
cisco
You make possible

# WAN Architectures and Design Principles

Dave Fusik, Systems Architect

BRKRST-2041

cisco Live!
Barcelona | January 27–31, 2020

# Cisco Webex Teams

## Questions?

Use Cisco Webex Teams to chat with the speaker after the session

## How

1. Find this session in the Cisco Events Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

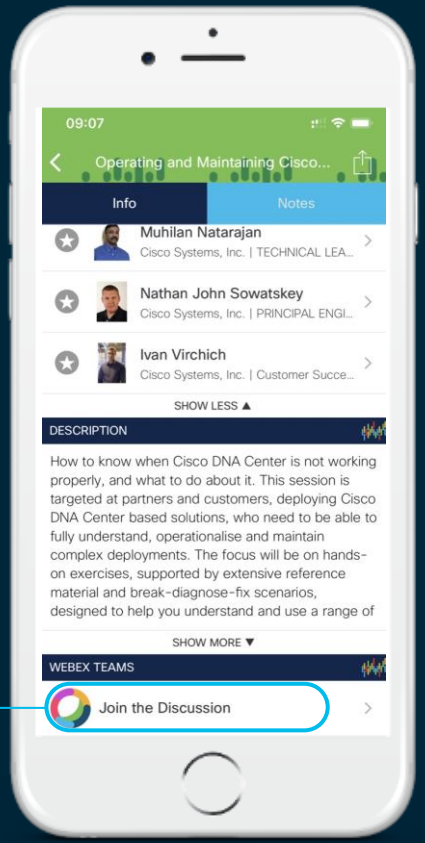

# Who is Dave Fusik?

22+ years at Cisco

3 years in TAC

CISCO CERTIFIED CCIE ROUTING AND SWITCHING #4768

5 years in CPOC

CISCO CERTIFIED CCIE SECURITY

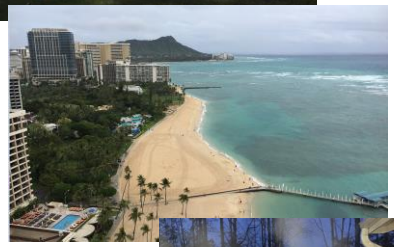

Systems Architect

14+ years in Sales

CISCO CERTIFIED CCDE DESIGN #2013::70

cisco Live!

# Agenda

- Introduction
- What is Wide Area Network (WAN) Architecture and Design?
- What to consider when designing a WAN
- Impacts of Evolving technology on WAN design
- WAN Designs moving Forward
- Conclusions

Main Message:

# *Foundational Design is key to WAN Architecture*

# The Challenge

- Allow the business to adopt changes rapidly and smoothly
- Quickly realize strategic advantage from new technologies
- Build a network that can gracefully adapt to an evolving technology landscape

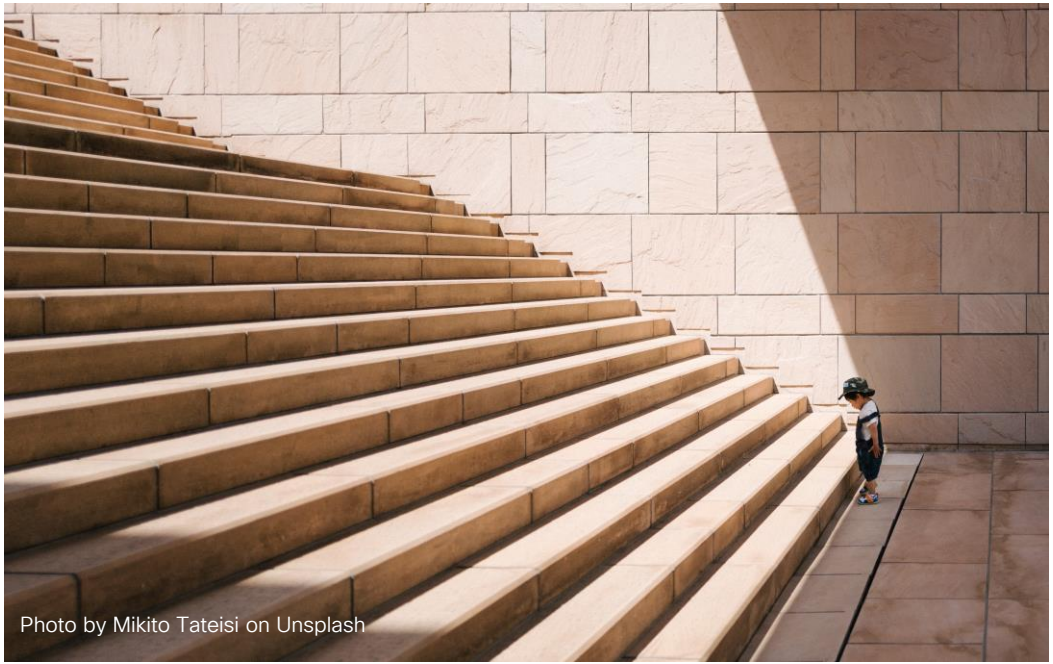


Cloud, SDN, IPv6, 5G, What's next?

# The WAN Technology Continuum

| Early Networking | Early-Mid 1990s | Mid 1990s-Late 2000s | Today |
|---|---|---|---|
| Flat/Bridged<br>Experimental Networks | Multiprotocol<br>Business Enabling | Large Scale<br>Mission Critical | Global Scale<br>IP Ubiquity<br>Cloud Connected |
| **Architectural Lessons** | **Architectural Lessons** | **Architectural Lessons** | **Planning** |
| Protocols required for Scale & Restoration | Route First, Bridge only if Must | Redundancy<br>Build to Scale | ? |

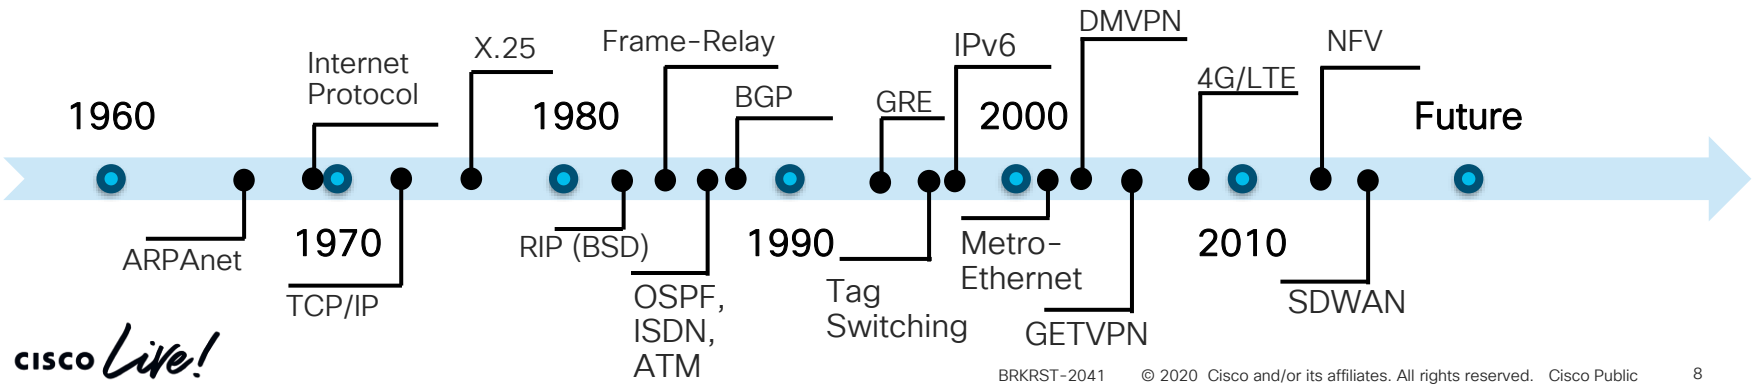

# What is WAN Architecture and Design?

# WAN Architecture and Design

- Network Architecture
  - The way network devices and services are structured or organized to serve and protect the connectivity needs of client devices
  - Depending on the place in the network, the requirements and the threats vary, so different frameworks are built
  - In the WAN, this means connecting users to applications, between LAN locations, sometimes over long distances
- Network Design
  - The process of translating business needs, budget, and operational constraints into a technological approach that addresses the architectural requirements
  - Includes documentation, such as implementation guides and topology diagrams
  - WAN designs need to minimize cost and enhance user experience when serving distributed applications to distributed users

# Architecture vs. Design

- Architecture looks toward strategy, structure and purpose
- Design drives toward practice and implementation
- Architecture goes nowhere without design
- Design may be too singularly focused without architecture

# Key Principles to WAN Design

***Simplicity*** can often be synonymous with elegance but must be paired with functional

***Modularity*** implies the use of building blocks that can be reused and fitted together to drive consistency

***Hierarchy*** creates vertical flow to horizontal expansion with natural points of aggregation

These are the tools to achieve ***Structure***

# Network Design Modularity

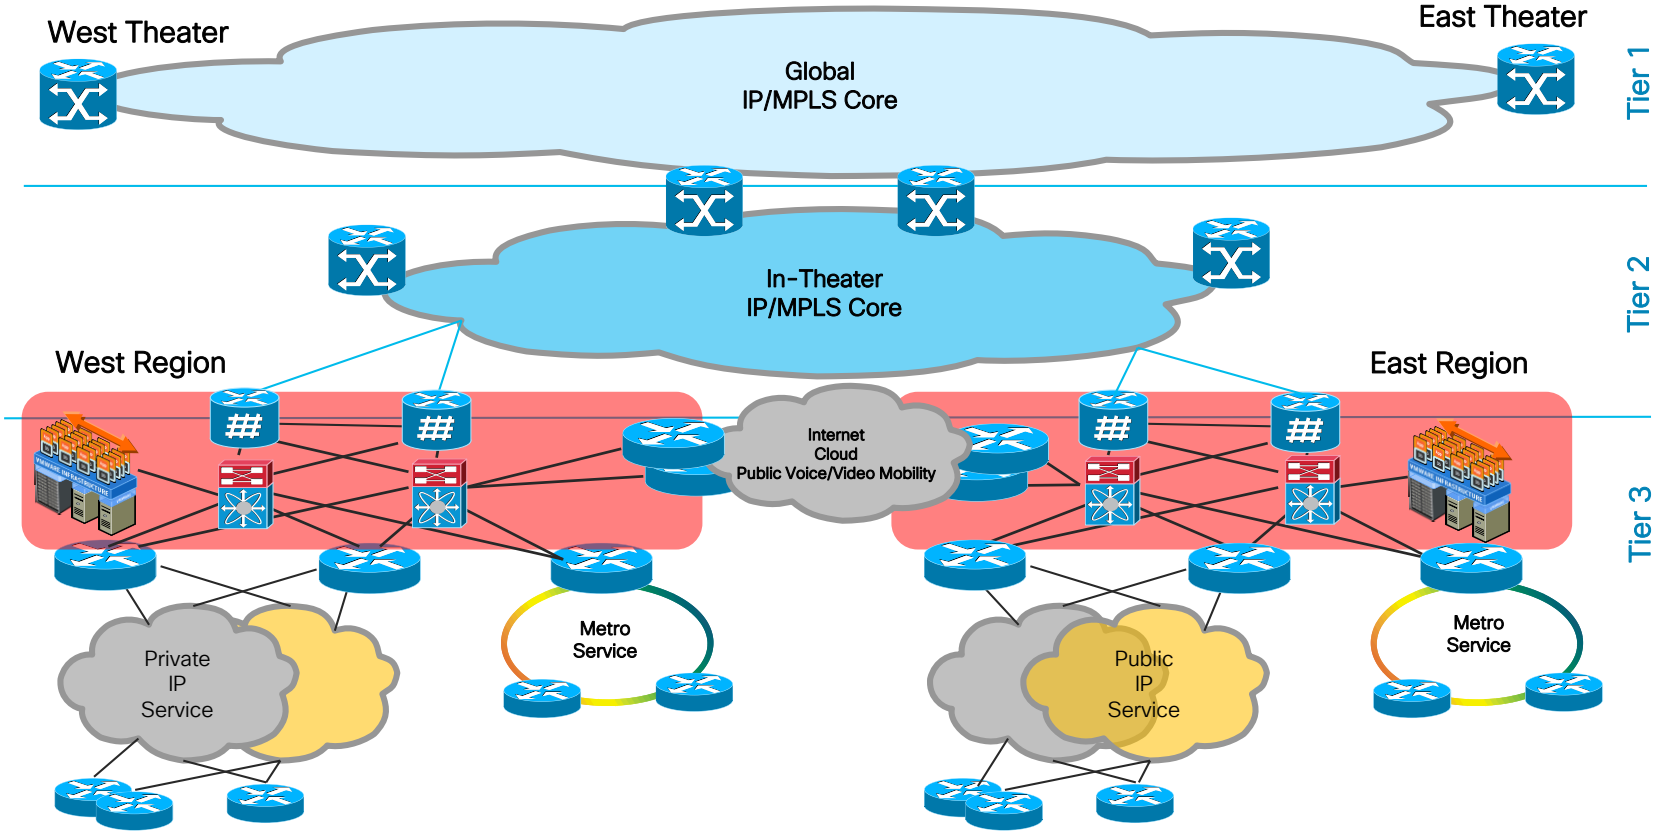

# Hierarchical Network Design

## Without a Rock Solid Foundation the Rest Doesn't Matter

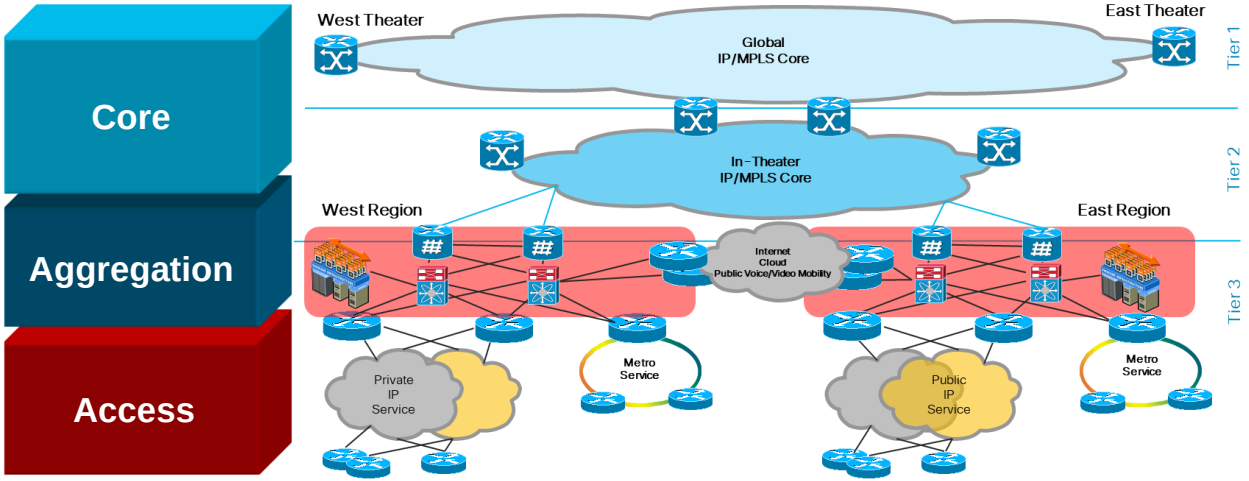


- Hierarchy—each layer has specific role
- Modular topology—building blocks
- Easy to grow, understand, and troubleshoot
- Creates small fault domains—clear demarcations and isolation
- Promotes load balancing and redundancy
- Promotes deterministic traffic patterns
- Incorporates balance of both Layer 2 and Layer 3 technology, leveraging the strength of both
- Utilizes Layer 3 routing for load balancing, fast convergence, scalability, and control

# Do I Need a Core Layer?

## It's Really a Question of Scale, Complexity, and Convergence

- No Core
- Fully-meshed distribution layers
- Physical cabling requirement
- Routing complexity

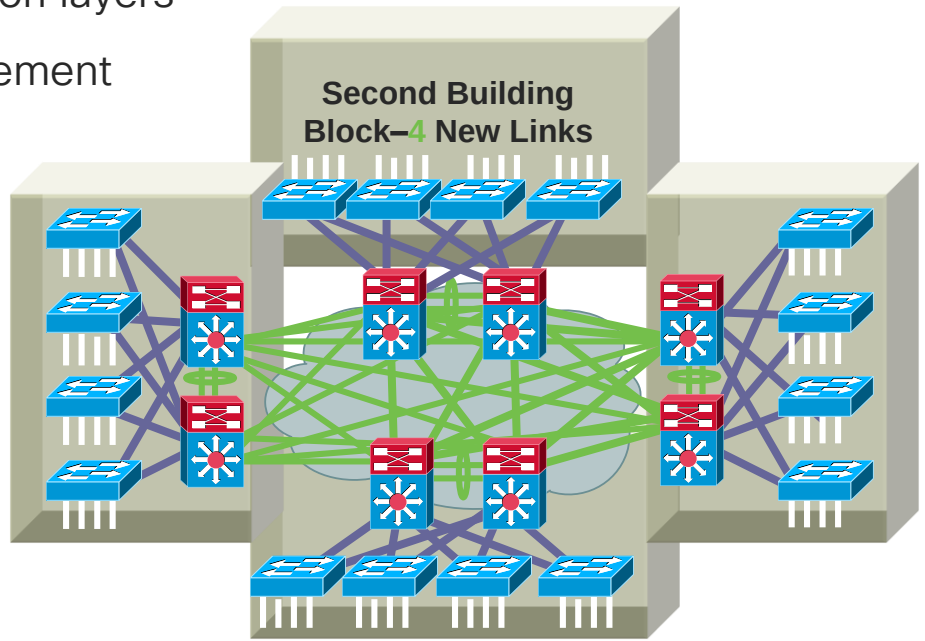

# What to consider when designing a WAN

# Business Requirements and Constraints

- Business Environment
  - Market transitions
  - Competitive pressures
  - Project goals
  - Mergers and acquisitions

- Workforce Productivity
  - User experience
  - Access to resources
  - Employee satisfaction

- Costs
  - OPEX and CAPEX
  - Lifecycle and ROI
  - IT Capabilities
  - Opportunity costs

- Compliance and Policy
  - Government and Industry Regulations
  - Security mandates
  - Reputation and perception

# Technical Requirements and Constraints

- Application requirements
  - Bandwidth, Latency, Jitter
  - Connectivity and Protocols
    - L2 or L3, IPv4 or IPv6, Multicast,
- Policy and Compliance
  - Security
  - Segmentation
  - Encryption
- Performance and Resiliency
  - Quality-of-Experience
  - High Availability
    - Convergence and Recovery
    - Device quantities and capabilities
- Existing Network Infrastructure
  - Greenfield or Brownfield
  - Available documentation
  - Current designs and technologies

# Physical Requirements and Constraints

- Company Locations
  - 10's, 100's, or 1000's of sites
  - Where in the world
  - Site diversity
    - retail store, campus, large manufacturing plant, etc.
- Topology Implications
  - Single or dual connected
  - Geographical dispersity
    - Local, Regional, Global
  - Network role
    - Data Center, Colo Facility, Branch, Remote access, Public/Guest access
- Operational requirements
  - Access to resources
  - Transport options
  - Available power
  - Size and quantity of equipment
- Risks associated with the Business and Technical requirements

# When Considering High Availability

- Assess system criticality
- How to measure availability
- Eliminate single points of failure
- Failure detection and recovery
- Environmental conditions

# Defining Availability

- System Availability: a ratio of the expected uptime to the experienced downtime over a period of time of the same duration
- Branch WAN High Availability: Between 99.99%(4) and 99.999%(5)
- Ultra High Availability: Between 99.9999%(6) and 99.999999%(8)

| Availability | Downtime / Year |
|---|---|
| 98.000000% | 7.3 Days |
| 99.000000% | 3.65 Days |
| 99.500000% | 1.825 Days |
| 99.900000% | 8.76 Hrs |
| 99.990000% | 52.56 Min |
| 99.999000% | 5.256 Min |
| 99.999900% | 31.536 Sec |
| 99.999990% | 3.1536 Sec |
| 99.999999% | .31536 Sec |

Branch WAN HA Targets

Ultra HA Targets

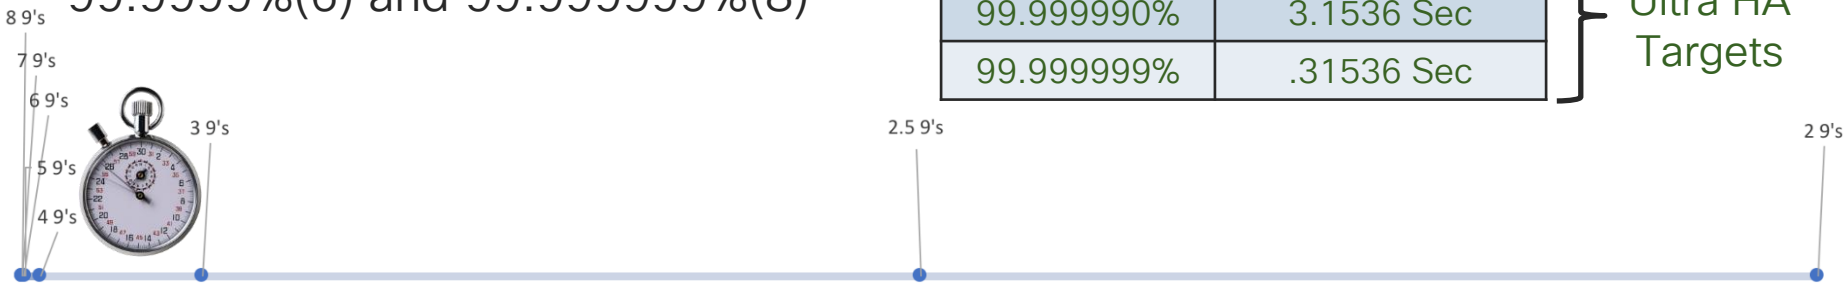

# Where Can Outages Occur?

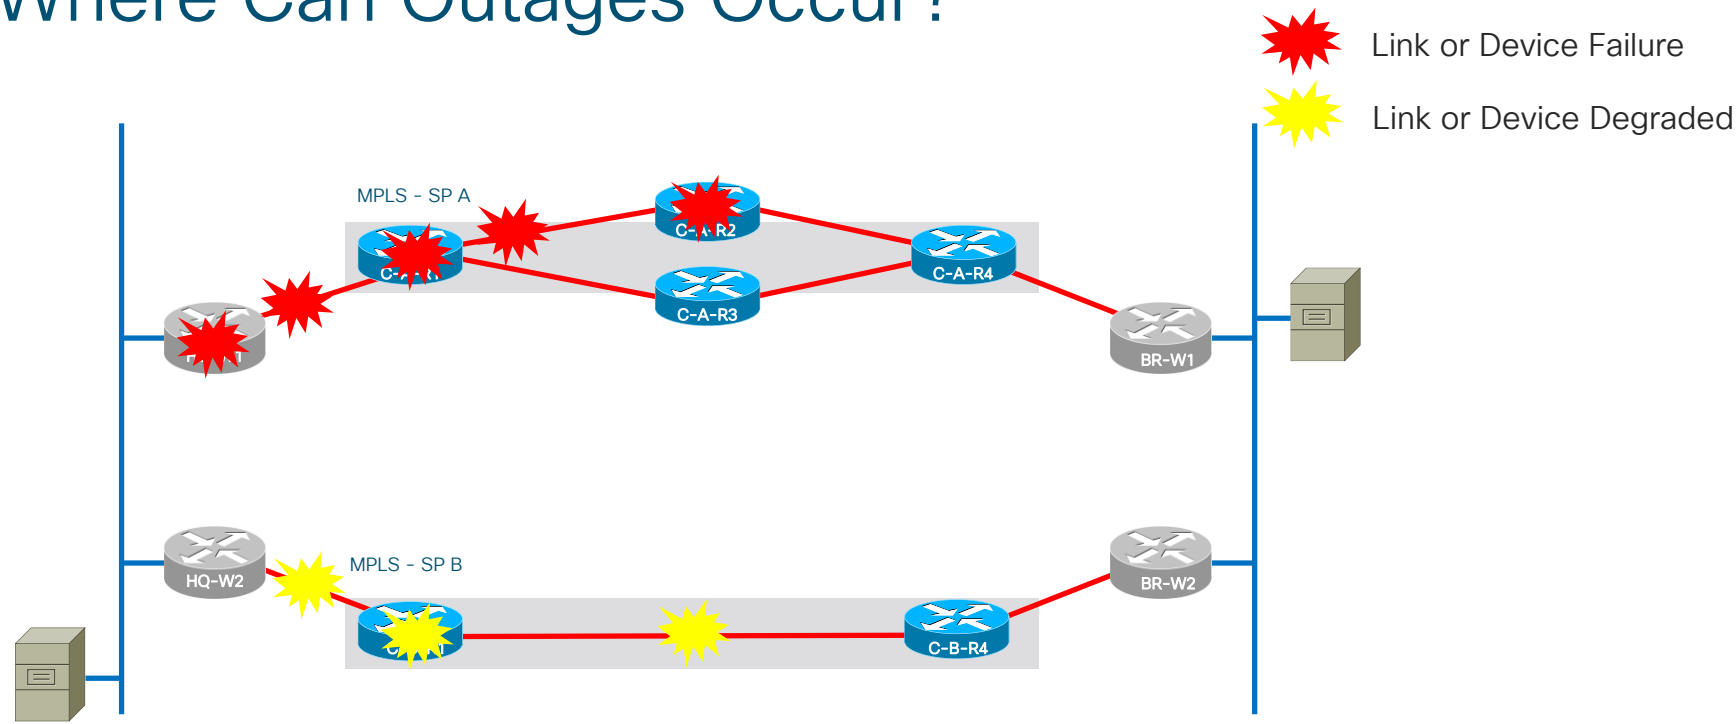


- How does outage manifest?
- How quickly can network detect?
- How long is bidirectional reconvergence?

