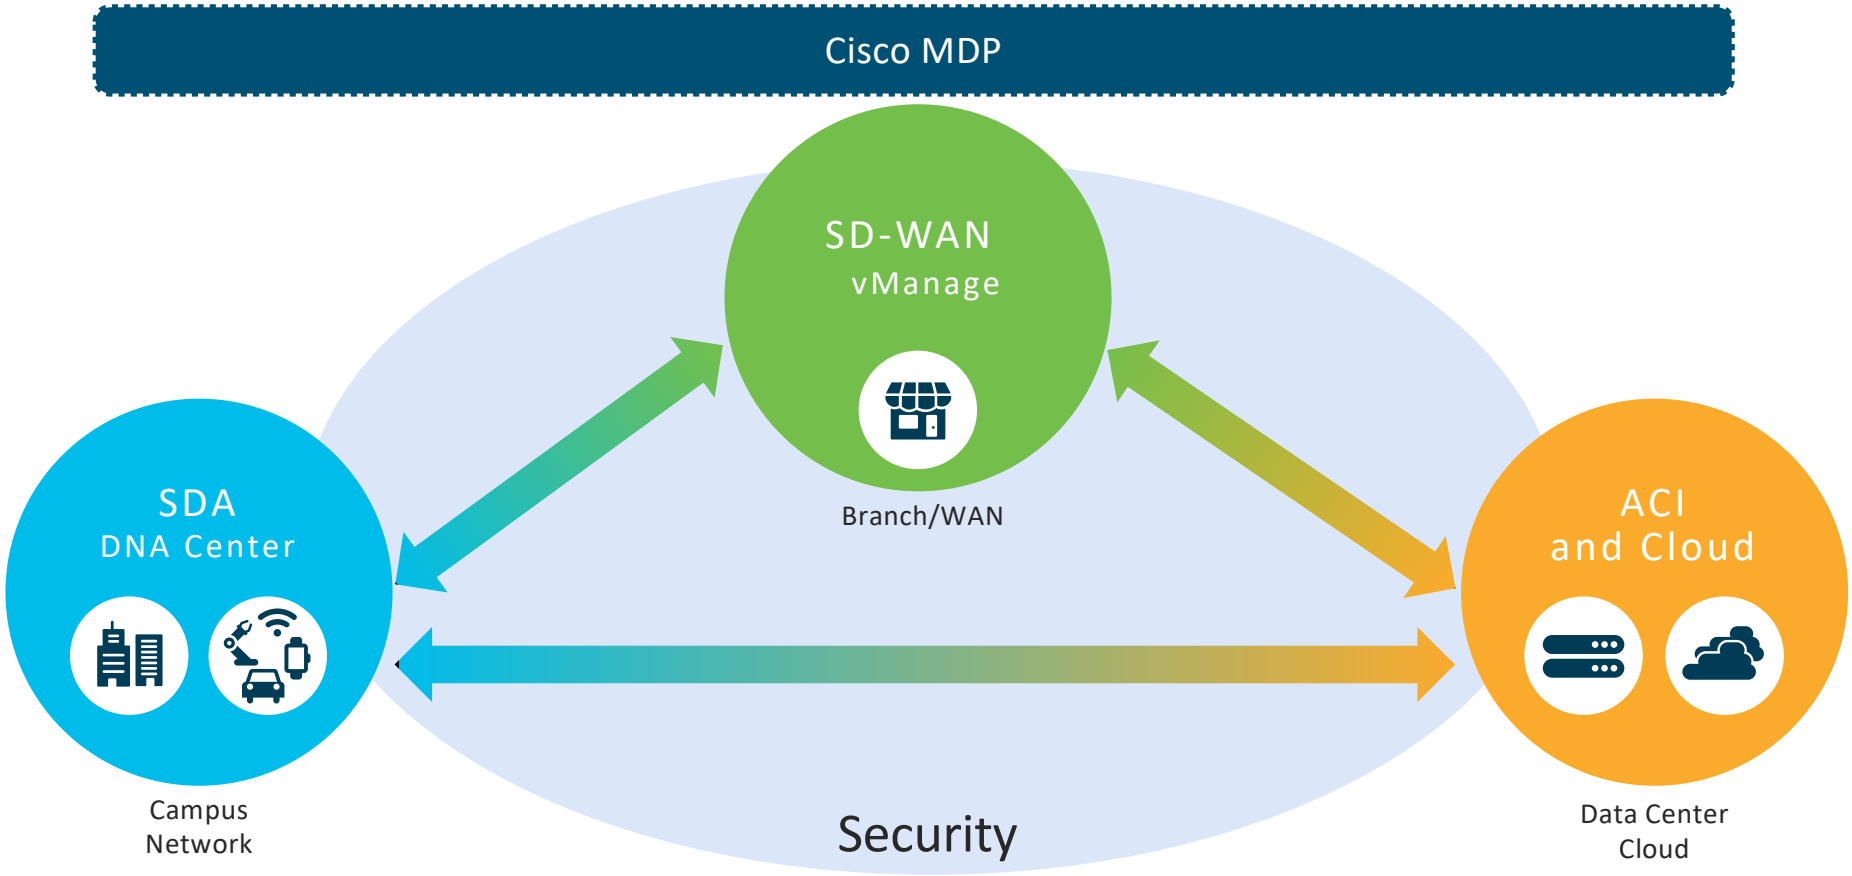

The Network Fabric

- The network based on the Fabric concept is made by two layers:
  - The Underlay Network
  - One or more Overlay Networks
- An Overlay Network is a logical topology used to virtually connect devices, built on top of the physical Underlay Network
- An Overlay Network often uses alternate forwarding attributes to provide additional services, not provided by the Underlay

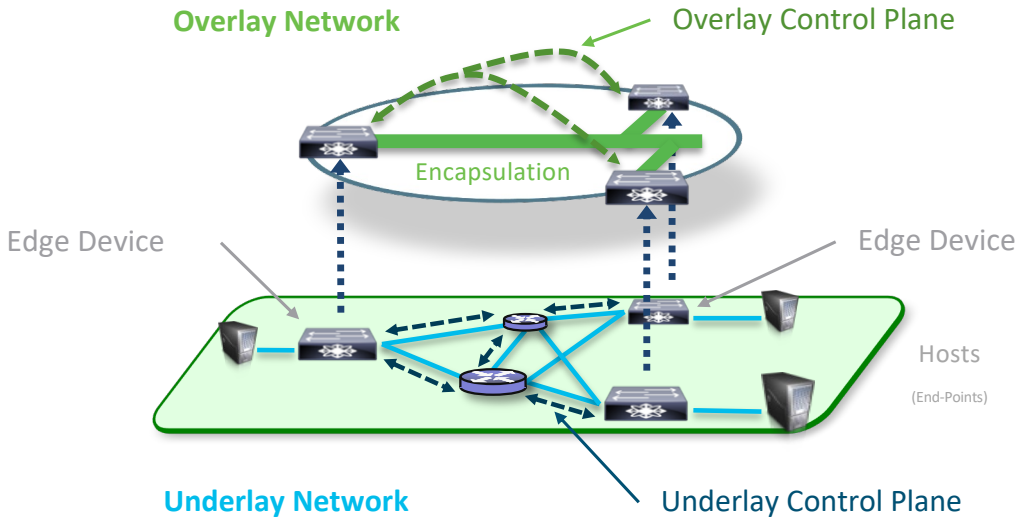

Network Controllers are Foundational to Intent-Based Networks

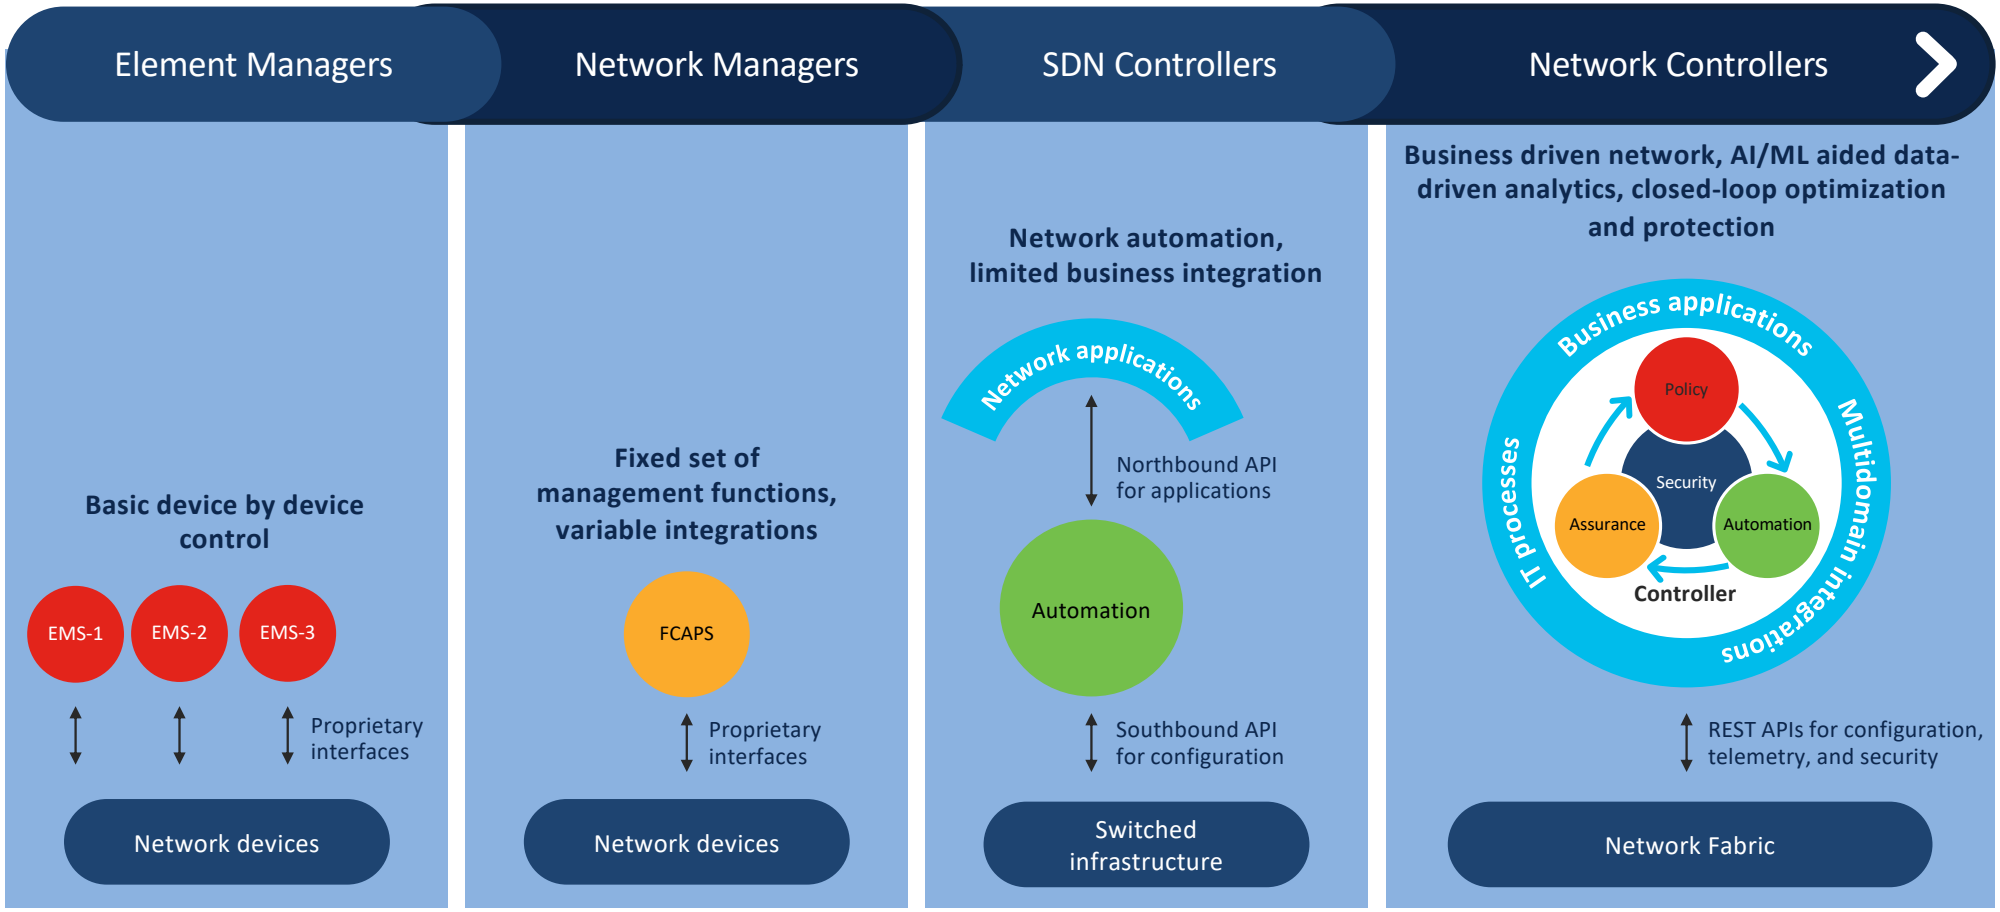

Agenda

- SDN Introduction
- Cisco Intent-Based Networking
- Cisco IBN for DC
- Cisco IBN for WAN
- Cisco IBN for Enterprise Network

Organizations are moving to Multi-Cloud and SaaS

93%

of Enterprises have embraced a Multicloud Strategy

- Flexera 2020 State of the Cloud Report

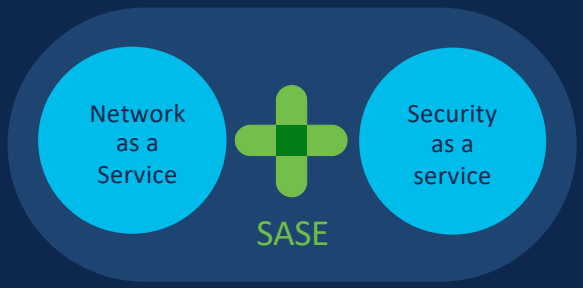

Over the next five, market for secure access service edge (SASE) will grow at a CAGR of 42%, reaching almost $11 billion by 2024

- Gartner Research

90%

of the organizations worldwide are using 1 or more SaaS Applications

- Markets & Markets Research

Historic traffic flows

Led to the age of perimeter-based security and networking

Network:
Centralized

Security:
Single, on-premise security stack

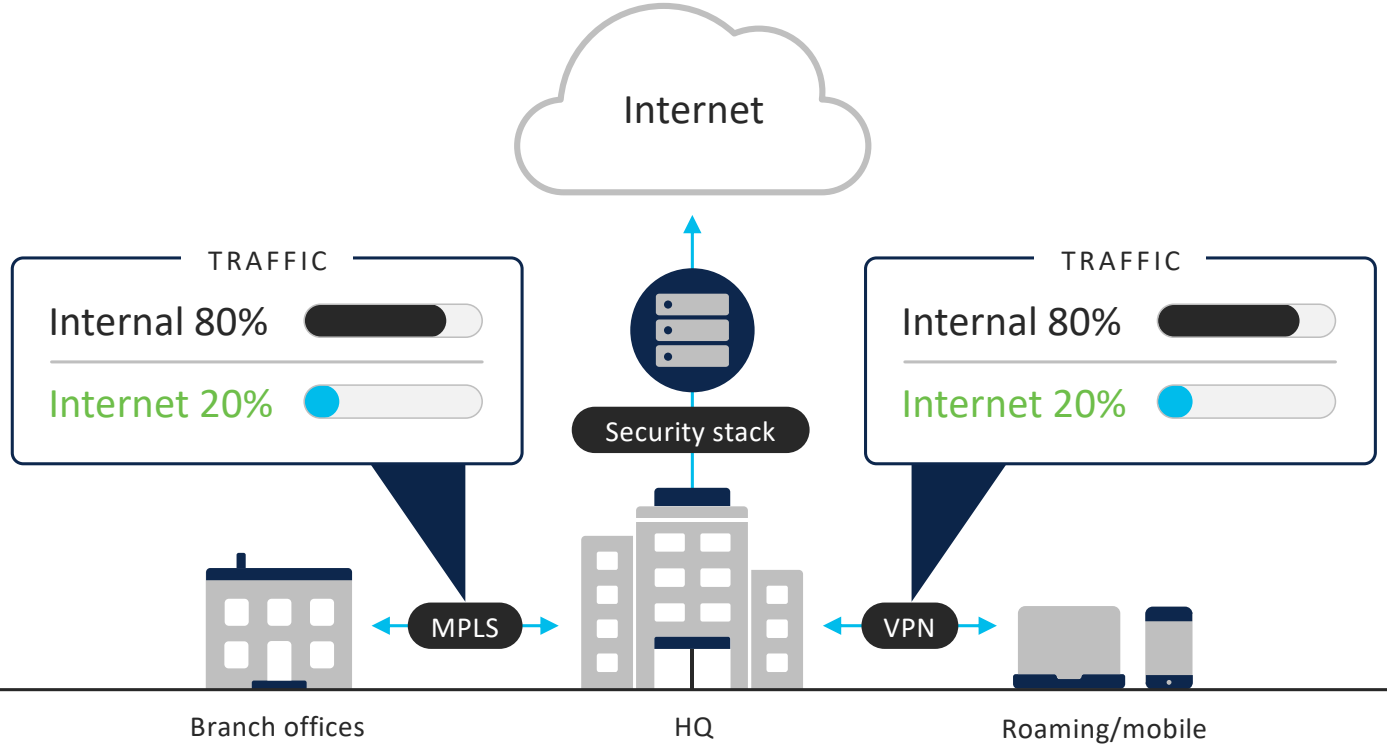

Changes in the types of traffic and destinations

Have inverted the traffic model

Problems:

- Costs
- Performance
- # Tools/vendors
- Integrations
- Maintenance

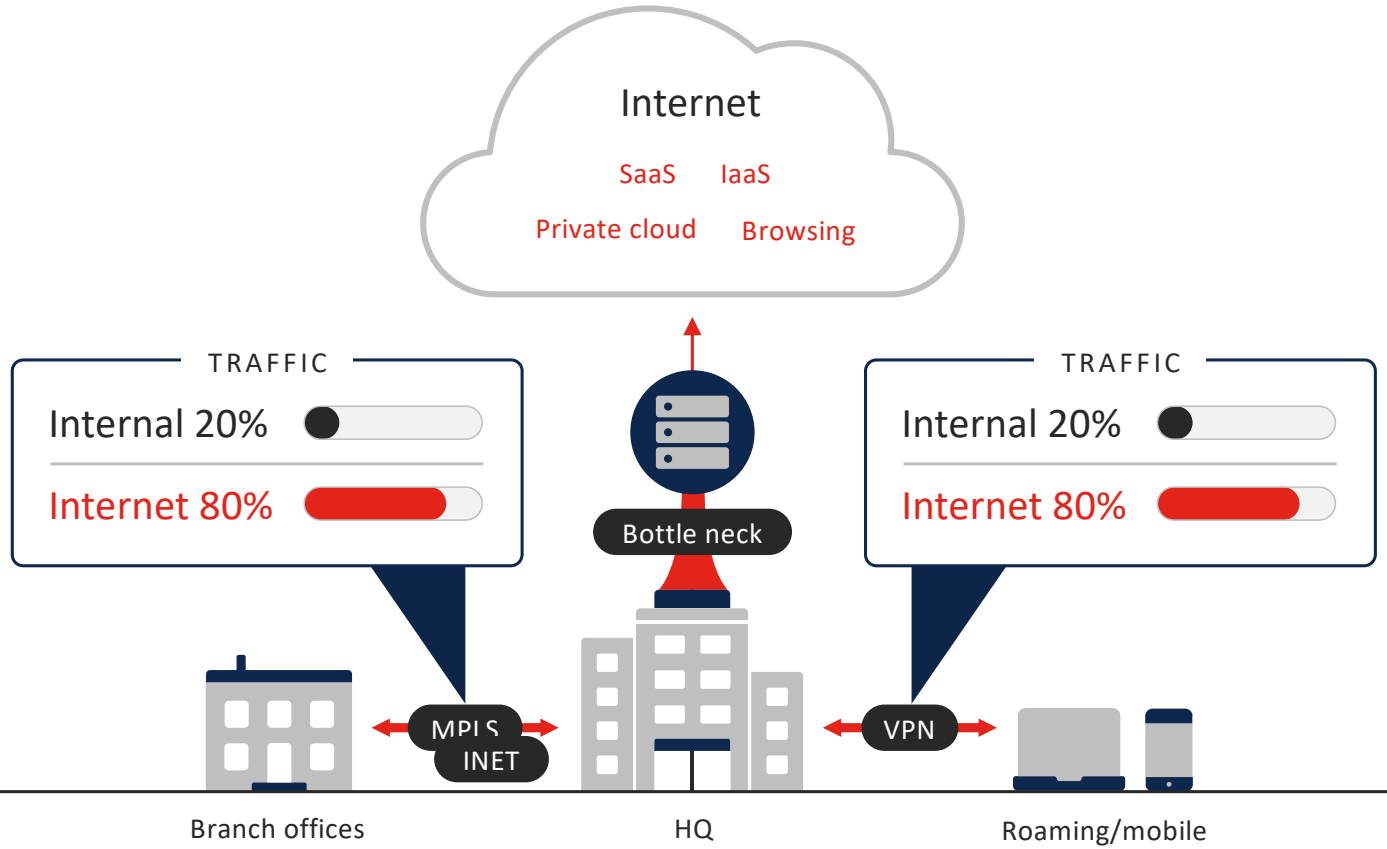

Networking and Security teams struggle to…

…connect users to applications and data

- Poor user experience when accessing cloud apps
- Complexity in connecting to multiple cloud providers
- Lack of end-to-end granular visibility of application performance

…protect against evolving threat vectors

- Gaps in security protection
- Inconsistent policies enforced across disparate locations
- Difficult to verify identity of users and devices

This requires a new approach to networking and security…

A more modern approach

Security:
Enforced at the cloud edge

Network:
Optimized routing from anywhere to the cloud

Architecture:
Shifting from DC-Centric to Internet/Cloud Centric

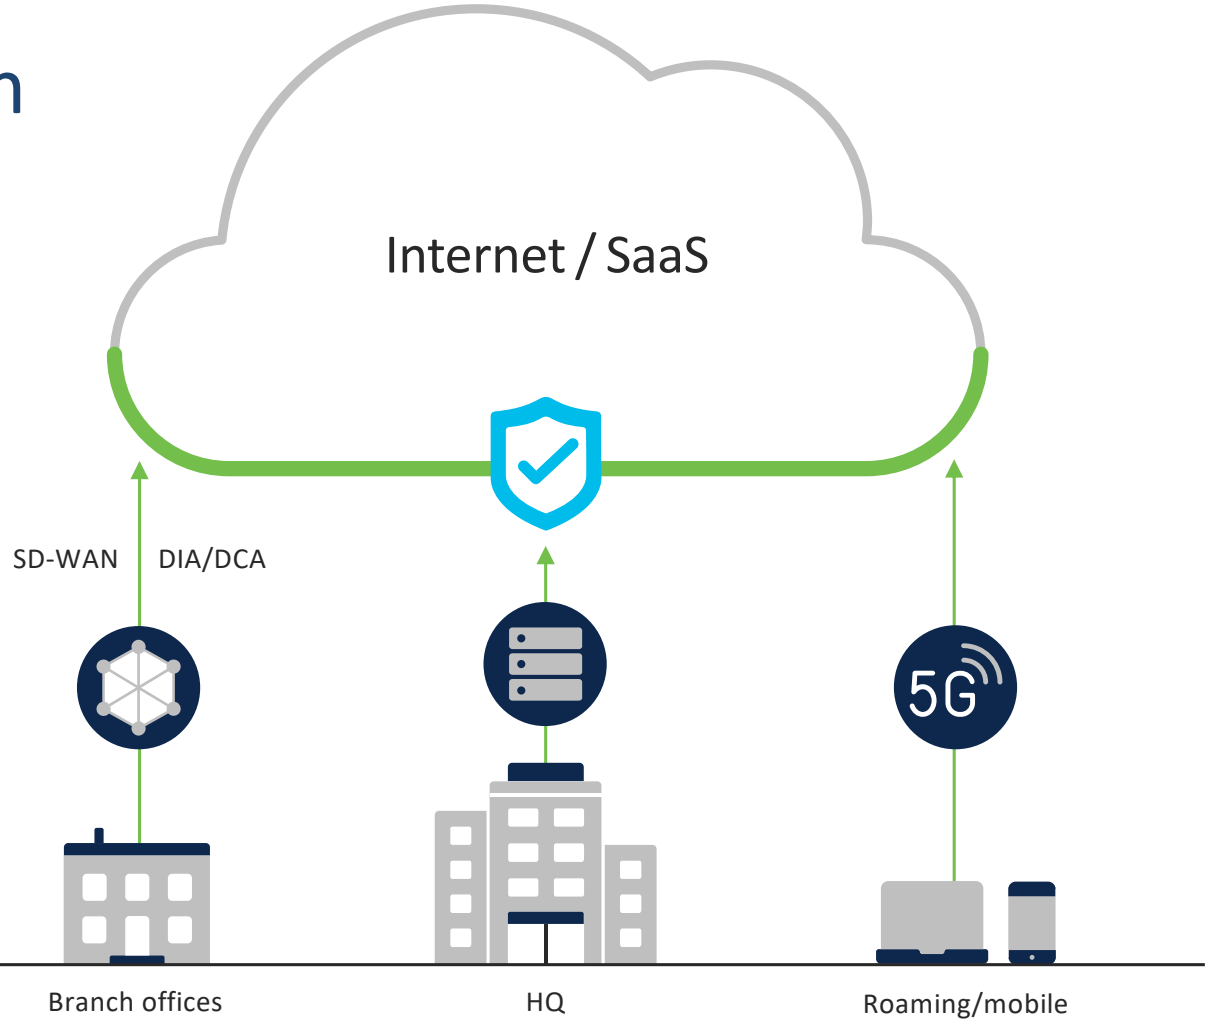

What is SD-WAN

- The Software-Defined Wide Area Network (SD-WAN) is a technology in which we can implement an Enterprise WAN based on Software-Defined Networking (SDN)
- SDWAN represents an evolution of networking from an older, hardware-based model to a secure, software-based, virtual IP fabric.
- It is called an overlay Network because forms a software overlay that runs over standard network transport services, including the public-internet, MPLS and broadband