# Building Highly Available WANs
## Redundancy and Path Diversity Matter

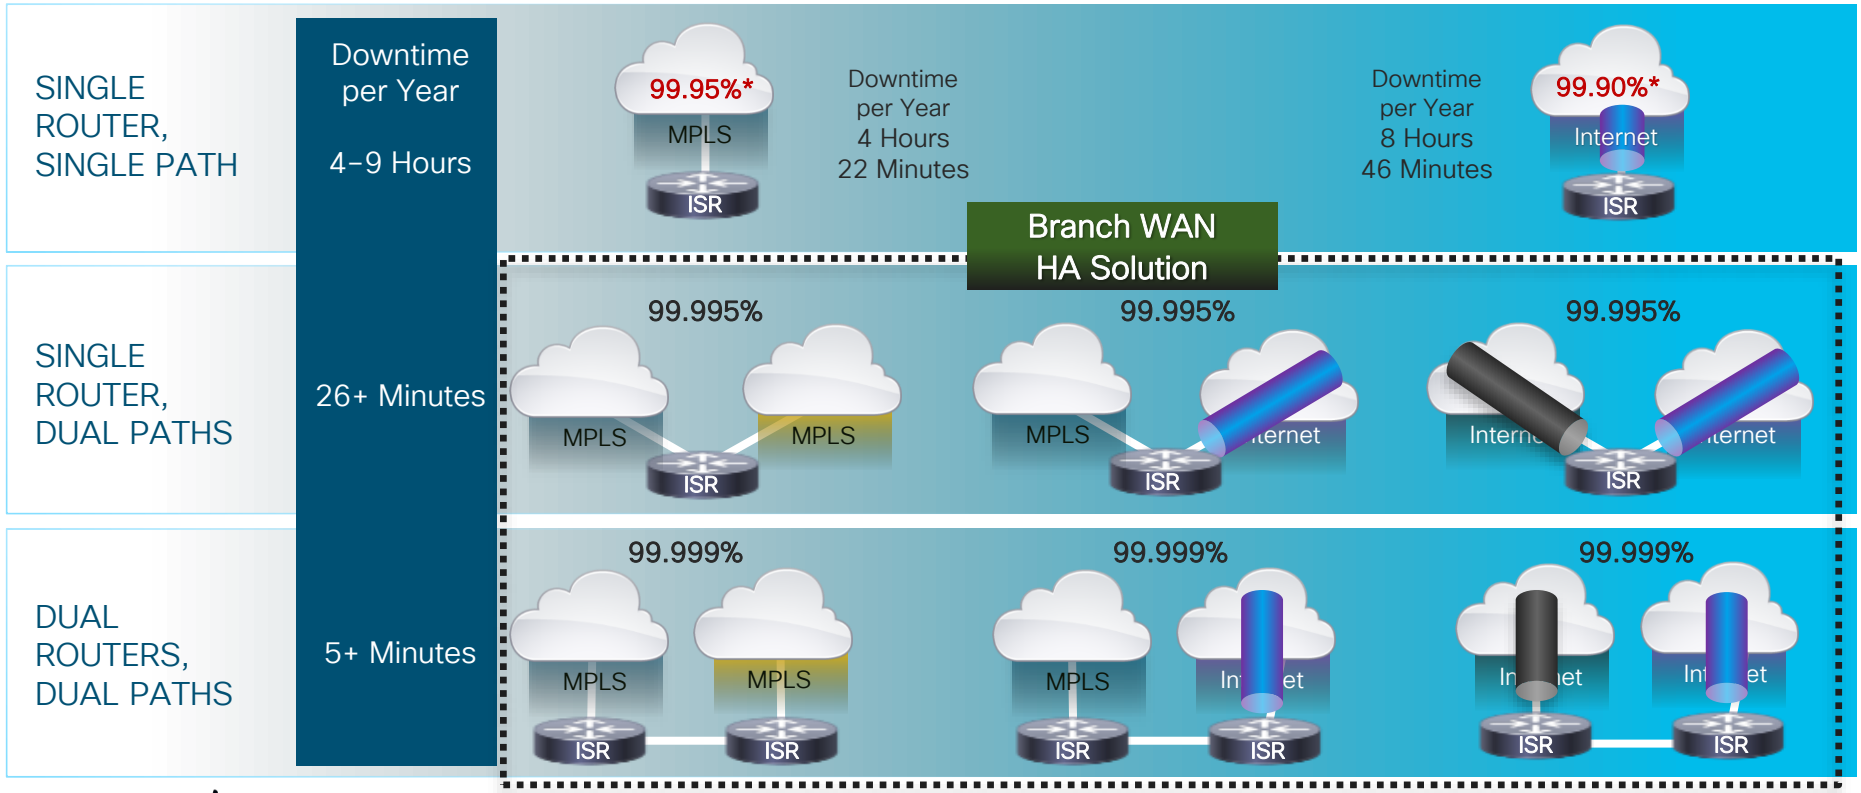


* Typical MPLS and Business Grade Broadband Availability SLAs and Downtime per Year, calculated with Cisco AS DAAP tool.

# Redundancy vs. Convergence Time
## More Is Not Always Better

- In principle, redundancy is easy
- Any system with more parallel paths through the system will fail less often
- The problem is a network isn't really a single system but a group of interacting systems
- Increasing parallel paths increases routing complexity, therefore increasing convergence times

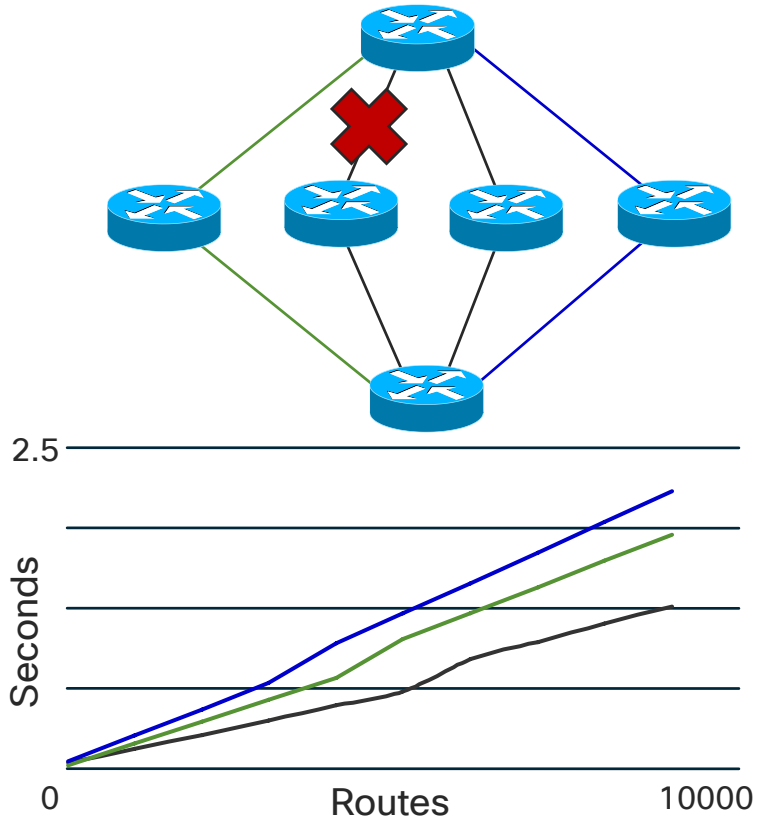

# Current and Evolving Technologies that impact WAN design

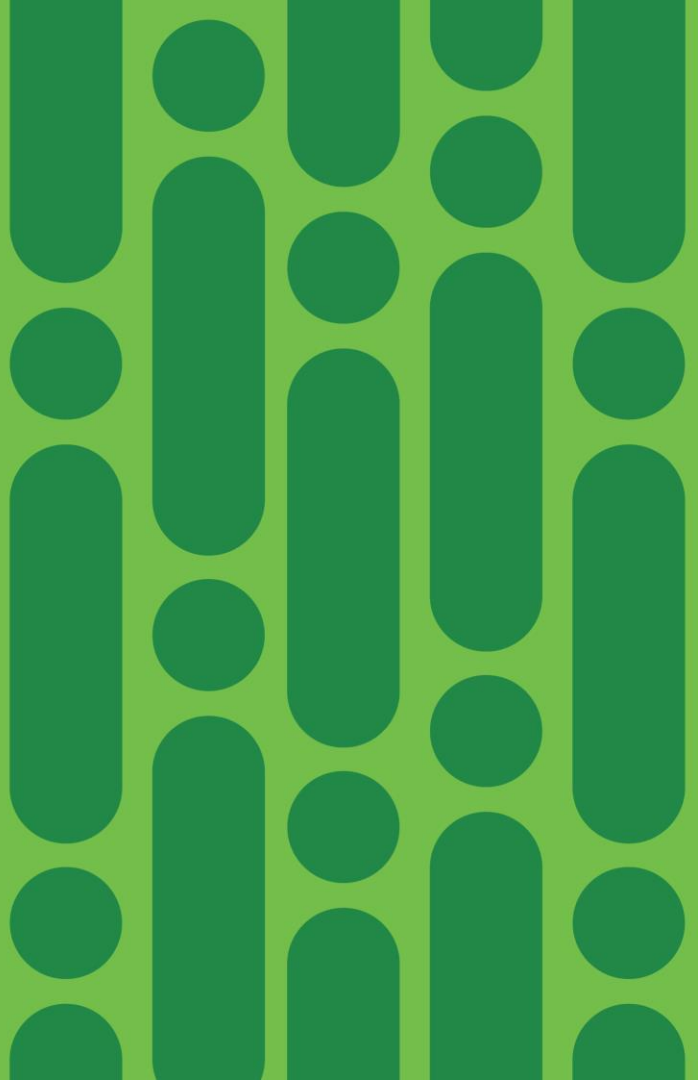

# WAN Locations and Devices

- Organization sites
  - Headquarters Campus
  - Branch Office
  - Retail store
  - Factory, etc.
- Remote Access
  - Mobile workers
  - Home office
- Cloud
  - Private Data Center
  - Public IaaS
  - SaaS
  - Colocation Facility

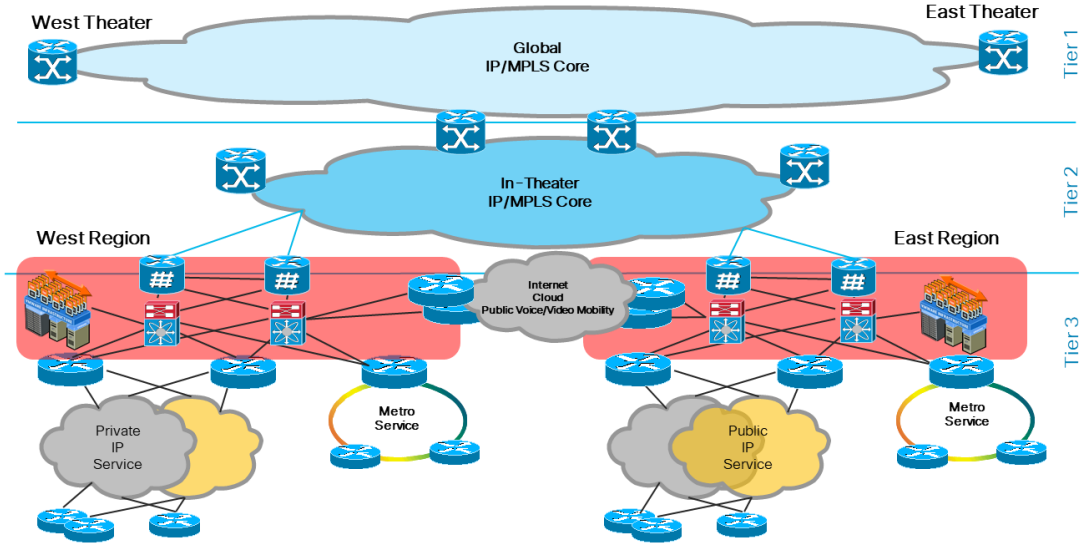


- Physical devices
  - Router/CPE
  - Firewall
  - Multi-purpose compute
  - Client devices
- Virtualized Network Functions
  - Virtual router
  - Virtual Firewall
  - etc…

# Cisco Enterprise Routing Portfolio

## Branch

### ISR 900

- Fixed and fanless
- IOS Classic based

### ISR 1000

- Integrated wired and wireless access
- PoE/PoE+

### ISR 4000

- WAN and voice module flexibility
- Compute with UCS E
- Integrated Security stack
- WAN Optimization

### vEdge 100

- 4G LTE & Wireless

### vEdge 1000 & 2000

- Fixed/Pluggable Module

## Aggregation

### ASR 1000

- Hardware and software redundancy
- High-performance service with hardware assist
- **Fixed Chassis**

### vEdge 5000

- Modular
- RPS

## Virtual and Cloud

### Cisco ENCS

- Service chaining virtual functions
- Options for WAN connectivity
- Open for 3rd party services & apps

### ISRv CSR 1000V

### vEdge Cloud

- Cisco DNA virtualization
- Extend enterprise routing, security & management to cloud

SD-WAN

# Cisco Cloud Services Router (CSR) 1000V

## Cisco IOS XE Software in a virtual network function form-factor

### Software
Same IOS XE software as the ASR1000 and ISR4000

### Infrastructure Agnostic
Runs on x86 platforms

**Supported Hypervisors:**
VMware ESXi, RHEL Linux KVM, Suse Linux KVM, Citrix Xen, Microsoft Hyper-V, Cisco NFVIS and CSP5000

**Supported Cloud Platforms:**
Amazon Web Services, Microsoft Azure, Google Cloud Platform

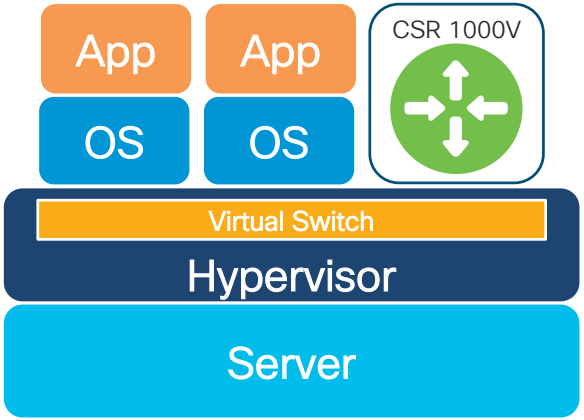


### Performance Elasticity
Available licenses range from 10 Mbps to 10 Gbps

CPU footprint ranges from 1vCPU to 8vCPU

### Programmability
NetConf/Yang, RESTConf, Guest Shell and SSH/Telnet

### License Options
Term based 1 year, 3 year or 5 year

Enterprise-class networking with rapid deployment and flexibility

# Cisco vEdge Cloud Router

## Cisco vEdge Software in a virtual network function form-factor

### Software
Same software as the physical vEdge router platforms

### Infrastructure Agnostic
Runs on x86 platforms

**Supported Hypervisors:**
VMware ESXi, RHEL Linux KVM, Suse Linux KVM, Citrix Xen, Microsoft Hyper-V, Cisco NFVIS and CSP5000

**Supported Cloud Platforms:**
Amazon Web Services, Microsoft Azure, Google Cloud Platform

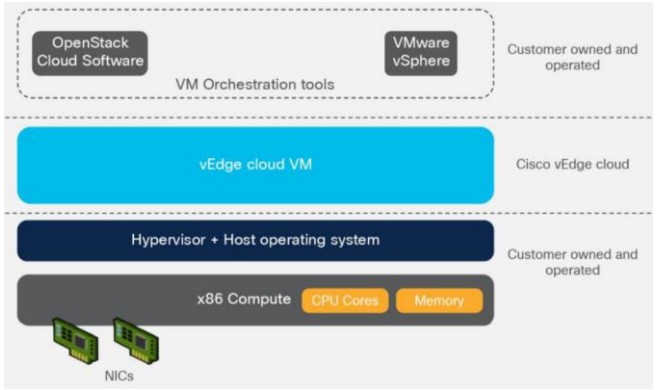


### Performance
Available licenses range from 10 Mbps to 100 Mbps

CPU footprint minimum 2vCPUs

### Positioning
Extends SD-WAN Overlay into Cloud Environments

### License Options
Term based 1 year, 3 year or 5 year

Enterprise-class networking with rapid deployment and flexibility

# Platform Built for Enterprise NFV

## ENCS 5000 Series for the Branch

Best of Routing & Compute

Complete Virtualized Services

Open for Third Party Services and Apps

**Enterprise Network Compute System**

ENCS 5100 Series


ENCS 5400 Series

8 Integrated LAN Ports with Optional POE

USB 3.0 Storage

Hardware Acceleration for VM Traffic

2 Onboard Gigabit Ethernet ports with SFP

Network Interface Module for LTE & legacy WAN

2 HDD or SSD RAID 0 & 1

cisco Live!

# What is Cisco SD-Branch?

## Network services in minutes, on any platform

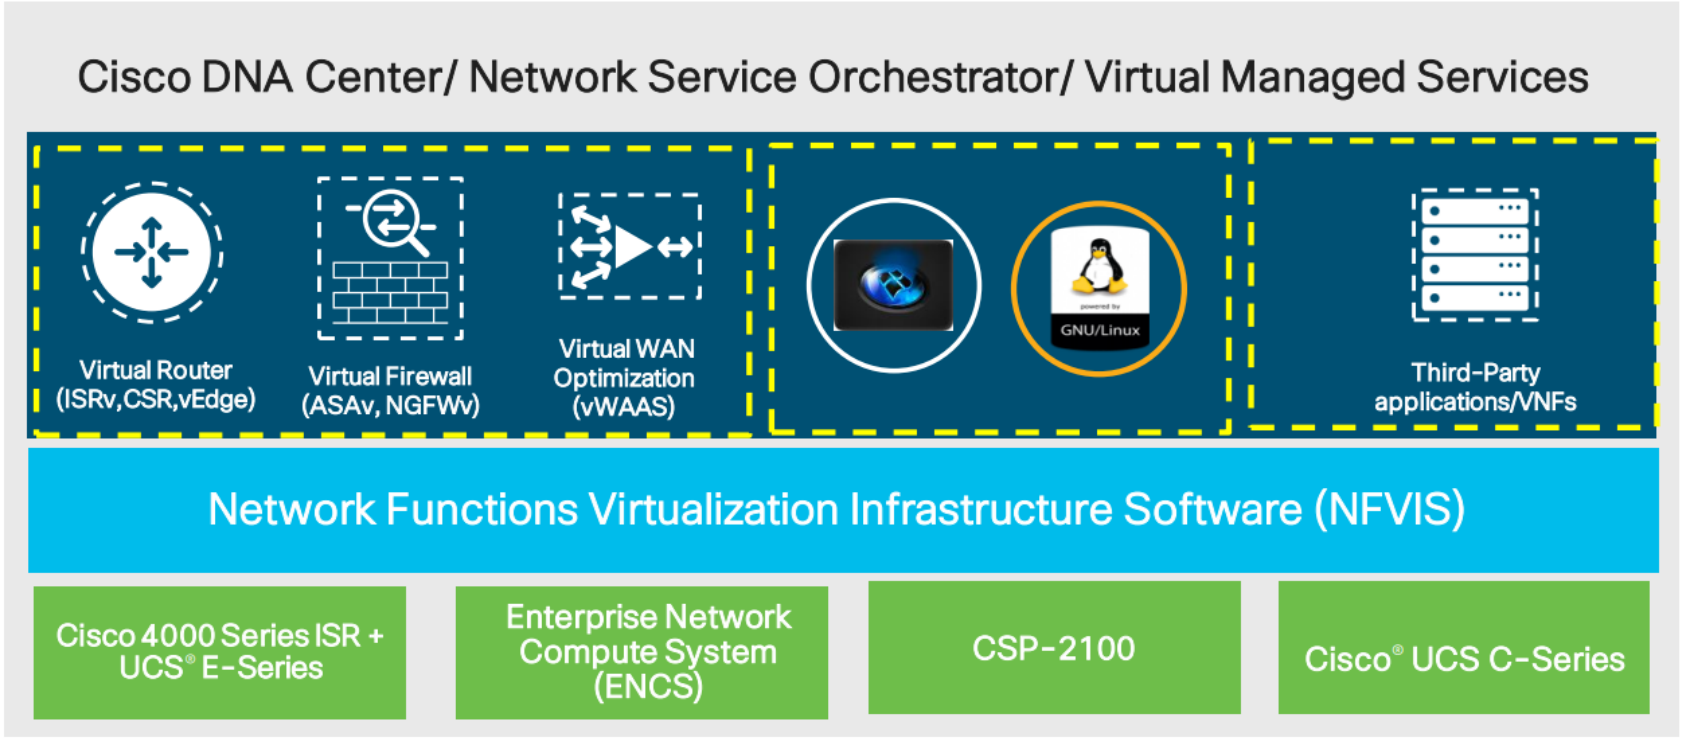

# What changes with Cisco SD-Branch?

## Before

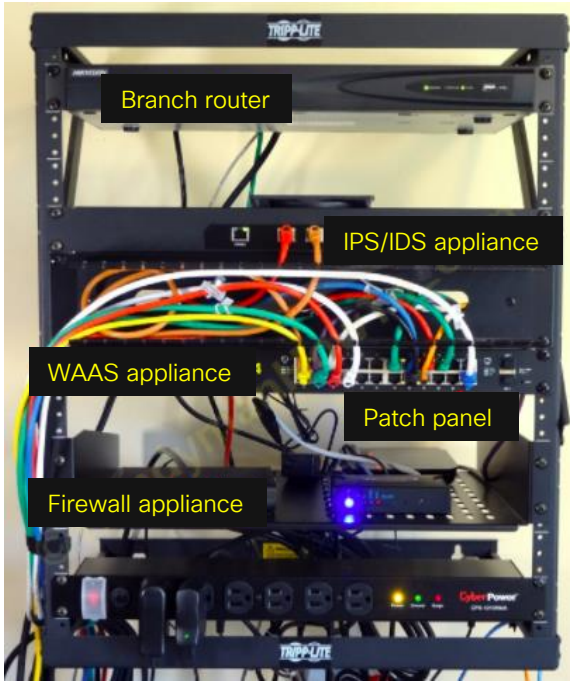


## After

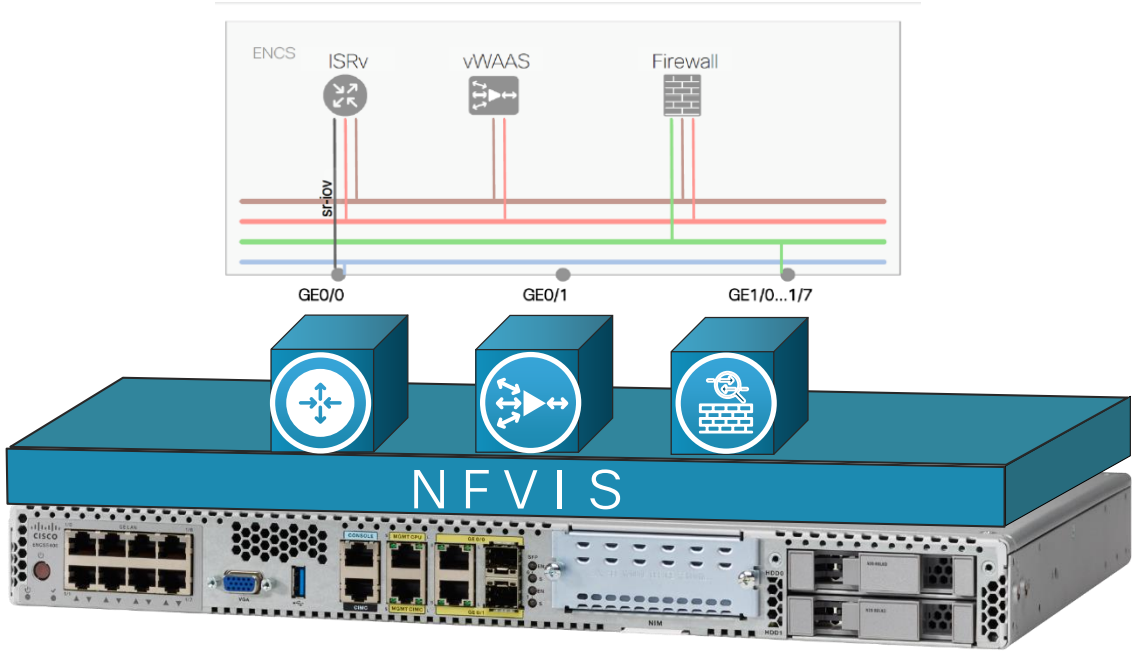


A single x86 compute platform housing multiple VNFs

# ISRv and CSR 1000V

**Integrated Services Router - Virtual**

- Packaged for NFVIS
- Branch-Specific Features
- Branch-Specific Pricing
- Look-and-feel of an ISR 4000
- Not available separately

**Cloud Services Router**

- Cloud and VDC Deployments
- Aggregation Use-Cases
- Flexible Pricing & Packaging
- Virtual ASR 1000 Series
- Available on multiple platforms

# WAN Connection and Transport Technologies

- Dark Fiber
  - Highest flexibility, control, and security but only point-to-point connectivity
  - Most costly unless owned by the organization
- MPLS
  - Widely available service with flexible bandwidth options
  - Provider manages complex WAN routing with QoS SLAs
  - Offers simplicity with global scale if the organization can afford it
- Metro Ethernet
  - Layer 2 Ethernet connectivity service between up to hundreds of locations within a specific geographic region
  - Organization manages its own routing and QoS policies but may offer higher bandwidth at less cost than MPLS

- Broadband
  - Lower cost, high bandwidth Internet connectivity
  - Organization manages a secure overlay VPN between sites but has no control over latency or QoS
  - Available as wired (DSL, Cable) or wireless (3G/4G/5G or satellite)
- Legacy T1
  - Last resort option but available anywhere
  - Cost comparable to Metro Ethernet but only 1.5Mbps bandwidth
  - Point-to-point layer 2 connectivity and requires non-Ethernet type port on router

# MPLS VPN Models

CE = Customer Edge router

PE = Provider Edge router

**Technology Options**

## MPLS Layer-2 VPNs

### Point-to-Point Layer-2 VPNs

- CE connected to PE via L2 connection (Eth, FR, ATM, etc.)
- CE-CE L2 p2p connectivity
- CE-CE routing
- No SP involvement

### Multi-Point Layer-2 VPNs

- CE connected to PE via Ethernet connection
- CE-CE L2 (Eth) mp connectivity
- CE-CE routing
- No SP involvement

## MPLS Layer-3 VPNs

- CE connected to PE via IP-based connection (over any layer-2 type)
  - –Static routing
  - –PE-CE routing protocol;

    eBGP, OSPF, IS-IS
- CE has peering relationship with PE
- PEs participate in customer routing
- PEs maintain customer-specific routing tables

cisco Live!

# Broadband Internet

- Widely available in wired or wireless
- Wired is generally an Ethernet handoff
- High bandwidth to the Internet so creates security vulnerability that must be managed
- Provides access to Public Cloud services such as IaaS and SaaS
- Does not support QoS or Multicast
- Overlay IP encapsulation with IPSec creates a secure VPN tunnel between Enterprise locations
- No service guarantee for critical applications but offers a low cost backup or bandwidth augmentation option

# Types of Overlay Service

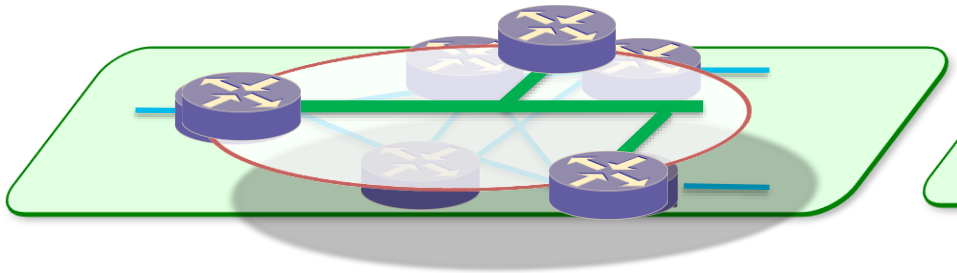

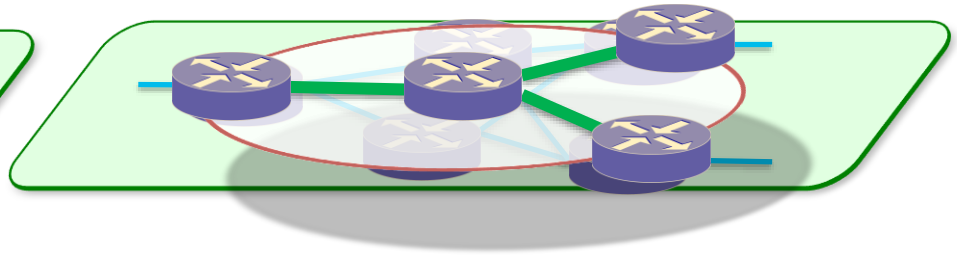

## Layer 2 Overlays

- Virtual Extensible LAN (VXLAN)
  - MAC-in-UDP encapsulation
  - 24-bit segment ID for up to 16M logical networks
- Other L2 overlay technologies
  - MPLS-over-GRE/mGRE, L2TPv3, OTV

## Layer 3 Overlays

- IPSec–Encapsulating Security Payload (ESP)
  - Strong encryption
  - IP Unicast only
- Generic Routing Encapsulation (GRE)
  - IP Unicast, Multicast, Broadcast
  - Multiprotocol support
- Other L3 overlay technologies
  - MPLS-over-GRE/mGRE, LISP



  

# GRE and IPSec Overlay Encapsulation Example

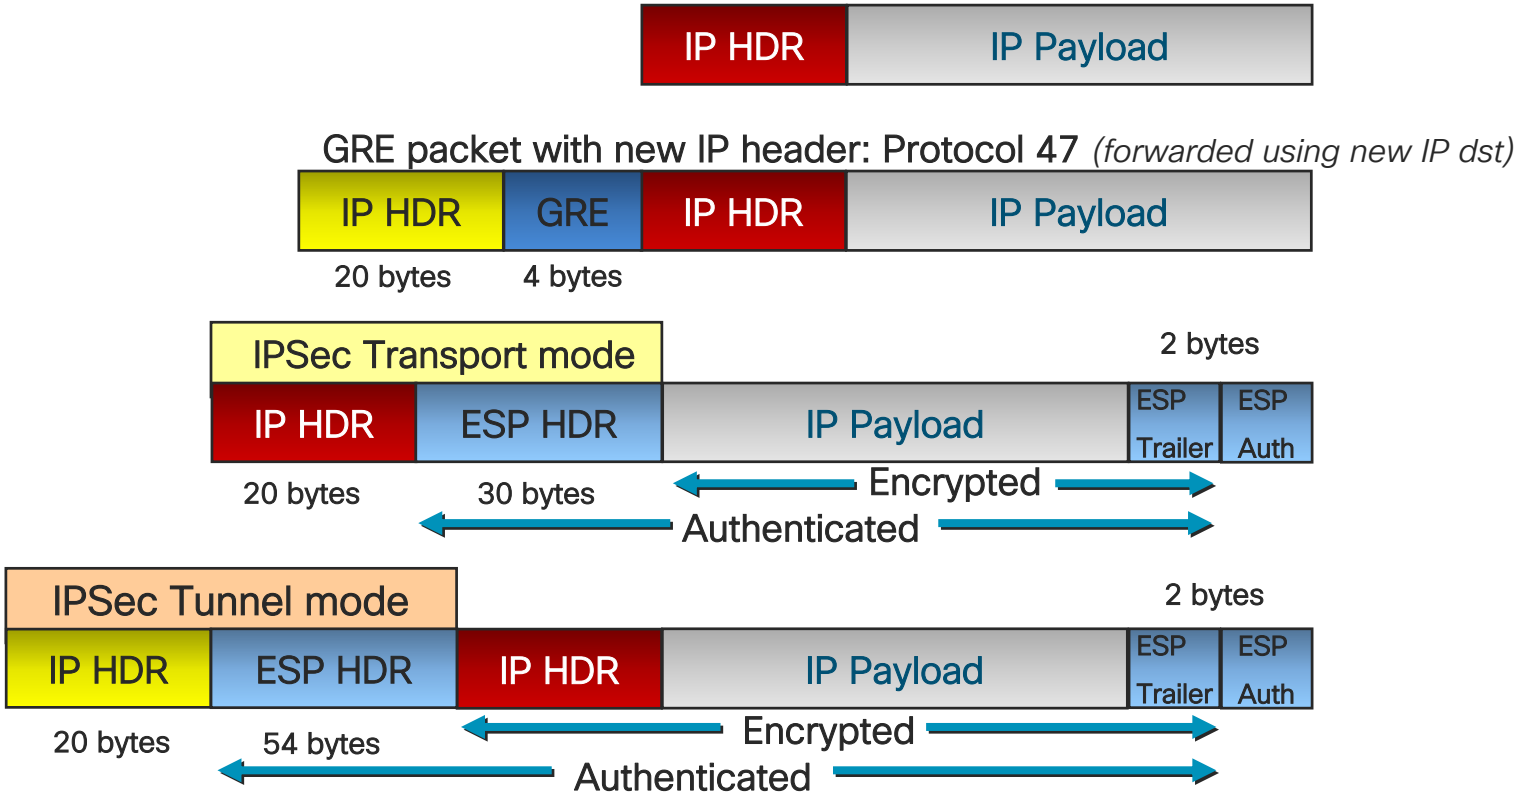

# Wide Area Network Design Trends

- Single Carrier Designs
  - Enterprise connects all sites to a single MPLS VPN carrier for L3 connectivity
    - Simple design with consistent features
    - Bound to single carrier for feature velocity
    - Vulnerable to MPLS cloud failure scenario

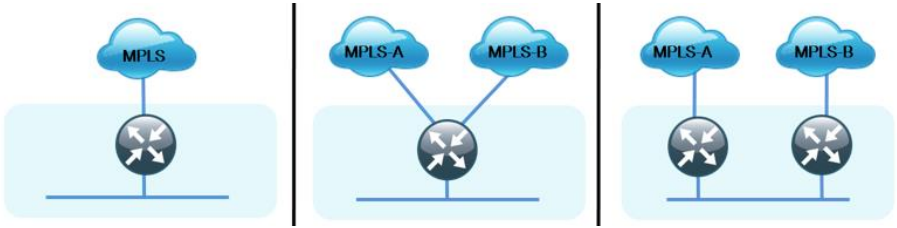


- Dual Carrier Designs
  - Enterprise single/dual connects sites into one/both MPLS VPN carriers
    - Protection against full MPLS cloud failure
    - Leverage for competitive services pricing
    - Complexity from service differences between carriers (QoS, BGP AS, etc.)
    - Must settle for least common denominator features

# Wide Area Network Design Trends (cont.)

- Hybrid and Overlay Designs
  - Tunneling/encryption enables transport agnostic design
    - + On-demand or permanent backup links
    - + Commodity broadband services offer lower cost, higher bandwidth
    - + Flexible overlay topology independent of physical underlay connectivity
    - – Two “layers” to support
    - – SLA over commodity transport
    - – Must consider potential for frag

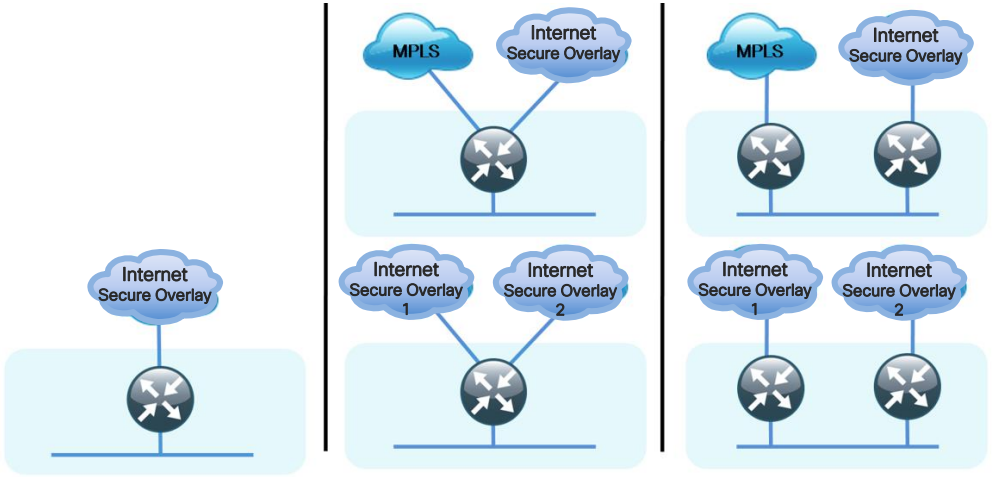

# Legacy IPsec VPN Technologies Comparison

| Features | DMVPN | FlexVPN | GET VPN |
|---|---|---|---|
| Infrastructure Network | ▪ Public or Private Transport<br>▪ Overlay Routing<br>▪ IPv4/IPv6 dual Stack | ▪ Public or Private Transport<br>▪ Overlay Routing | ▪ Private IP Transport<br>▪ Flat/Non-Overlay IP Routing |
| Network Style | ▪ Large Scale Hub and Spoke with dynamic Any-to-Any | ▪ Converged Site to Site and Remote Access | ▪ Any-to-Any; (Site-to-Site) |
| Failover Redundancy | ▪ Active/Active based on Dynamic Routing | ▪ Dynamic Routing or IKEv2 Route Distribution<br>▪ Server Clustering | ▪ Transport Routing<br>▪ COOP Based on GDOI |
| Scalability | ▪ Unlimited<br>▪ 3000+ Client/Server | ▪ Unlimited<br>▪ 3000+ Client/Server | ▪ 8000 GM total<br>▪ 4000 GM/KS |
| IP Multicast | ▪ Multicast replication at hub | ▪ Multicast replication at hub | ▪ Multicast replication in IP WAN network |
| QoS | ▪ Per Tunnel QoS, Hub to Spoke | ▪ Per SA QoS, Hub to Spoke<br>▪ Per SA QoS, Spoke to Spoke | ▪ Transport QoS |
| Policy Control | ▪ Locally Managed | ▪ Centralized Policy Management | ▪ Central or Local Management |
| Technology | ▪ Tunneled VPN<br>▪ Multi-Point GRE Tunnel<br>▪ IKEv1 & IKEv2 | ▪ Tunneled VPN<br>▪ Point to Point Tunnels<br>▪ IKEv2 Only | ▪ Tunnel-less VPN<br>▪ Group Protection<br>▪ IKEv1 & IKEv2 |

# Link Speeds Out-Pacing IP Encryption

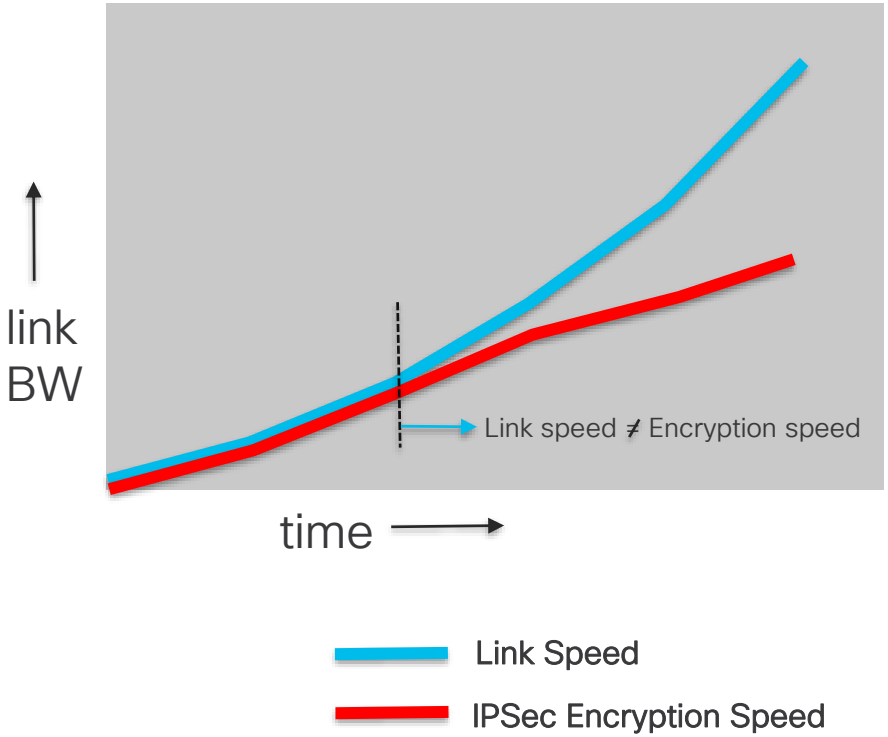


- Bandwidth application requirements out-pacing IP encryption capabilities
- Bi-directional and packet sizes further impact encryption performance
- IPSec engines dictate aggregate performance of the platform (much lower throughput)
- Cost per bit for IPSec much more expensive
- Encryption must align with link speed (100G+) to support next-generation applications

# What is MAC Security (MACsec)?

## Hop-by-Hop Encryption via IEEE 802.1AE

- Hop-by-Hop Encryption model
  - Packets are decrypted on ingress port
  - Packets are in the clear in the device
  - Packets are encrypted on egress port
- Supports 1/10G, 40G, 100G encryption speeds
- Data plane (IEEE 802.1AE) and control plane (IEEE 802.1x-Rev)
- Transparent to IPv4/v6, MPLS, multicast, routing
- Encryption aligns with Link PHY speed (Ethernet)

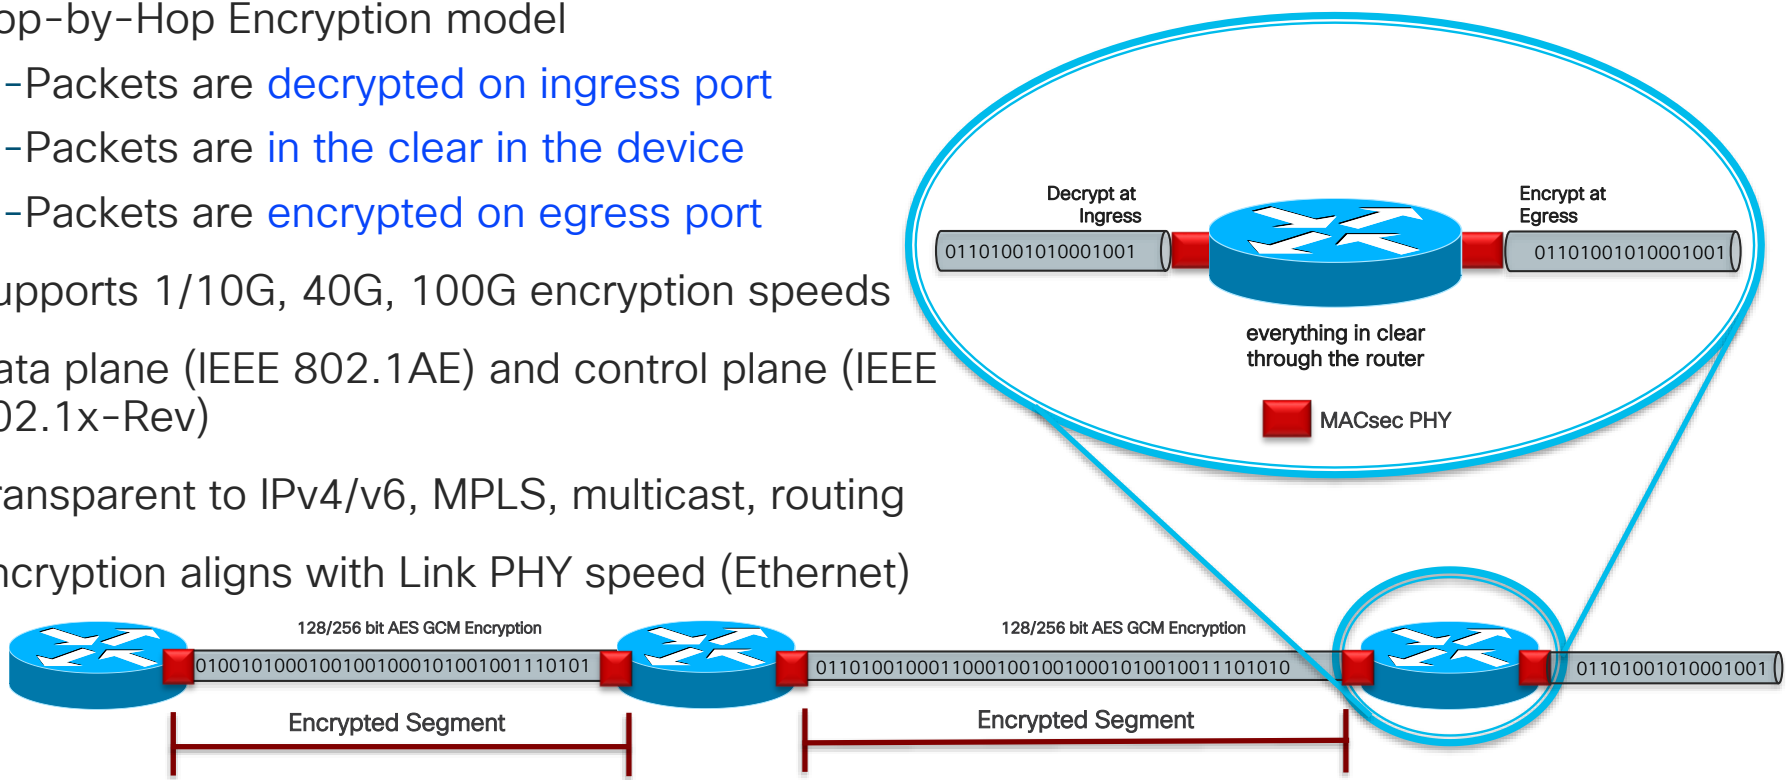

# What is “WAN MACsec?

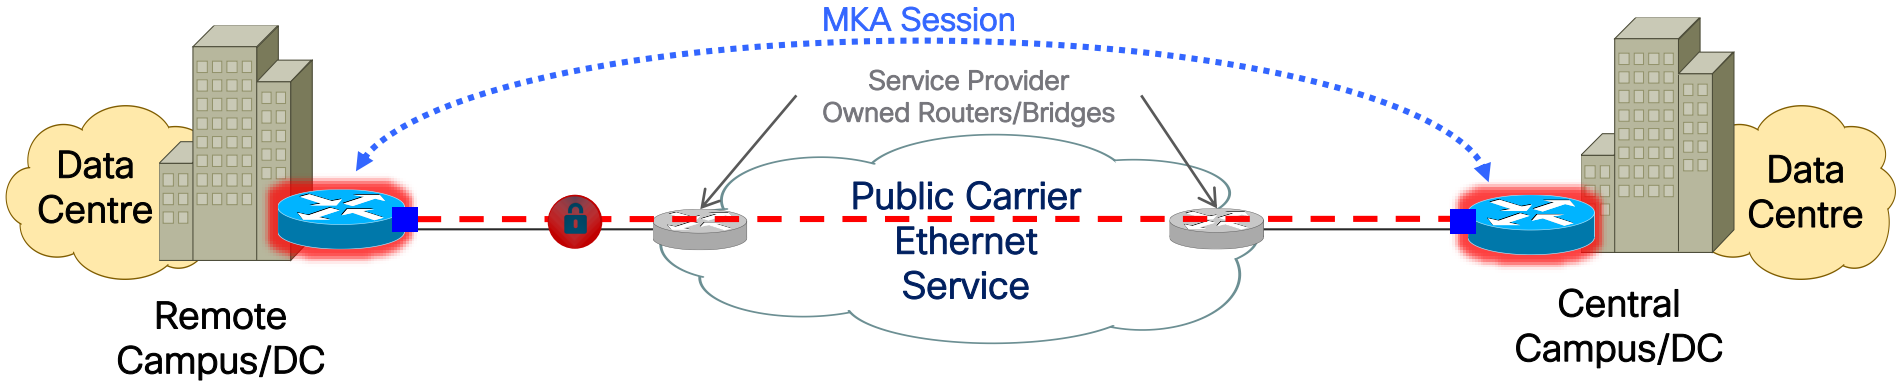


- Leverage MACsec over “public” standard Ethernet transport
- Optimize MACsec + WAN features to accommodate running over public Ethernet transport
- Target “line-rate” encryption for high-speed applications
  - Inter DC, MPLS WAN links, massive data projects
- Targets 100G, but support 1/10/40G as well

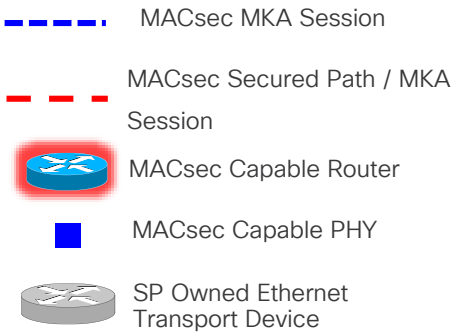

# WAN MACsec Use Cases

## Most Common Use Cases Leveraging WAN MACsec in the Enterprise

- 10GE → 100GE High speed Site to Site
  - Campus, WAN, DC→DC, Metro E
- Data Centre Interconnect
  - High Speed replication and storage transfers
- IP/MPLS core/edge links (PE-P, P-P, PE-PE)
  - MPLS labels, VPN, Segment Routing is transparent to MACsec encryption
  - No GRE, simple. Encryption = Link BW
- High Speed hub-and-spoke
  - Leverage low-cost/high-speed Metro E transport
  - Simple configuration, no GRE tunnels
- Hybrid Encryption Design Options
  - Ability to leverage BOTH MACsec and IPSec at various network points

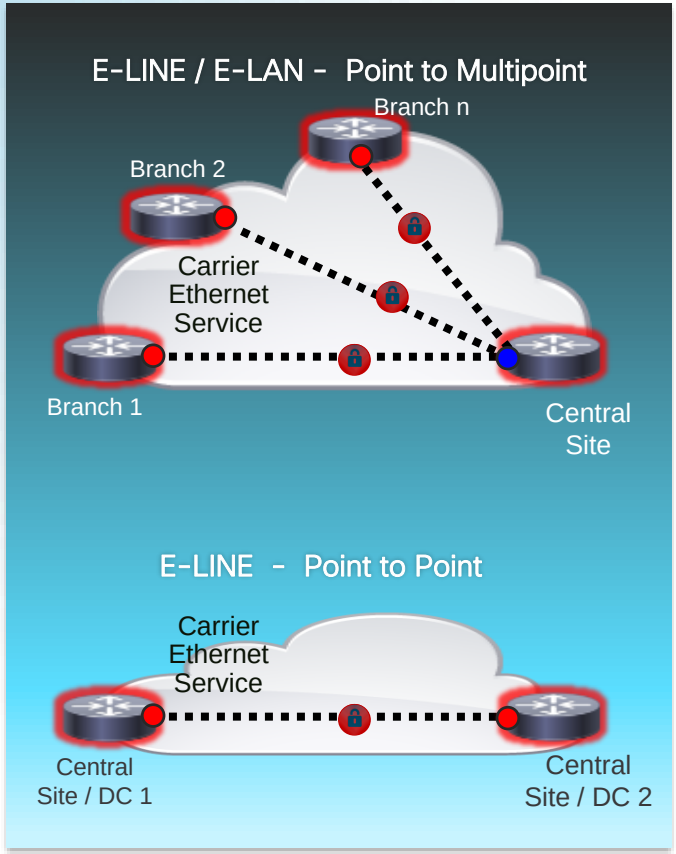

# What Is Enterprise L3 “Network” Segmentation?

- Giving One physical network the ability to support multiple L3 virtual networks
- End-user perspective does not change
- Maintains Hierarchy, Virtualizes devices, data paths, and services