Cisco SD-WAN Architecture

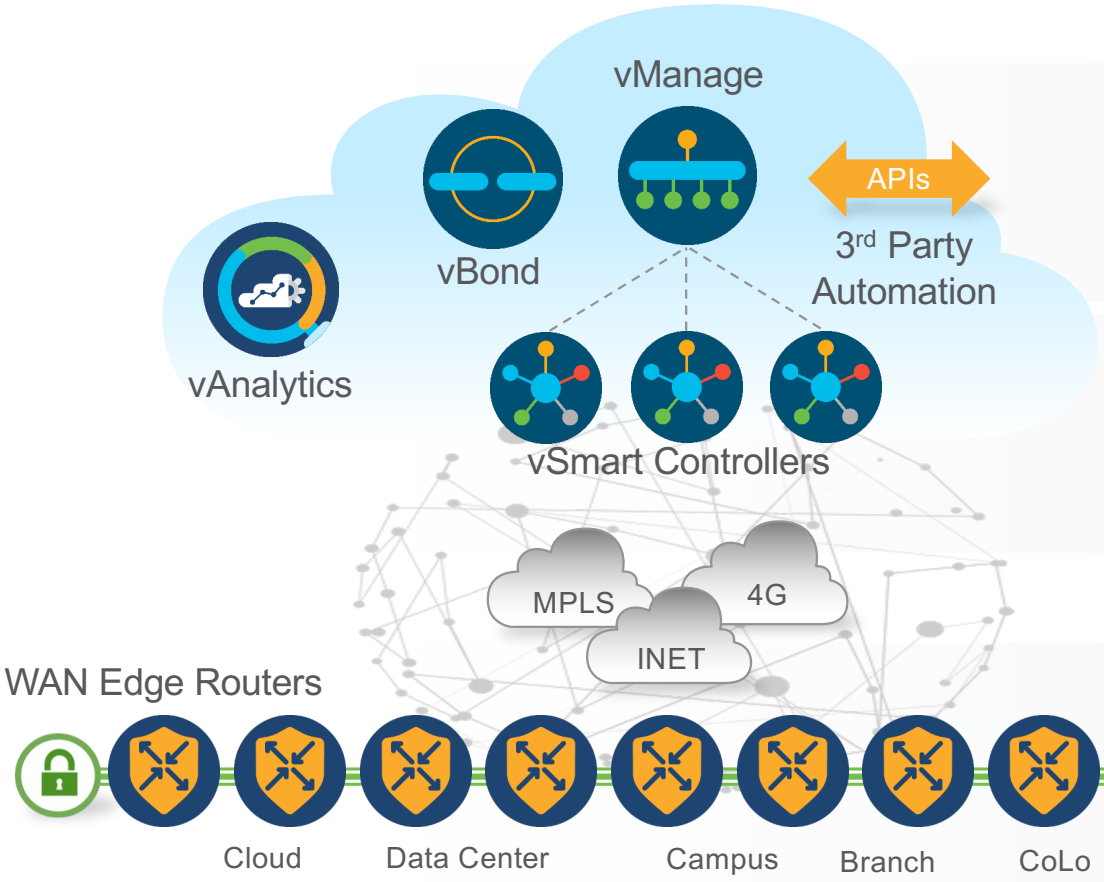

Management and Orchestration Plane

Single pane of glass for provisioning, troubleshooting and monitoring
First point of authentication
Distributes list of vSmarts/vManage to all WAN Edge routers
Facilitates NAT traversal

Control Plane

Dissimilates control plane information between WAN Edge Routers
Distributes data plane and app-aware routing policies to the WAN Edge routers
Implements control plane policies, such as service chaining, multi-topology and multi-hop
Dramatically reduces control plane complexity

Data Plane

Physical or virtual form factor
Zero Touch Provisioning
Establishes secure SD-WAN fabric
Leverages traditional routing protocols like OSPF, BGP, and EIGRP

Cisco SD-WAN Solution Elements
Orchestration Plane

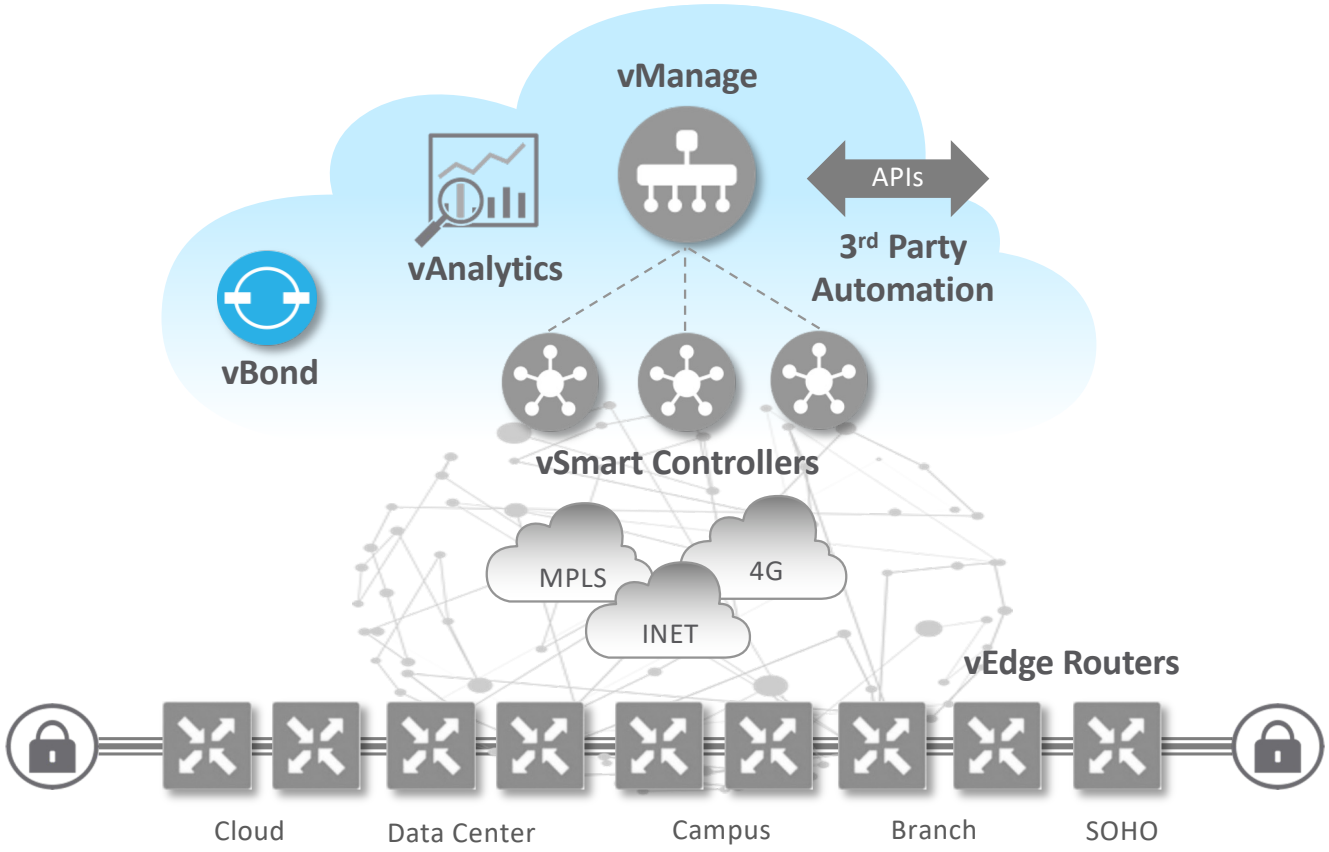

Orchestration Plane

Cisco vBond

- Orchestrates control and management plane
- First point of authentication (white-list model)
- Distributes list of vSmarts/ vManage to all vEdge routers
- Facilitates NAT traversal
- Requires public IP Address [could sit behind 1:1 NAT]
- Highly resilient

Cisco SD-WAN Solution Elements
Management Plane

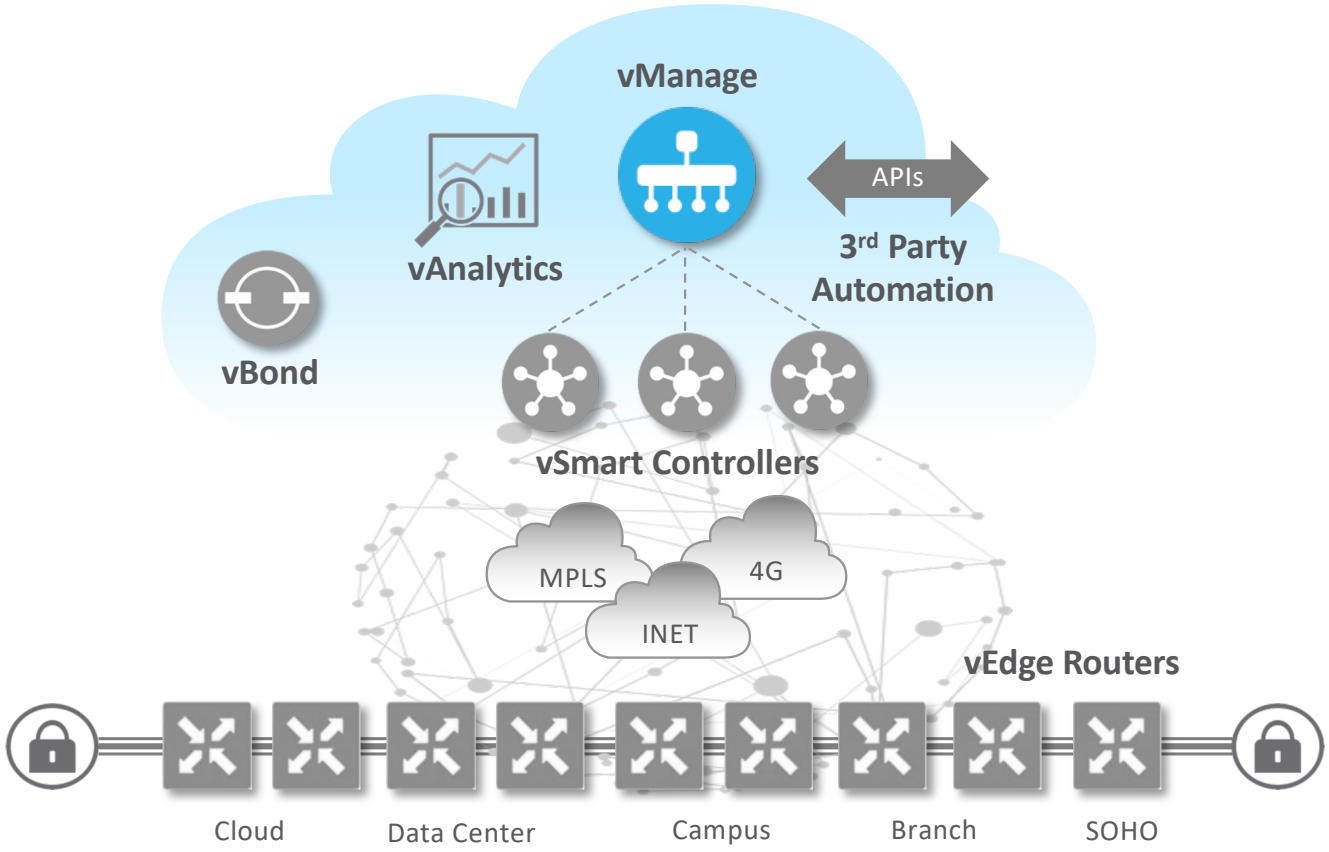

Management Plane

Cisco vManage

- Single pane of glass for Day0, Day1 and Day2 operations
- Multitenant with web scale
- Centralized provisioning
- Policies and Templates
- Troubleshooting and Monitoring
- Software upgrades
- GUI with RBAC
- Programmatic interfaces (REST, NETCONF)
- Highly resilient

Cisco SD-WAN Solution Elements
Control Plane

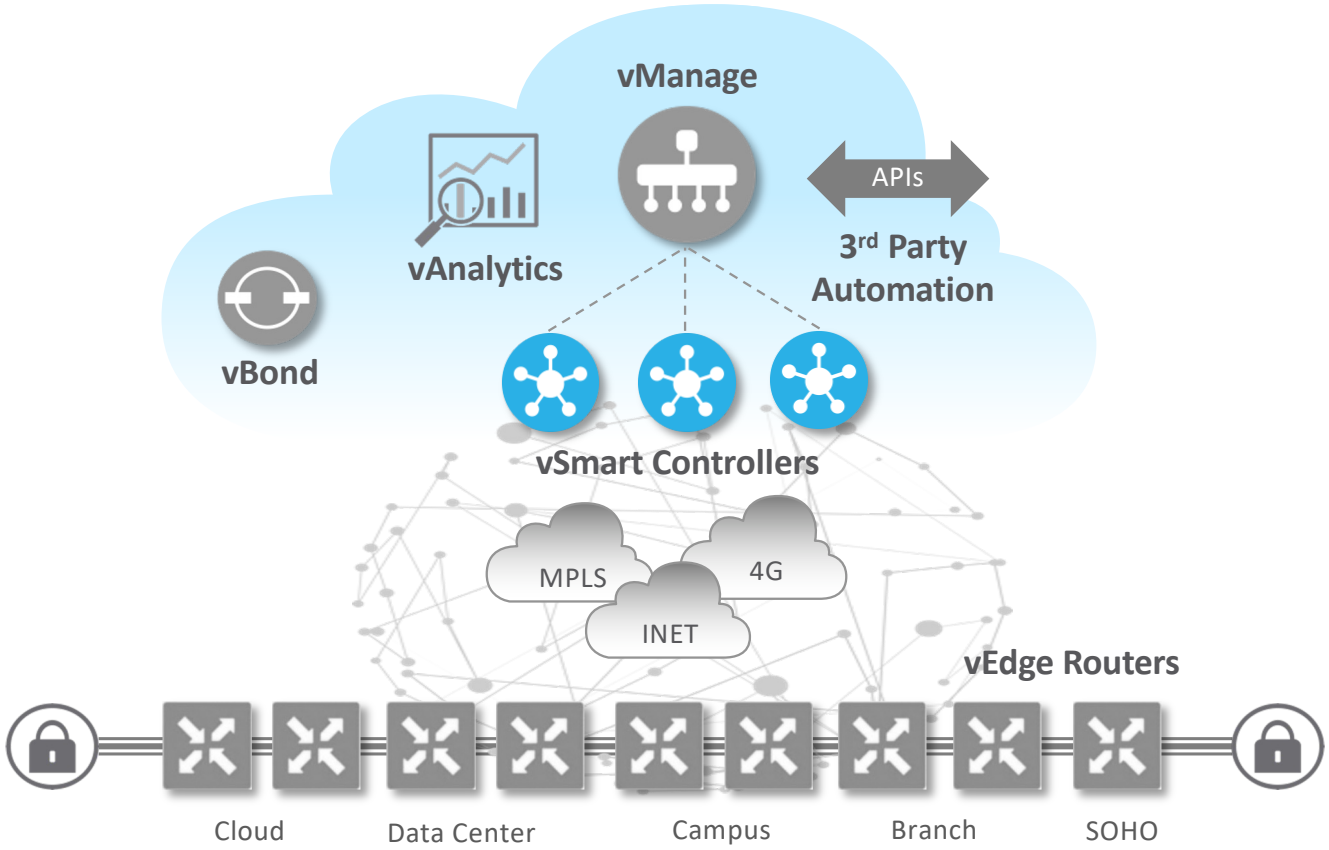

Control Plane

Cisco vSmart

- Facilitates fabric discovery
- Dissimilates control plane information between vEdges
- Distributes data plane and app-aware routing policies to the vEdge routers
- Implements control plane policies, such as service chaining, multi-topology and multi-hop
- Dramatically reduces control plane complexity
- Highly resilient

Cisco SD-WAN Solution Elements
Data Plane

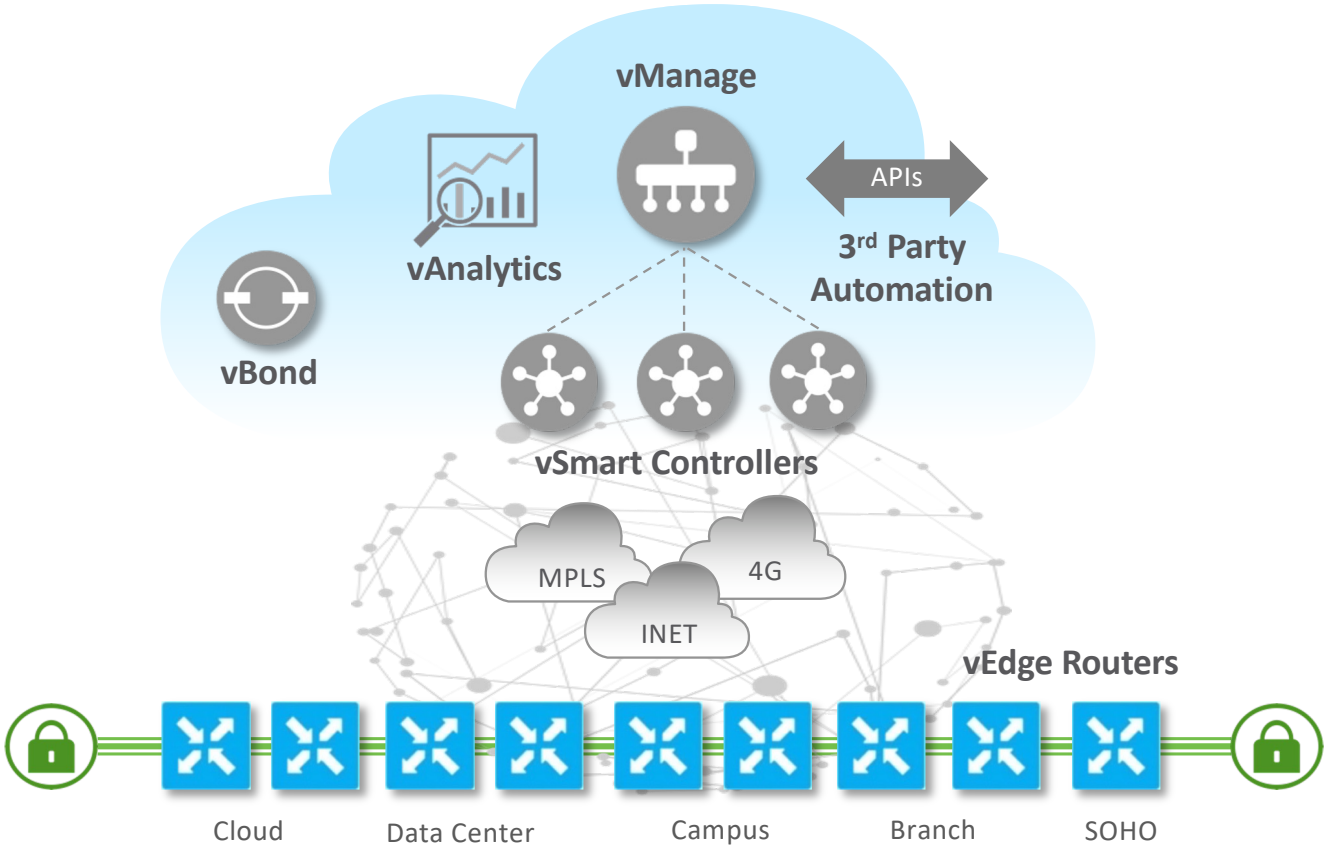

Data Plane
Physical/Virtual

Cisco vEdge

- WAN edge router
- Provides secure data plane with remote vEdge routers
- Establishes secure control plane with vSmart controllers (OMP)
- Implements data plane and application aware routing policies
- Exports performance statistics
- Leverages traditional routing protocols like OSPF, BGP and VRRP
- Support Zero Touch Deployment
- Physical or Virtual form factor (100Mb, 1Gb, 10Gb)

Controllers
Deployment Methodology

On-Premise

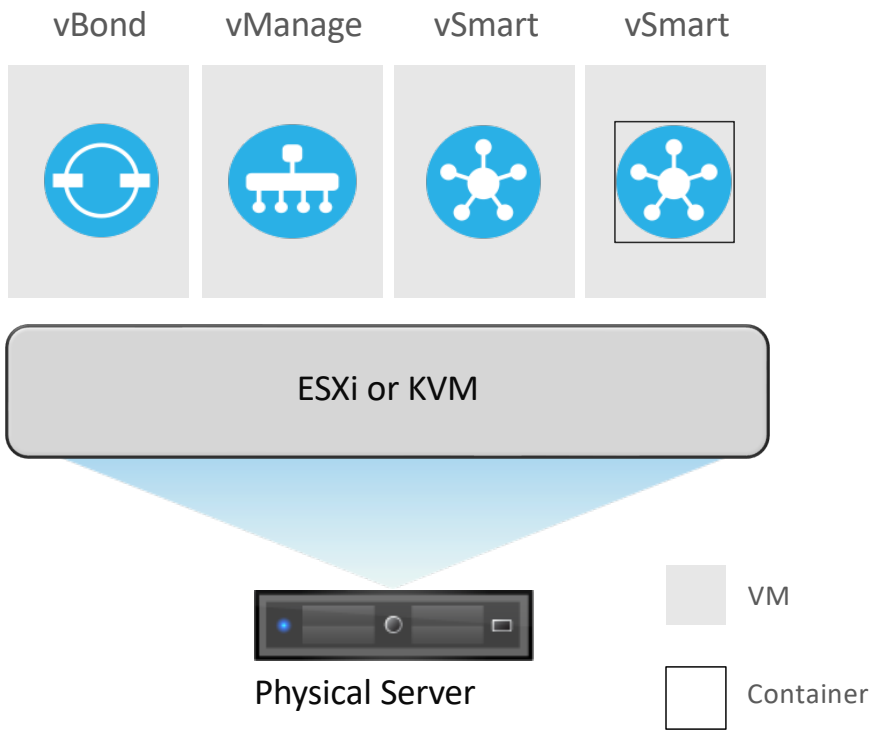

Hosted

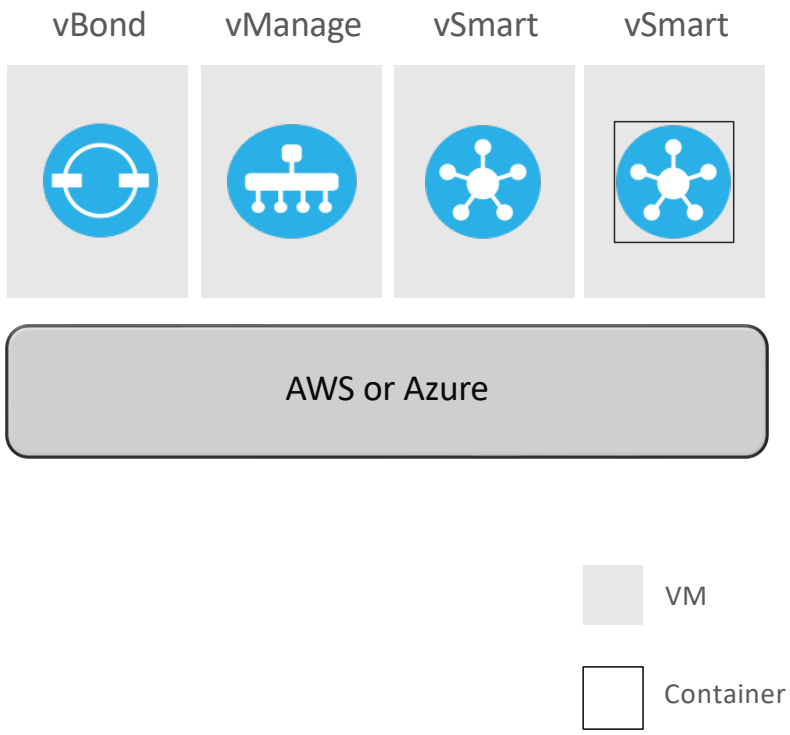

Cisco SD-WAN Fabric Terminology

- **Overlay Management Protocol** – Control plane protocol distributing reachability, security and policies throughout the fabric
- **Transport Locator (TLOC)** – Transport attachment point and next hop route attribute
- **Color** – Control plane tag used for IPSec tunnel establishment logic
- **Site ID** – Unique per-site numeric identifier used in policy application
- **System IP** – Unique per-device (Cisco WAN Edge and controllers) IPv4 notation identifier. Also used as Router ID for BGP and OSPF.
- **Organization Name** – Overlay identifier common to all elements of the fabric
- **VPN** – Also known as VRF in IOS-XE. Used for device-level and network-level segmentation.

Overlay Management Protocol (OMP)
Unified Control Plane

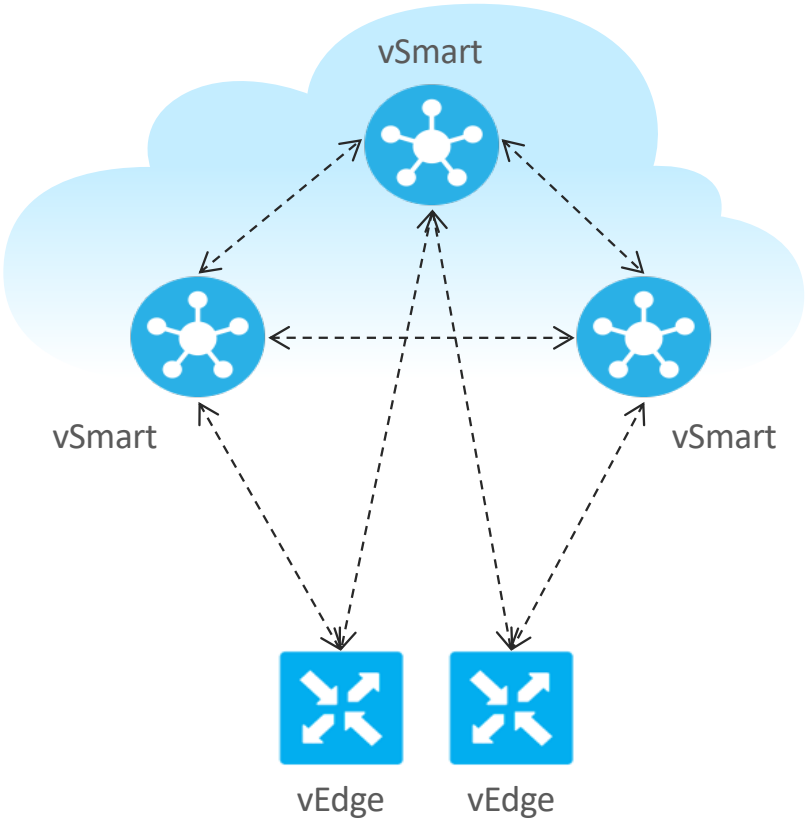

Note: vEdge routers need not connect to all vSmart Controllers

- TCP based extensible control plane protocol
- Runs between vEdge routers and vSmart controllers and between the vSmart controllers
  - Inside TLS/DTLS connections
- Advertises control plane context
- Dramatically lowers control plane complexity and raises overall solution scale

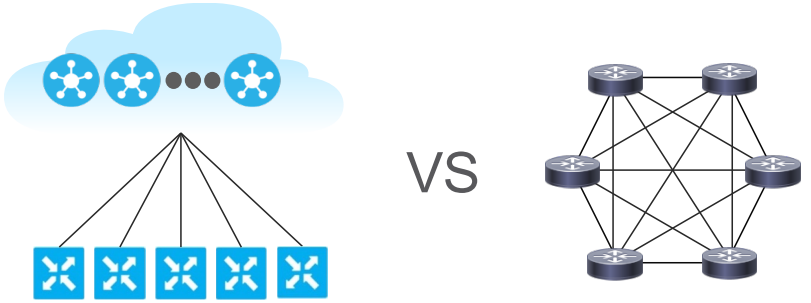

Control Plane Sessions

- Secure Channel to SD-WAN Controllers operates over DTLS/TLS authenticated and secured tunnels.
- OMP between vEdge routers and vSmart controllers and between the vSmart controllers
- NETCONF – Provisioning from vManage. Access via admin credentials over authenticated tunnel.

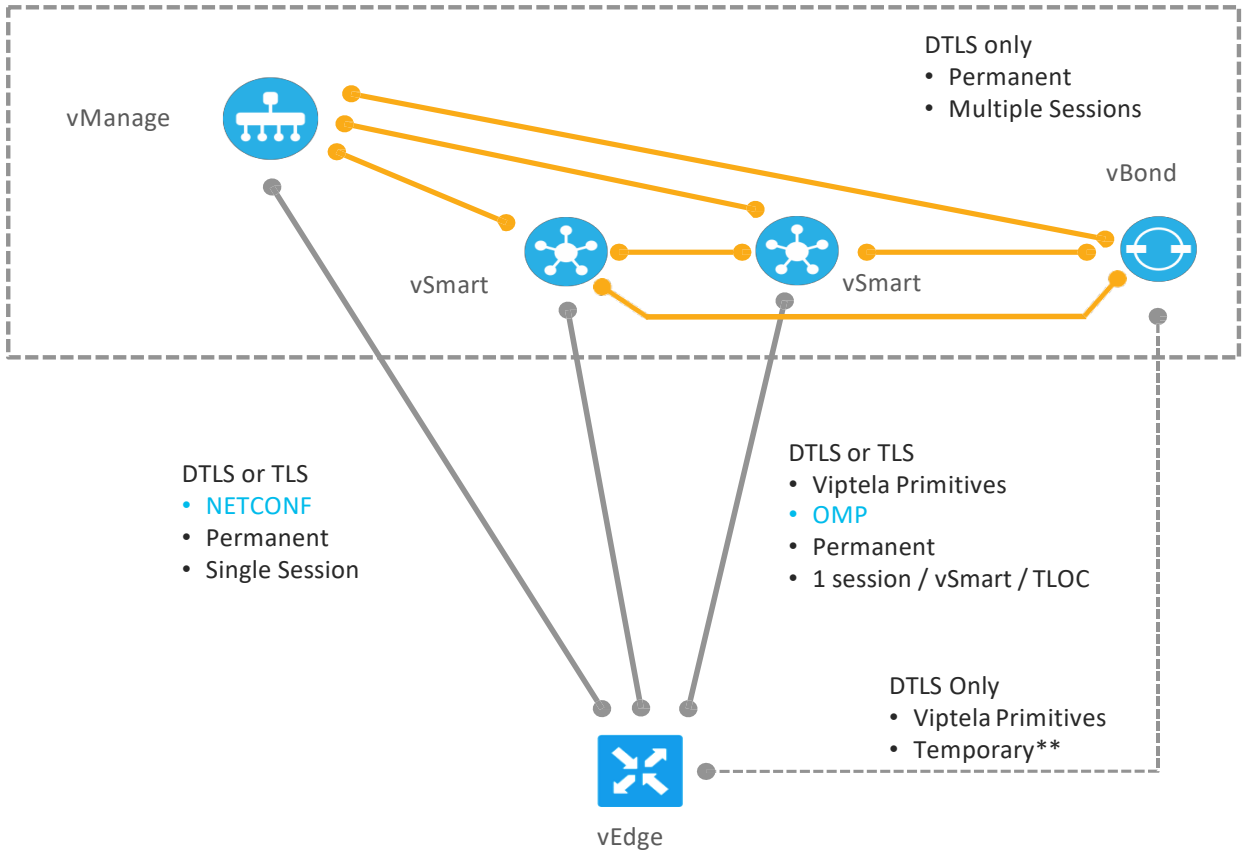

Transport Locators (TLOCs)

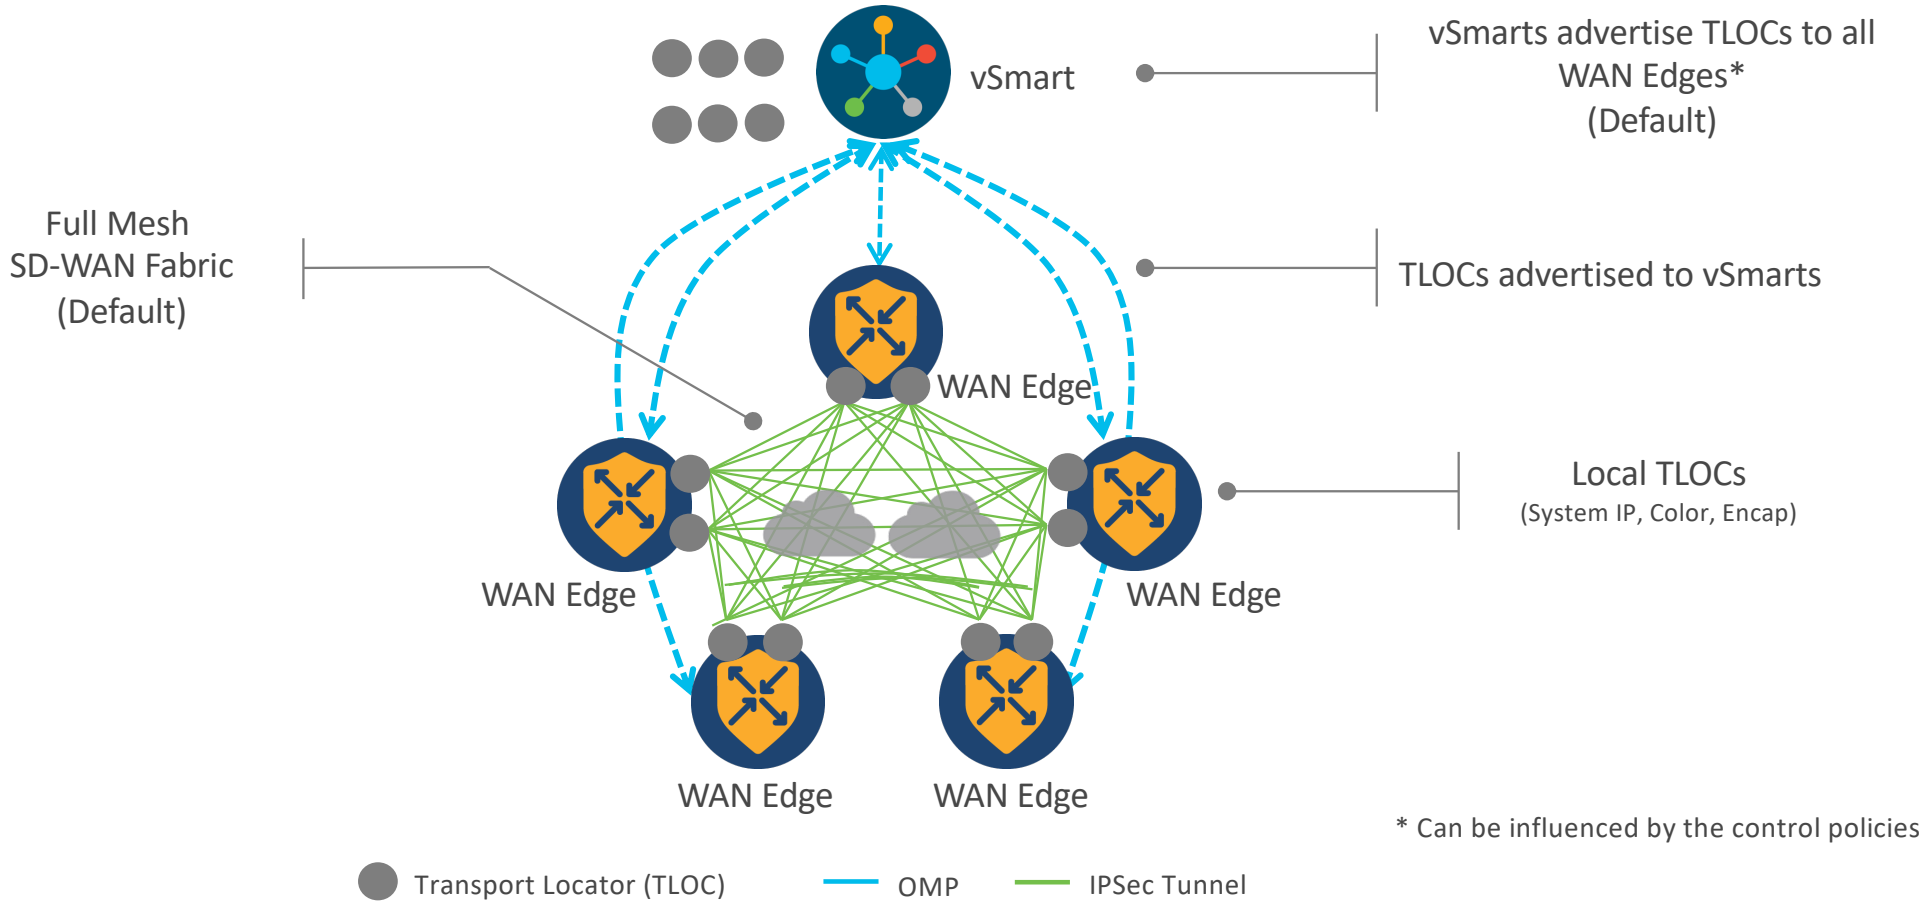

End-to-End Segmentation

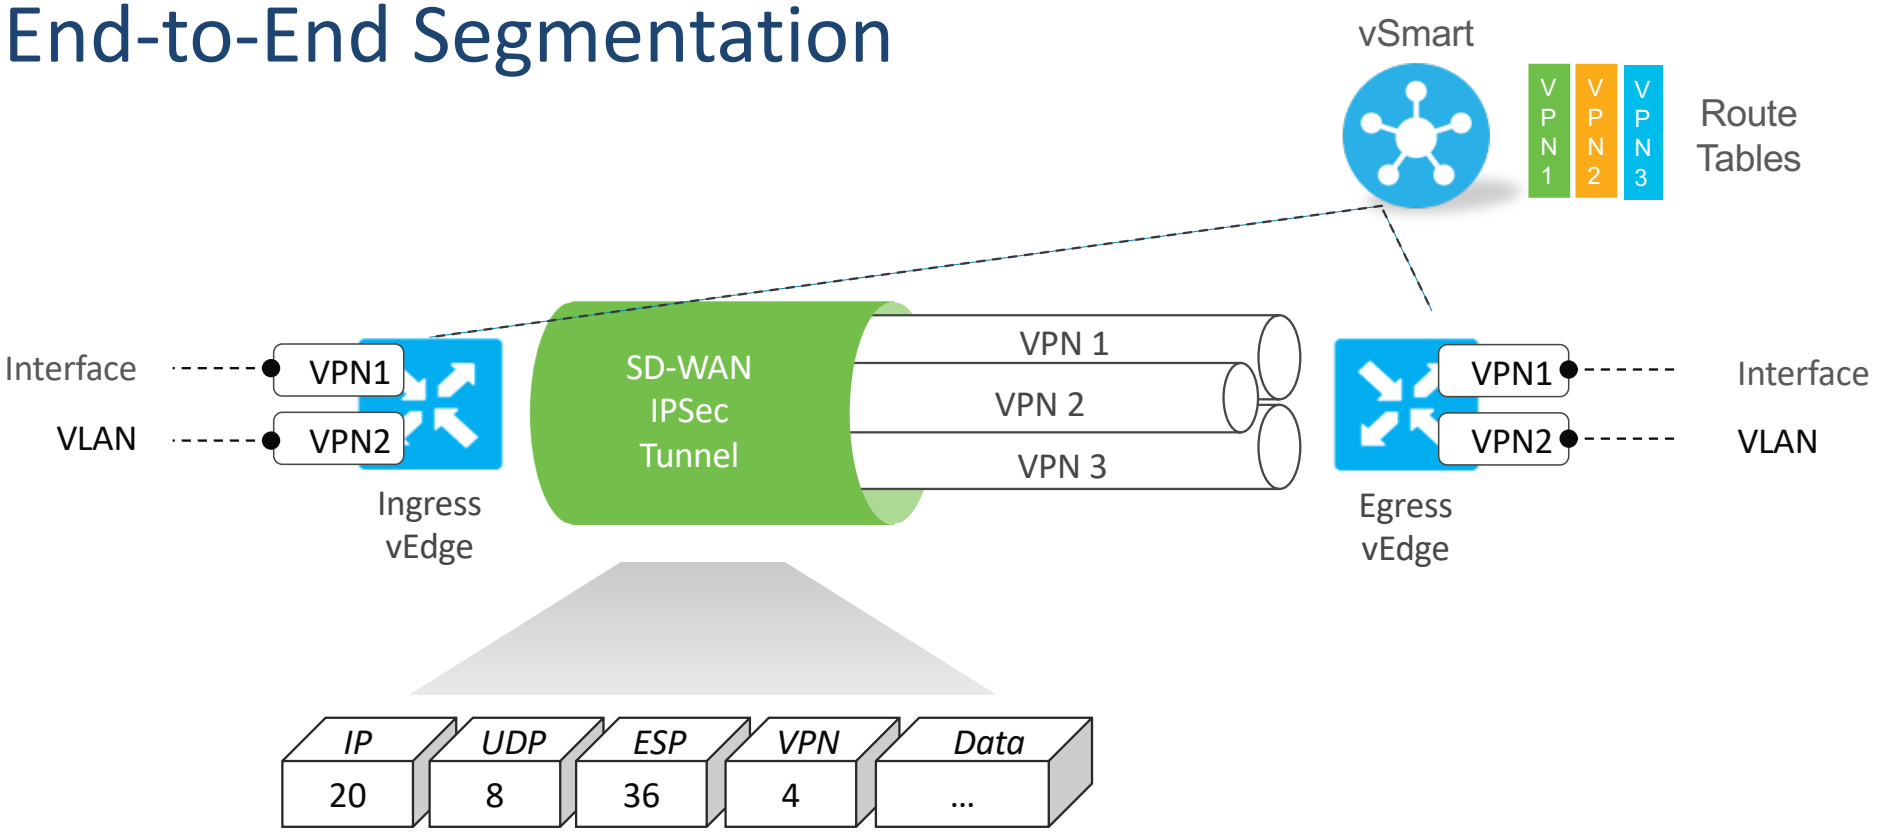

- Segment connectivity across fabric w/o reliance on underlay transport
- vEdge routers maintain per-VPN routing table
- Labels are used to identify VPN for destination route lookup
- Interfaces and sub-interfaces (802.1Q tags) are mapped into VPNs

Labels

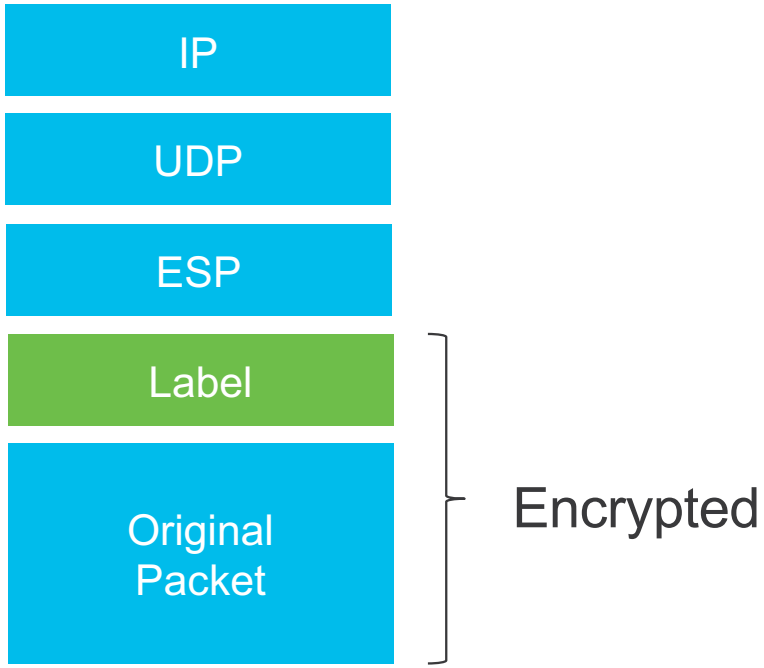

- Labels identity VPN route table on vEdge router
  - Per-VPN
  - Locally significant on each vEdge
- Pushed on the ingress vEdge, popped on the egress vEdge
- Appear in encrypted part of the IPSec packet
- Exchanged through the OMP routes
- Used for segmentation

https://www.ietf.org/rfc/rfc4023.txt

Cisco SD-WAN VPNs (VRFs)

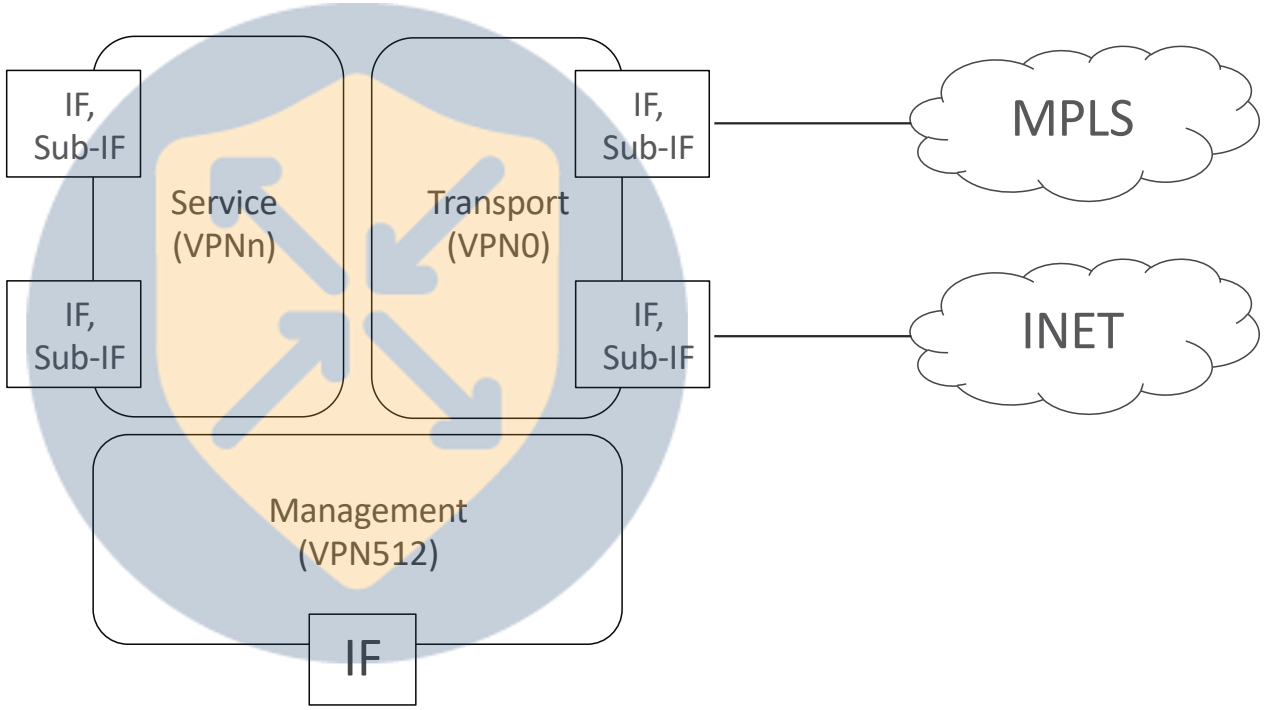

- VPNs are isolated from each other, with each VPN has its own forwarding table
- Reachability within VPN is advertised by OMP
- VPN0 is reserved for WAN uplinks (Transport)
- VPN512 is reserved for Management interfaces
- VPNn represents user-defined LAN segments (Service)