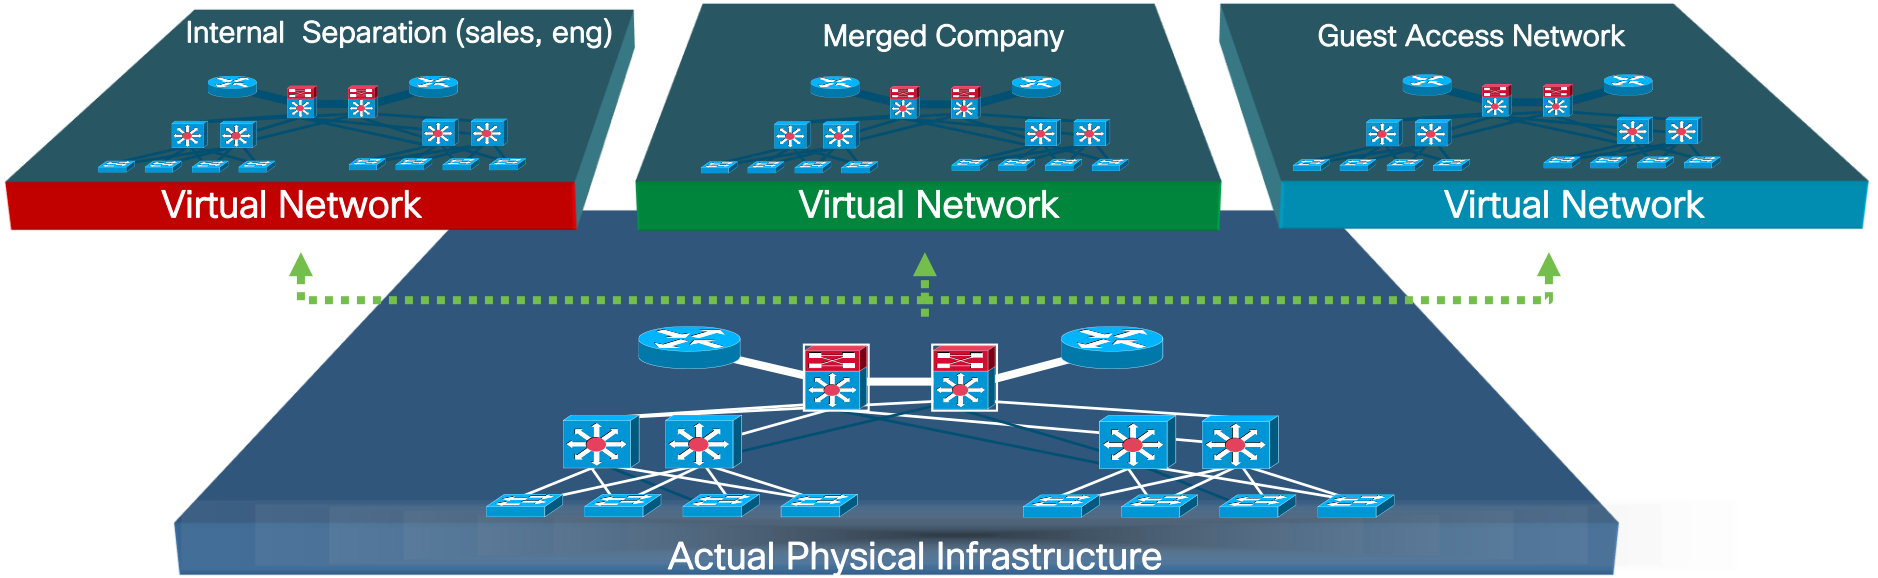

# Virtual Routing and Forwarding Instance – VRF

## Virtual Routing Table and Forwarding Separate to Customer Traffic

- Logical routing context within the same PE device
- Unique to a VPN
- Allows overlapping customer IP addresses
- Deployment use cases
  - Business VPN services
  - Network segmentation
  - Data Center access

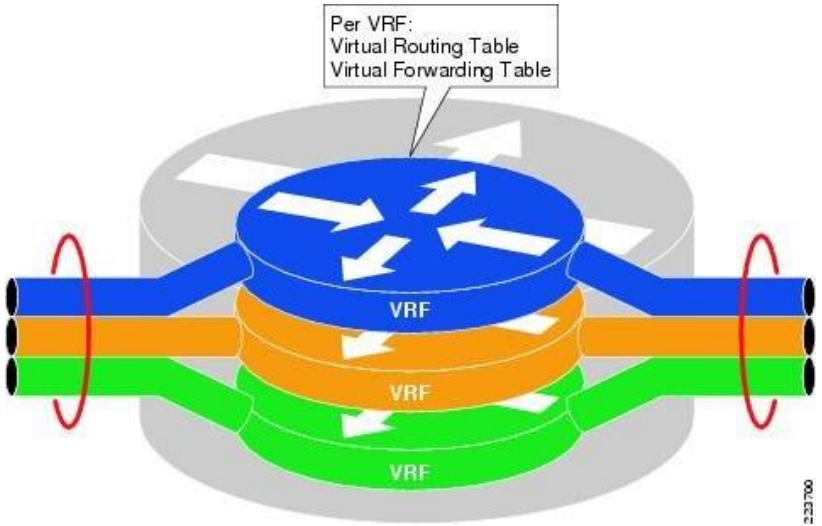

# Enterprise Network Segmentation over the WAN

## The Building Blocks – Example Technologies

### Device Partitioning

VLAN
VRF
VXLAN
Virtual Device Context (VDC)
Cloud Services Router (CSR)
IOS-XRv 64-bit

### WAN Segmentation Interconnect

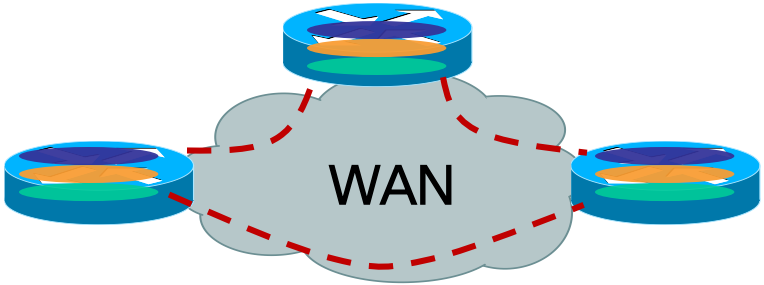


| L2 VPNs | L3 VPNs |
|---|---|
| EVPN/VxLAN | MPLS BGP L3 VPN |
| PW/VPLS | L3 VPN over IP |
| OTV | BGP EVPN (VXLAN, SR) |
| | VXLAN to MPLS Integration |

### Device Pooling

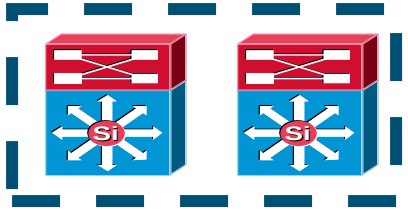


StackWise Virtual (SVL)
Virtual Port Channel (vPC)
Stackwise
Inter-Chassis Control Protocol (ICCP)
HSRP/GLBP

# Why L3 Network Segmentation?

## Key Drivers and Benefits

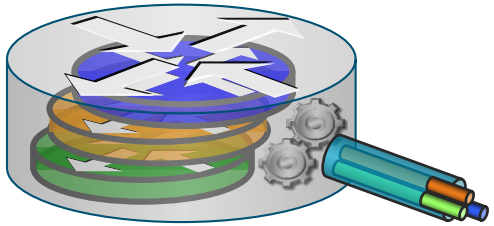

- Cost Reduction
  - Allowing a single physical network the ability to offer multiple virtual networks to tenants
- Simpler OAM
  - Reducing the physical network devices that need to be managed and monitored
- Security
  - Maintaining segmentation of the network for different departments over a single device/Campus/WAN
- Agility
  - Accelerates adding network segments (virtual) over same physical networks
- High Availability
  - Leverage segmentation through clustering devices that appear as one (vastly increased uptime)
- Data Center Applications
  - Offer per/multi-tenant segmentation from the DC into the WAN/campus/Branch and cloud
  - End-to-end Segmentation from-server-to-campus-to-WAN

# Why L3 Network Segmentation?

## Current and Evolving Use Cases

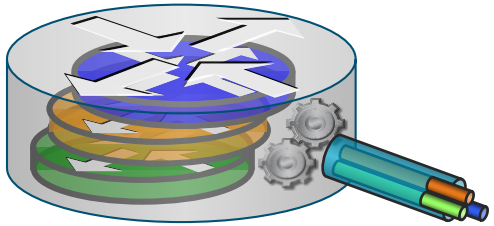

- Multi-Tenant Dwelling Separation
  - Airports – United, Delta, etc…
  - Government Facilities – agencies sharing single building/campus
  - Intra Organization segmentation – Sales, Engineering, HR, LoB
  - Company mergers – allowing slow migration for transition, overlapping addressing
  - IoT Device Isolation – segment from the user data (IP cameras, badge readers)
- Regulation requirements
  - Health Care – HIPPA
  - Financial and Transactional – Sarbanes-Oxley
  - PCI Compliance
- Security for Isolation
  - Key Fundamental element for Zero Trust Security framework
  - Quarantine Zone – Honey Pot, Steered Traffic as result of DDoS, Anomaly Enforcement
  - Mandates to logically separate varying levels of security (e.g. enclaves)
- Public Cloud and Key Component of Policy Construct
  - L3 segmentation for “per tenant” - GBP, and leveraged in Intent-based network policies

# WAN Segmentation Trends

Segmentation Domain

## MPLS Backbone

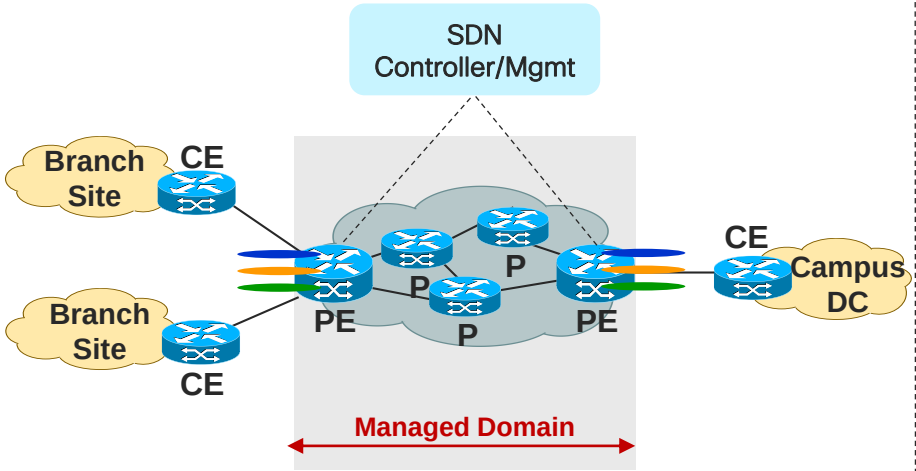


- Targets “Service Provider like” customers who need to control SLA’s, rapid service turn up times, tighter granular service options, end-to-end control, provisioning, and visibility
- Segment Routing, SR-TE, Centralized WAN controller

## Enterprise SD-WAN

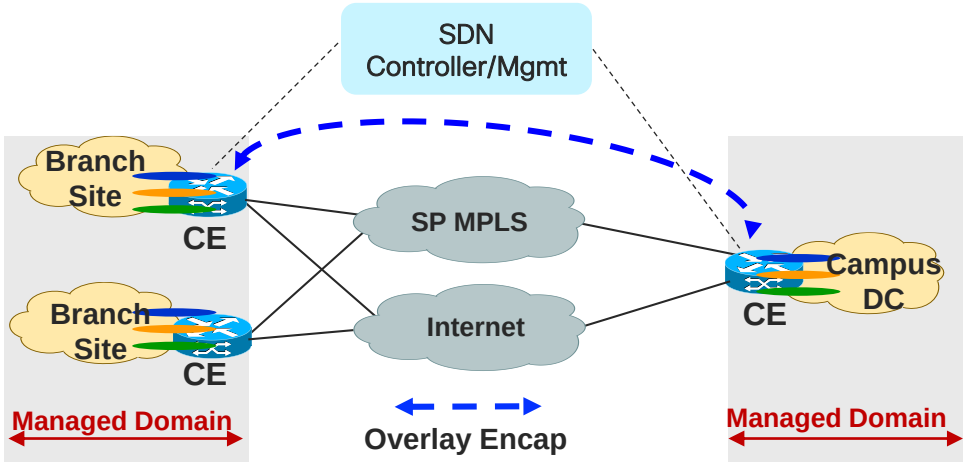


- Targets enterprise customers looking to consume secure WAN transport, with central mgmt., control, and application visibility
- Cisco SD-WAN, MPLS VPN over IP (central controller and/or open tools for automation)

# Quality of Service (QoS) Operations

## How Does It Work and Essential Elements

| Classification and Marking | Queuing and Dropping | Post-Queuing Operations |
|---|---|---|
| IDENTIFY & PRIORITIZE | MANAGE & SORT | PROCESS & SEND |

- **Classification and Marking:**
  - The first element to a QoS policy is to classify/identify the traffic that is to be treated differently. Following classification, marking tools can set an attribute of a frame or packet to a specific value.
- **Policing:**
  - Determine whether packets are conforming to administratively-defined traffic rates and take action accordingly. Such action could include marking, remarking or dropping a packet.
- **Scheduling (including Queuing and Dropping):**
  - Scheduling tools determine how a frame/packet exits a device. Queuing algorithms are activated only when a device is experiencing congestion and are deactivated when the congestion clears.

# Enabling QoS in the WAN
## Traffic Profiles and Requirements

### Voice

- Smooth
- Benign
- Drop sensitive
- Delay sensitive
- UDP priority

Bandwidth per call depends on codec, Sampling-Rate, and Layer 2 Media

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 1%
- Bandwidth (30-128Kbps)
- One-Way Requirements

### SD Video Conf

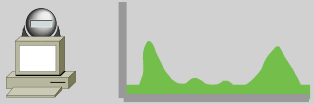

- Bursty
- Greedy
- Drop sensitive
- Delay sensitive
- UDP priority

SD/VC has the same requirements as VoIP, but traffic patterns and BW varies greatly

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 0.05%
- Bandwidth (1Mbps)
- One-Way Requirements

### Telepresence

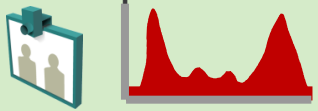

- Bursty
- Drop sensitive
- Delay sensitive
- Jitter sensitive
- UDP priority

HD/VC has tighter req's than VoIP for jitter and BW varies based on the resolutions

- Latency ≤ 200 ms
- Jitter ≤ 20 ms
- Loss ≤ 0.10%
- Bandwidth (5.5-16Mbps)
- One-Way Requirements

### Data

- Smooth/bursty
- Benign/greedy
- Drop insensitive
- Delay insensitive
- TCP retransmits

Traffic patterns for Data vary across applications

Data Classes:

- Mission-Critical Apps
- Transactional/Interactive Apps
- Bulk Data Apps
- Best Effort Apps (Default)

# Getting Started with QoS design

## Relevant

- Needed to support the core business objective
- Applications should be understood, marked and treated in accordance to best practice

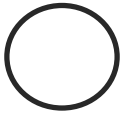

## Business as usual

- May or may not support business objectives directly
- The traffic can be grouped to qos class queues with proper marking or just tied to single qos class or default queues

## Not Important

- Consumer oriented traffic type
- Treated less than best class effort

# WAN Edge Bandwidth Allocation Models

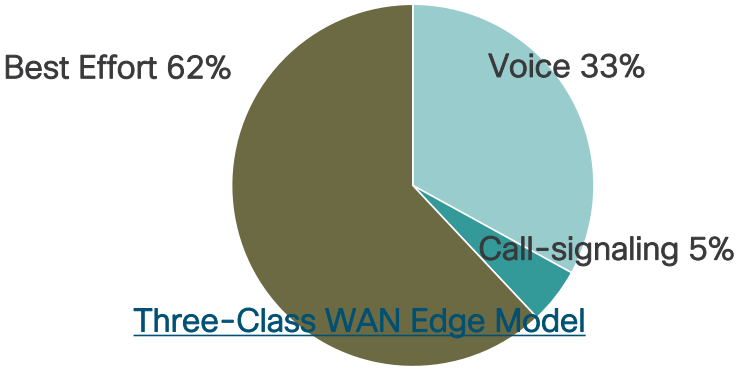


Three-Class WAN Edge Model

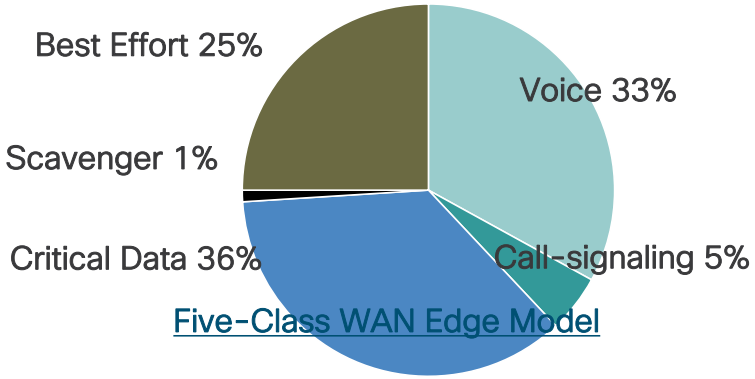


Five-Class WAN Edge Model

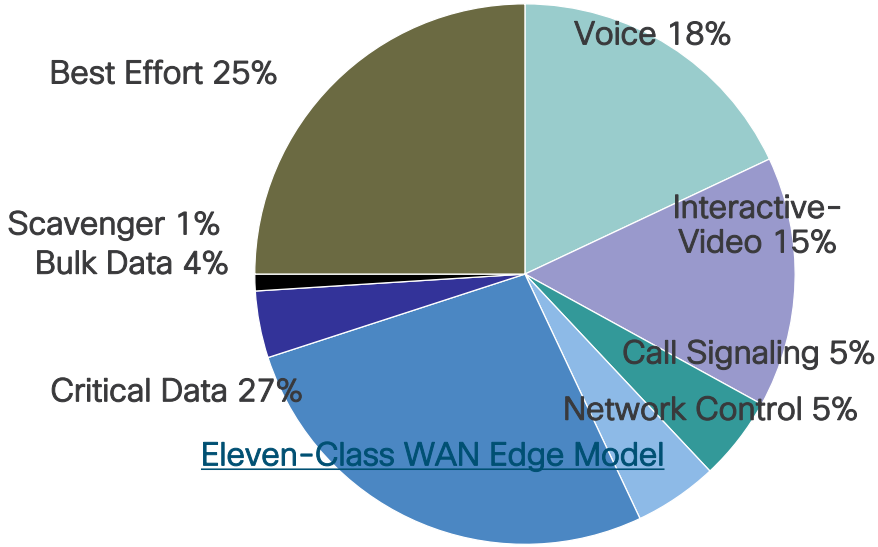


Eleven-Class WAN Edge Model

# QoS Tools and Techniques

## Classifying and Marking

- Network Based Application Recognition (NBAR2)
- Application Visibility and Control (AVC)
- Layer 2 or 3 marking of CoS/EXP or DSCP/IP precedence

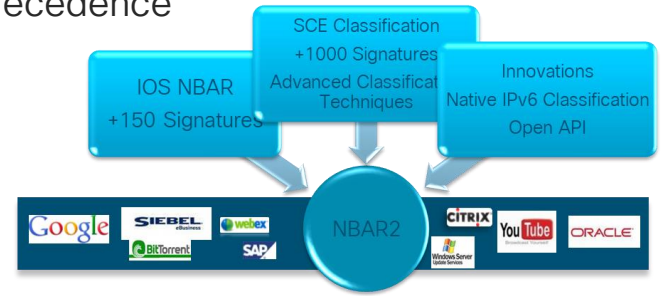


- New DPI engine provides Advanced Application Classification and Field Extraction Capabilities from Service classification engine

## Policing and Markdown

- Define traffic metering contracts
- Markdown out-of-contract flows
- Conform, Exceed, Violate actions

## Scheduling

- Re-order and selectively drop during congestion
- Class Based Weighted Fair Queuing (CBWFQ)
- Low Latency Queuing (LLQ) and Multi-LLQ

## Link-specific tools

- Traffic Shaping and Hierarchical QoS (HQoS)
- Compression
- Fragmentation and Interleaving

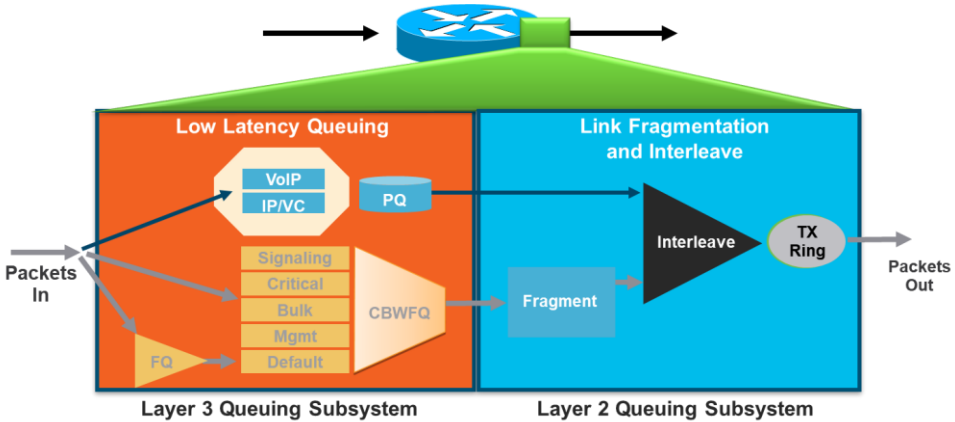

# GRE/IPSec QoS Consideration

## ToS Byte Preservation

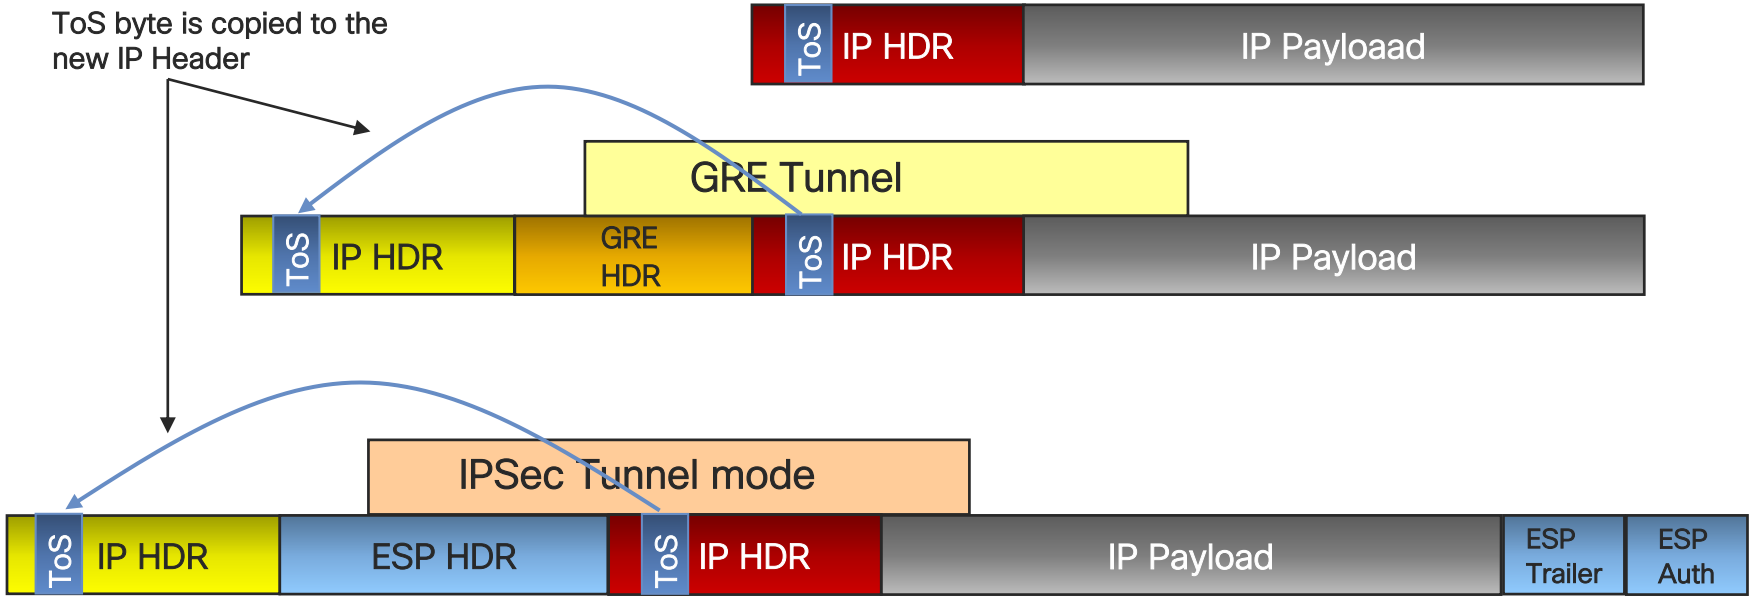

# QoS for IPv6

- The IPv6 implementation of DiffServ is identical to IPv4
- The same classifiers can be used to differentiate both IPv6 and IPv4 packets
  - Source IP address, destination IP address, IP Protocol field, source port number, and destination port number
  - IP precedence or DSCP values
  - TCP/IP header parameters, such as packet length
  - Source and destination MAC addresses
- The match precedence and match dscp commands filter IPv4 and IPv6 traffic