Fabric Operation Walk-Through

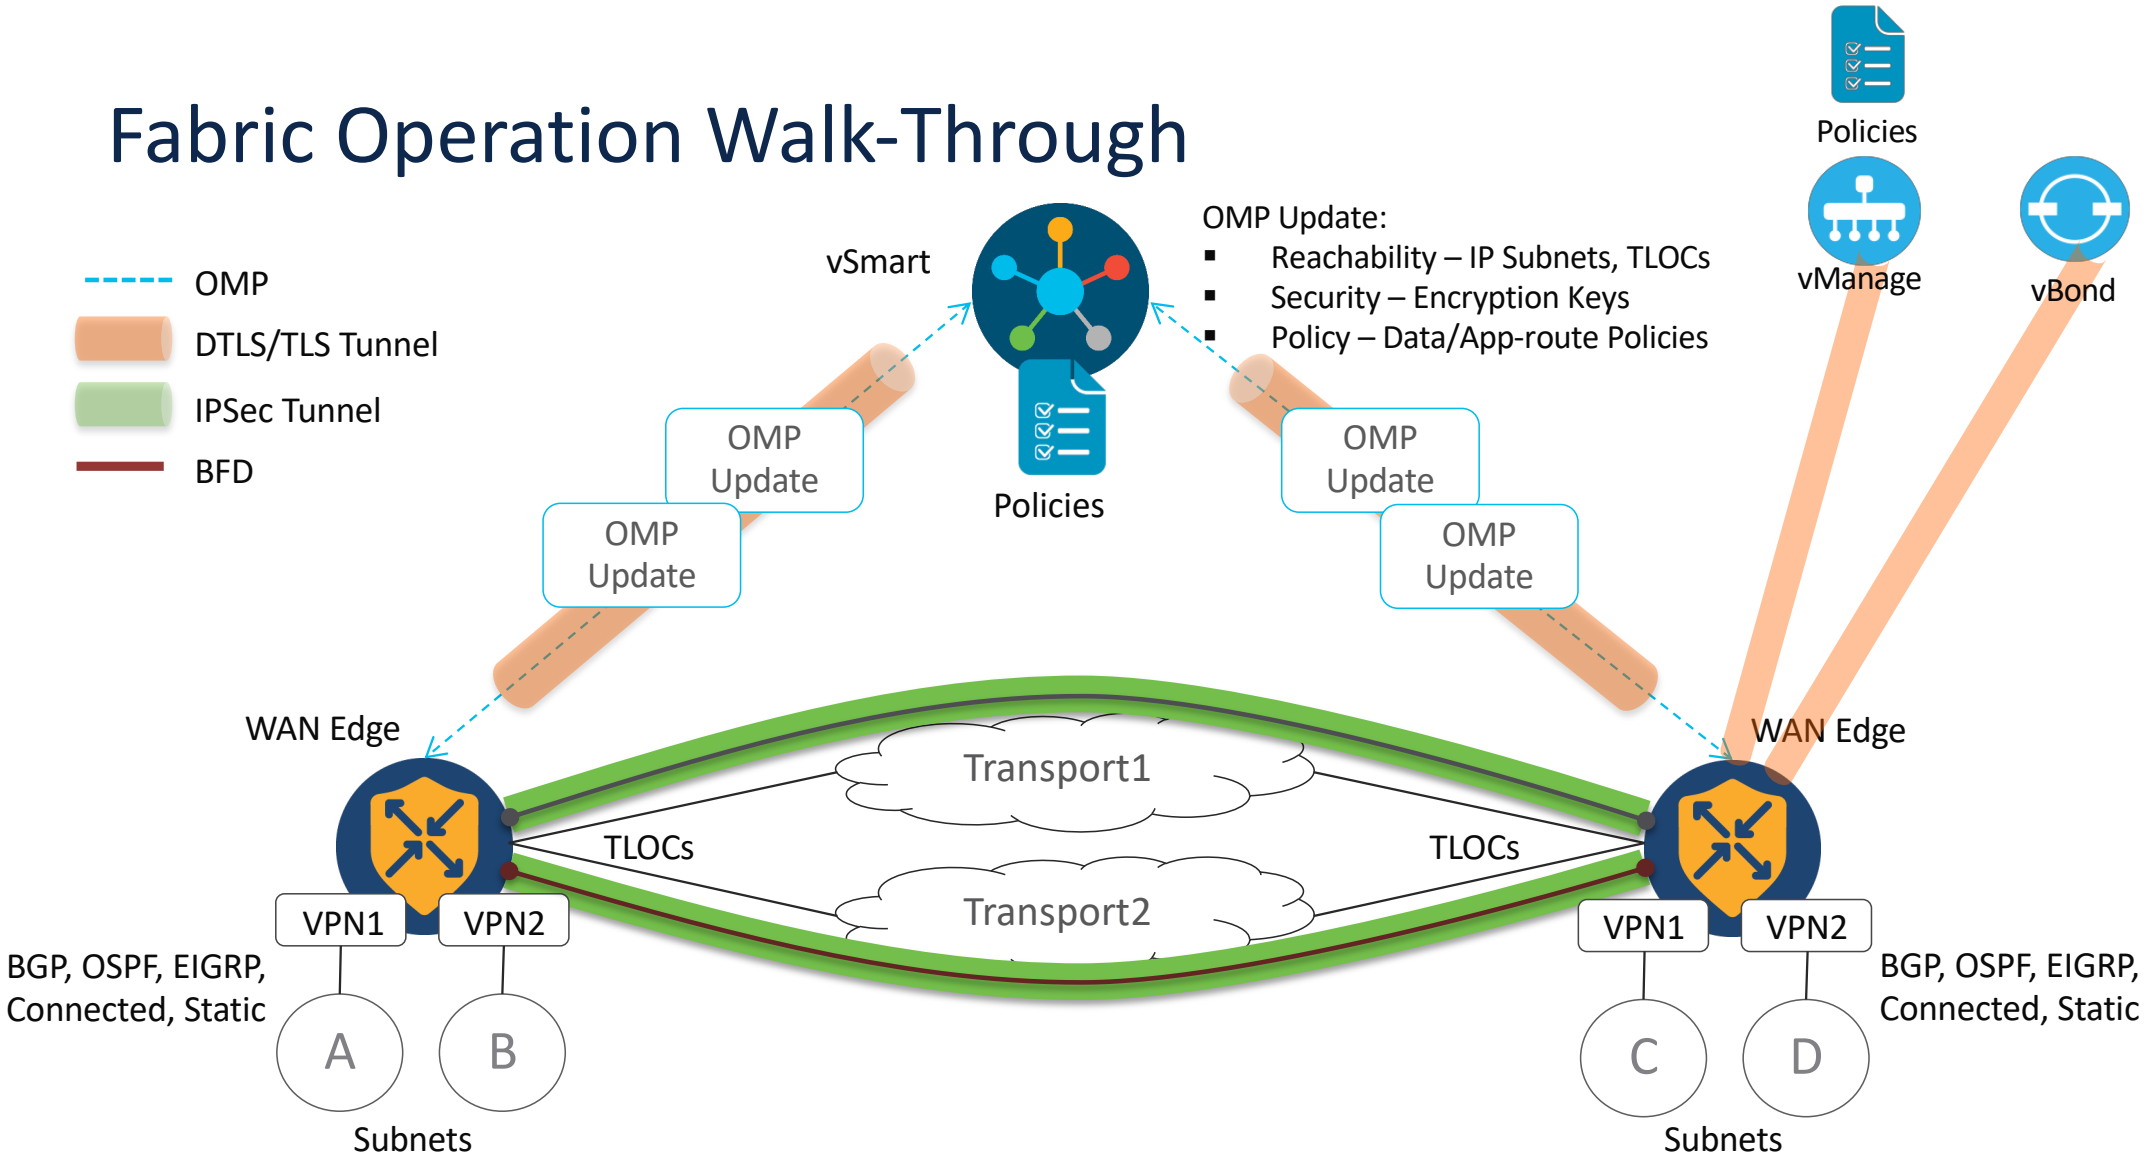

Cisco SD-WAN Policy Architecture

Policy Categories

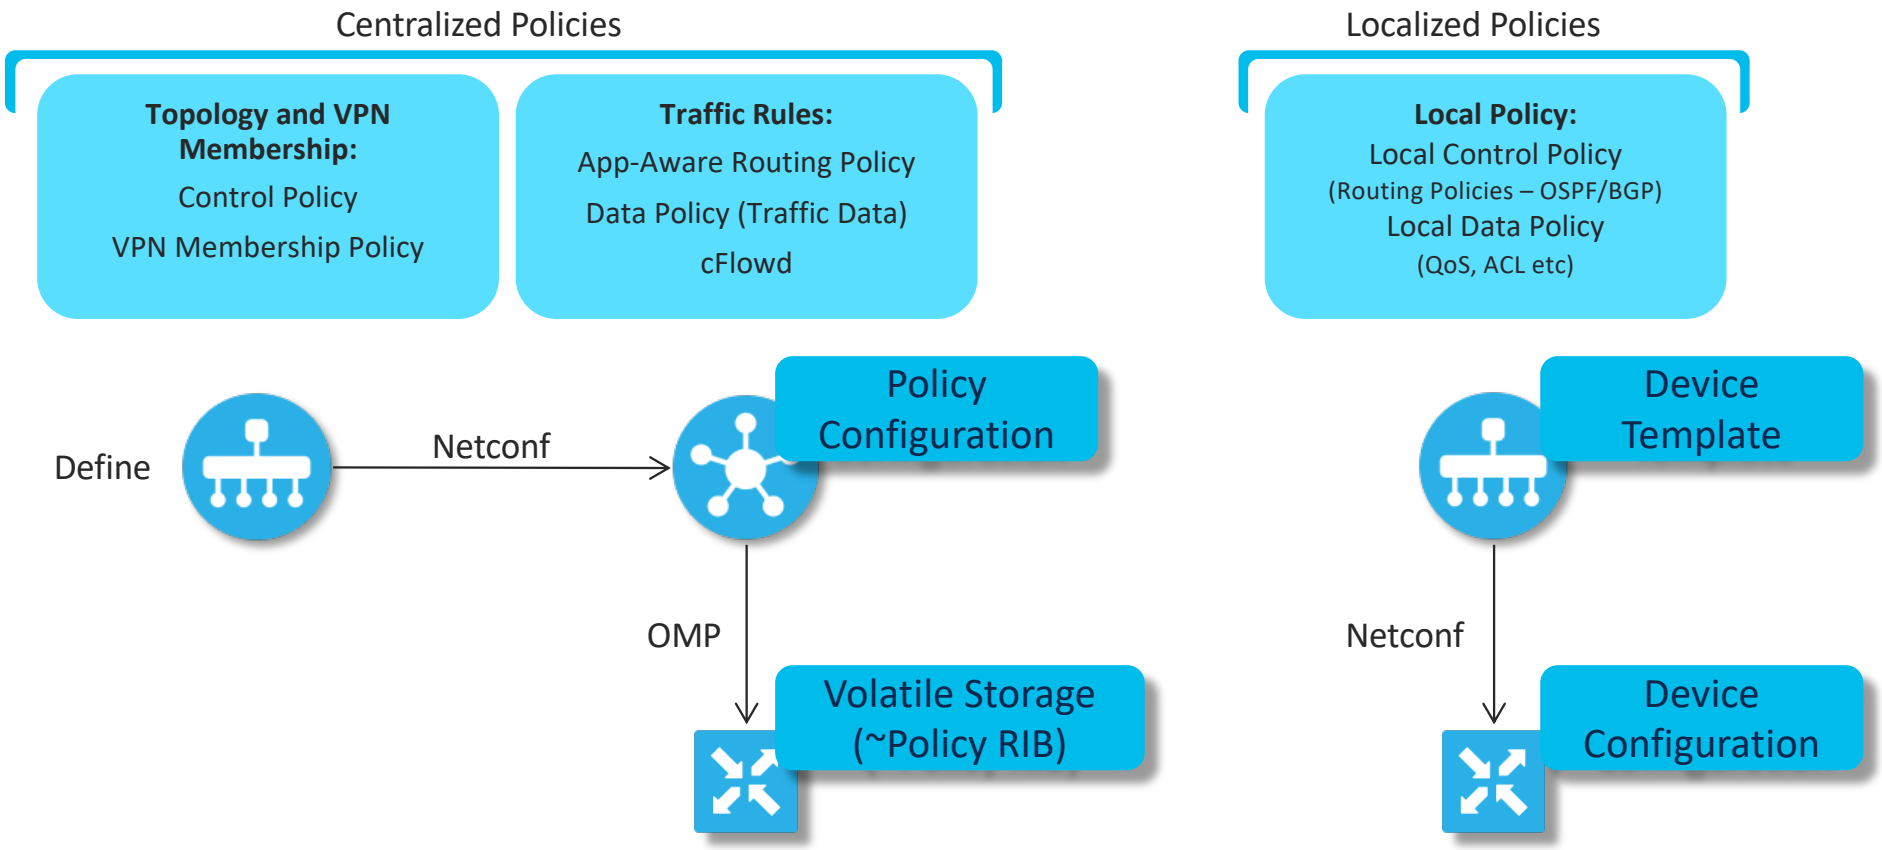

Cisco SD-WAN Policy Architecture

Suite of Policies to address different functional domains

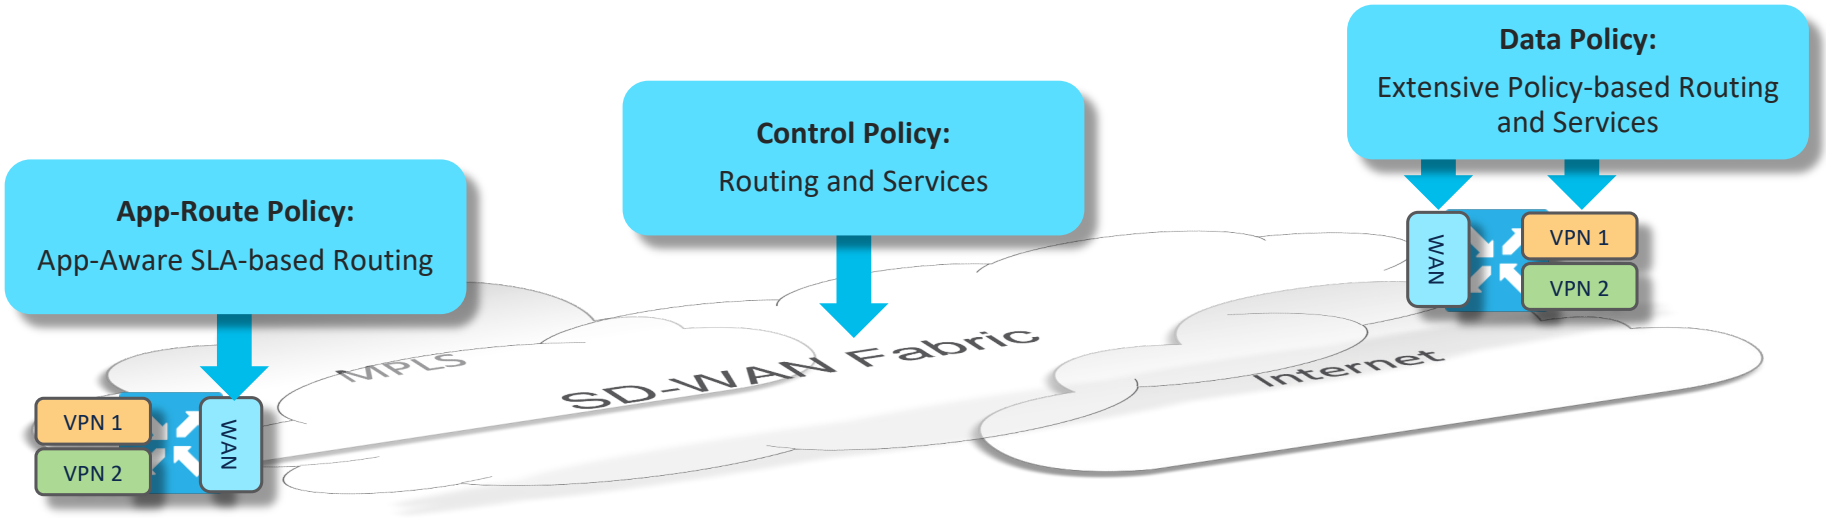

- Control Policies are applied at vSmart: Tailors routing information advertised to WAN endpoints
- App-Route Policies are applied at WAN Edge: SLA-driven path selection for applications
- Data Policies are applied at WAN Edge: Extensive Policy driven routing

Control Policies

Overlay Management Protocol Routing Policies

- Control policies are applied and executed on vSmart to influence routing in the Overlay domain
- Control policies filter or manipulate OMP Routing information to:
  - Enable services
  - Influence path selection
- Control Policies controls the following services:
  - Service Chaining
  - Traffic Engineering
  - Extranet VPNs
  - Service and Path affinity
  - Arbitrary VPN Topologies
  - and more …
- The Control Policy is one of the most powerful tools in the Cisco SD-WAN toolbox

Data Policies

Policy-driven Routing and Service Enablement

- Data policies:
  - Applied on vSmart
  - Advertised to and executed on WAN Edge
- A Data policy acts on an entire VPN and is not interface-specific
- Different Data Policies can be applied to different VPNs
- Data Policies are used to enable the following functions and services:
  - Application Pinning
  - NAT/DIA
  - Classification, Policing and Marking
  - and more …
- Use a Data Policy for any type of data plane centered traffic management

App-Route Policies

Centralized Policy for enabling SLA-driven routing on WAN Edge endpoints

- App-route policies:
  - Applied on vSmart
  - Advertised to and executed on vEdge
- Monitors SLAs for active overlay paths to direct Applications along qualified paths
- Allows for the use of L3/L4 keys or DPI Signatures for application identification
- Delivers a fully distributed SLA-driven routing mechanism

Typical SDWAN Deployment Architecture

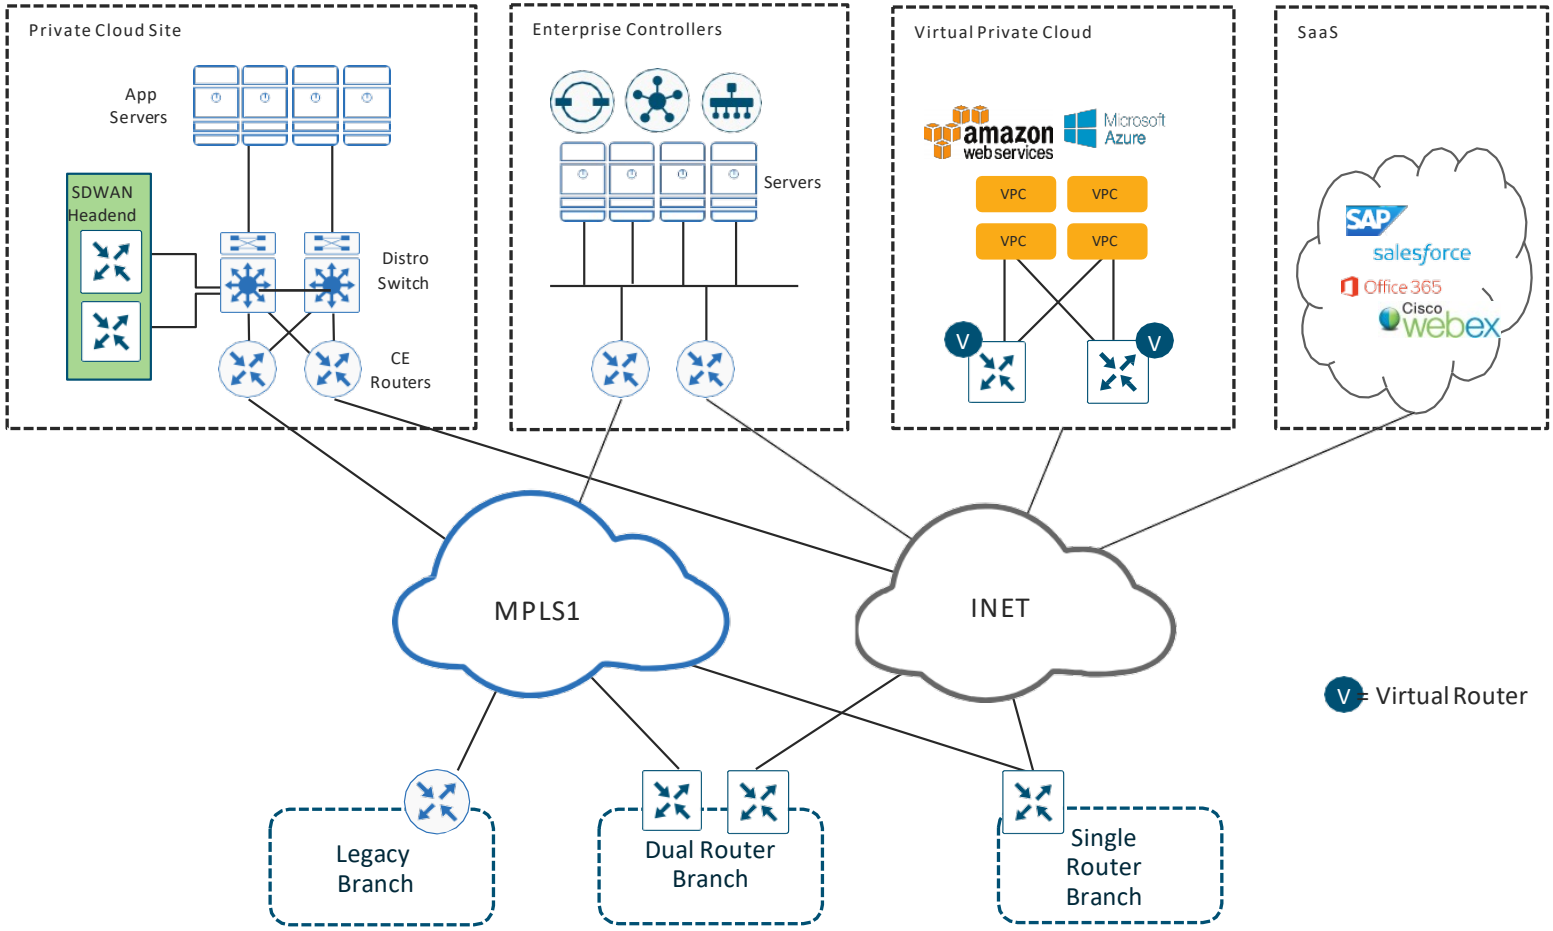

Cisco SD-WAN use cases

Aggregating features and capabilities to deliver business needs

Secure Automated WAN	Secure connectivity between remote offices, data centers, and public/private cloud over a transport-independent network
Application Performance Optimization	Improves the application experience for users at remote offices
	Locally offloads Internet traffic at the remote office
Cloud Branch Multicloud Access	Connects Cloud (IaaS/SaaS) applications to remote offices over optimal path
Regional Hub Branch Multicloud Access	Aggregates regional remote offices that utilize cloud applications with better security control and management