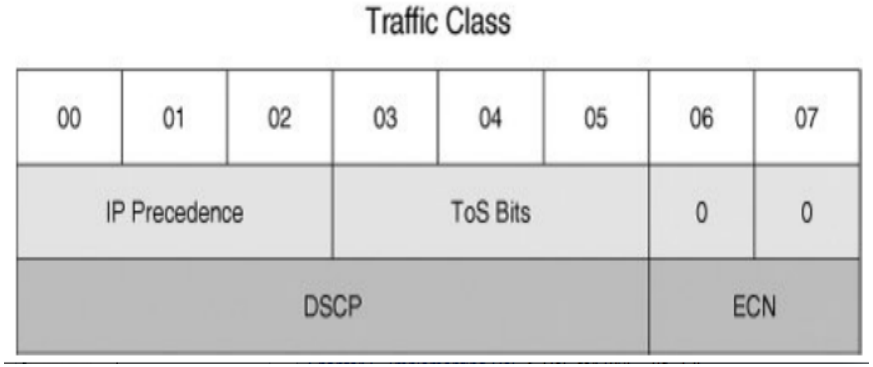


**To match packets on both IPv4 and IPv6 protocols:**

```
class-map match-all ipv6+ipv4forprec5
 match precedence 5
```

**To match packets for IPv6 protocols only:**

```
class-map match-all ipv6onlyprec5
 match protocol ipv6
 match precedence 5
```

# What Are the QoS Implications of MPLS VPNs?

Bottom Line:

- Enterprises must Co-manage QoS with Their MPLS VPN Service Providers
- Their Policies must be both consistent and complementary

# IP Multiservice VPN Service Providers
## Service-Level Agreements

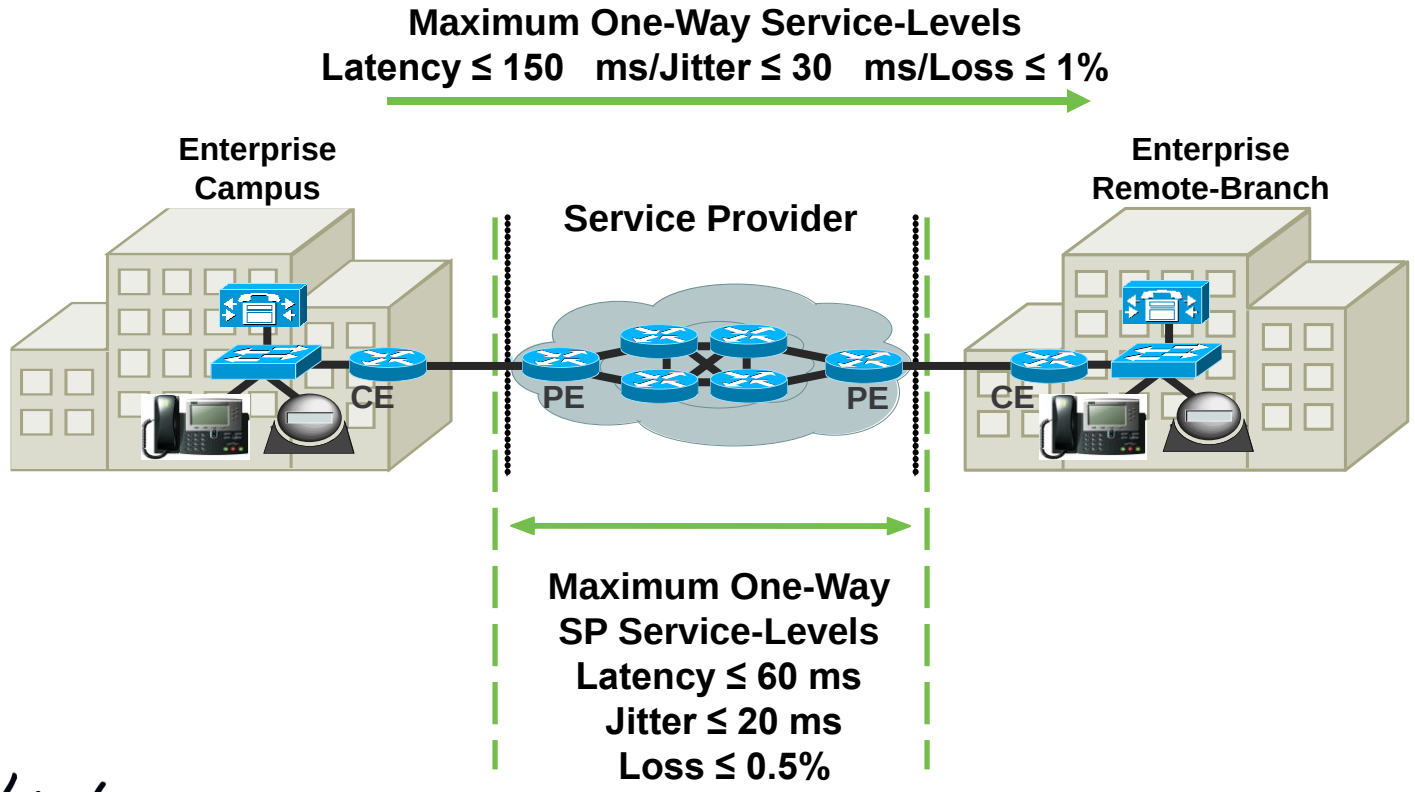

# Enterprise-to-Service Provider Mapping
## Five-Class Provider-Edge Model Remarking Diagram

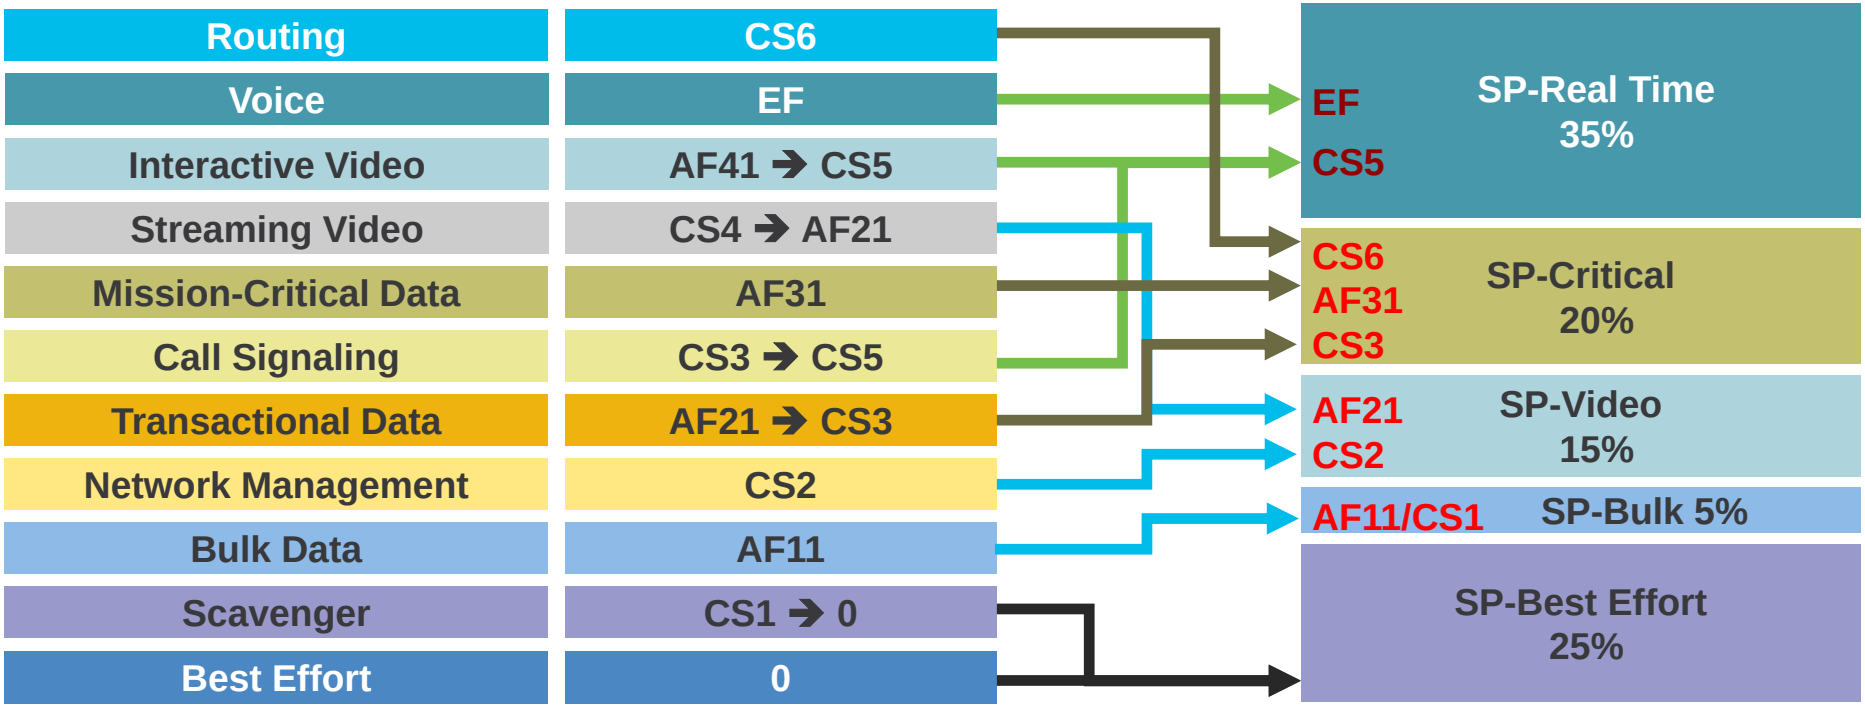

# IP Multicast in the Enterprise WAN

- IPs: 224.0.0.0 – 239.255.255.255
- Group destination IP, never a source
- Single source transmission efficiently delivered to a group of receivers
- Protocol-Independent Multicast (PIM) relies on unicast routing to build a loop-free, hop-by-hop, path
- PIM must be enabled along the entire end-to-end path
- Not supported over the Internet
- Service Providers offer MPLS VPN with Multicast capabilities

- L2 WAN transport allows Enterprise to fully manage the Multicast domain
- Can operate in Overlay but may require head-end replication limiting overall efficiency

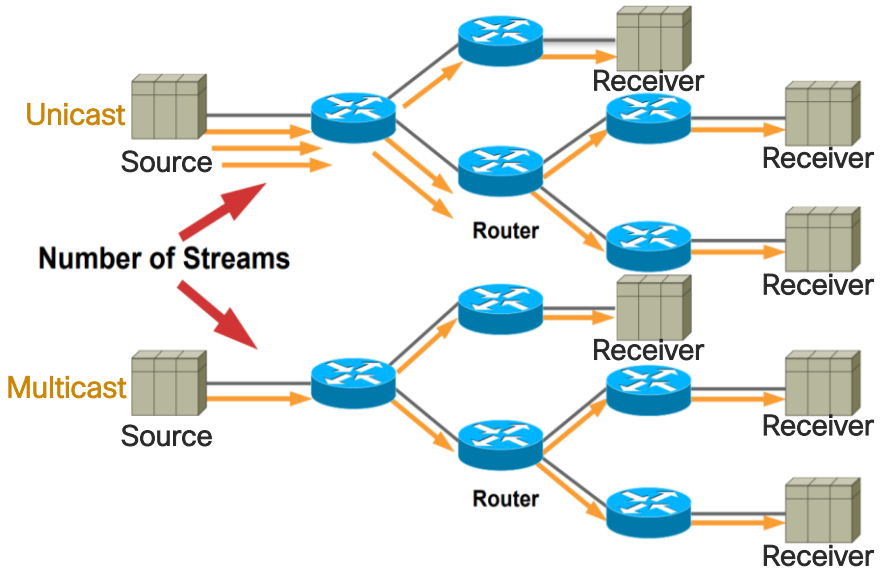

# Securing the WAN

- Perimeter security required at all Enterprise Internet connections points
- Private connections (eg. MPLS) provide a relative level of security
- Backhauling Internet traffic to data centers with appropriate perimeter security creates latency, congestion, and cost
- Deploying perimeter security at every location for DIA even more costly and difficult to manage
- The goal is a single security policy enforced across the entire WAN

***Security Tools***

- ✓ Firewalls
- ✓ Intrusion Prevention
- ✓ Visibility
- ✓ URL Filtering
- ✓ Advanced Malware Protection
- ✓ DNS Security
- ✓ Transport Security
- ✓ DDoS Protection
- ✓ etc…

# Cloud Connectivity Challenges

- **Complexity & Dependency** – Need a simple and scalable way to securely extend the private network across Multicloud environments
- **Inconsistent security policies between private & public** – Need to apply consistent security policies
- **Degraded application performance and ambiguity for best path to reach the cloud** – Need to enhance application experience

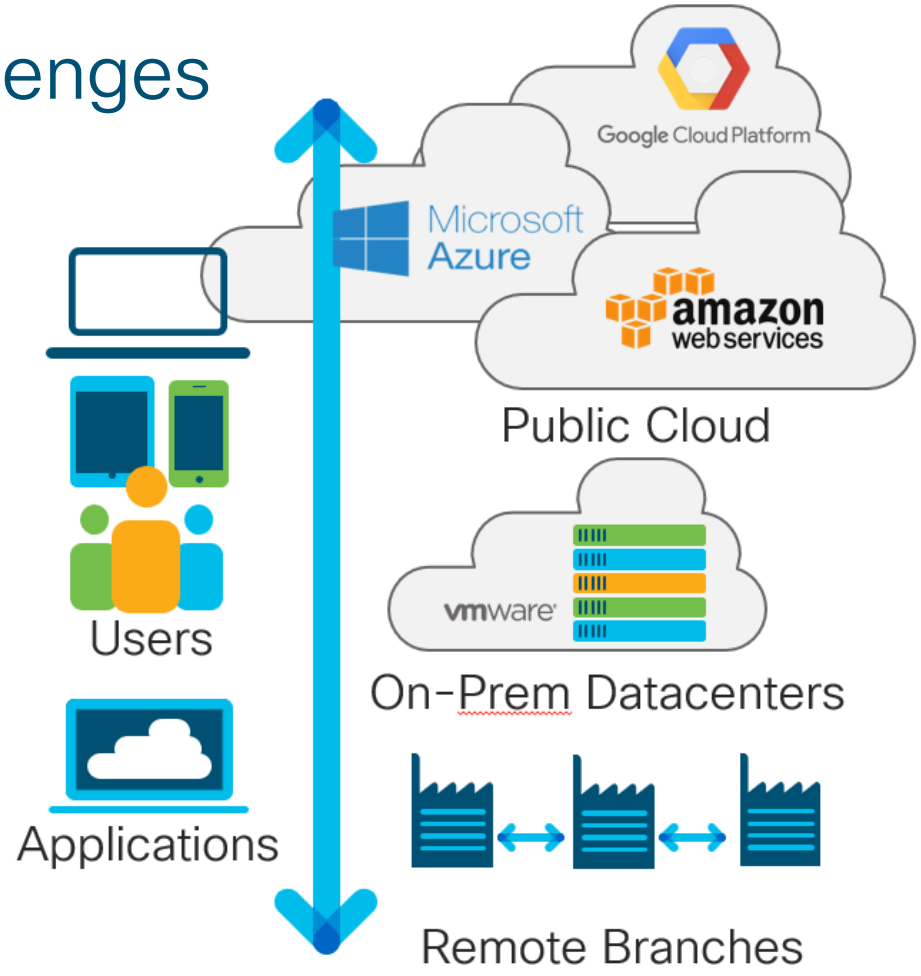

# Public Cloud Deployment Models

## Application VPC Gateway

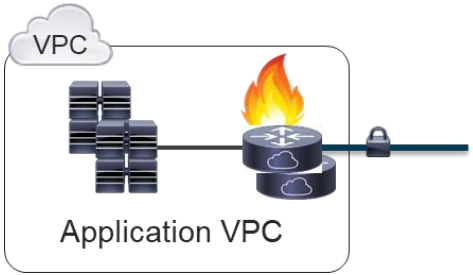


- CSR deployed in application VPC
- Provide IPsec gateway for entire VPC
- Need high availability

## Transit VPC

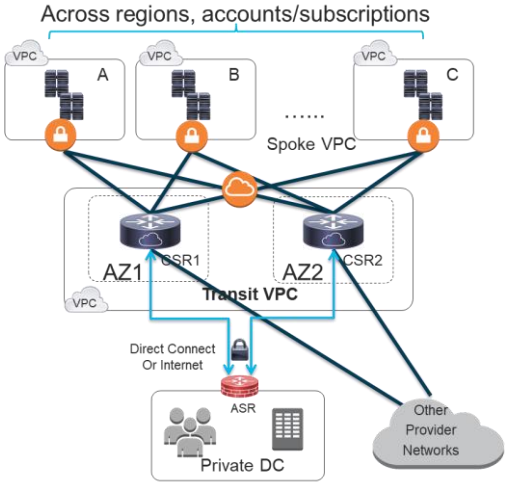


- CSR deployed in dedicated Transit Hub
- High speed traffic routing for spoke VPC
- High availability is built-in natively

## Auto-scale

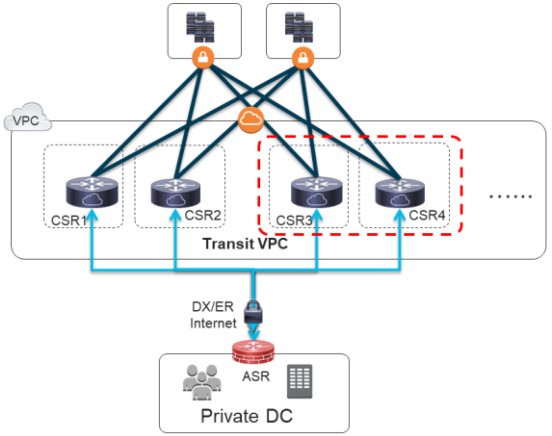


- Add pair of CSRs to scale out
- Remote end (VGW) has multiple tunnels and do L3 ECMP (Equal Cost Multiple Path)
- Monitors CSR real-time throughput and spin up new CSRs on demand

# Connecting to Public Cloud

## Internet connection

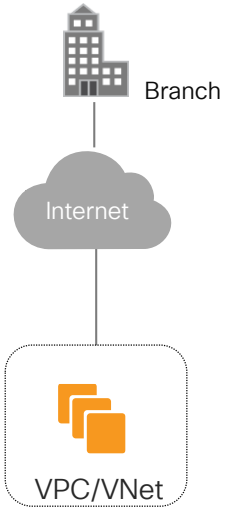


Internet only for connectivity

## IPSec tunnel from DC

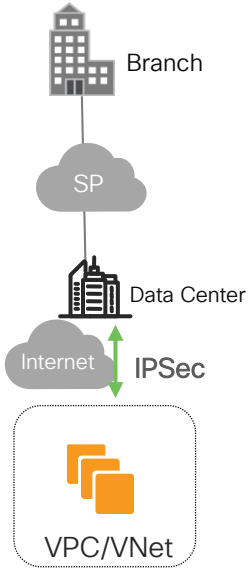


IPsec Tunnel from customer DC to the cloud

## DX / ER to Public Cloud through SP

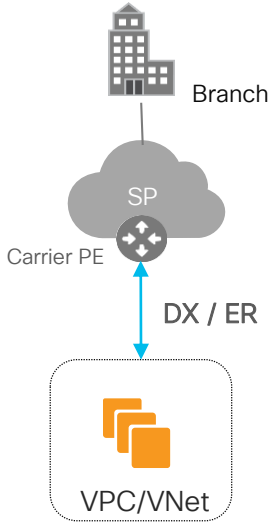


MPLS carriers (L3 VPN carrier) offers DX/ER as SP Managed Service

## Direct Connect to Public Cloud through co-locations

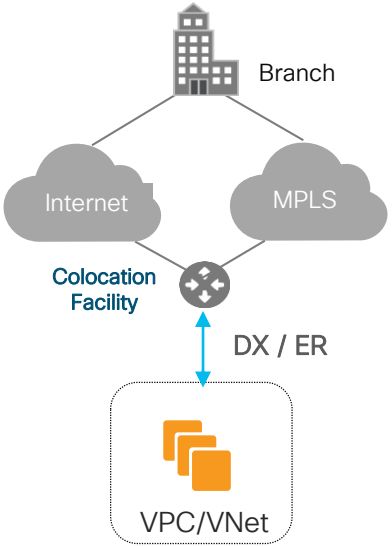


DX/ER from the co-location to the cloud

Remember the Main Message:

*Foundational Design is key to WAN Architecture*

# WAN Designs moving Forward

# Common WAN Topologies

## Design and Deployment Considerations

### Design Challenges with Growing Needs and New Innovation

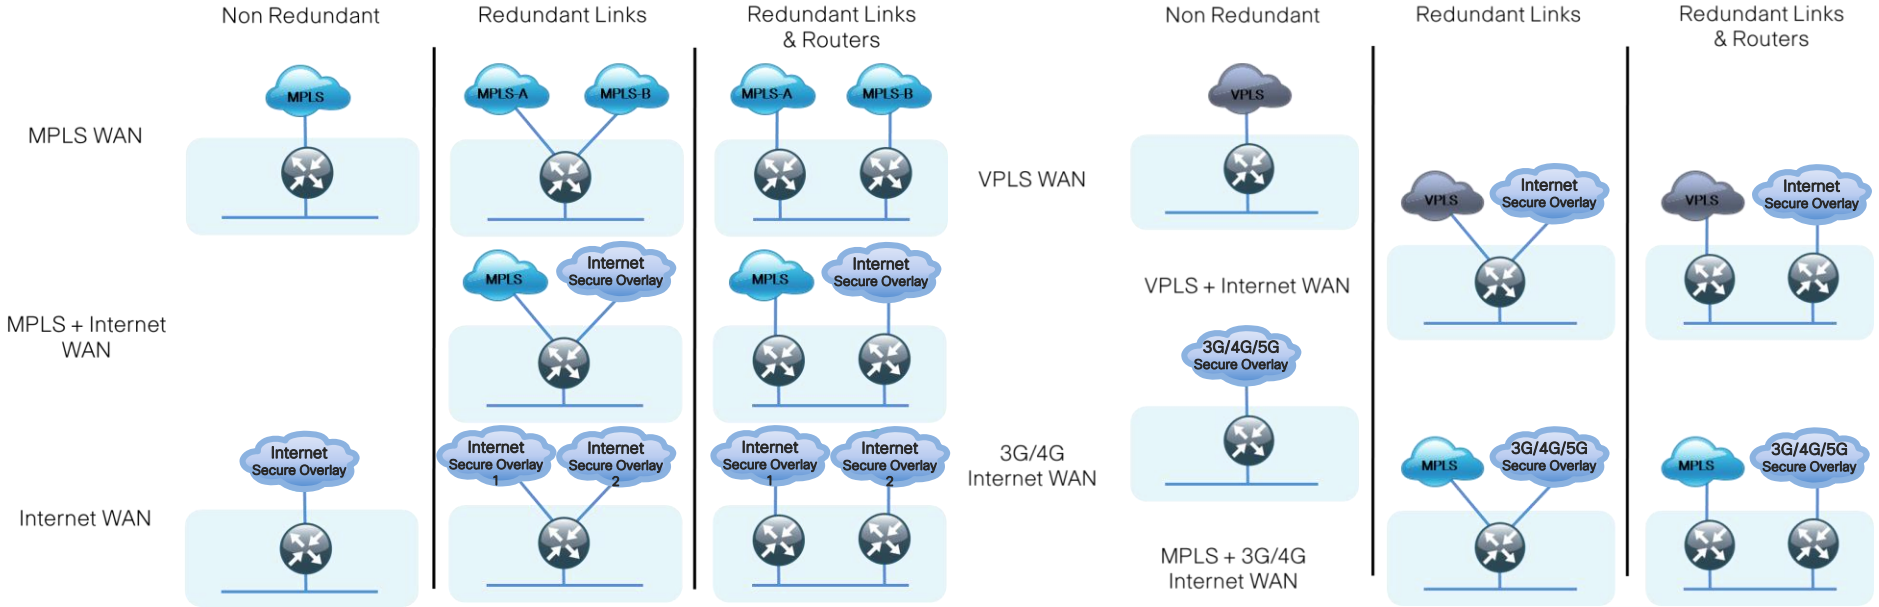

# Common WAN Topologies

## Growing Complexity – Scale, Policy, Segmentation

### Complexity Grows with Scale and Changing Business Requirements

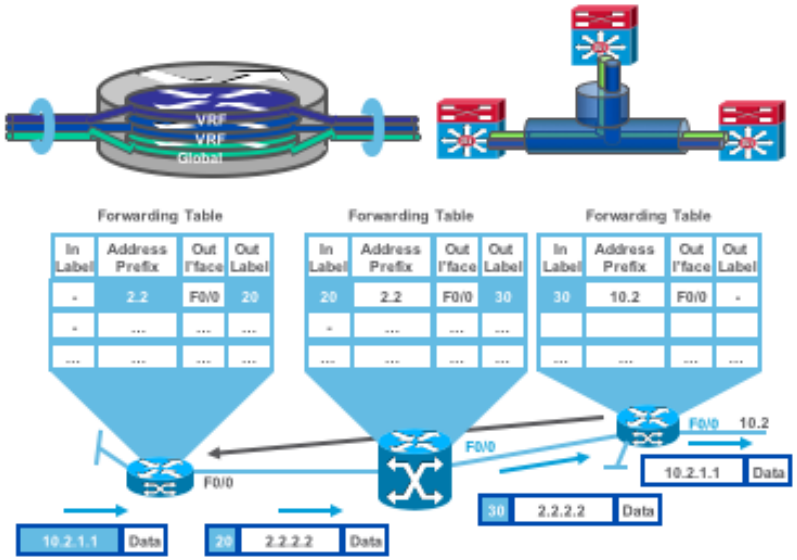


| Application | DSCP | CS6 Sent Unchanged | SPP3 (6-Class Model) |
|---|---|---|---|
| Internetwork Control | CS6 | | EF — SP-VOICE PQ-10% BW |
| VoIP | EF | | |
| Broadcast Video | CS5 → AF41 | | AF41 — SP-VIDEO 34% BWR |
| Multimedia Conferencing | AF41 | | |
| Real-Time Interactive | CS4 → AF41 | | |
| Multimedia Streaming | AF31 | | AF31 — SP-CLASS1DATA (UDP) 10% BWR |
| Signaling | CS3 → AF21 | | AF21 — SP-CLASS2DATA (TCP) 25% BWR |
| Transactional Data | AF21 | | |
| Network Management | CS2 → AF21 | | |
| Bulk Data | AF11 → AF21 | | AF11 — SP-CLASS3DATA—1% BWR |
| Scavenger | CS1 → AF11 | | DF — SP-DEFAULT 30% BWR |
| Best Effort | DF | | |

# Drivers for Change

- Today, large majority of application traffic on private network is destined off-network
- Some is critical traffic, not all, destined to SaaS, IaaS (e.g. O365, Salesforce.com, or Azure)
- Includes regular browsing traffic from each location
- MPLS can be an expensive conduit to a centralized Internet breakout point
- Enterprise pays for private bandwidth and then again for Internet bandwidth
- This change in traffic impacts capacity planning, application performance, and ultimately user satisfaction
- Major challenge to use traditional WAN features to deliver a cohesive solution and to troubleshoot

# A New Era in Network Architectures

**1st Wave – TDM**
TDM rigidity limits new services, forces architectural shift

TDM Era

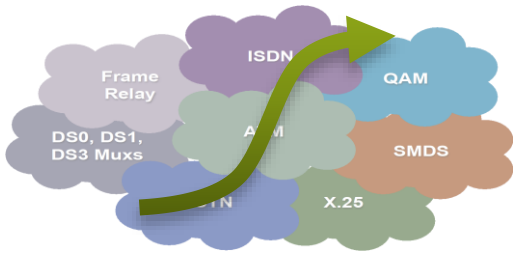


**2nd Wave – MPLS**
Commoditization of IP services plus high traffic growth limits profitability, forces architectural shift

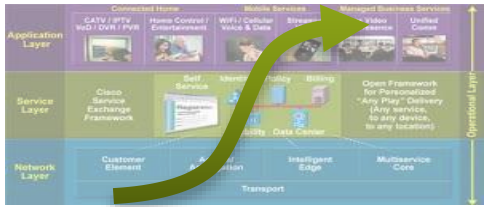

IP unleashes new wave of innovation and service revenues

**3rd Wave – EPN**
Evolved Programmable Network Era, Digital Transformation

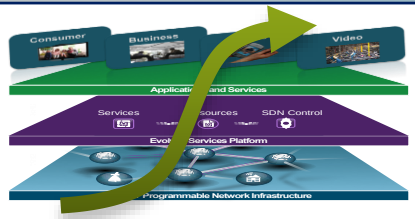


Network Function Virtualization, Software Defined Networking, and Service Orchestration enable
- Open and Dynamic
- Optimal resource utilization
- Accelerated innovation
- New services & revenues
- Reduced costs
- Reduced complexity

~5-10 Year Transition

~2-10 Years?