SD-WAN

Secure Automated WAN

Application Performance Optimization

Secure Direct Internet Access

Branch Multicloud Access

Regional Hub Multicloud Access

Secure Automated WAN

Hybrid branch: Remote office consuming apps from private and public clouds

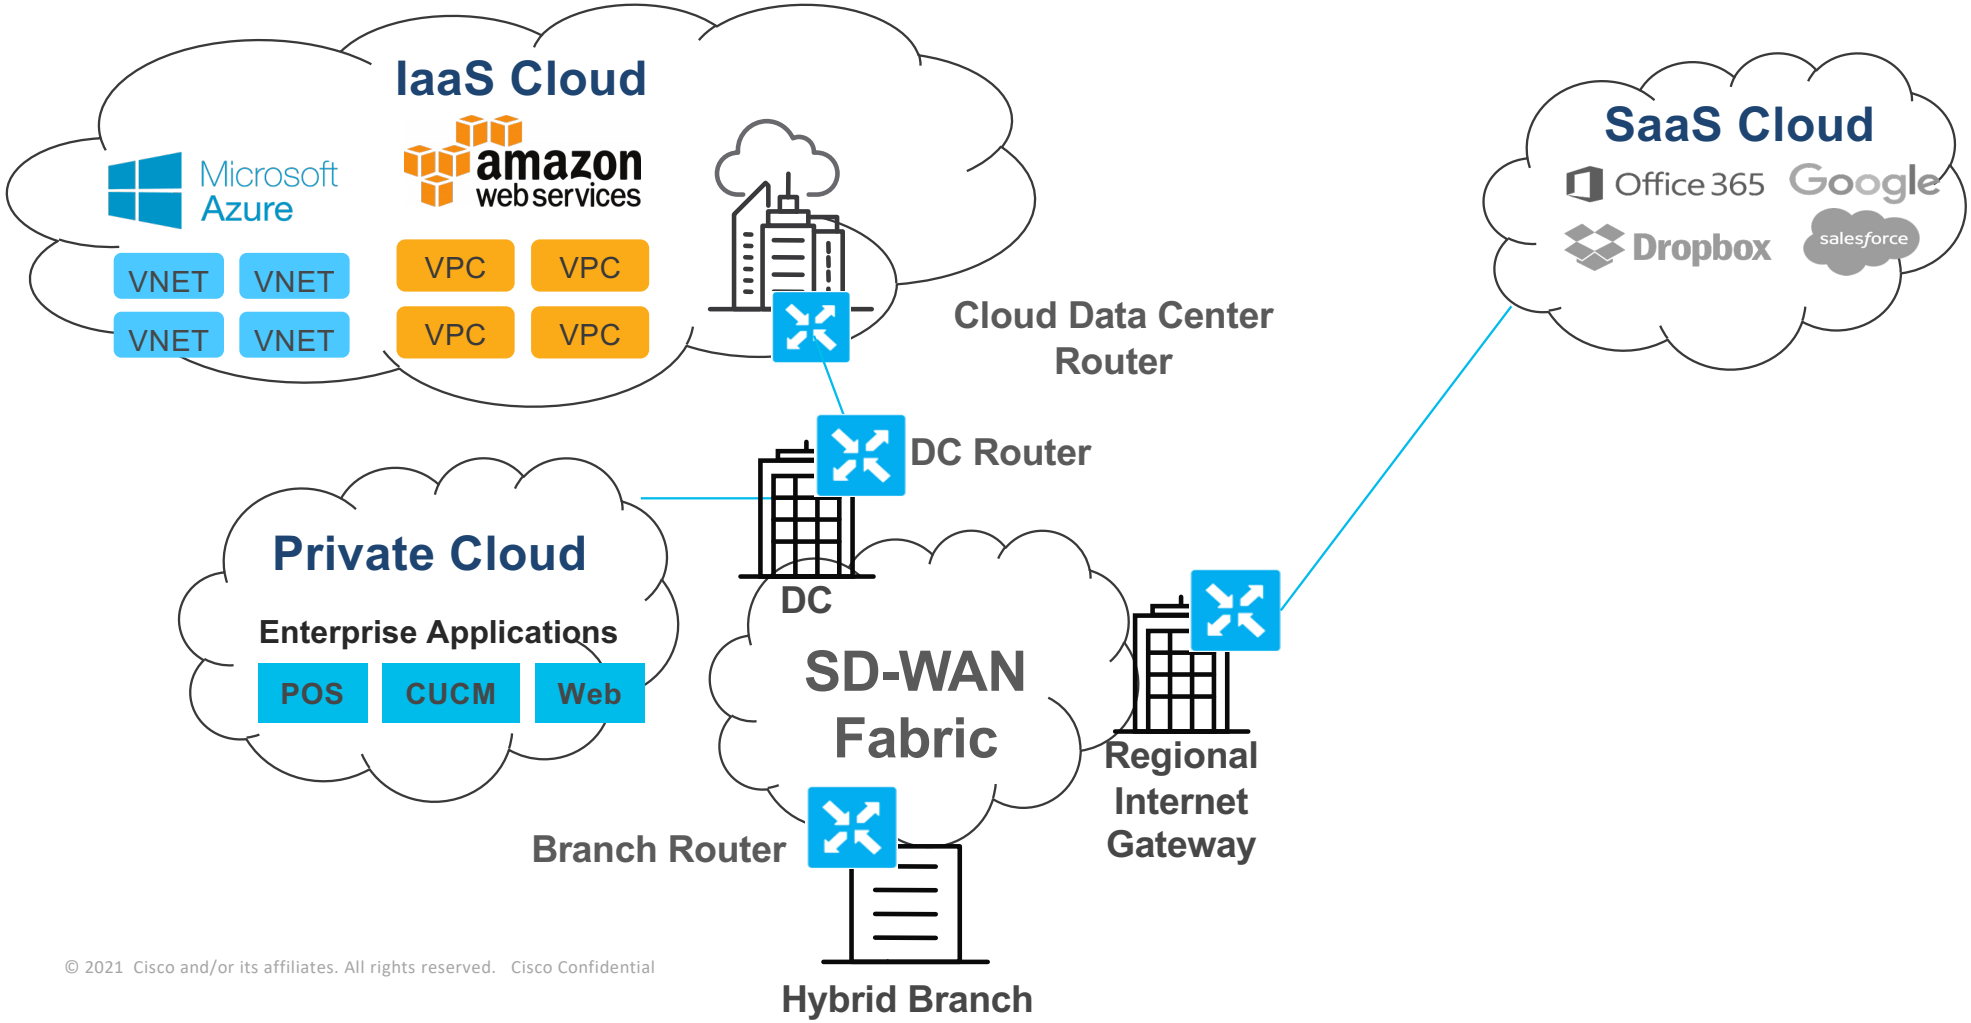

Zero Touch Provisioning

Automated configuration and fabric discovery

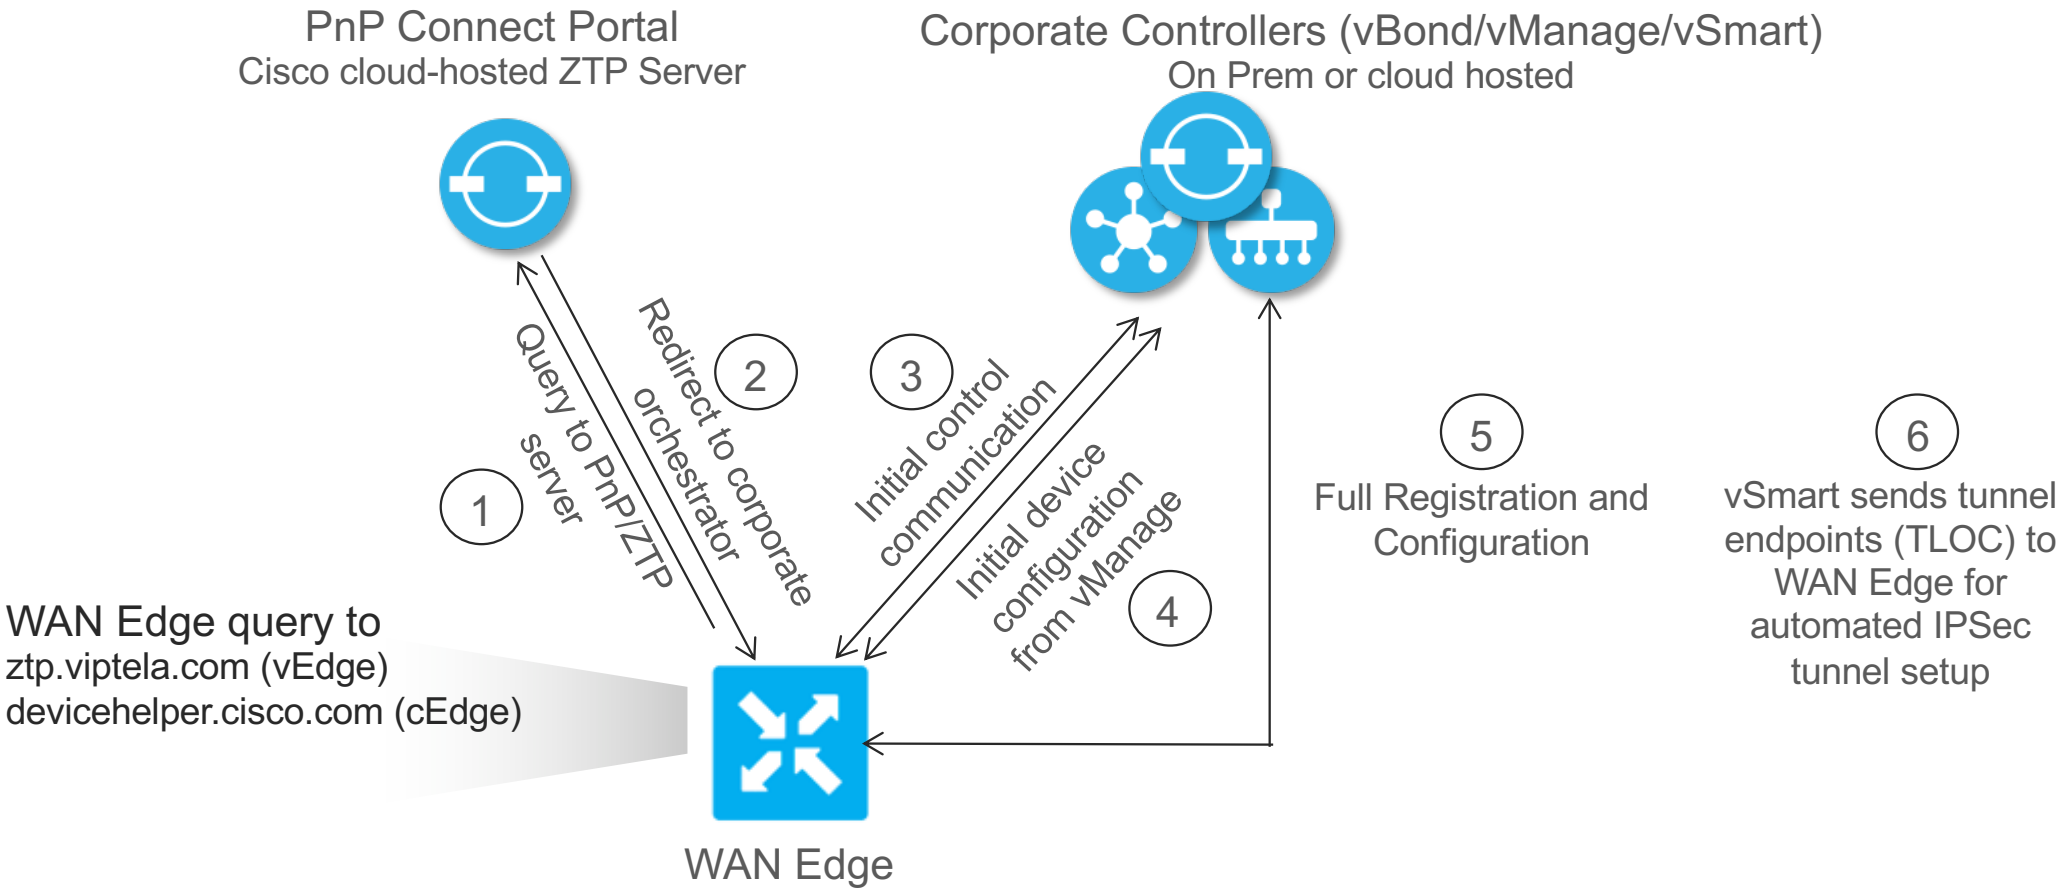

Bandwidth augmentation

Per-flow load sharing across transport-independent overlay

Equal Cost Load Sharing across multiple tunnels/transports

- Out of the box behavior, no policy required
- Can modify Tunnel preference and weight to influence traffic flows on specific links at a site
- Configure Circuit of last resort on 4G LTE if needed as backup only

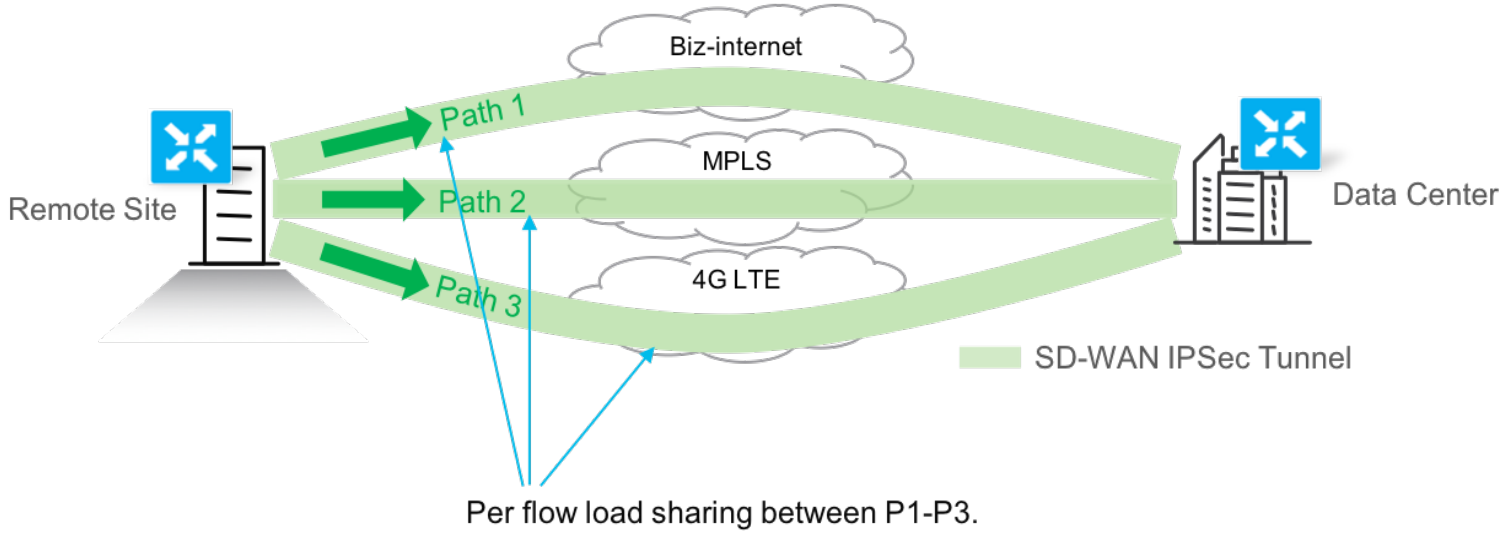

Bandwidth augmentation
Secure Automated WAN

Load sharing across preferred links only with Control Policy

- vSmart Control policy to raise preference on Paths (TLOCs) to make more desirable
- Similar to modifying IGP/BGP metrics

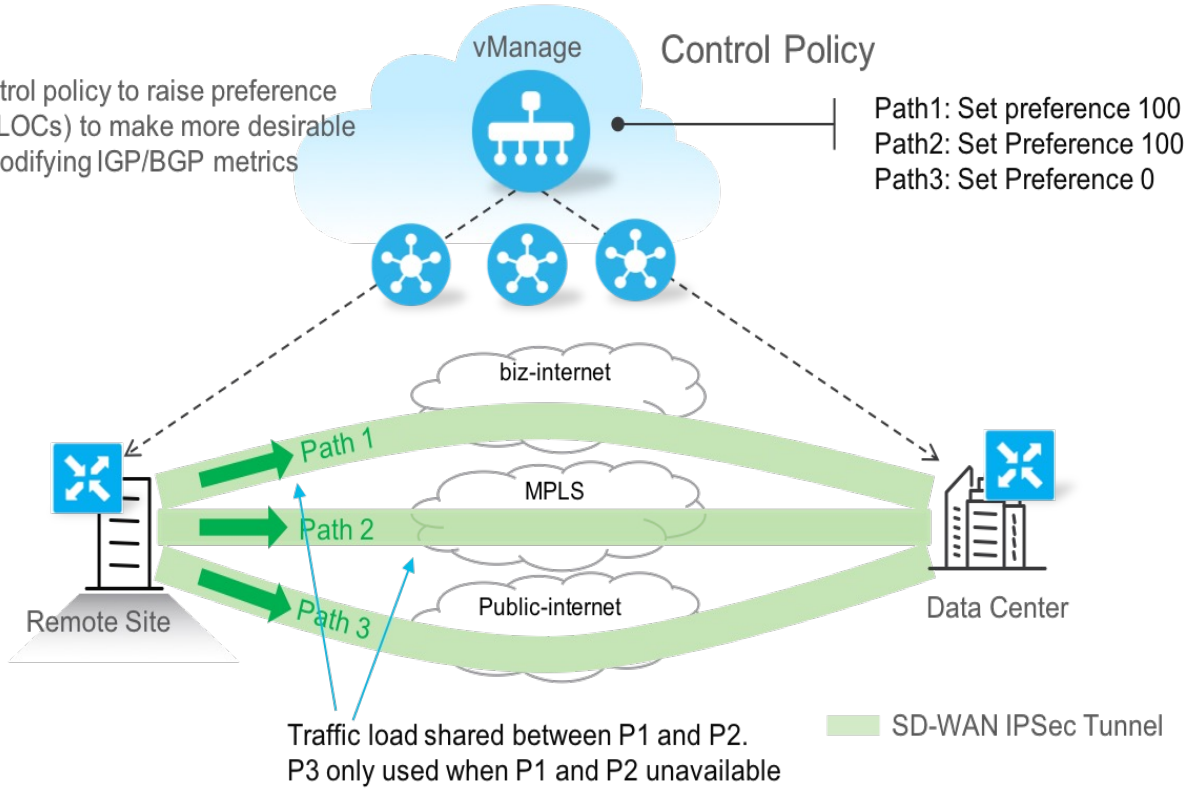

Secure Automated WAN

VPN segmentation

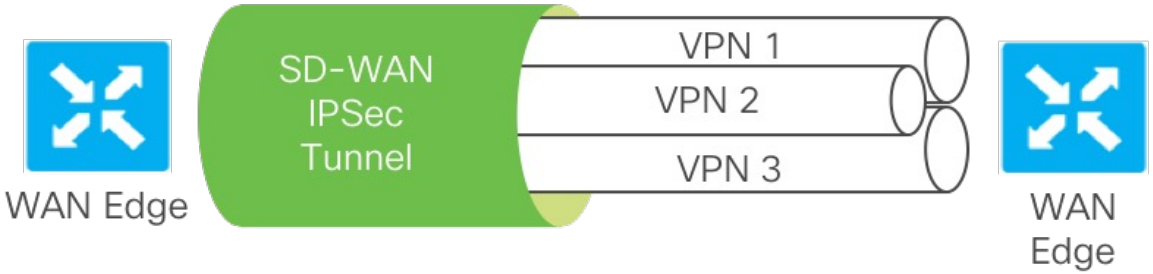

- Security Zoning
- Compliance
- Guest Wi-Fi
- Multi-Tenancy
- Extranet

Per-VPN Topology

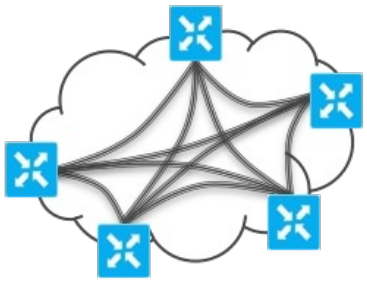

Full-Mesh

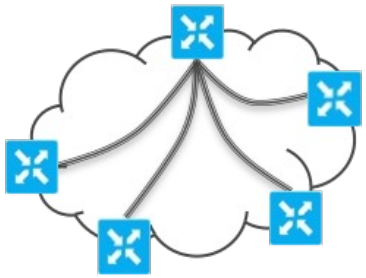

Hub-and-Spoke

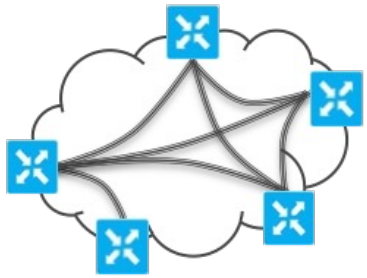

Partial Mesh

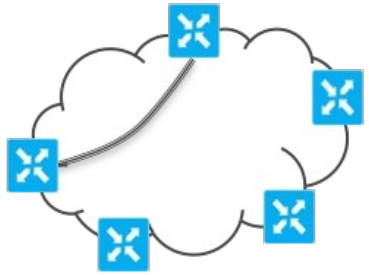

Point-to-Point

SD-WAN

Secure Automated WAN

Application Performance Optimization

Secure Direct Internet Access

Branch Multicloud Access

Regional Hub Multicloud Access

Quality of Service

Mitigating congested WAN links with traffic prioritization, queue management, and link-conditioning features

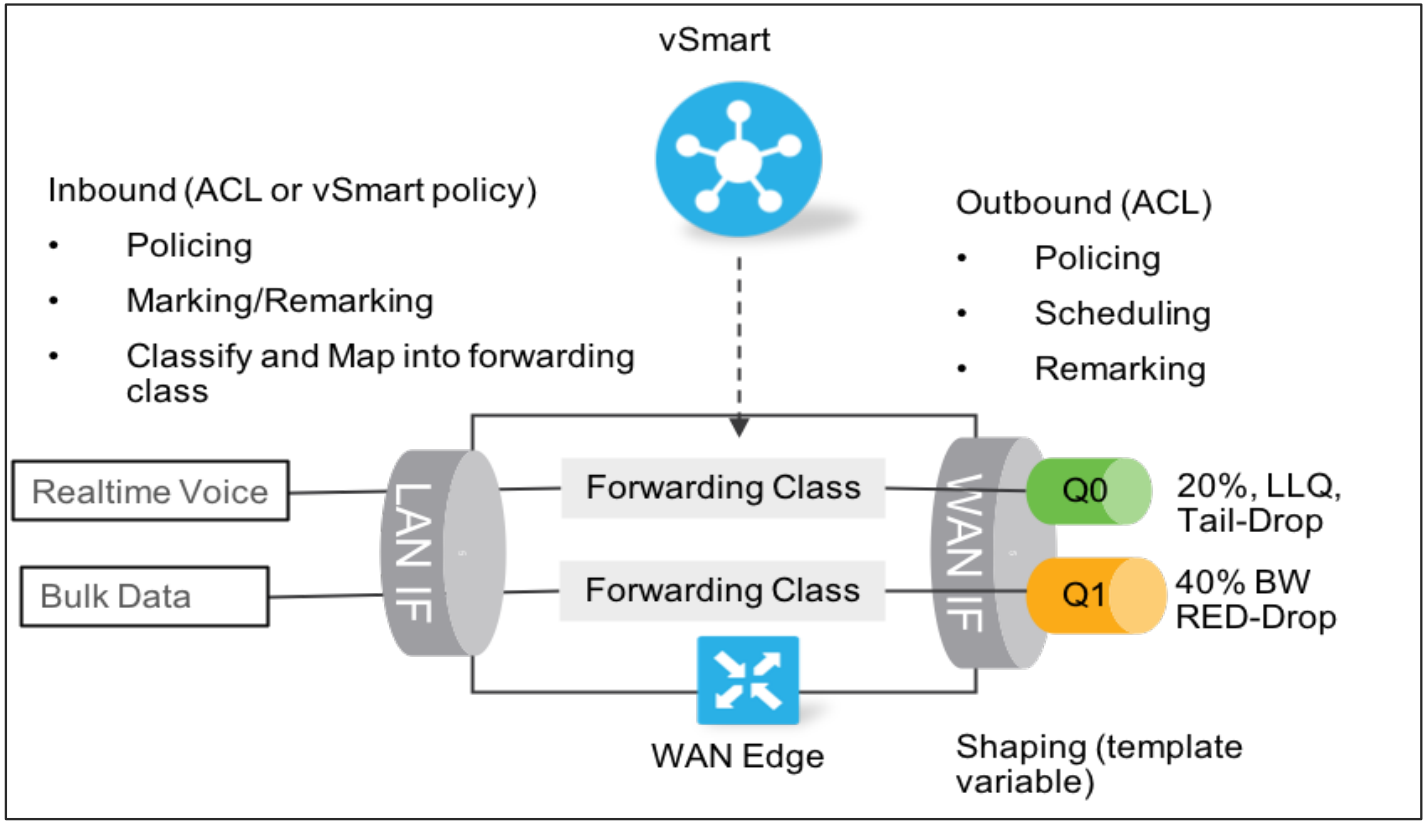

Per-Tunnel QoS with Adaptive Shaping

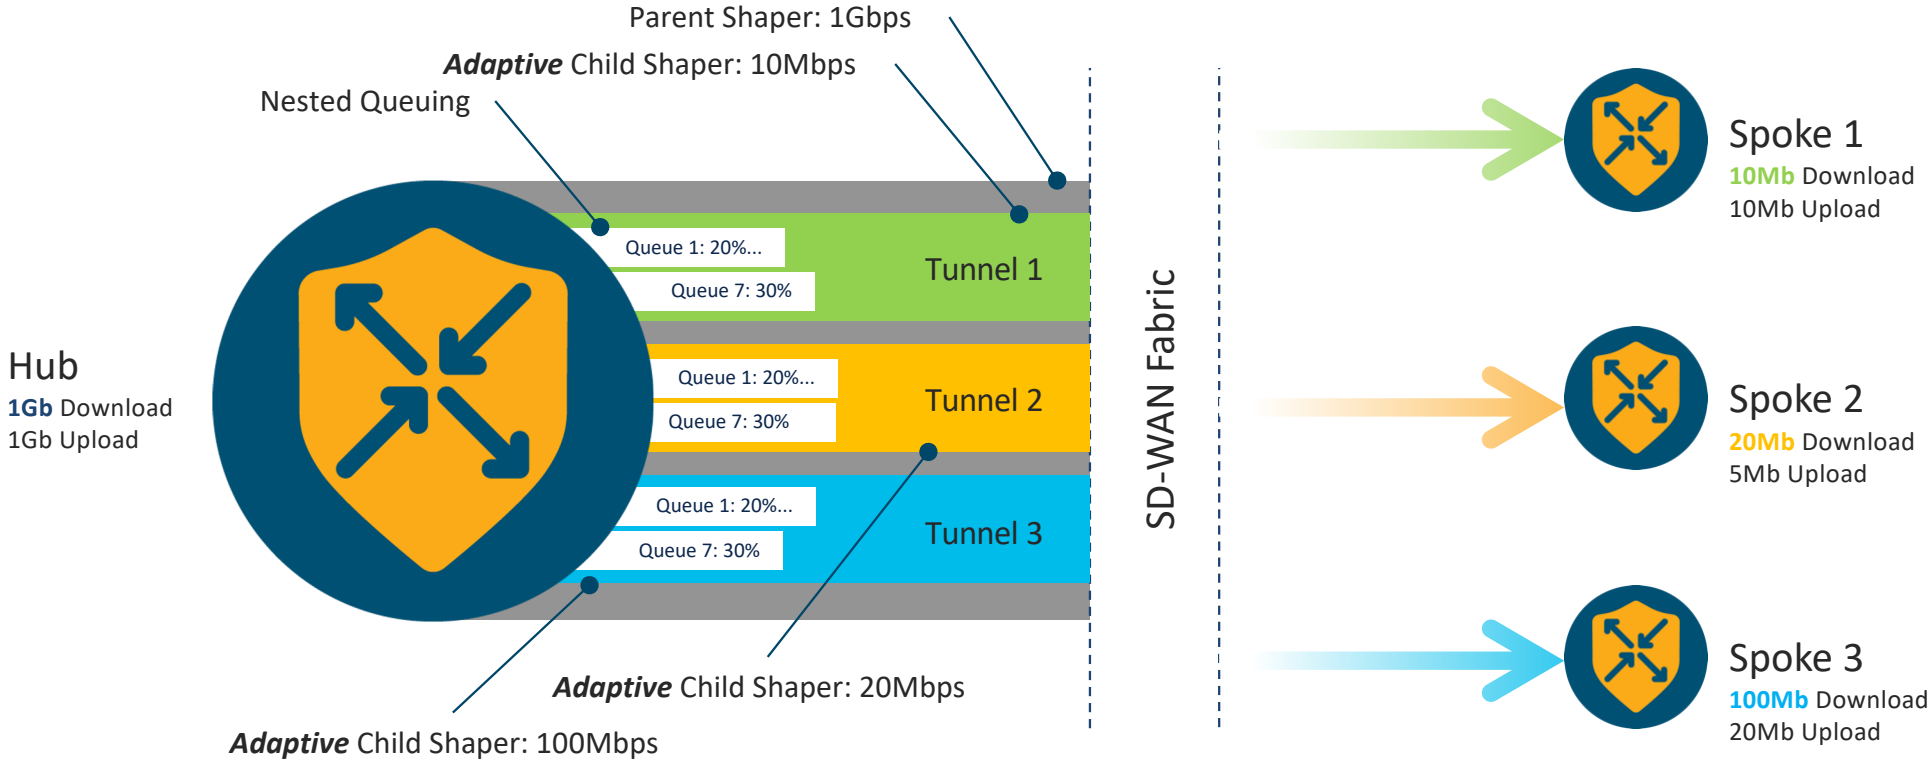

Per-Tunnel QoS allows the Hub site to dynamically adjust the sending rate of its traffic to accommodate lower bandwidth circuits at remote locations. Adaptive shapers measure the ***true*** circuit capacity at any given moment – rather than relying on static configuration.