# 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**

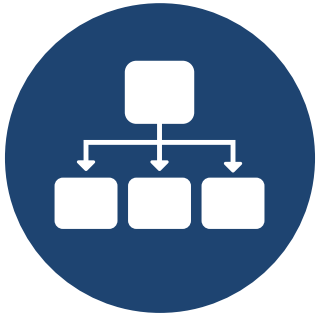

**Leverage the Power** of Existing **Distributed Systems**

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

# Cisco Digital Network Architecture

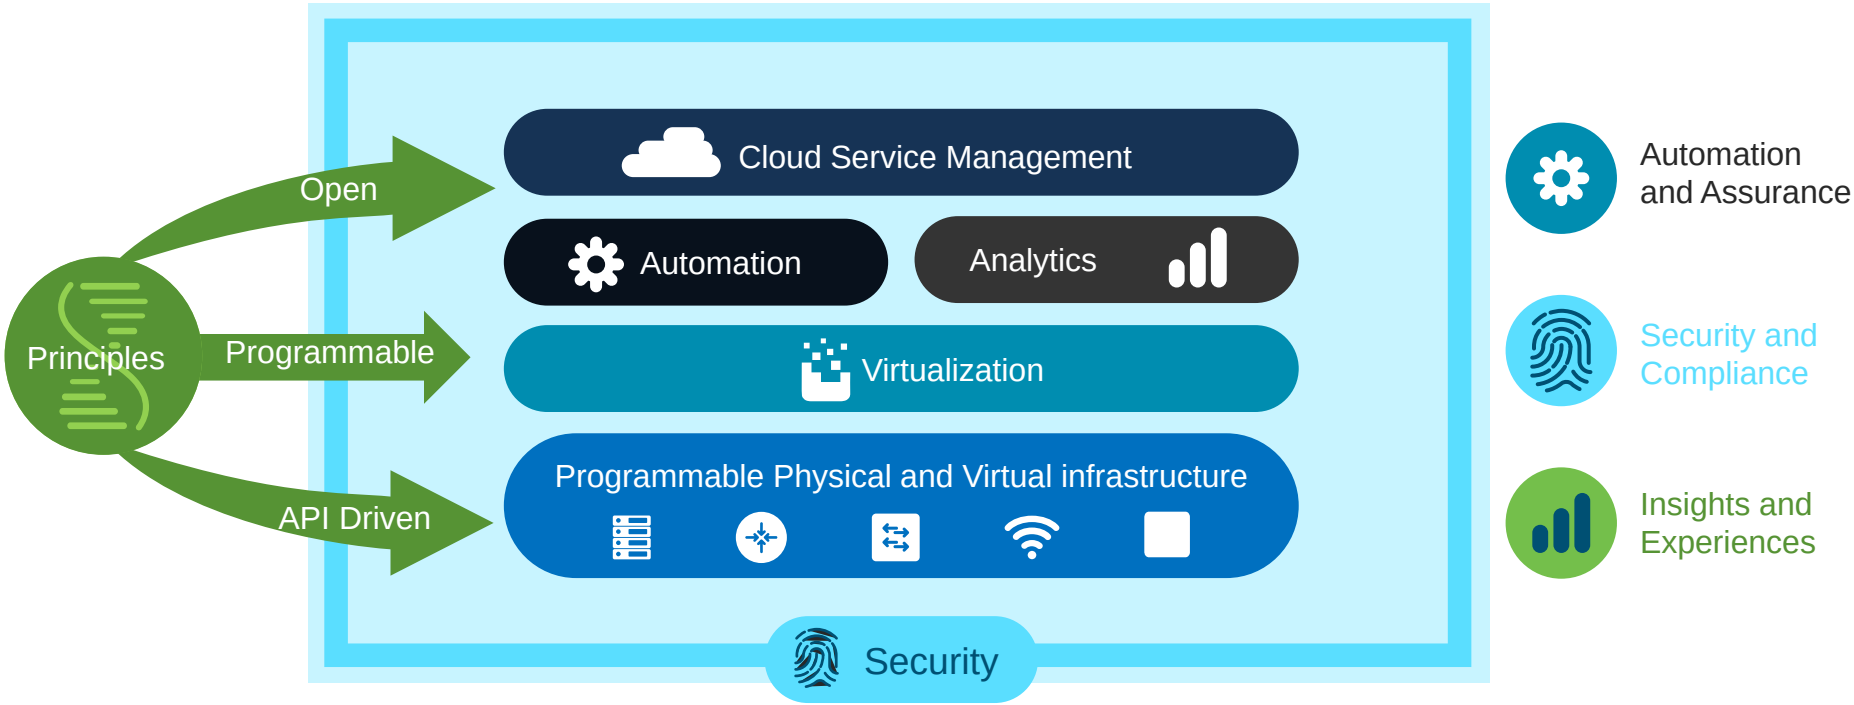

# Cisco Digital Network Architecture

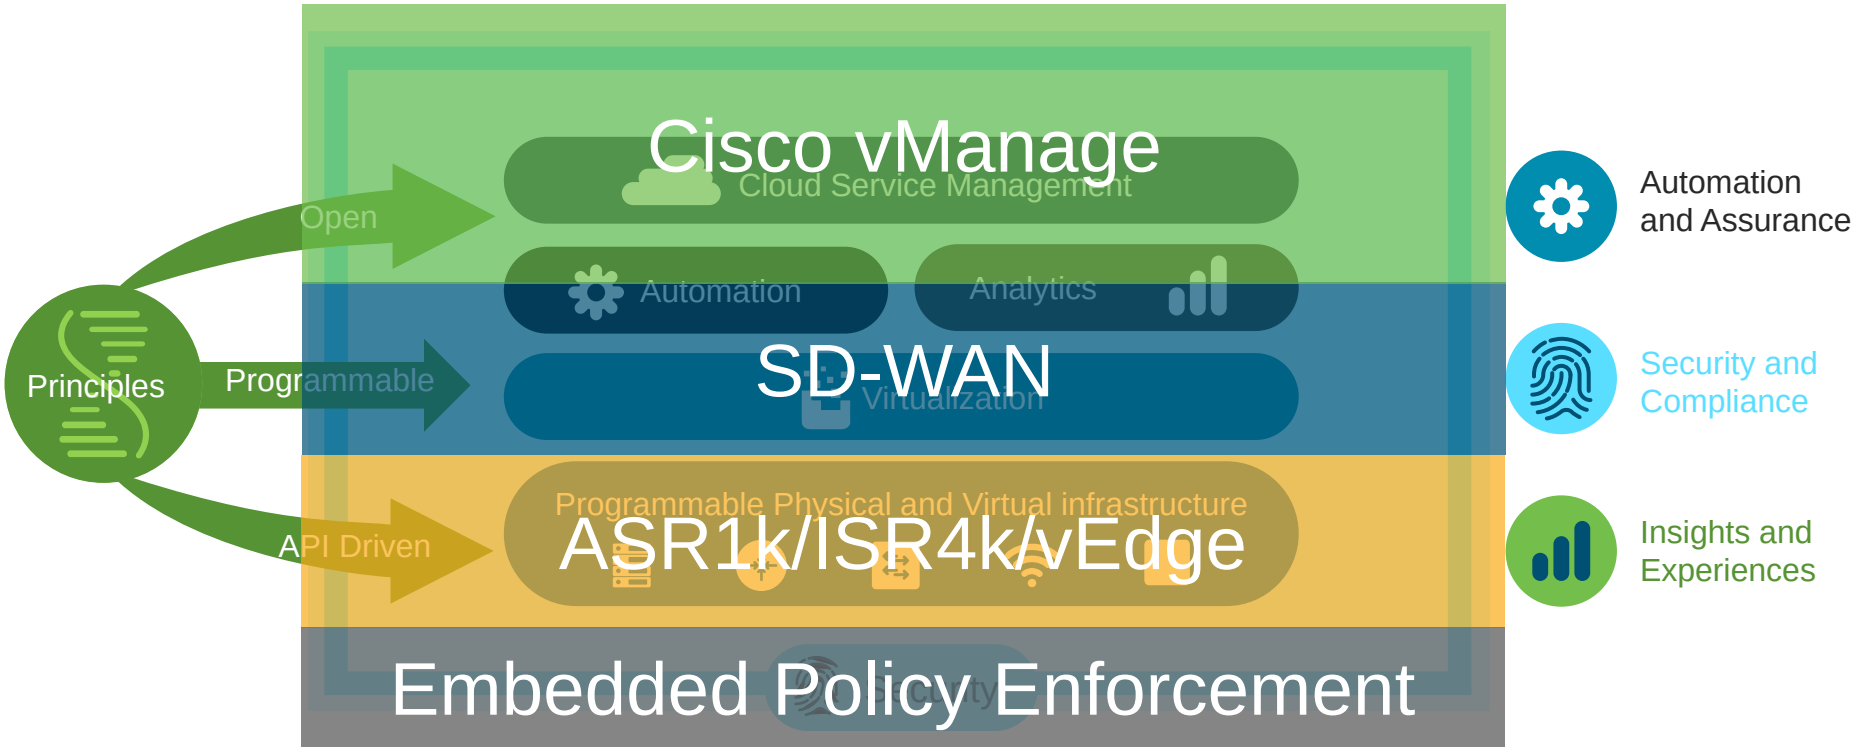

# SDWAN

# Network Transformation

## The Era of Digital Transformation

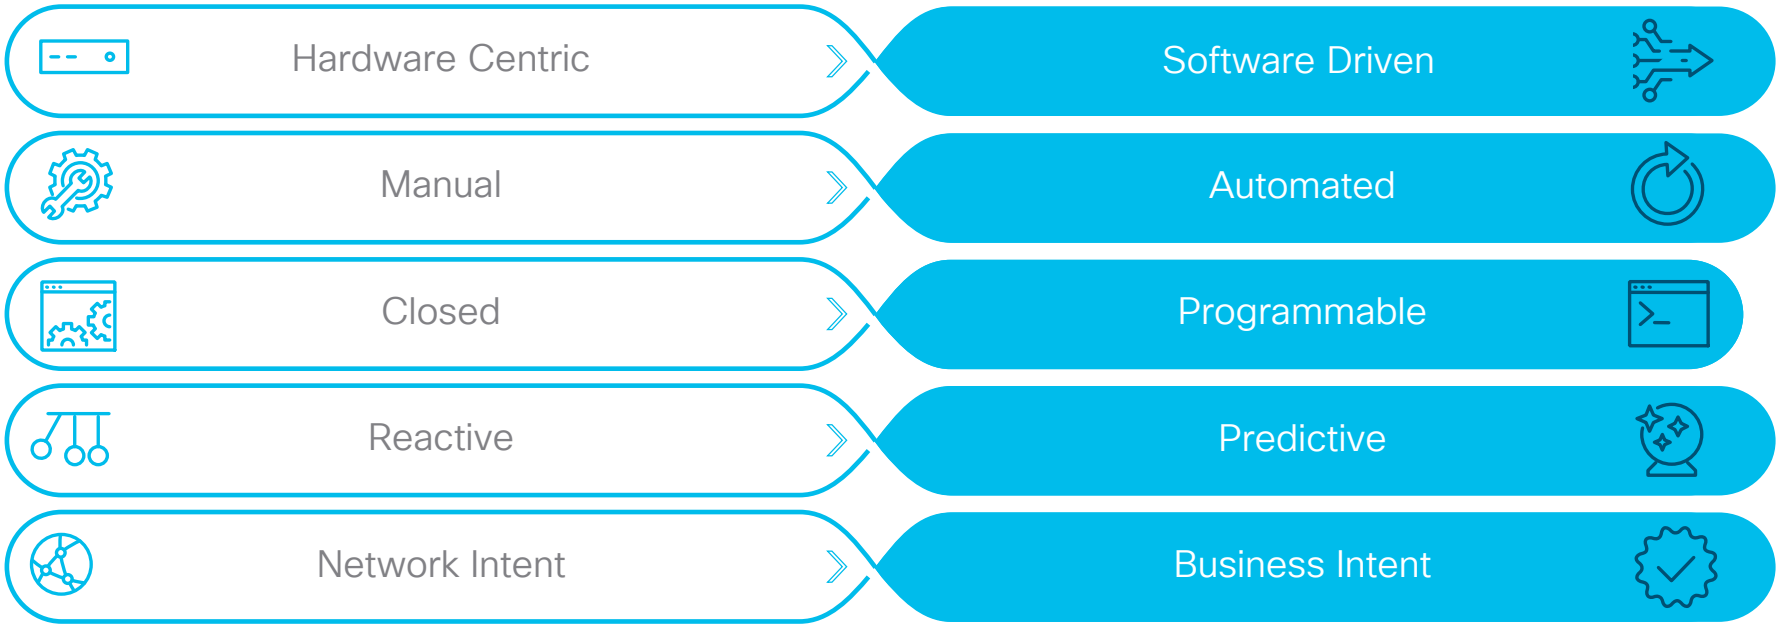


**CLOUD & ON-PREM**
Hosted, delivered, managed

**AUTOMATION & SCALE**
Speed, flexible, zero-touch, policy driven

**SECURITY & COMPLIANCE**
Segmentation, threat mitigation

**ASSURANCE & ANALYTICS**
Users, applications, devices

# Business Driven SD-WAN Infrastructure Design and Deploy for Impact Objectives

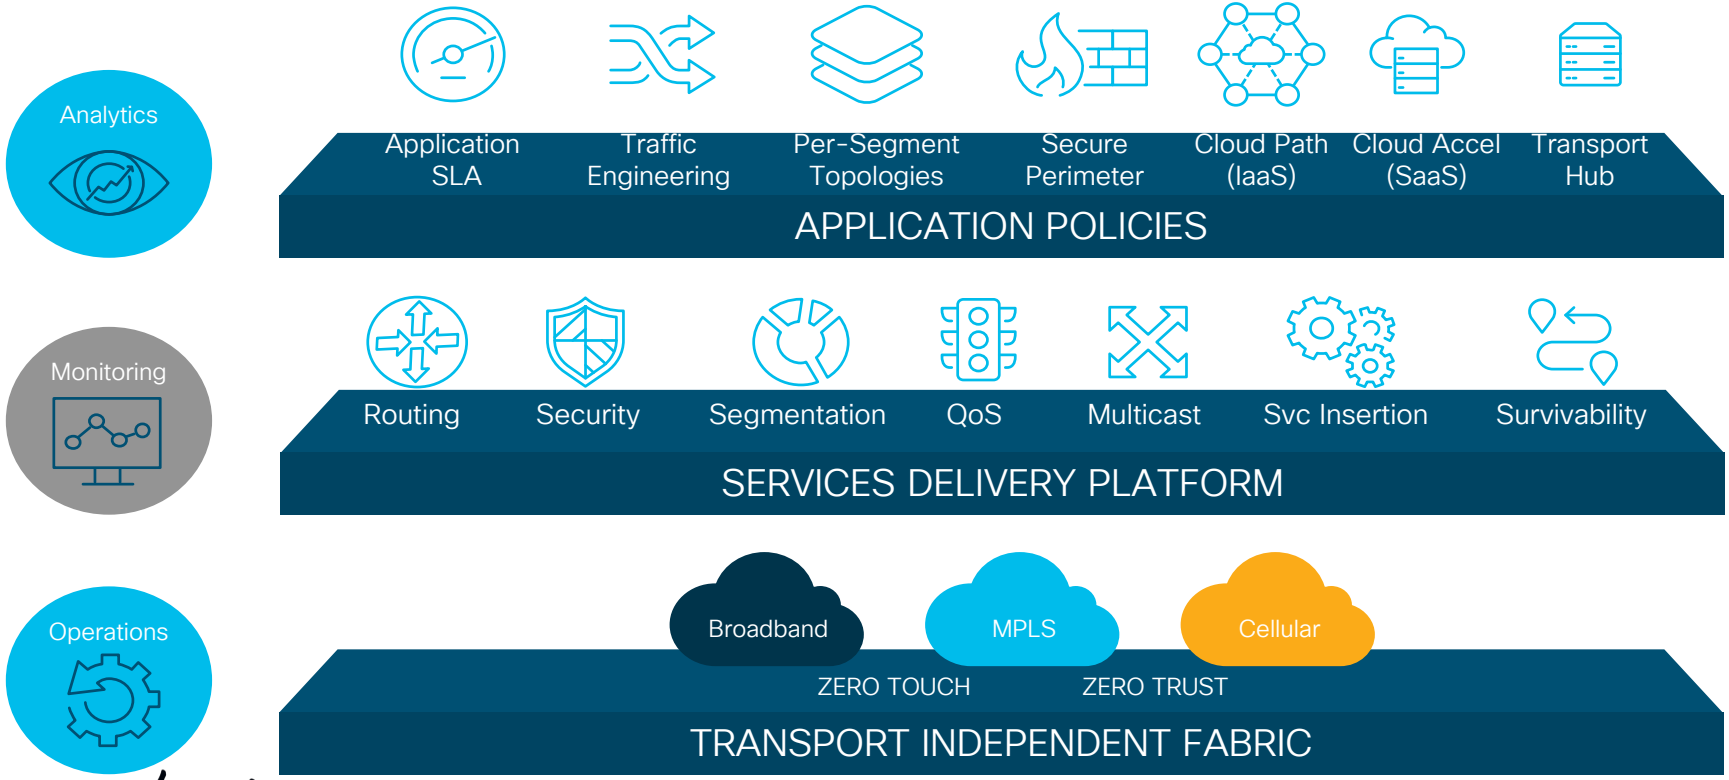

# Reinventing the WAN
# The Four Pillars and Focus Areas of Cisco SDWAN

- Security
- Connectivity
- Application Services
- Operations

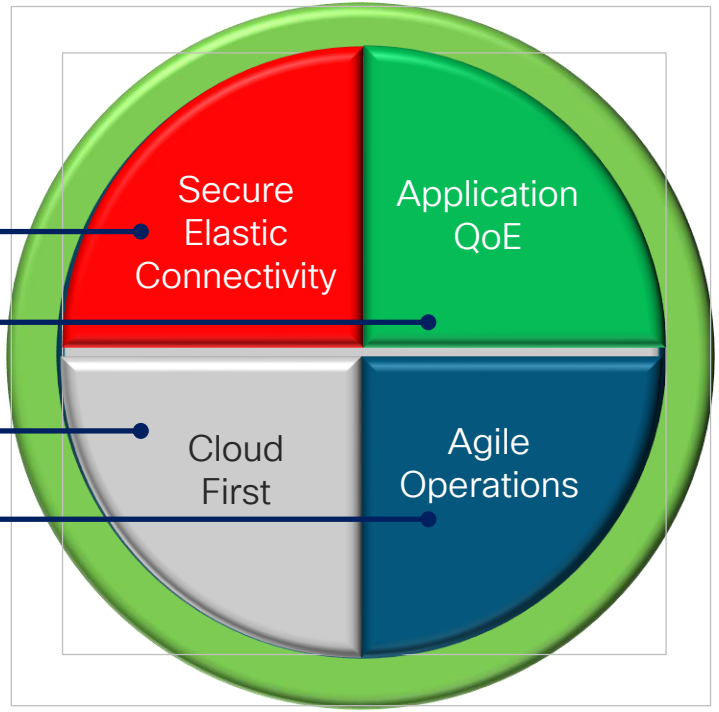

# Reinventing the WAN

## Security

Embedded Security

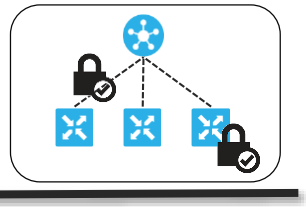

Centralized Device Auth-DB

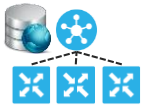

Authenticated/Encrypted Control Plane

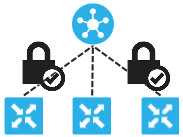

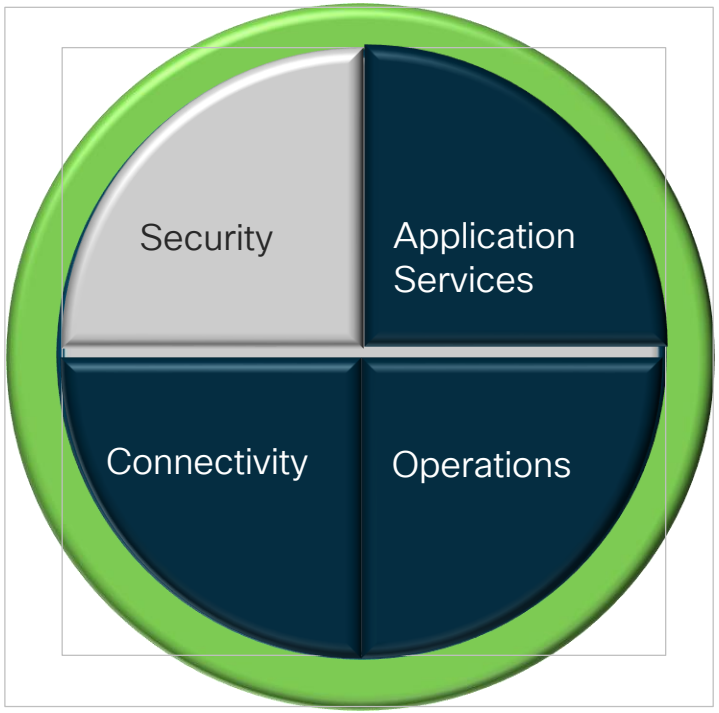


Secure Bring-up

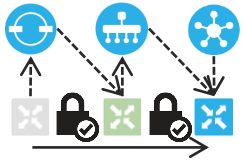

Scalable Data-Plane Encryption

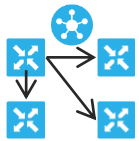

Automatic Key Rollover

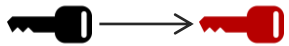

# Reinventing the WAN

## Connectivity

Hybrid WAN

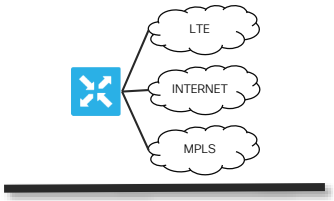


Segmentation/VPNs

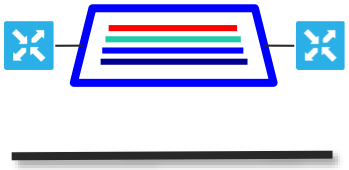

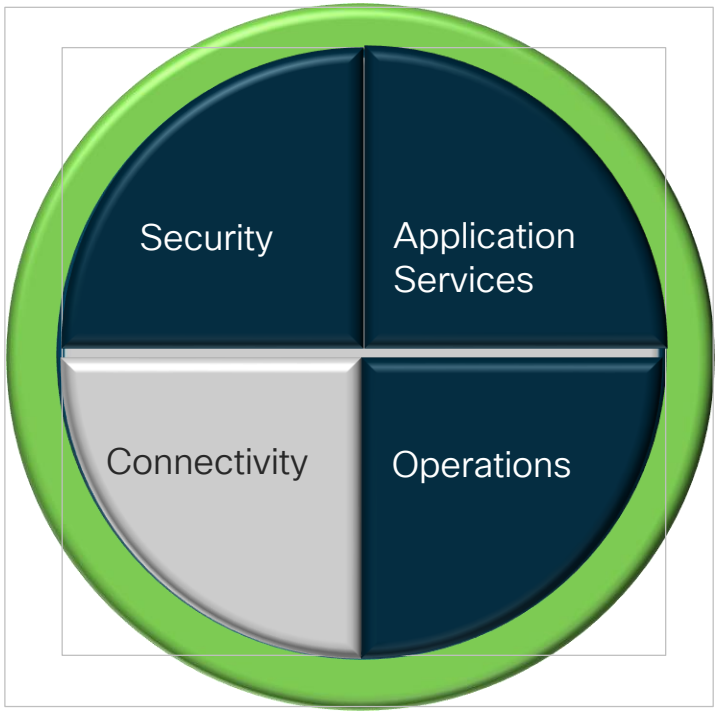


Provider/Transport Agnostic

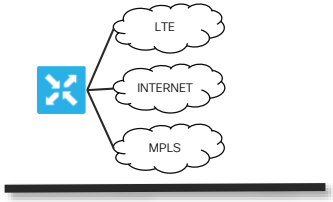


Dynamic Per-VPN Topologies

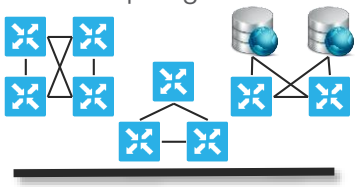

# Reinventing the WAN

## Application Services

Deep Packet Inspection

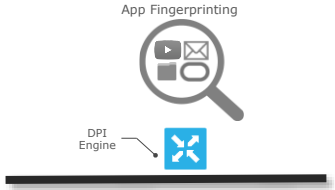


Transport SLA Monitoring

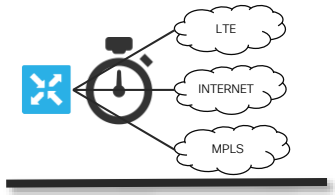


Application-Aware Routing

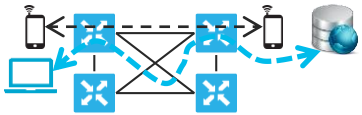

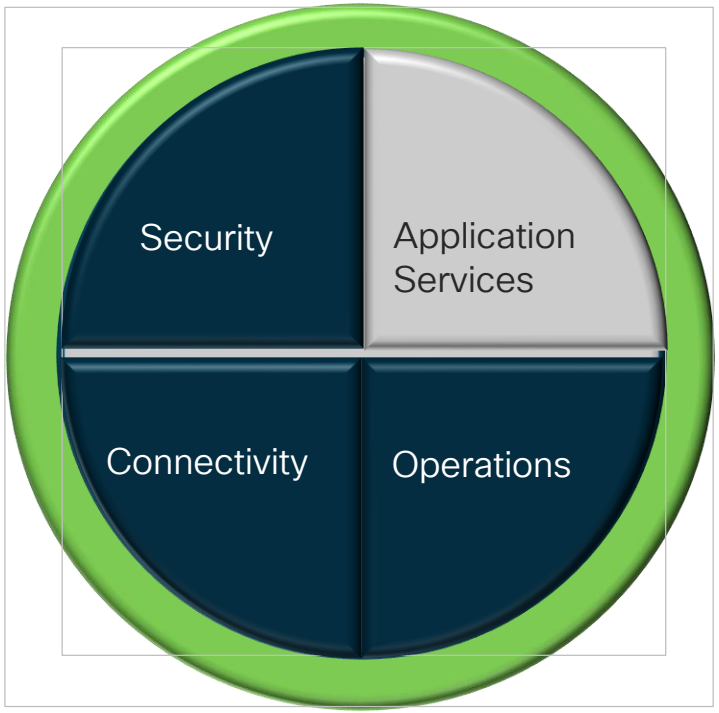


Central Orchestration

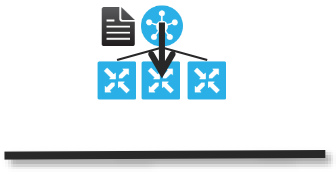

Application Layer Analytics

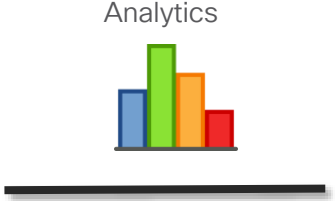

Cloud Services Integration

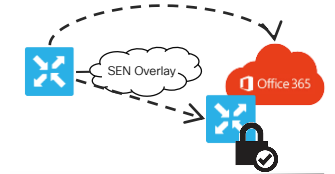

# Reinventing the WAN

## Operations

Centralized Operations
Distributed Execution

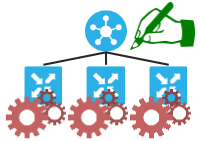

Template-based
Configurations

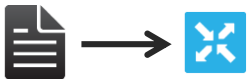

Programmatic APIs
Open Object Model
NetConf

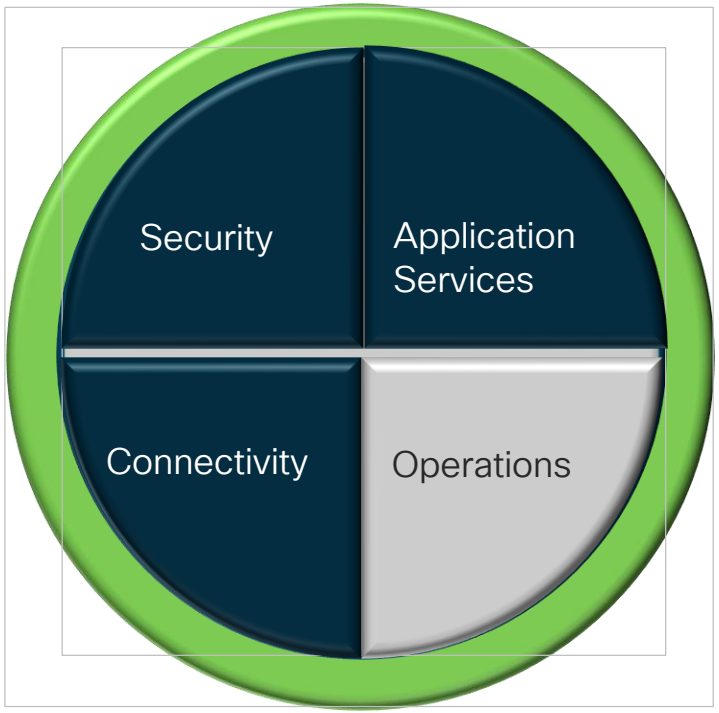


Centralized
Policy Orchestration

Zero Touch Provisioning

Ad-Hoc
Adds/Moves/Changes

# Cisco SDWAN Solution Overview

## Applying SDN Principles To The Wide Area Network

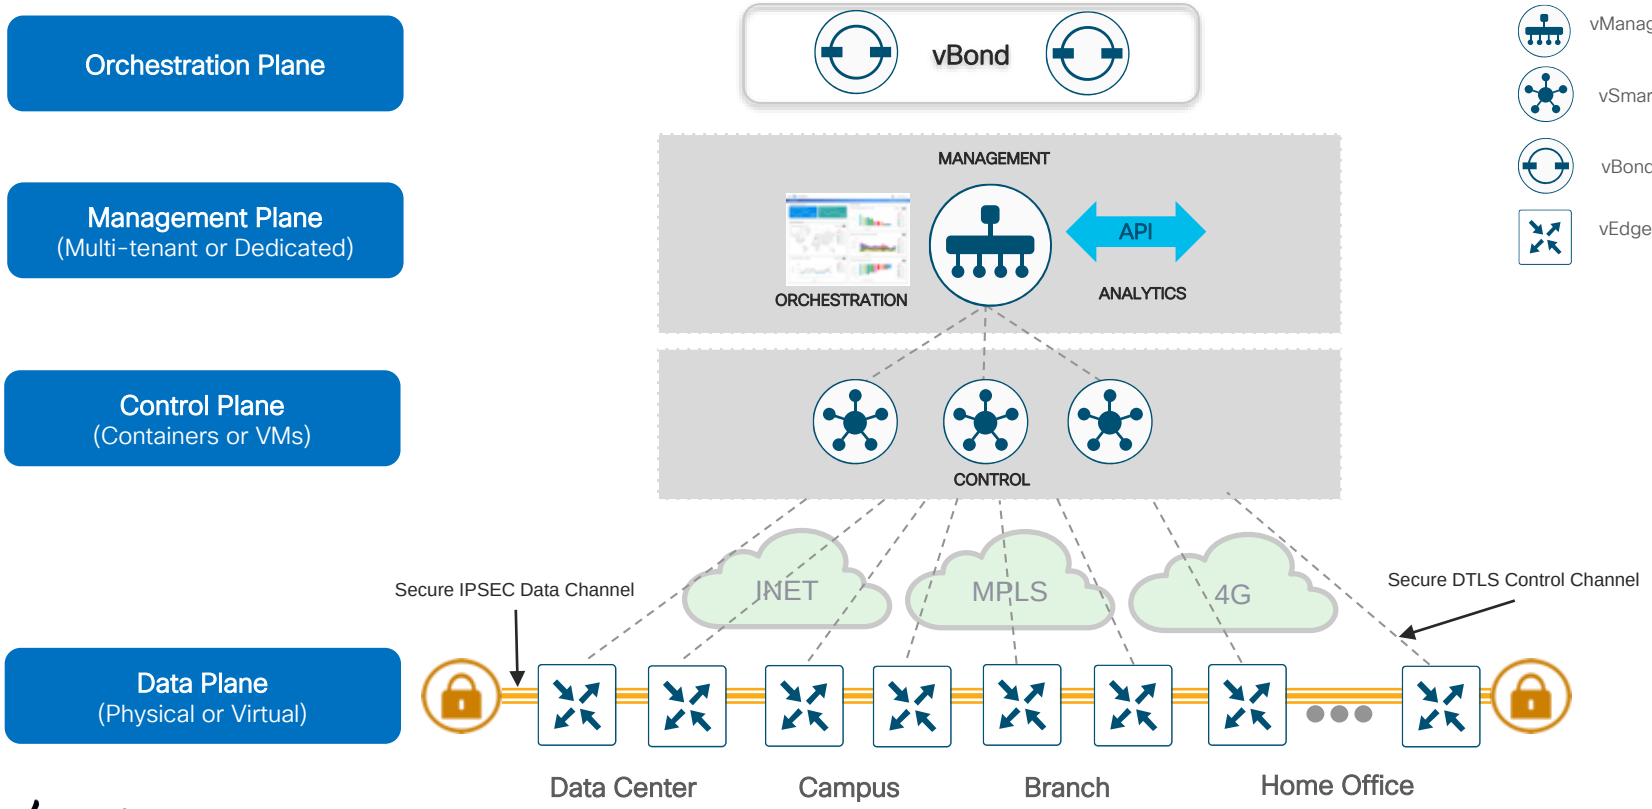

# Cisco SDWAN Typical Architecture

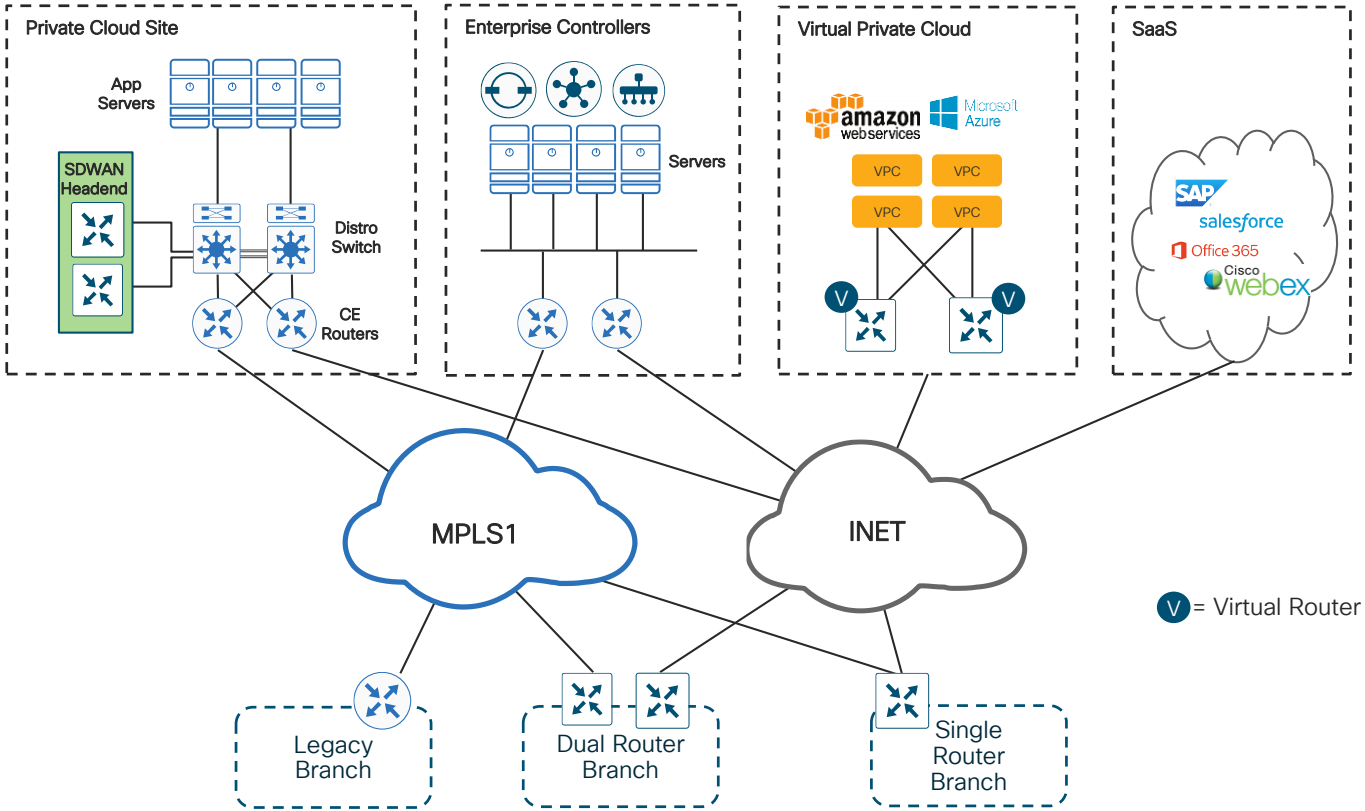

# Cloud-Delivered SDWAN Control

## Flexible Deployment Options

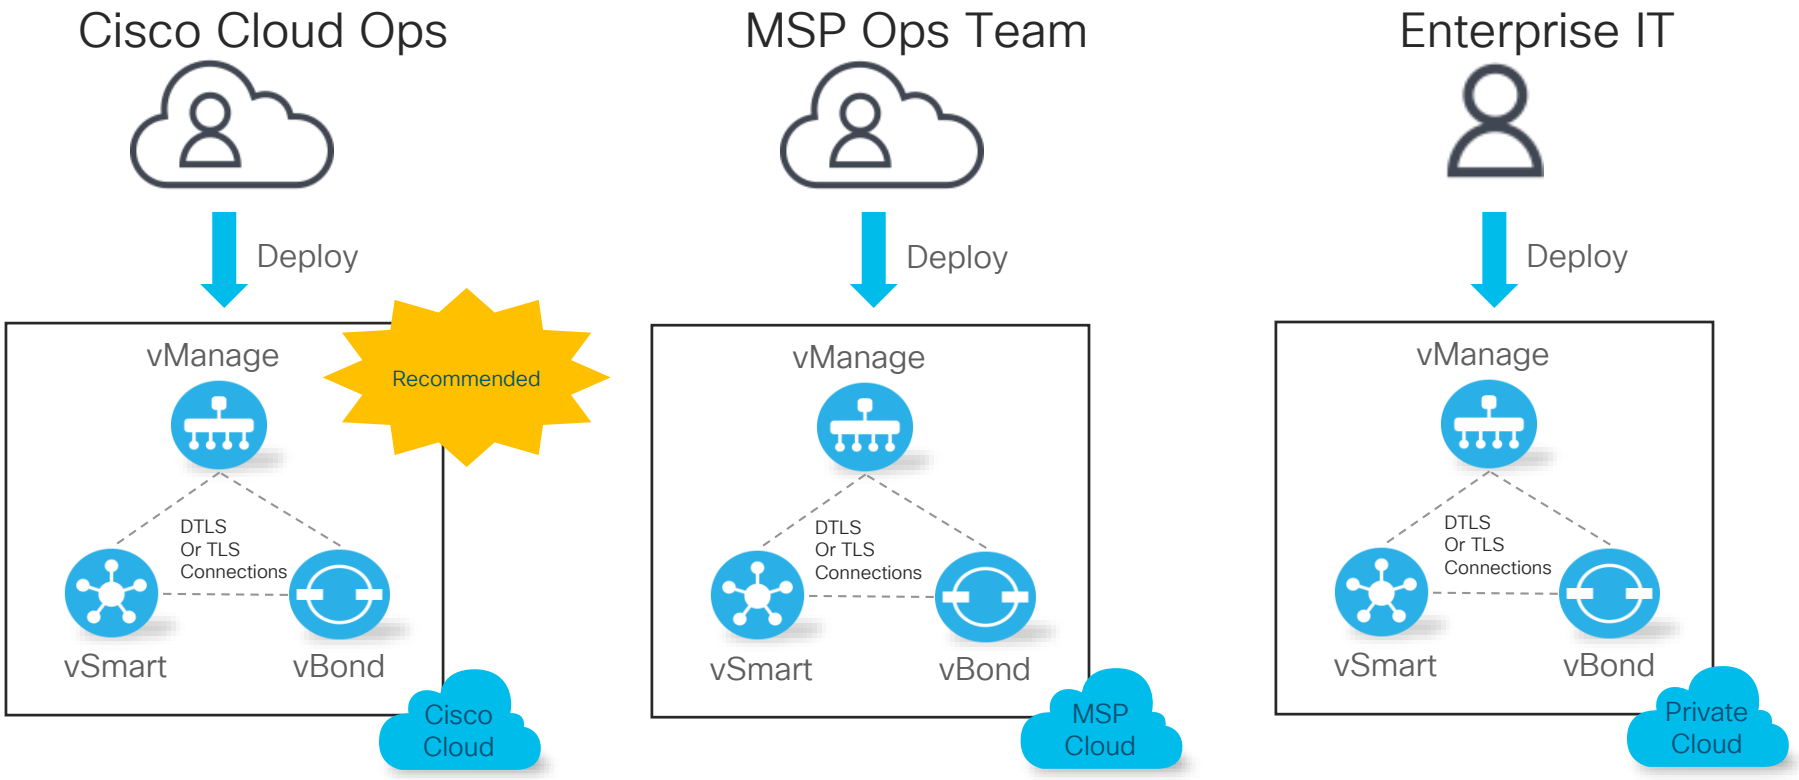

# Multi-Path Multi Destination – Per SLA

**App Aware Routing Policy**
App A path must have:
Latency < 150ms
Loss < 0%
Jitter < 5ms

**Path1: 10ms, 0% loss, 5ms jitter**
Path2: 90ms, 3% loss, 10ms jitter
Path3: 200ms, 1% loss, 10ms jitter
Path4: 180ms, 1% loss, 5ms jitter

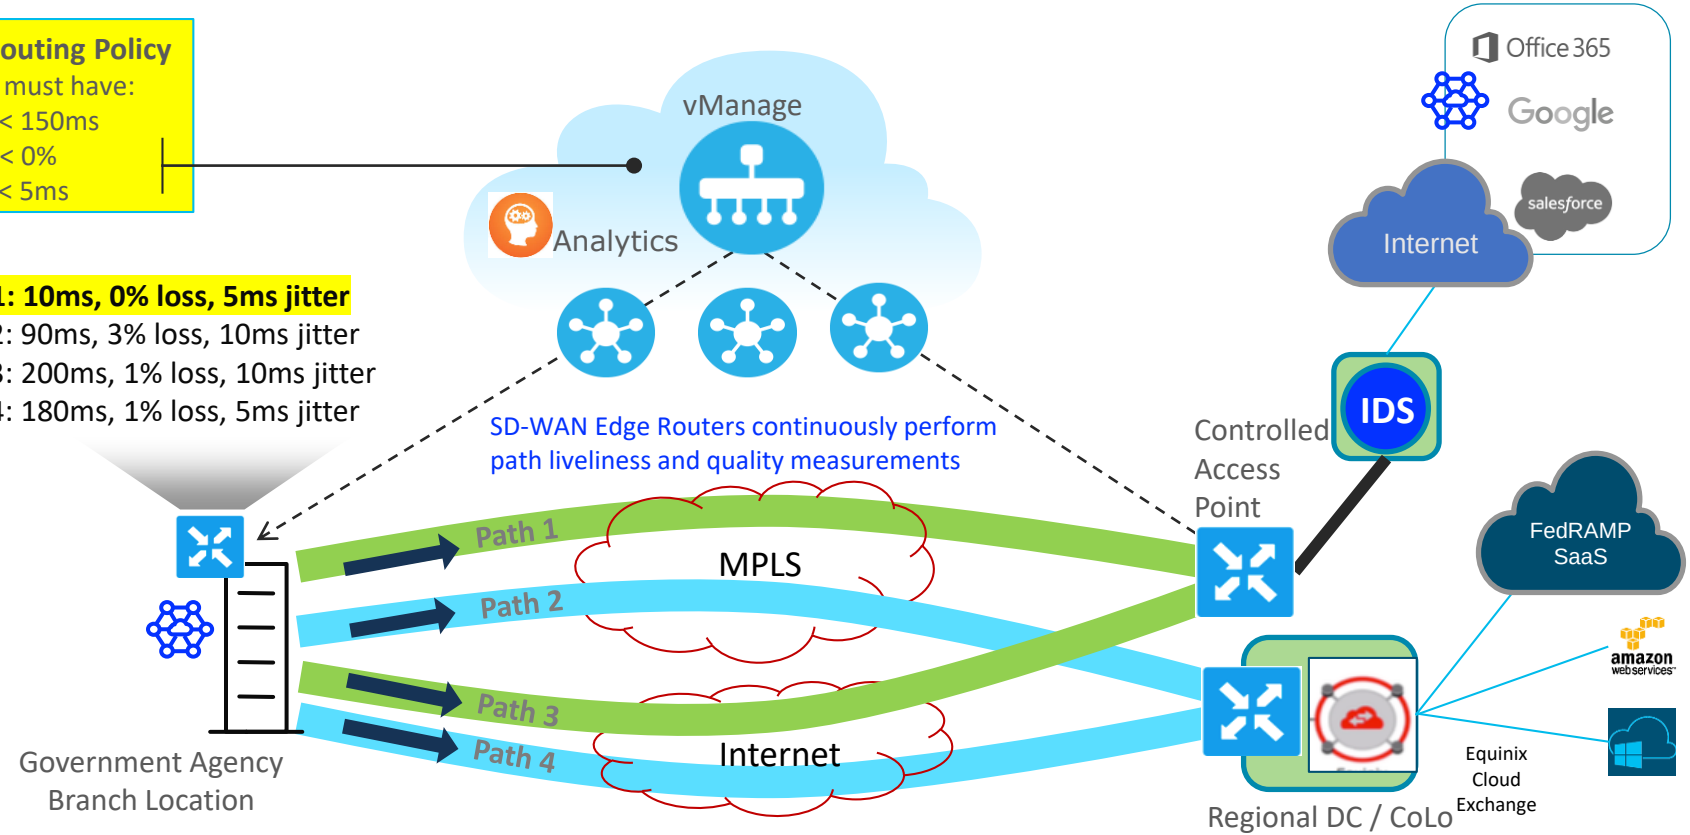


Application Aware Probe
Controlled Access Path
Regional CoLocation Center

# Cisco SD-WAN – Cloud OnRamp for SaaS

- User designates Cloud onRamp gateways which can be remote DMZs or local CPE (DIA case)
- App-Aware routing to SaaS end-point from gateway routers
- SLA metrics are computed by using httping based probes to the SaaS endpoint through the Cloud onRamp gateway
- Per application SLA metrics include loss and latency
- Path experiencing better SLA for the application is chosen