Application-aware routing

Protecting critical traffic with performance-based path selection

- vEdge Routers continuously perform path liveliness and quality measurements

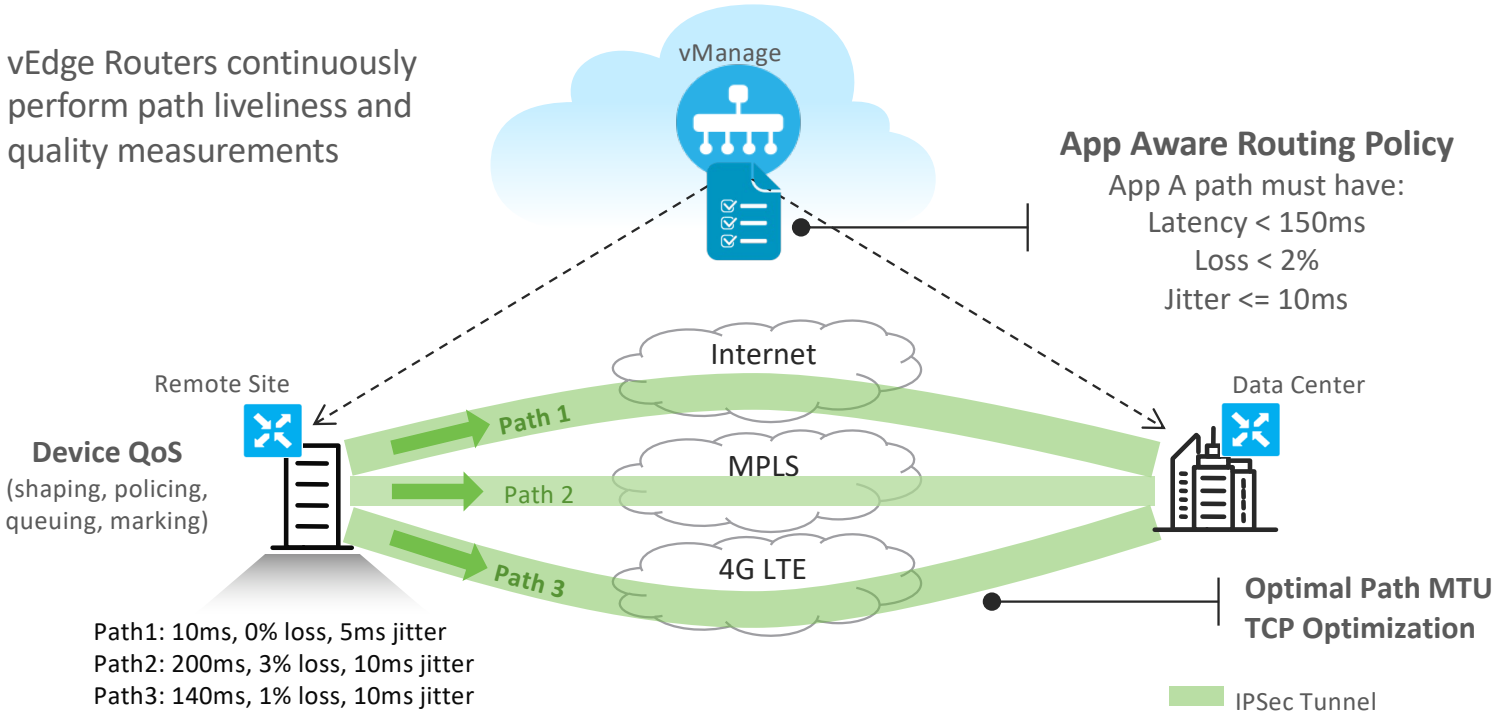

Latency and TCP throughput optimization

TCP optimization and session persistence

High latency and bad throughput can be improved with TCP optimization and session persistence
Examples: transcontinental or long-haul links and high-latency satellite links

With **TCP optimization**, a WAN Edge router acts as a TCP proxy between a client that is initiating a TCP flow and a server that is listening for a flow:

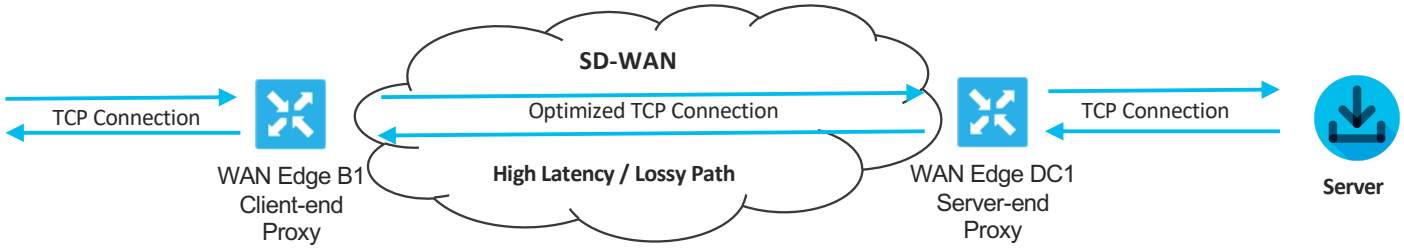

Session Persistence is an additional option to improve latency and throughput:

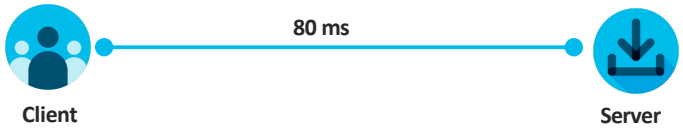

New connection for every request/response pair

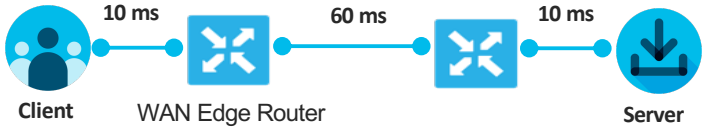

Single TCP connection to send and receive multiple requests/responses

Forward Error Correction

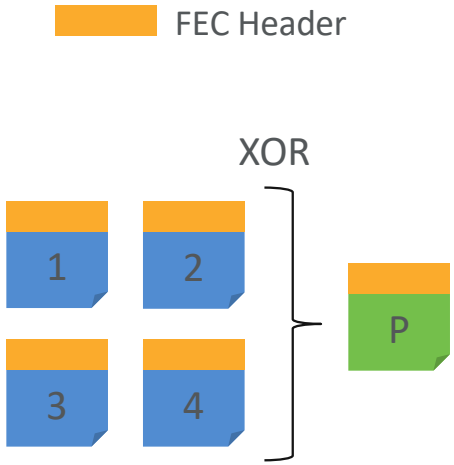

Highlights:

- Protects against packet loss for critical applications
- Protocol agnostic (TCP/UDP)
- Dynamically invoked
- Operates per-tunnel

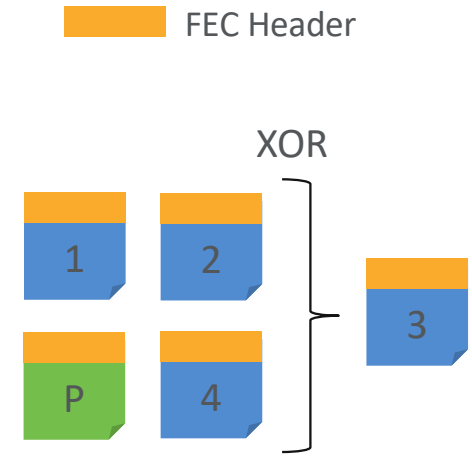

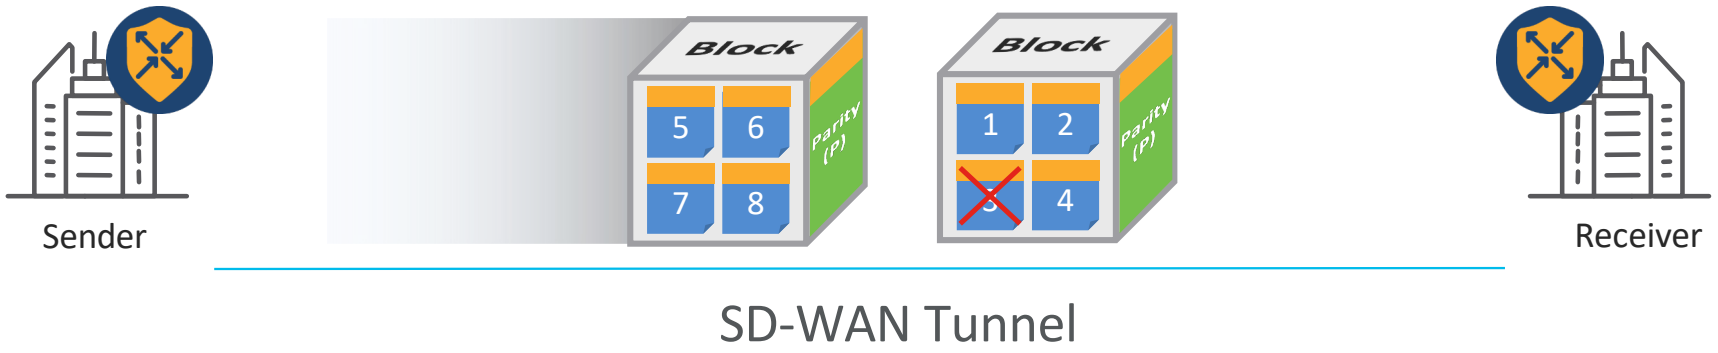

Packet Duplication

Highlights:

- Protects against packet loss for critical applications
- Protocol agnostic (TCP/UDP)
- Works only over multiple tunnels
- Duplicates are discarded on receiver

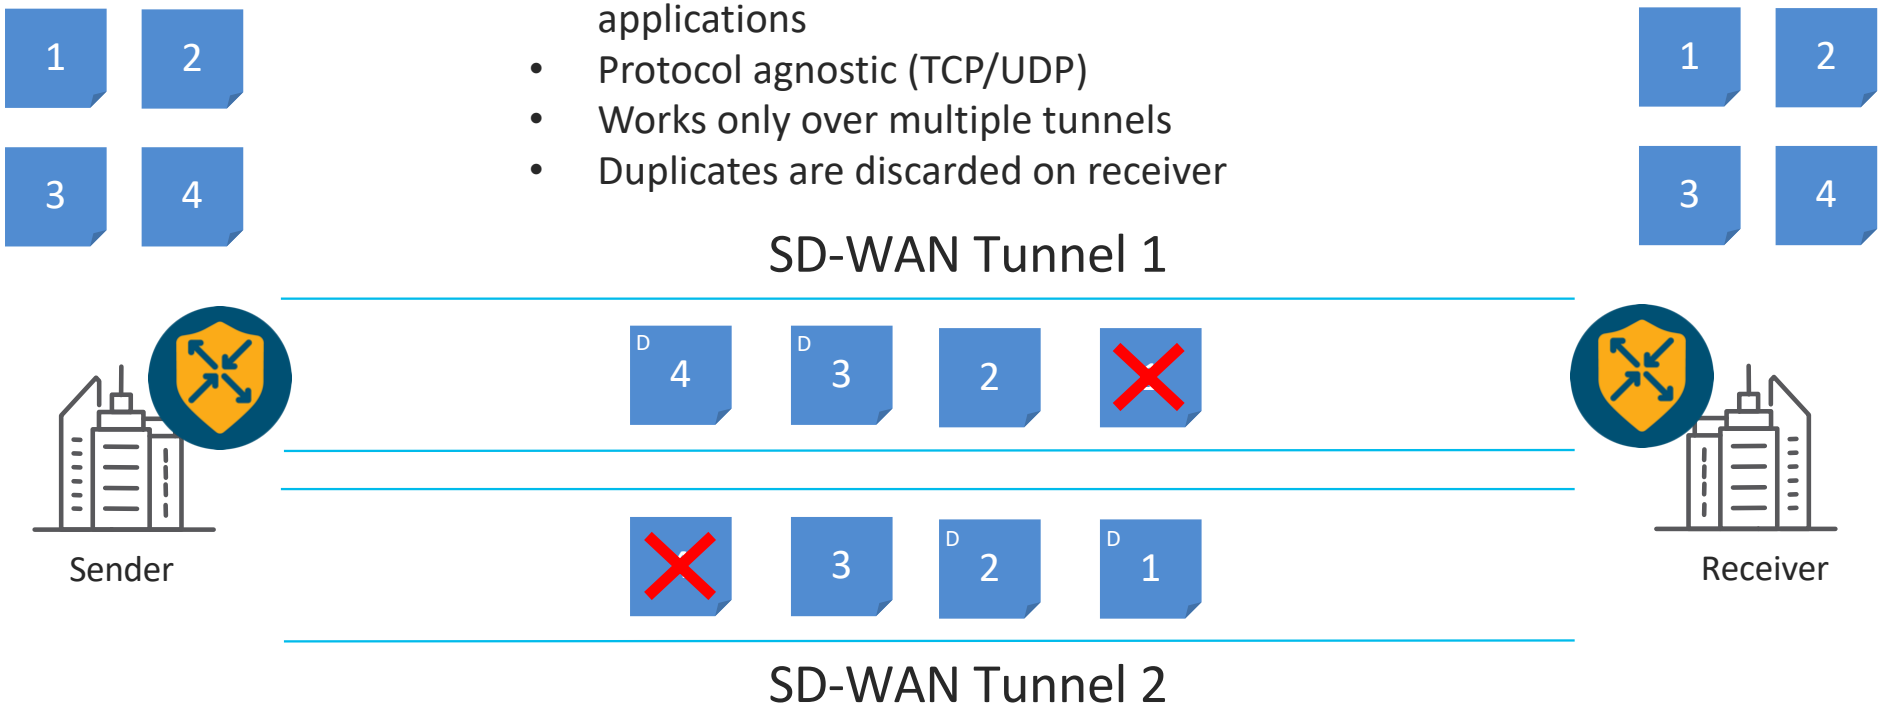

Management and Analytics Architecture

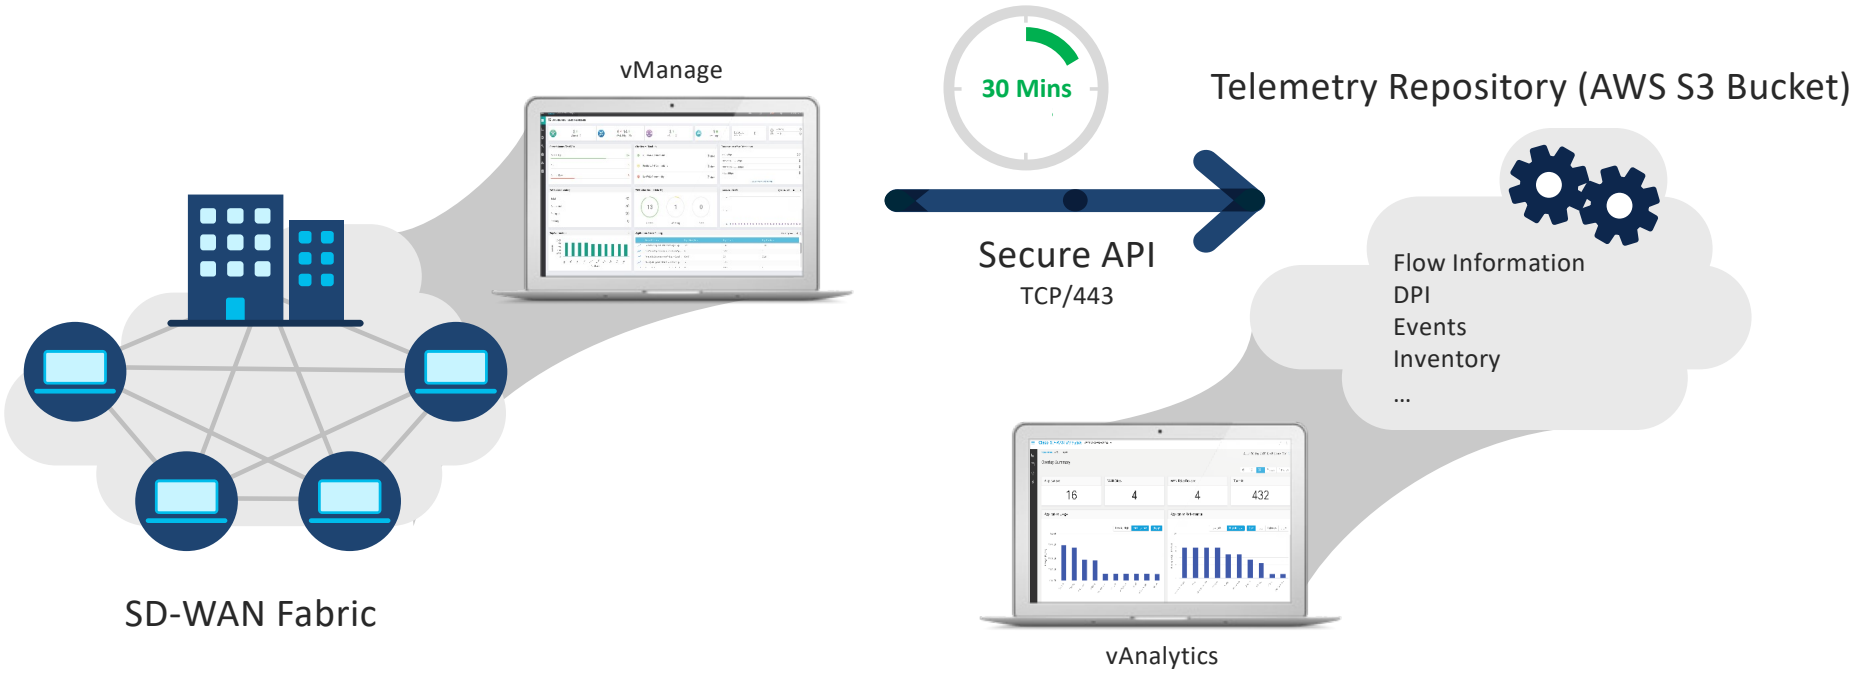

On-Prem or Cloud-Hosted SD-WAN (vManage)

Cloud-Hosted Analytics

Track Digital Experience with ThousandEyes Dashboards

Application availability

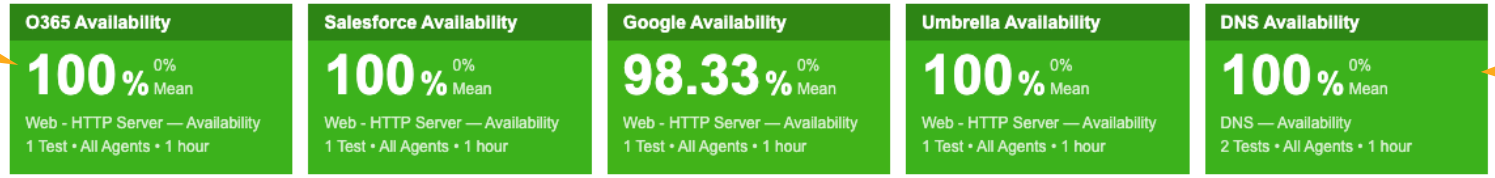

SASE & service availability

Per-office availability

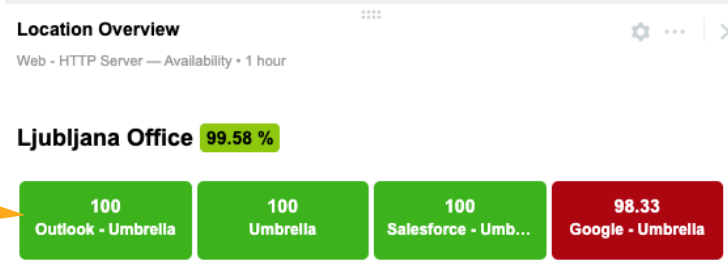

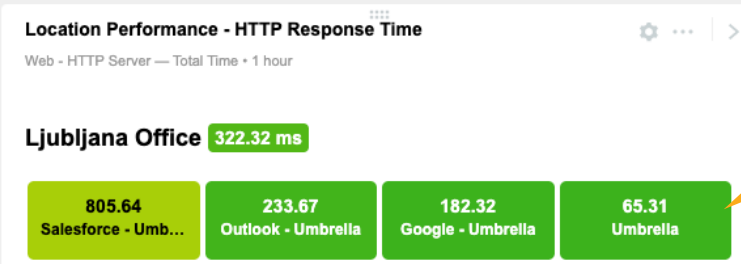

Per-office performance

Application performance

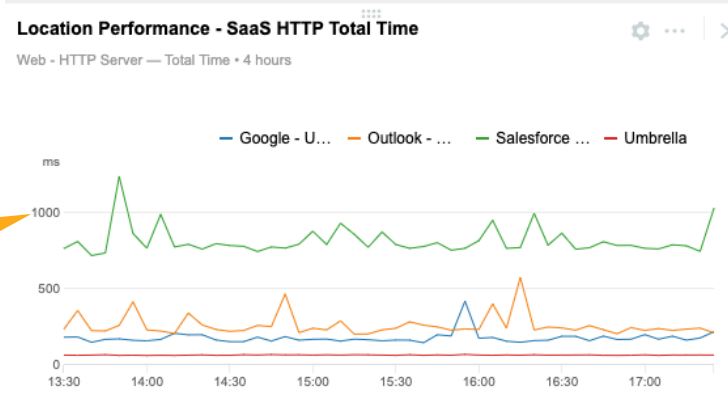

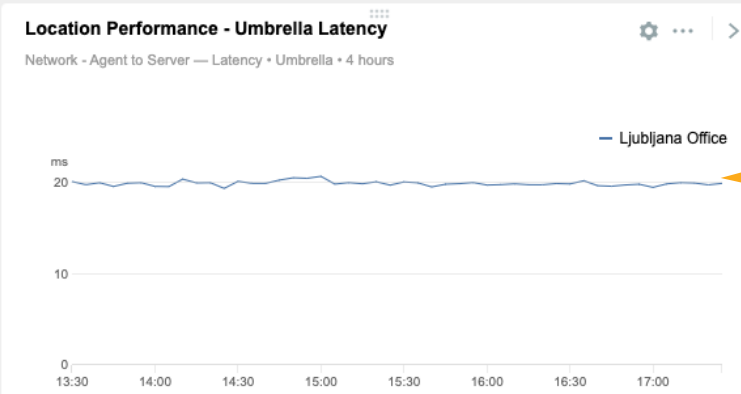

SASE performance

Multi-layer Correlation and Monitoring

When service degradation occurs, quickly identify where the problem is.

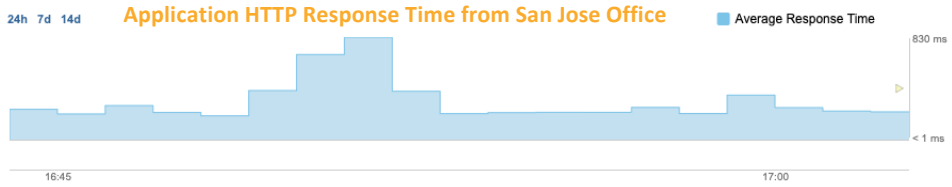

1 Increase in the service response time

2 Due to an increase in network latency

SD-WAN Overlay Path Visualization

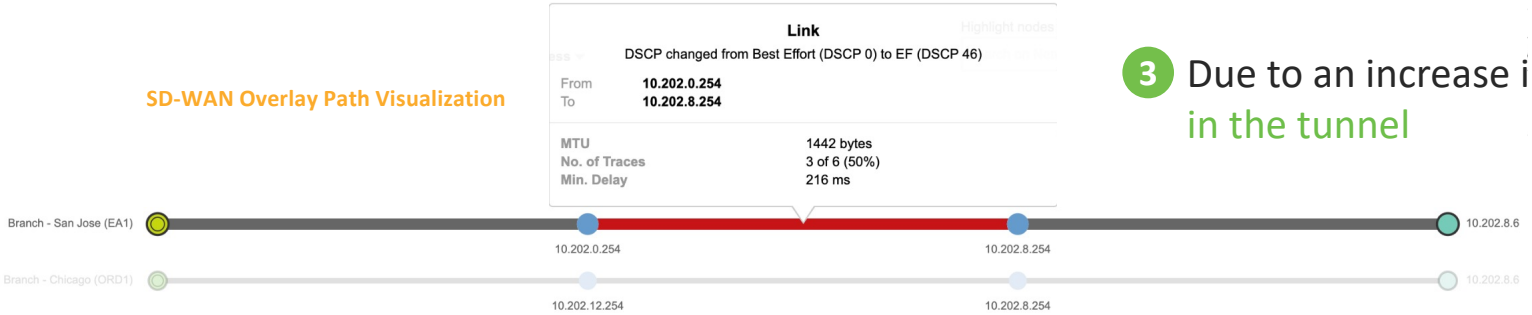

3 Due to an increase in network latency in the tunnel

SD-WAN Underlay Path Visualization

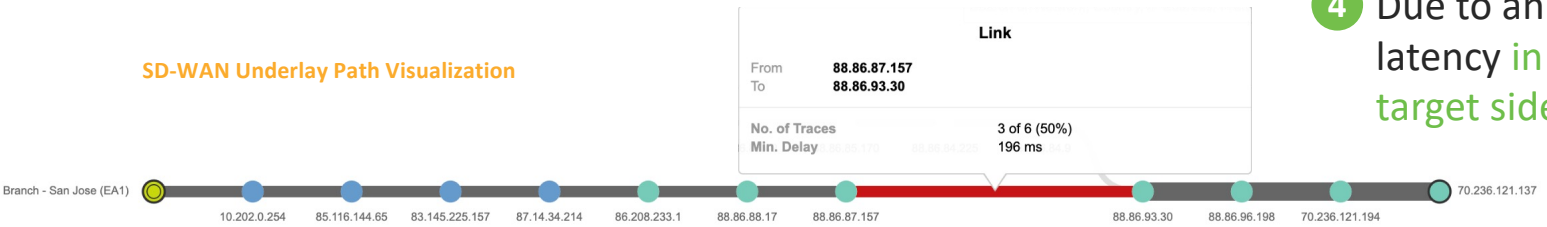

4 Due to an increase in network latency in a specific link on the target side of the underlay network

Collect performance data from every perspective

Enterprise Agents

- ✓ Deploy it in your office or data center.
- ✓ Test services in your DC, cloud or SaaS services.
- ✓ Test through Umbrella, SD-WAN or over the Internet.
- ✓ Deploy it on your Cisco devices (Cat 9K, Cat 8K, ISR 4K), or as a VM, Docker container, Linux package, Intel NUC or Raspberry Pi2 appliance.

Cloud Agents

- ✓ Globally distributed agents installed and managed by Cisco ThousandEyes in 200+ cities.
- ✓ Use them as a reference points to understand how the performance your employees are experiencing compares to performance of SaaS and other public cloud services from the public Internet.

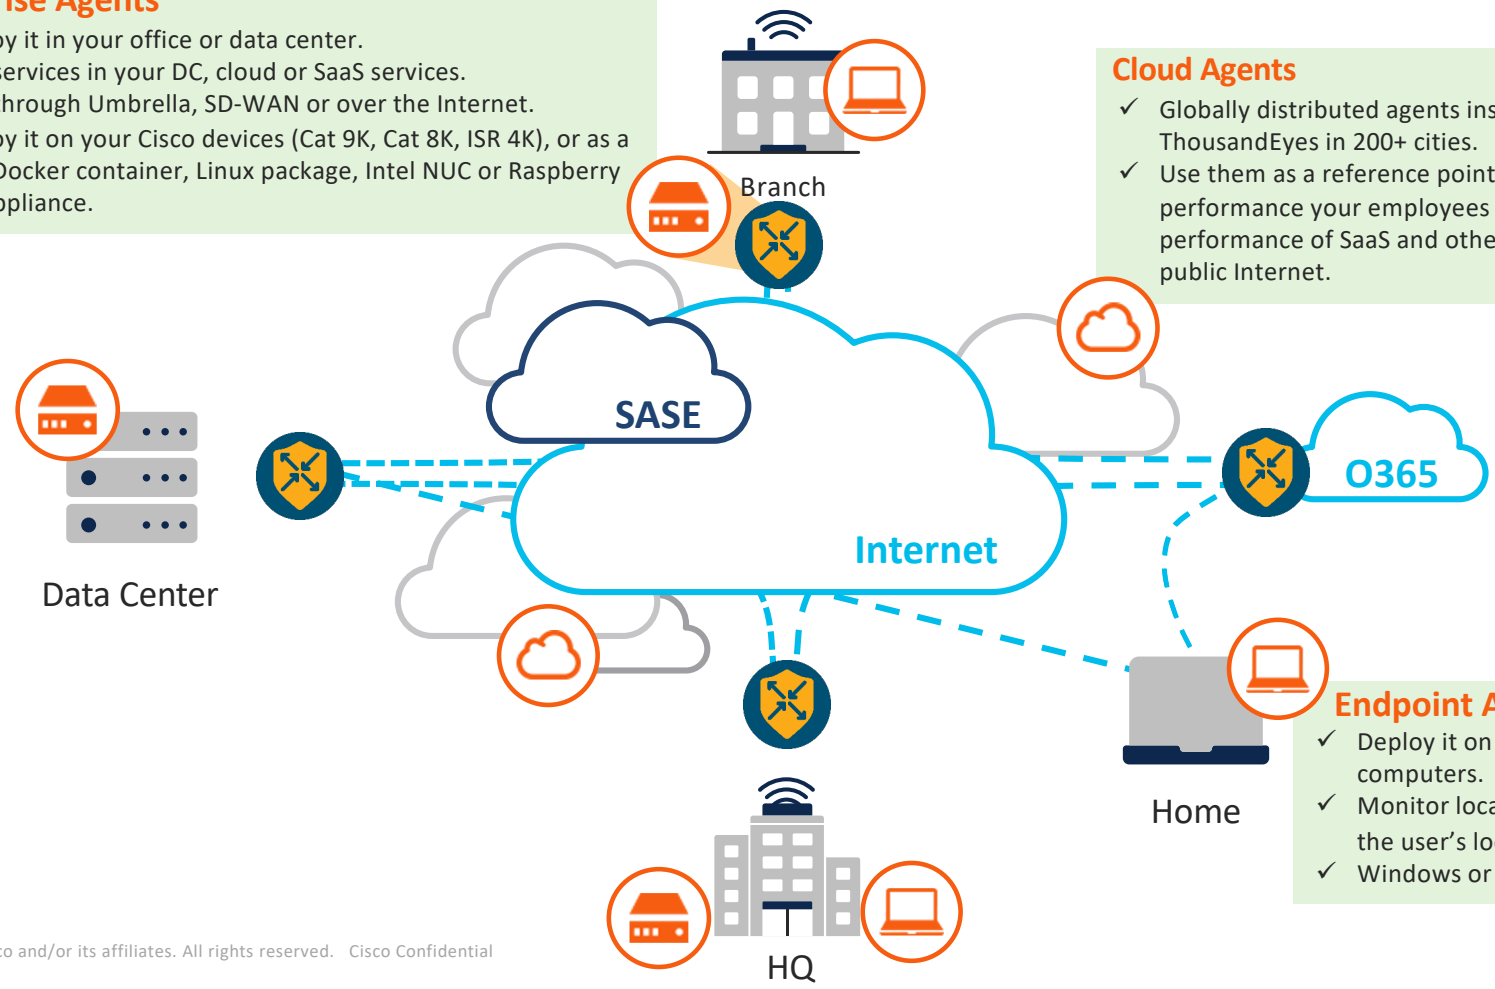

Endpoint Agents

- ✓ Deploy it on your employee's desktop & laptop computers.
- ✓ Monitor local network conditions regardless of the user's location (office, work-from-home).
- ✓ Windows or Mac

SD-WAN

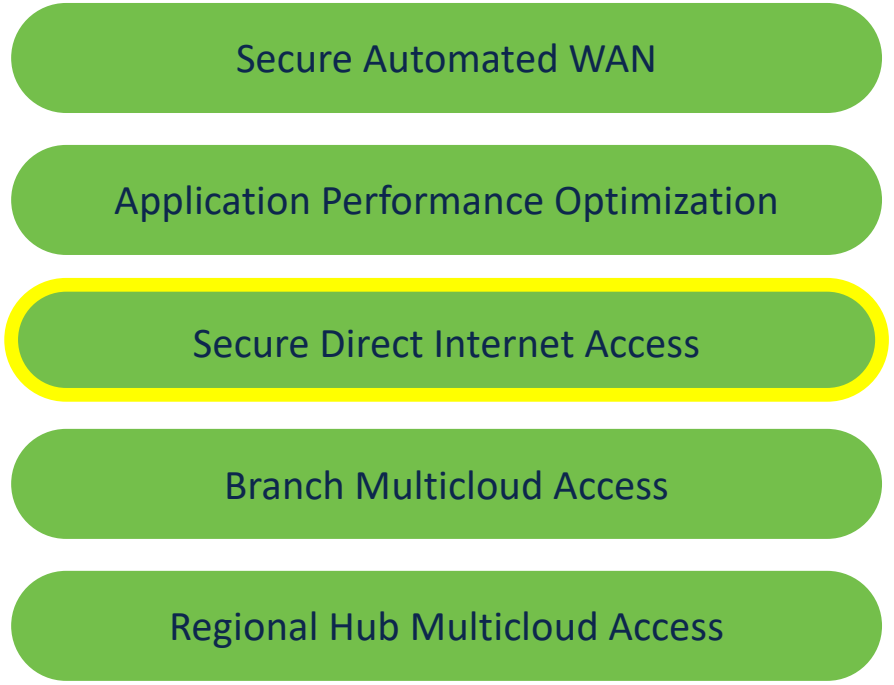

Internet access from SD-WAN

Regional Internet Exit and Direct Internet Access

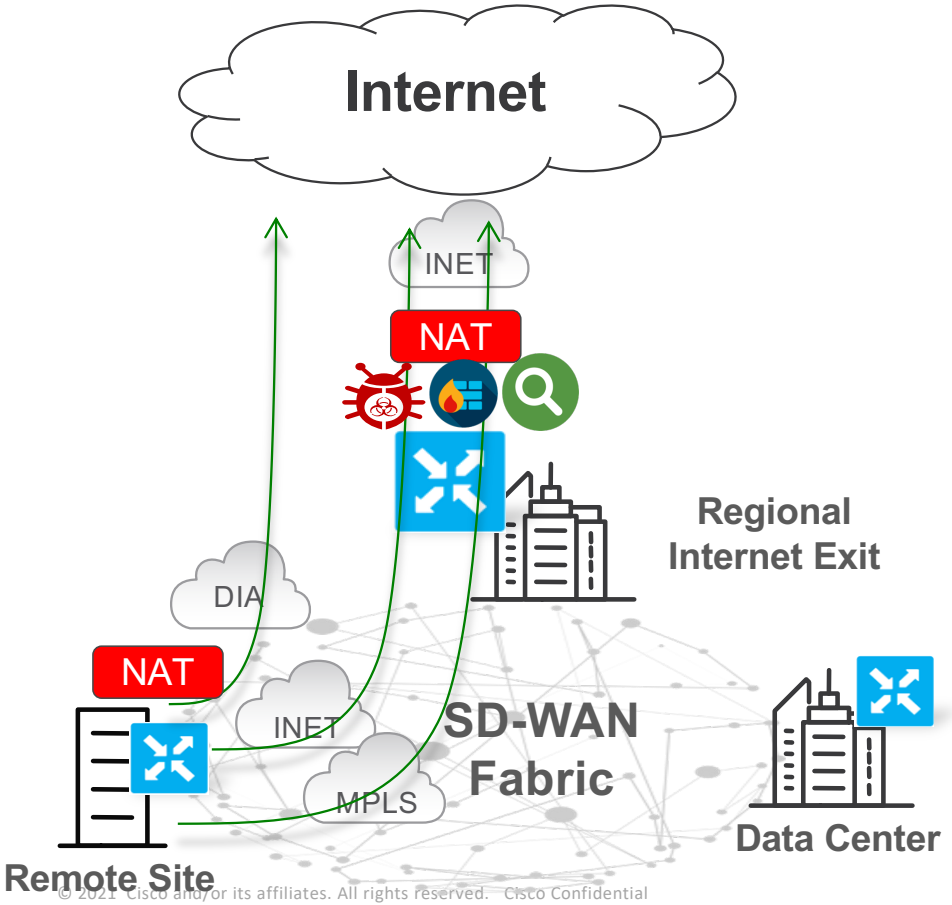

- **Regional Internet Exit:** Internet-bound traffic backhauled over SD-WAN tunnels to Data Center or designated Regional Internet Exit
  - Pros: Centralized Internet security services in a DMZ, with nothing additional needed at remote site
  - Cons: Additional latency with backhaul through DC, and additional traffic on SD-WAN fabric and centralized INET circuits

- **Direct Internet Access (DIA):** Internet-bound traffic from some or all VPNs leaves local Internet exit at remote site
  - Pros: Optimal path to Internet with no added latency or traffic on SD-WAN overlay
  - Cons: Poses a security challenge, as remote sites need local FW, IPS, AMP, etc.

Securing Direct Internet Access

Option 1: Leverage embedded SD-WAN security features

VPN1

Employee

VPN2

Guest

Enterprise FW App Aware

IDS/IPS

DNS/web layer security

URL Filtering

Advanced Malware Protection (AMP)

XE SD-WAN Router

- Define Inside and Outside Zones
- Inspect policy allows only return traffic
- Drop traffic by application or category

- Detect/protect against malicious traffic based on known signatures with Snort IPS engine
- Leverage TALOS cloud service for signature updates

- Integrate with Umbrella Cloud
- Block based on DNS Requests

- Enforce acceptable use controls
- Block based on Web Reputation scores
- On box and/or cloud updates

- Integrate with AMP for file reputation
- Integrate with Threat Grid for malware sandboxing
- On box and/or cloud updates

Securing Direct Internet Access

Option 2: Secure Internet Gateway (SIG)

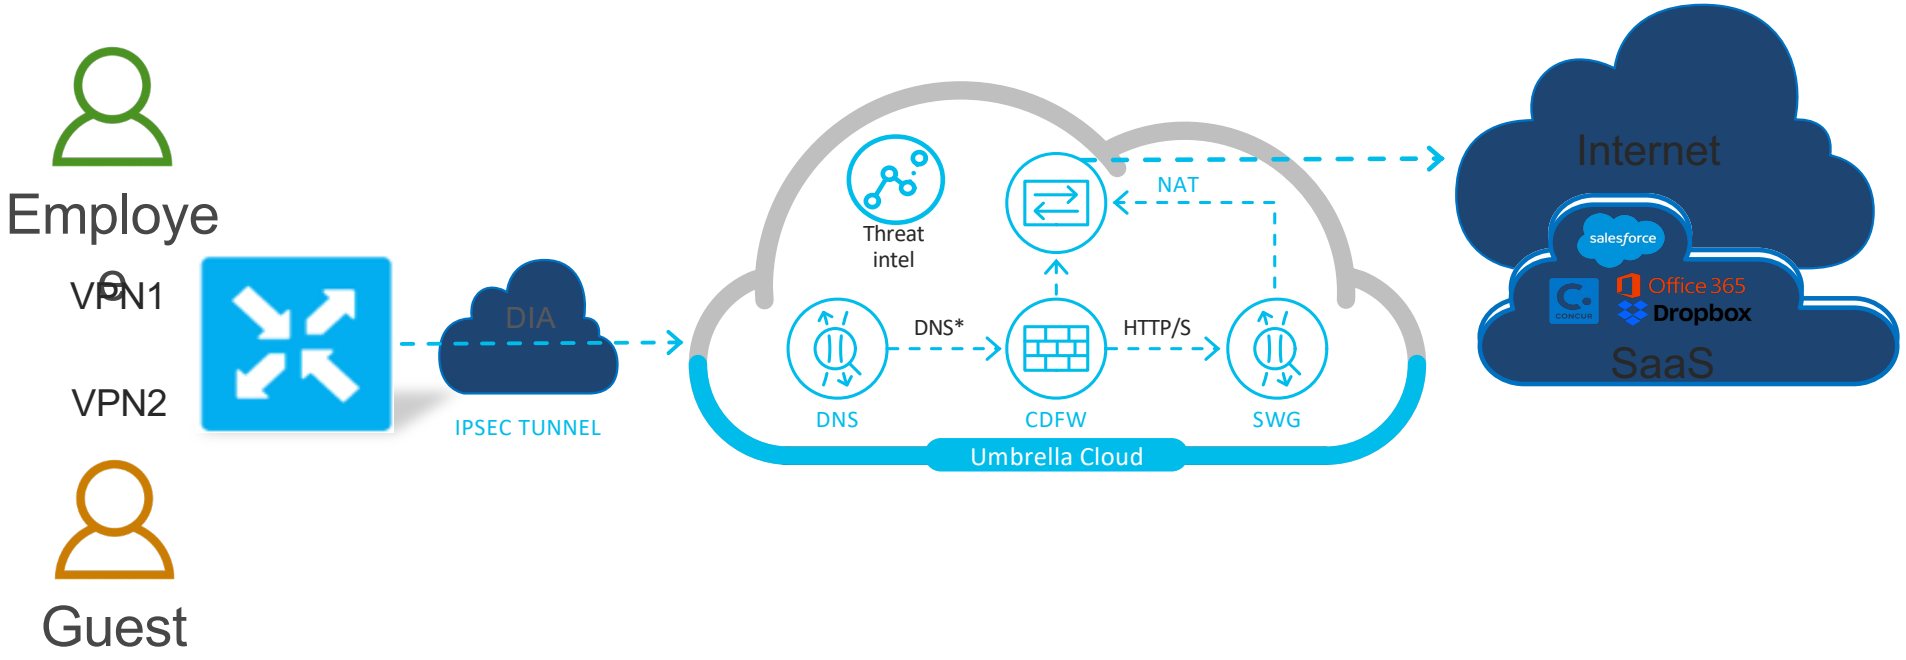

SD-WAN

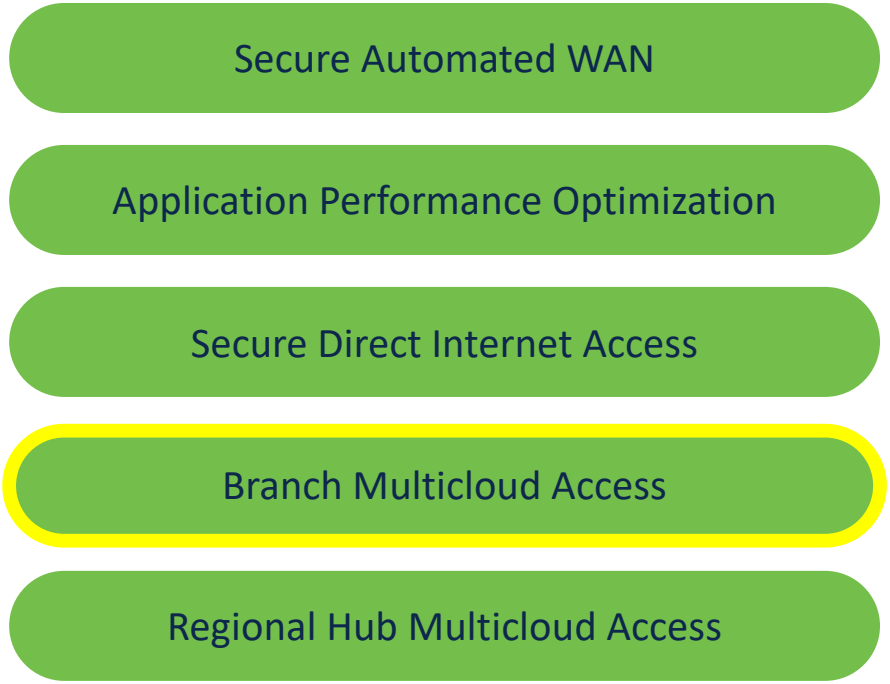

Evolutionary SaaS cloud adoption with SD-WAN

New Complications:

- Which way is cloud?
- What is cloud performance?
- Where is security?

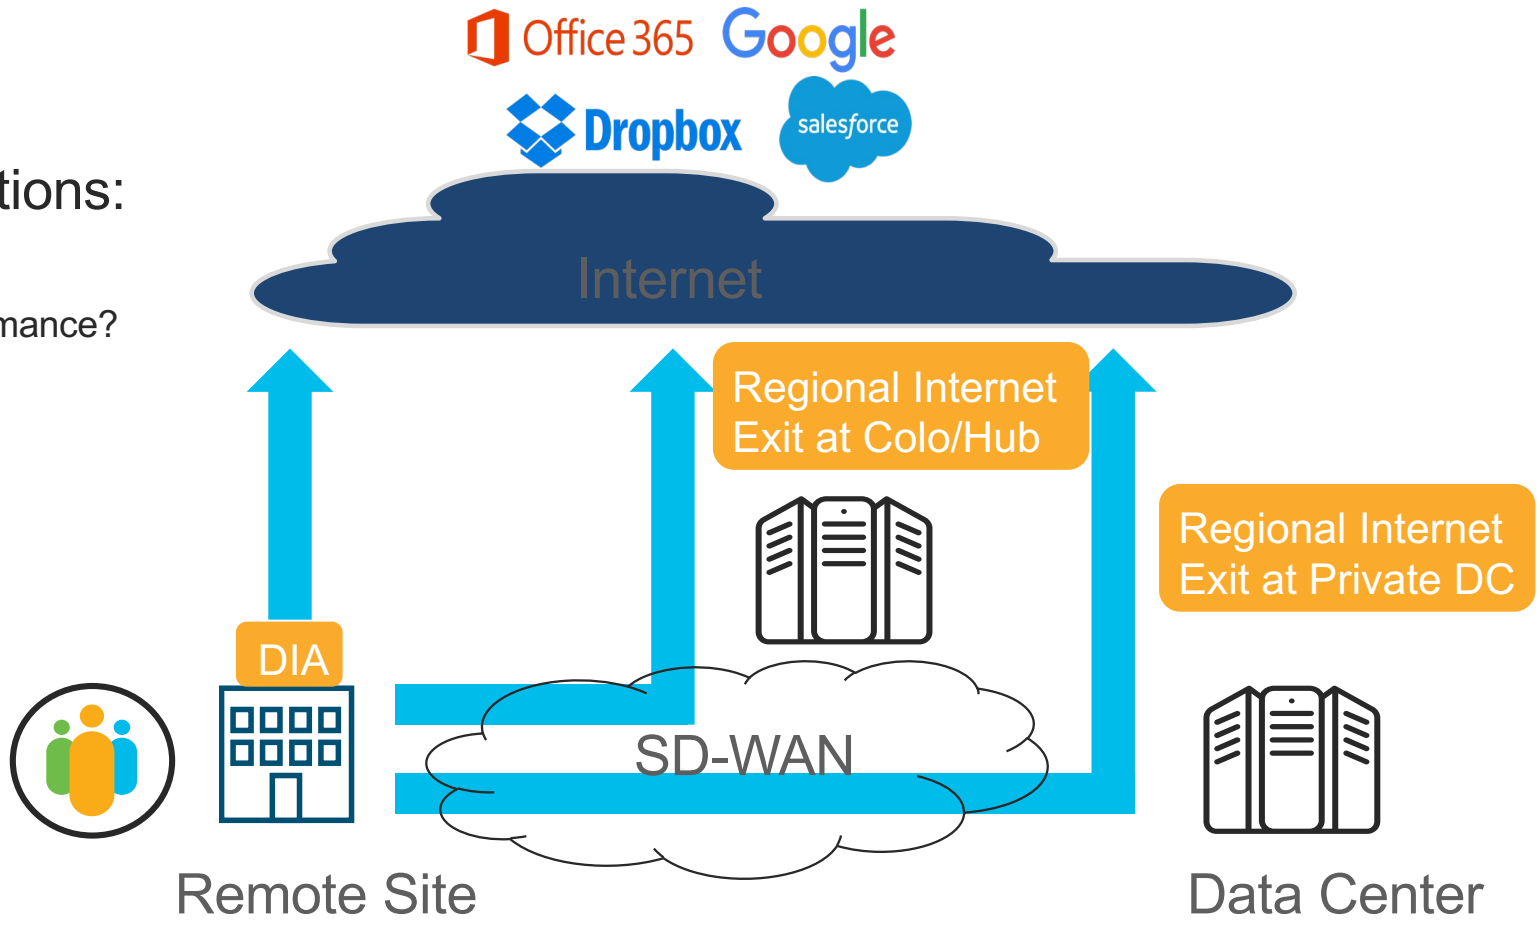

Cisco SD-WAN solution - Cloud onRamp for SaaS

Performance-based path selection (DIA and GW)

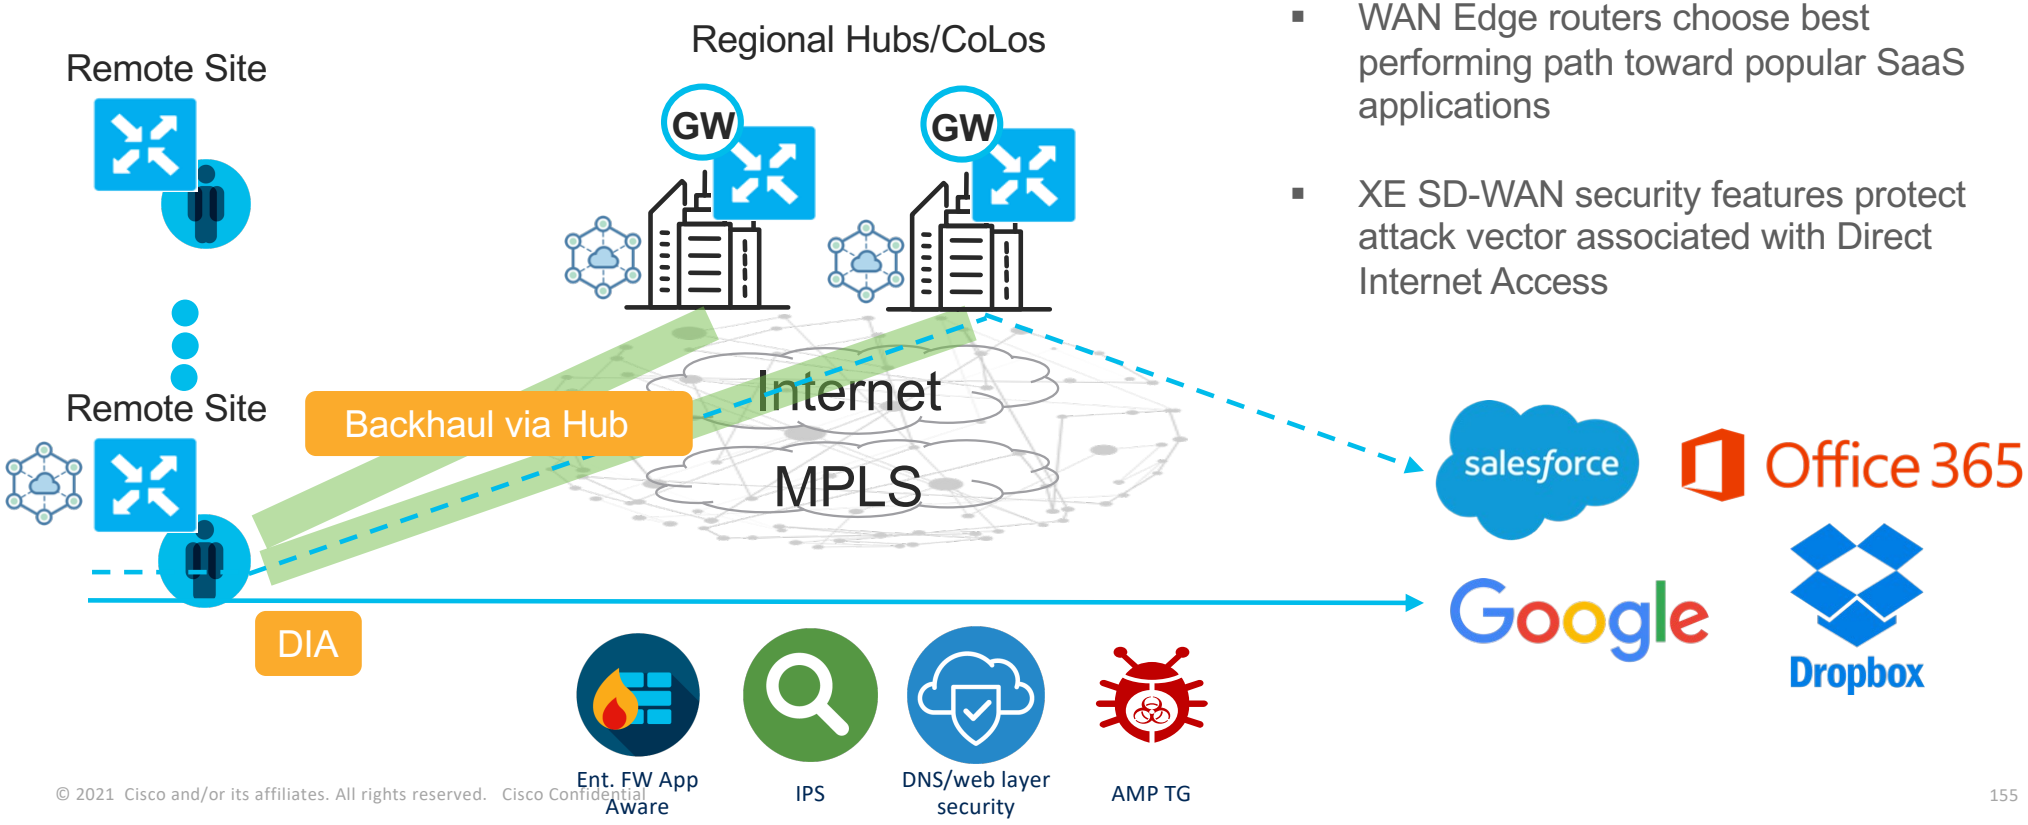

- Quality probing toward popular SaaS applications from all configured exits
- WAN Edge routers choose best performing path toward popular SaaS applications
- XE SD-WAN security features protect attack vector associated with Direct Internet Access

Cloud OnRamp for Microsoft 365

Improving the user experience

Use case

- How to optimize only certain M365 categories?
- How to gain M365 telemetry view to gain insights into application performance?
- When a specific path is having performance issues, how to automatically re-route traffic?

Features

- Dynamic URL/IP Categorization.
  - Distinct URLs for different applications (can be mapped to different traffic precedence and service-area); M365 traffic divided into 3 categories based on sensivity.
- Microsoft Informed Routing.
  - End-to-end telemetry using Application Infused Path Feedback (AIPF); import and export telemetry from/to Microsoft for best path selection.

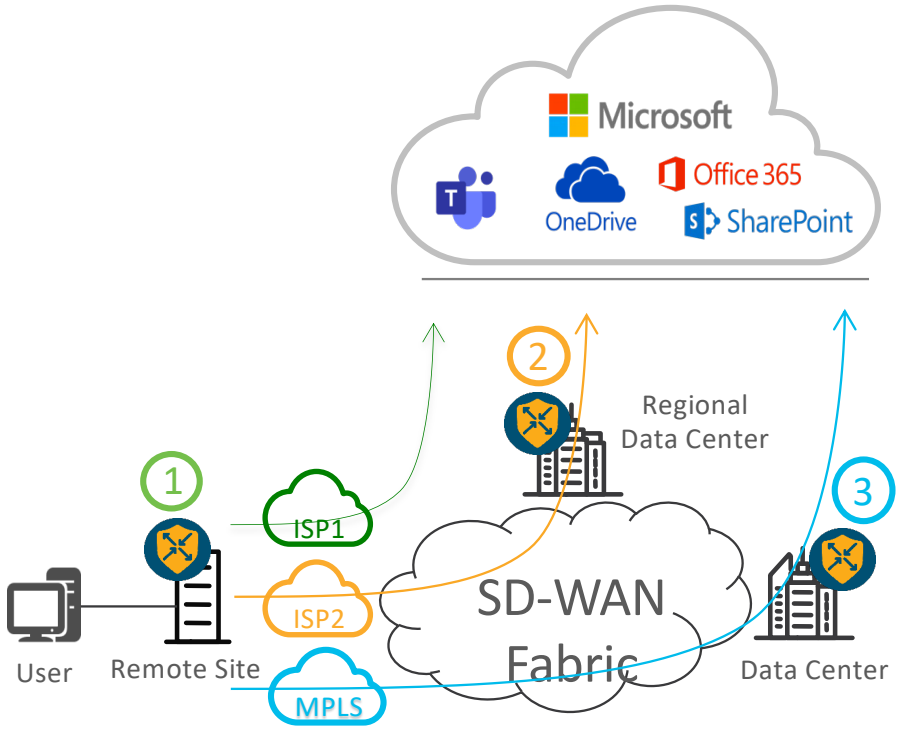

Traditional IaaS access

- No direct access to Public Cloud DC
- Limited segmentation and QoS
- Dependent on underlying technology

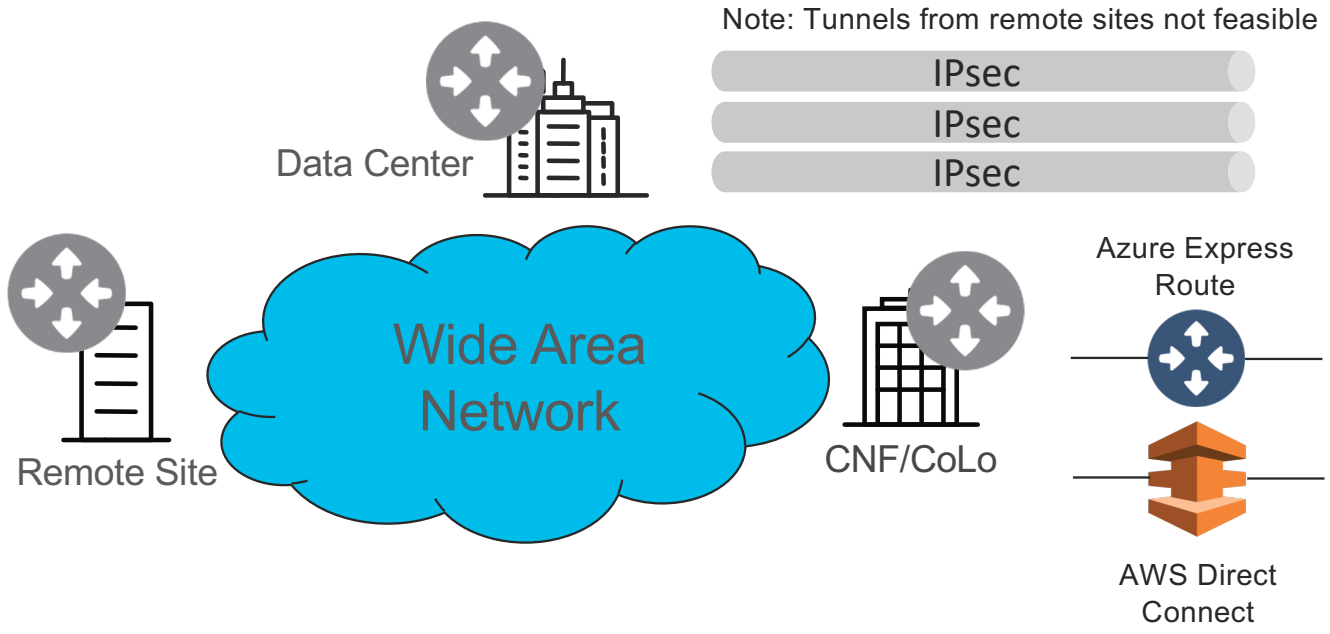

Public Cloud Data Centers

Microsoft Azure

VNET VNET
VNET VNET

amazon web services

VPC VPC
VPC VPC

Cloud onRamp for IaaS

Extending SD-WAN fabric to the Cloud DC (two options)

Using Cloud Provider Marketplace (DIY)

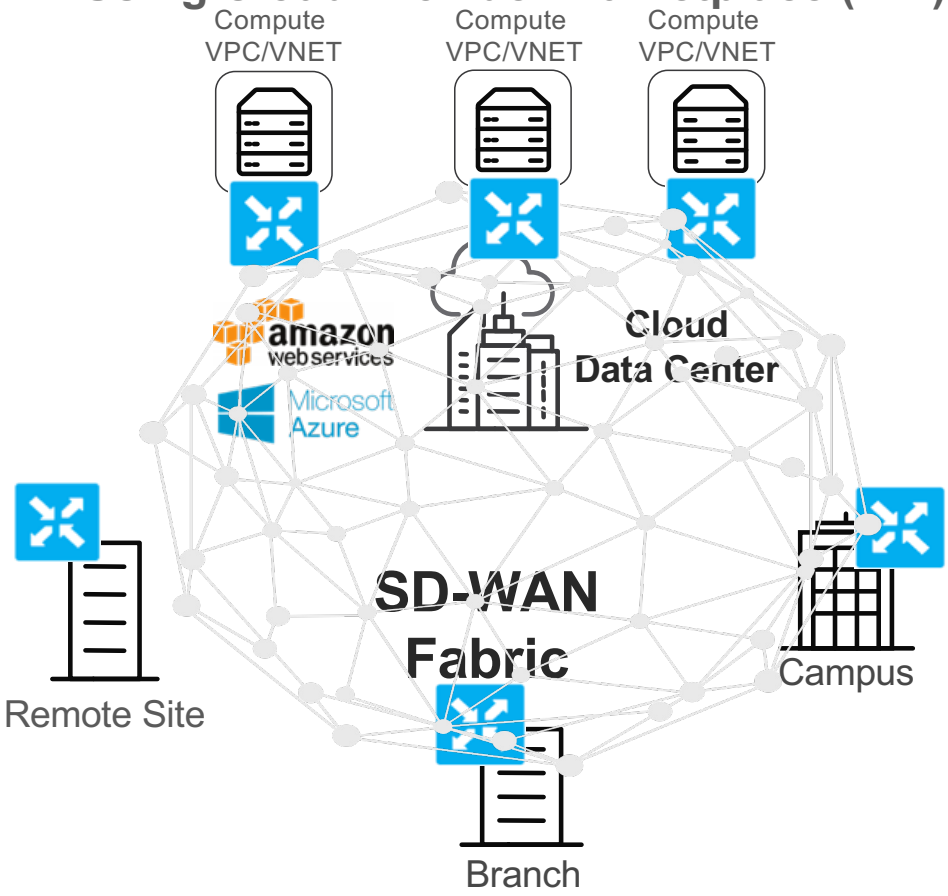

Fully Automated (CoR for IaaS)

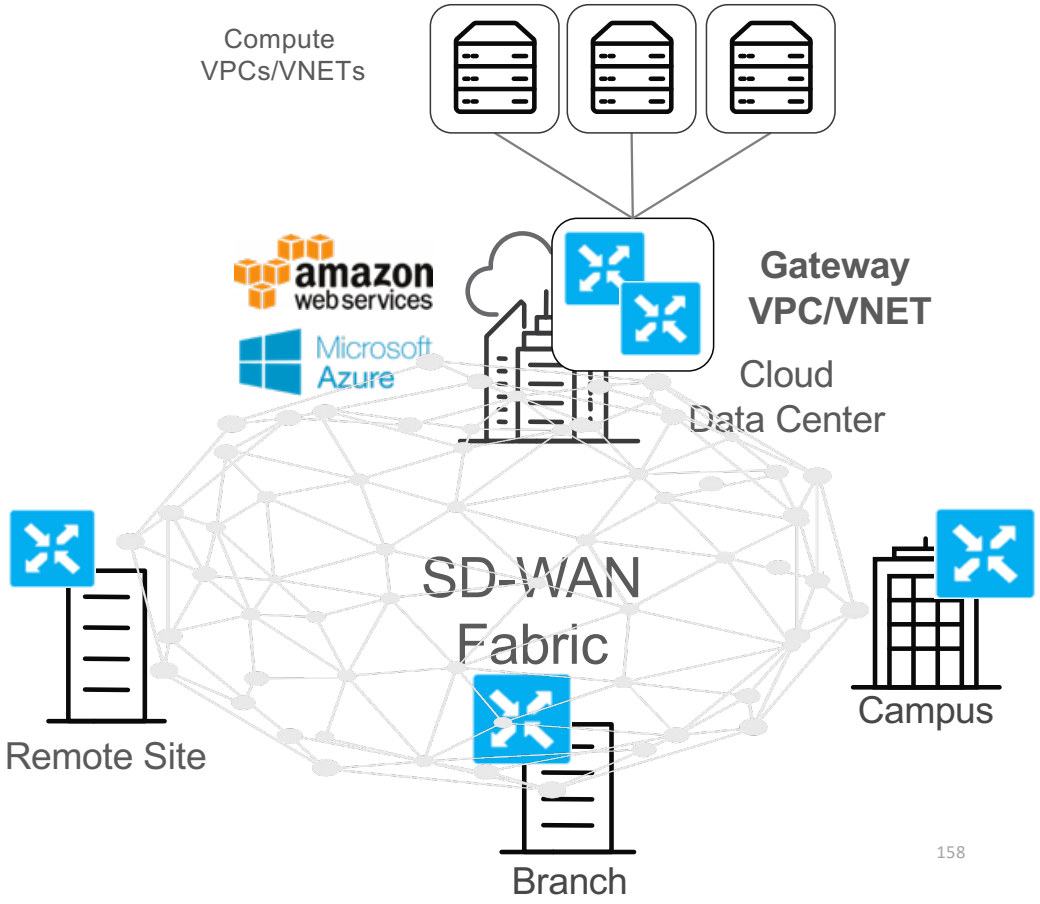

SD-WAN

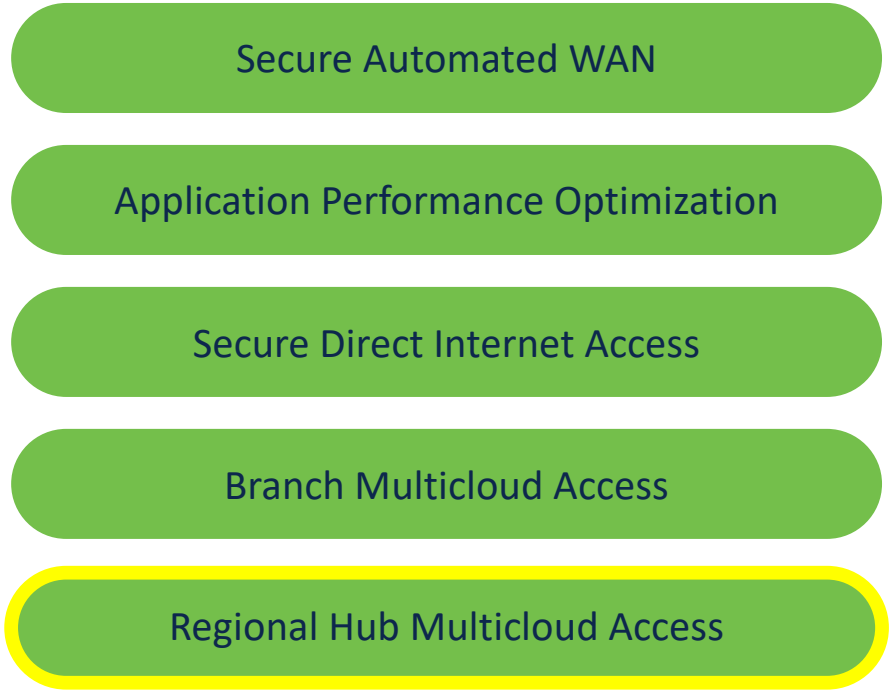

Challenges of providing Multicloud access to disparate user groups

Business Challenges:

- Optimizing IaaS and SaaS access
- Defining and maintaining service-level agreements (SLAs)
- Managing distributed Internet access
- Providing appropriate level of security for various user groups
- Operational efficiency

Multicloud access from the SD-WAN branch

Design options leveraging centralized security

Backhaul through DC/Regional Internet Exit

Pro: Simple networking and centralized branch security

Con: Backhaul latency may affect user experience at some sites

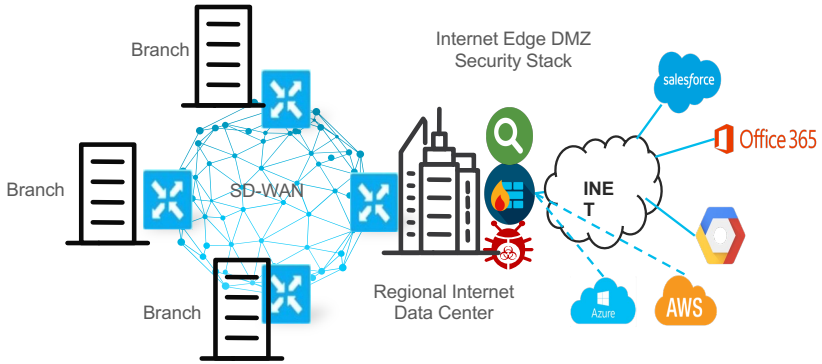

IPSec over DIA to Cloud Security Internet Gateway (SIG)

Pro: User experience improved for branches in close proximity to POPs

Con: Network design and level of security control constrained by provider

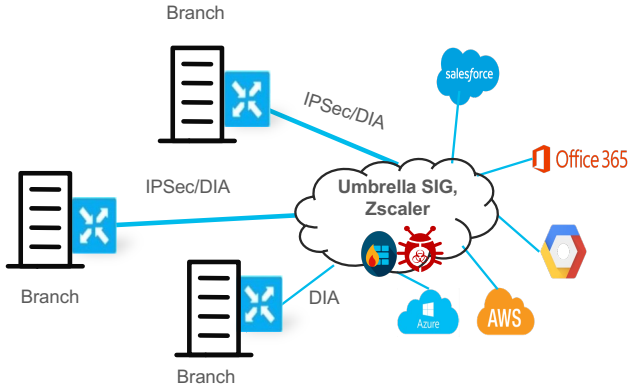

Multicloud access from the SD-WAN branch

Direct Internet Access designs leveraging local branch security

DIA with Security Appliance at each branch

Pro: User experience improved with branch security

Con: Increased CAPEX complexity with UTM appliance

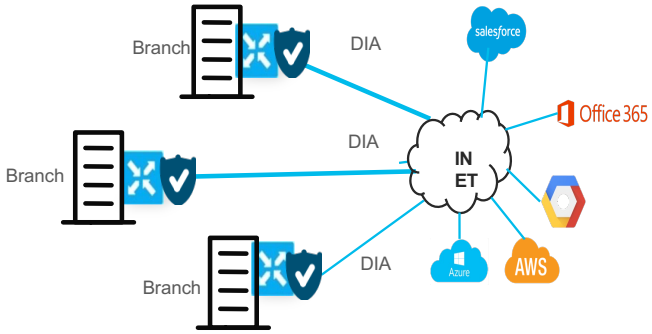

DIA with Branch Router embedded security features

Pro: User experience improved with full branch security stack

Con: Security feature availability dependent on platform

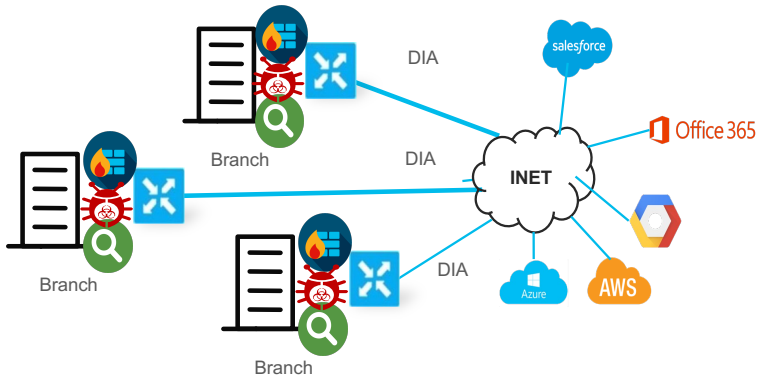

Multicloud access from the branch

Regional hub design options

Secure Agile Exchange (SAE)
Nexus 9K + CSP 5K with NSO orchestration

Pros:

- Improved user experience with removed latency of DC backhaul
- Maximum flexibility of networking and security services

Cons:

- Increased deployment complexity, requiring NSO orchestrator
- Separate SD-WAN and SAE management domains

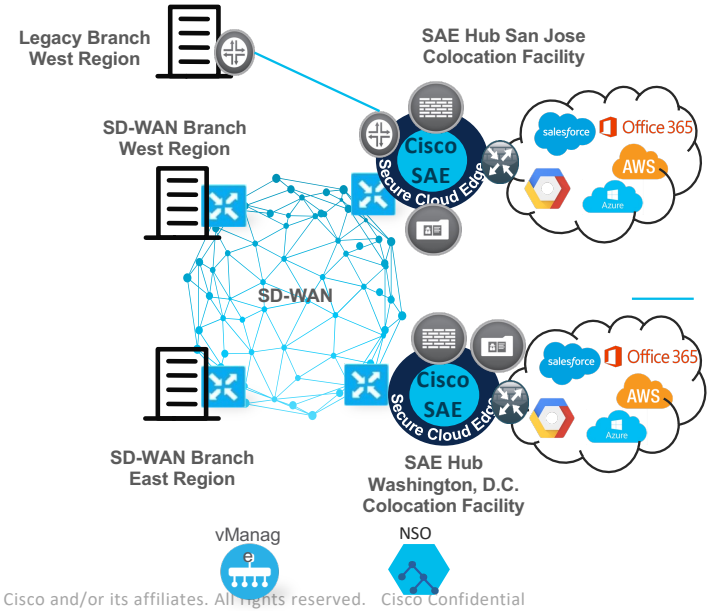

Cloud onRamp for Colocation (CoR4Colo)
Catalyst 9K + CSP 5K with vManage orchestration

Pros:

- Improved user experience with removed latency of DC backhaul
- Prescriptive solution, from equipment to cabling
- Solution integration with SD-WAN, with vManage orchestration

Con:

- Prescriptive solution less flexible for customization

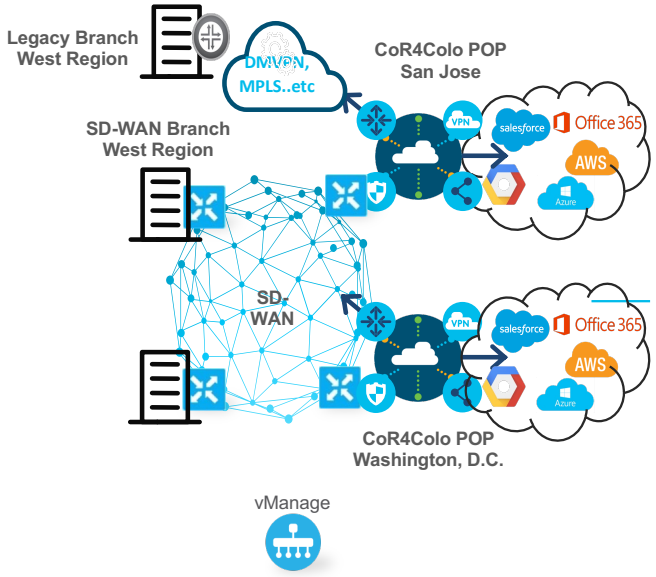

The bridge to possible