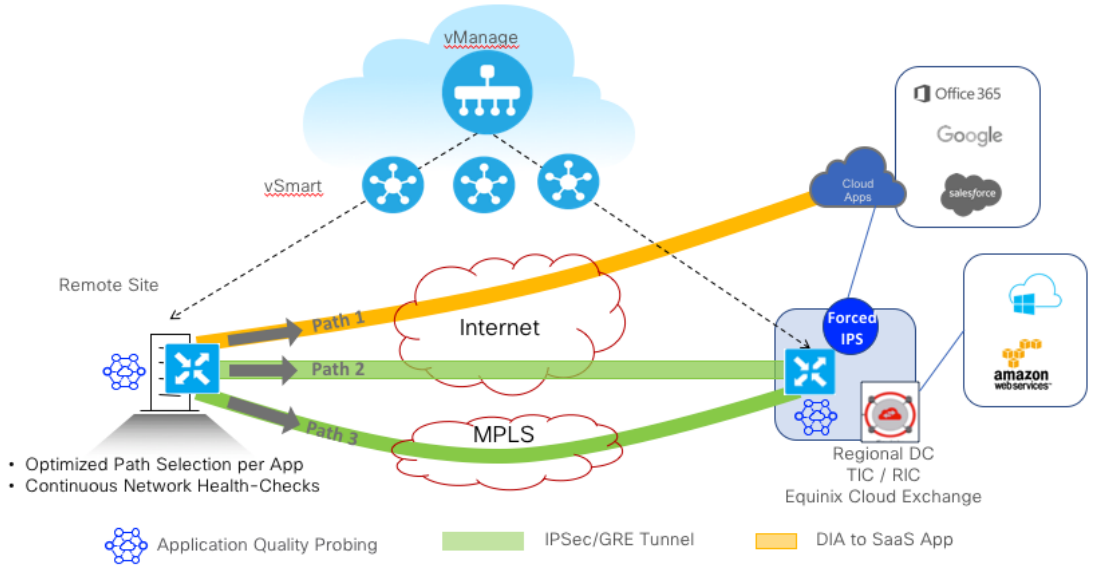

# Cloud Ready WAN Architecture

# Centralized Data Center Architecture (Legacy)
## Hosted Applications in the Agency Owned Data Center

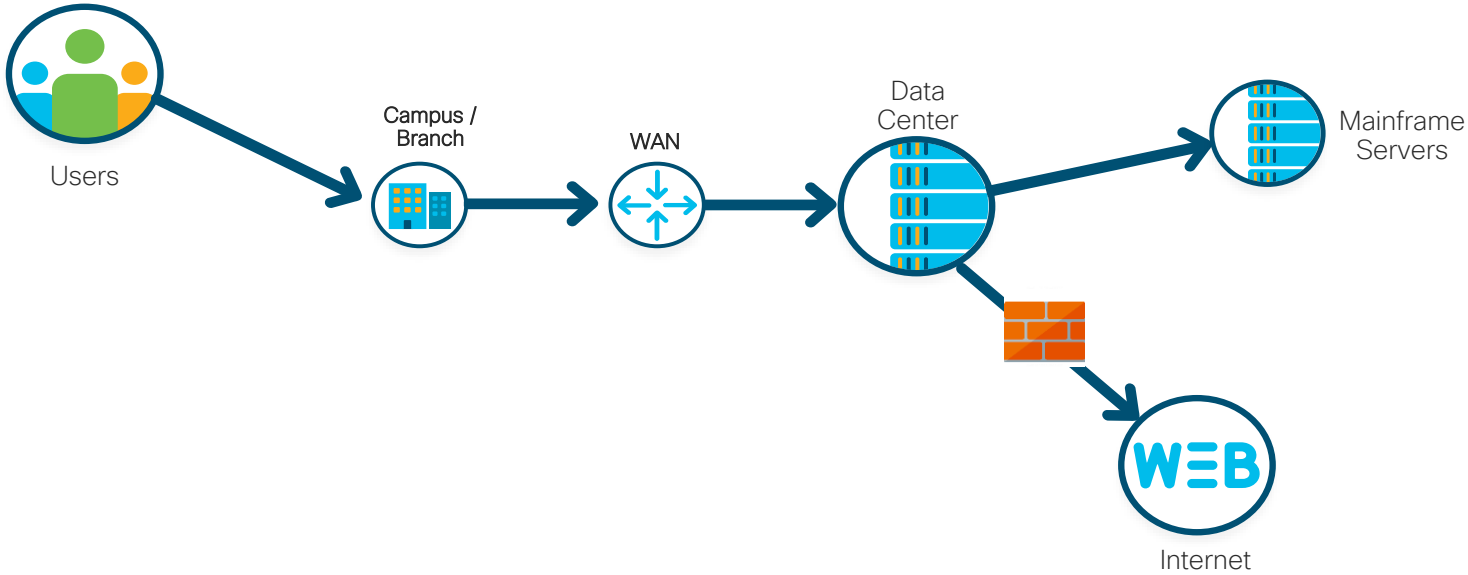


Full Security Stack

# Next Generation Enterprise Architecture

## Network Architecture Transition in a Multi-Cloud World

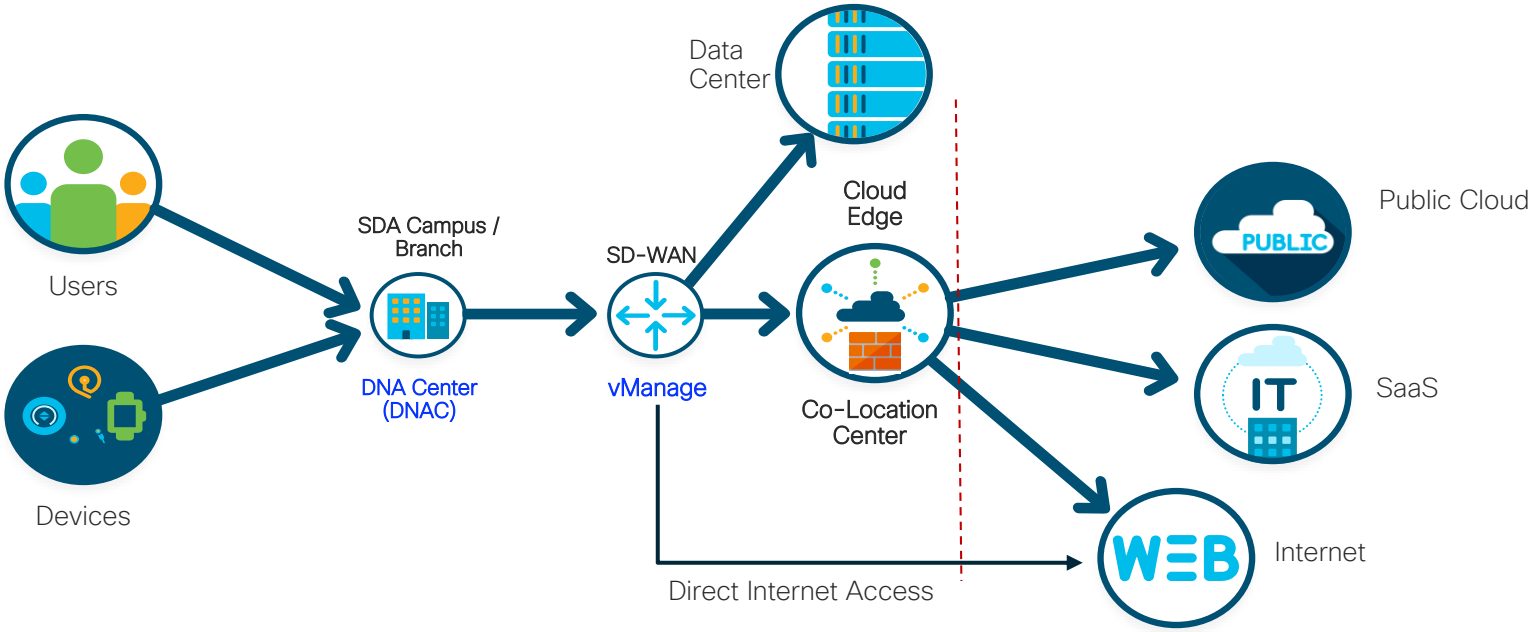


Deliver Segmentation, Security, Automation, anytime, anywhere, Any transport

# Cloud Ready Network Architecture

## Aligning WAN Design w/ Applications and Perimeter DMZ in Co-Location Centers

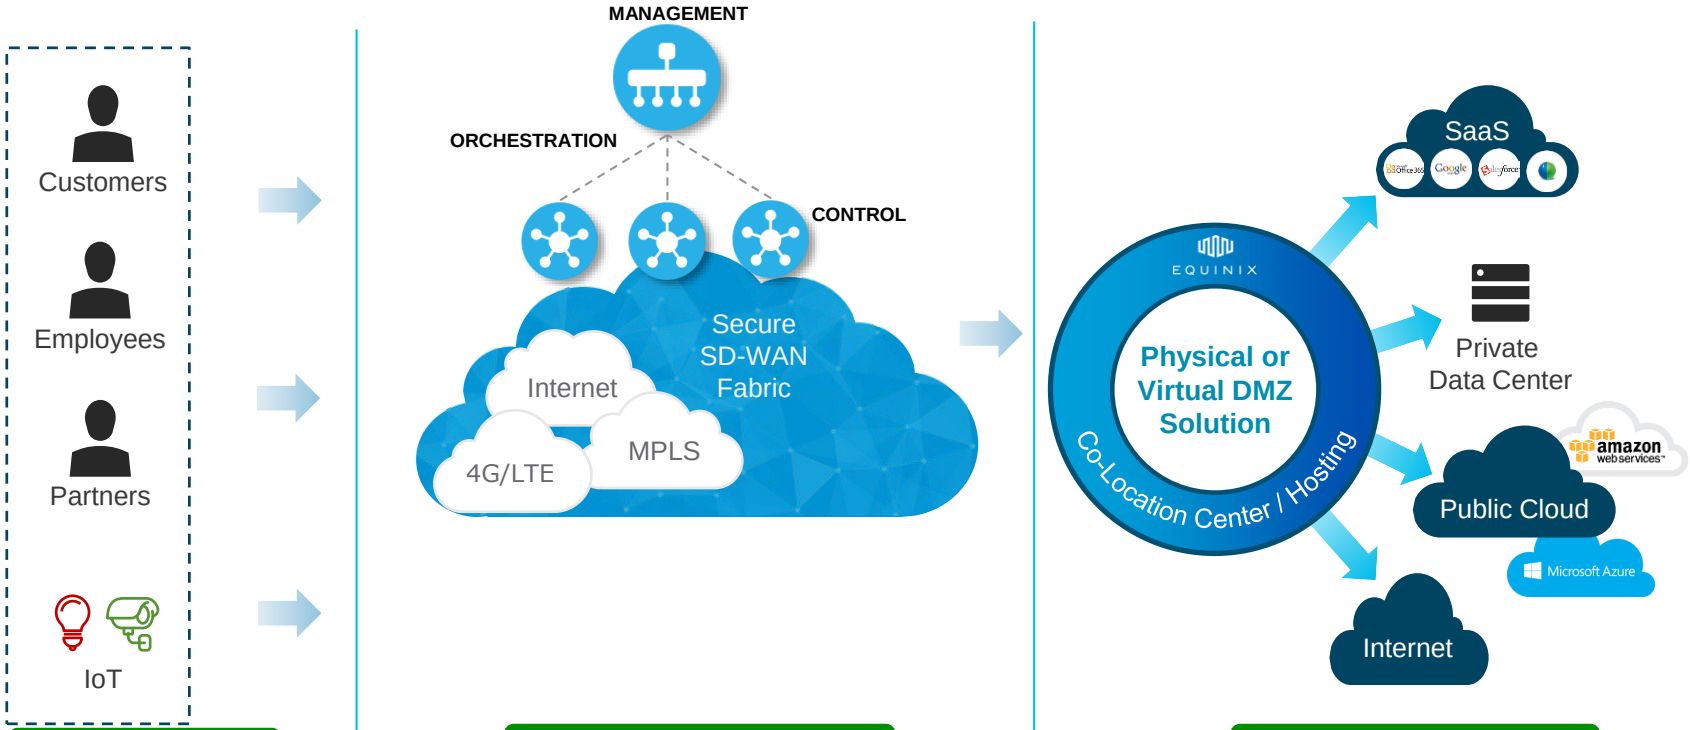


Management | Security | Policy | Orchestration | Analytics

# Cisco SD-WAN Cloud On-Ramp for CoLo

## Securely Connecting Users Cloud and Application Providers

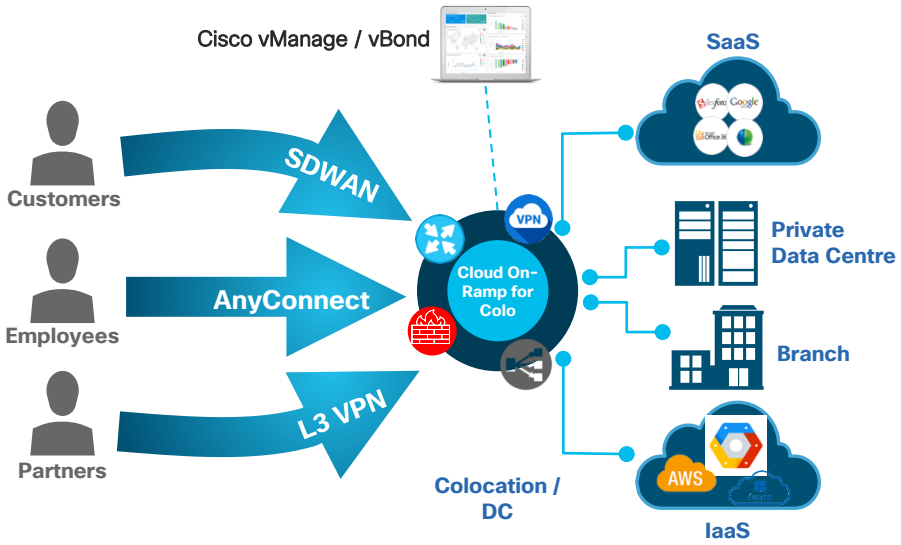


**Security**

Central policy enforcement

**Agility & Performance**

Rapid provisioning, change control and scale-out architecture via NFV fabric. Speed of software with the performance of hardware.

**Cost Savings**

Lower OpEx and CapEx through NFV. Reduce circuit costs and number of circuits.

Turn-key orchestration and automation of enterprise WAN Service-Chains!

# Modern Hierarchical Global WAN Design

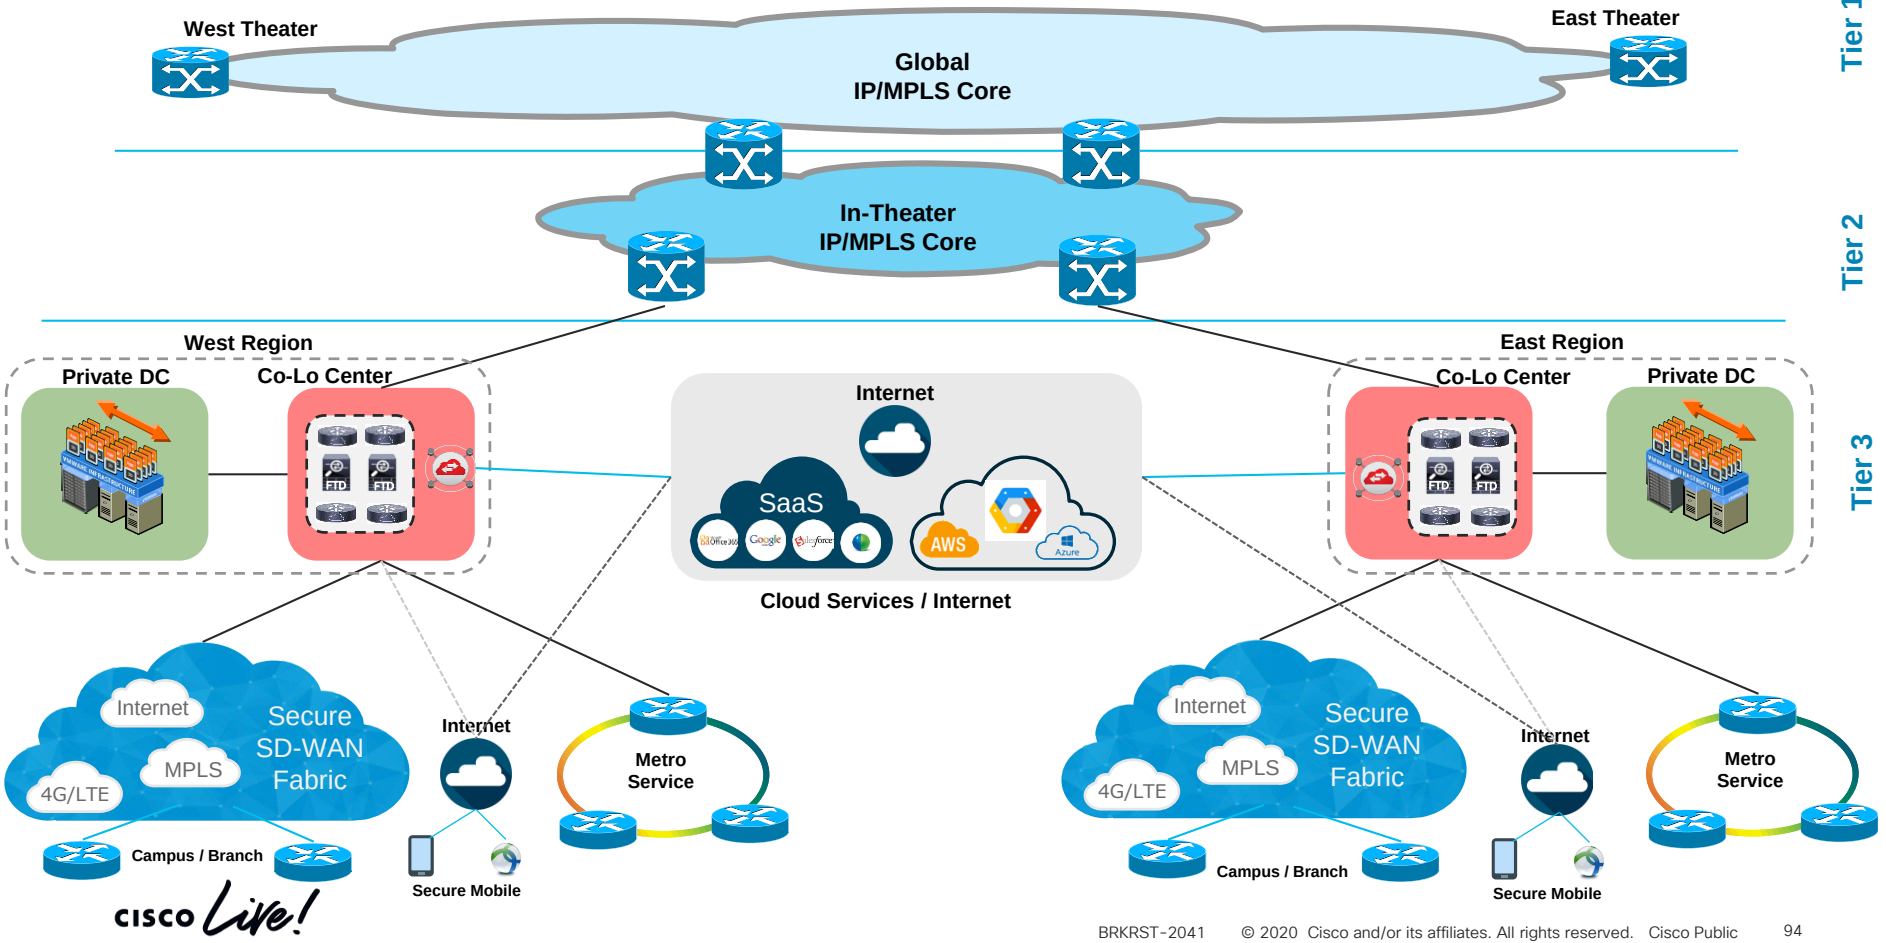

# Summary

# The WAN Technology Continuum

| Early Networking | Early-Mid 1990s | Mid 1990s-Late 2000s | Today |
|---|---|---|---|
| Flat/Bridged<br>Experimental Networks | Multiprotocol<br>Business Enabling | Large Scale<br>Mission Critical | Global Scale<br>IP Ubiquity<br>Cloud Connected |
| **Architectural Lessons** | **Architectural Lessons** | **Architectural Lessons** | **Planning** |
| Protocols required for Scale & Restoration | Route First, Bridge only if Must | Redundancy<br>Build to Scale | ? |

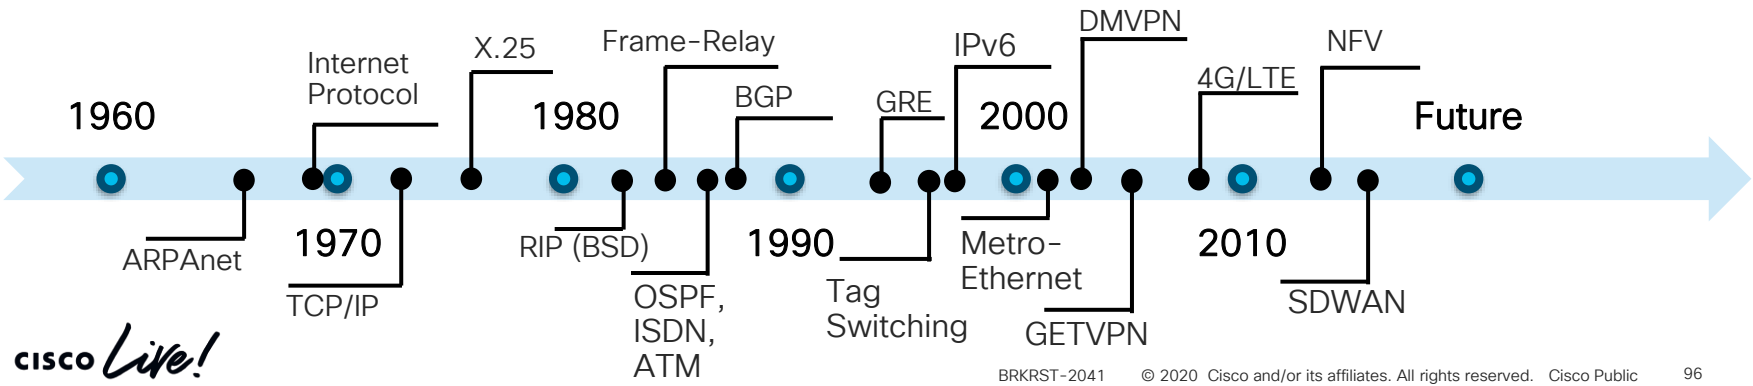

# The WAN Technology Continuum

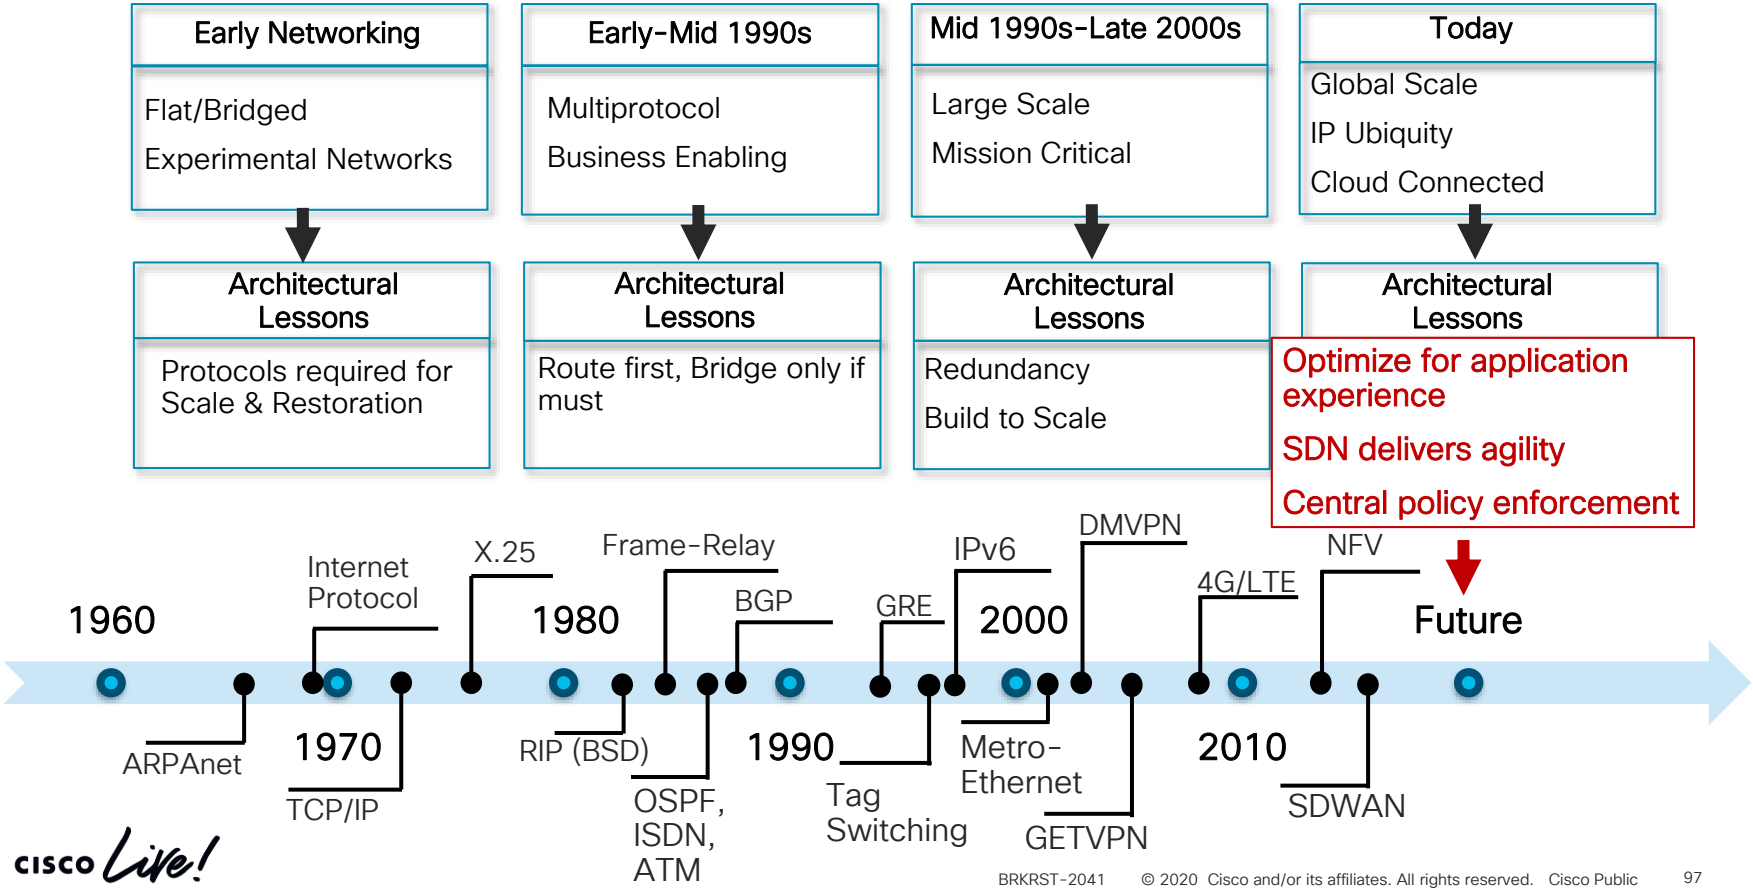

# The WAN of Yesterday, Today and Tomorrow

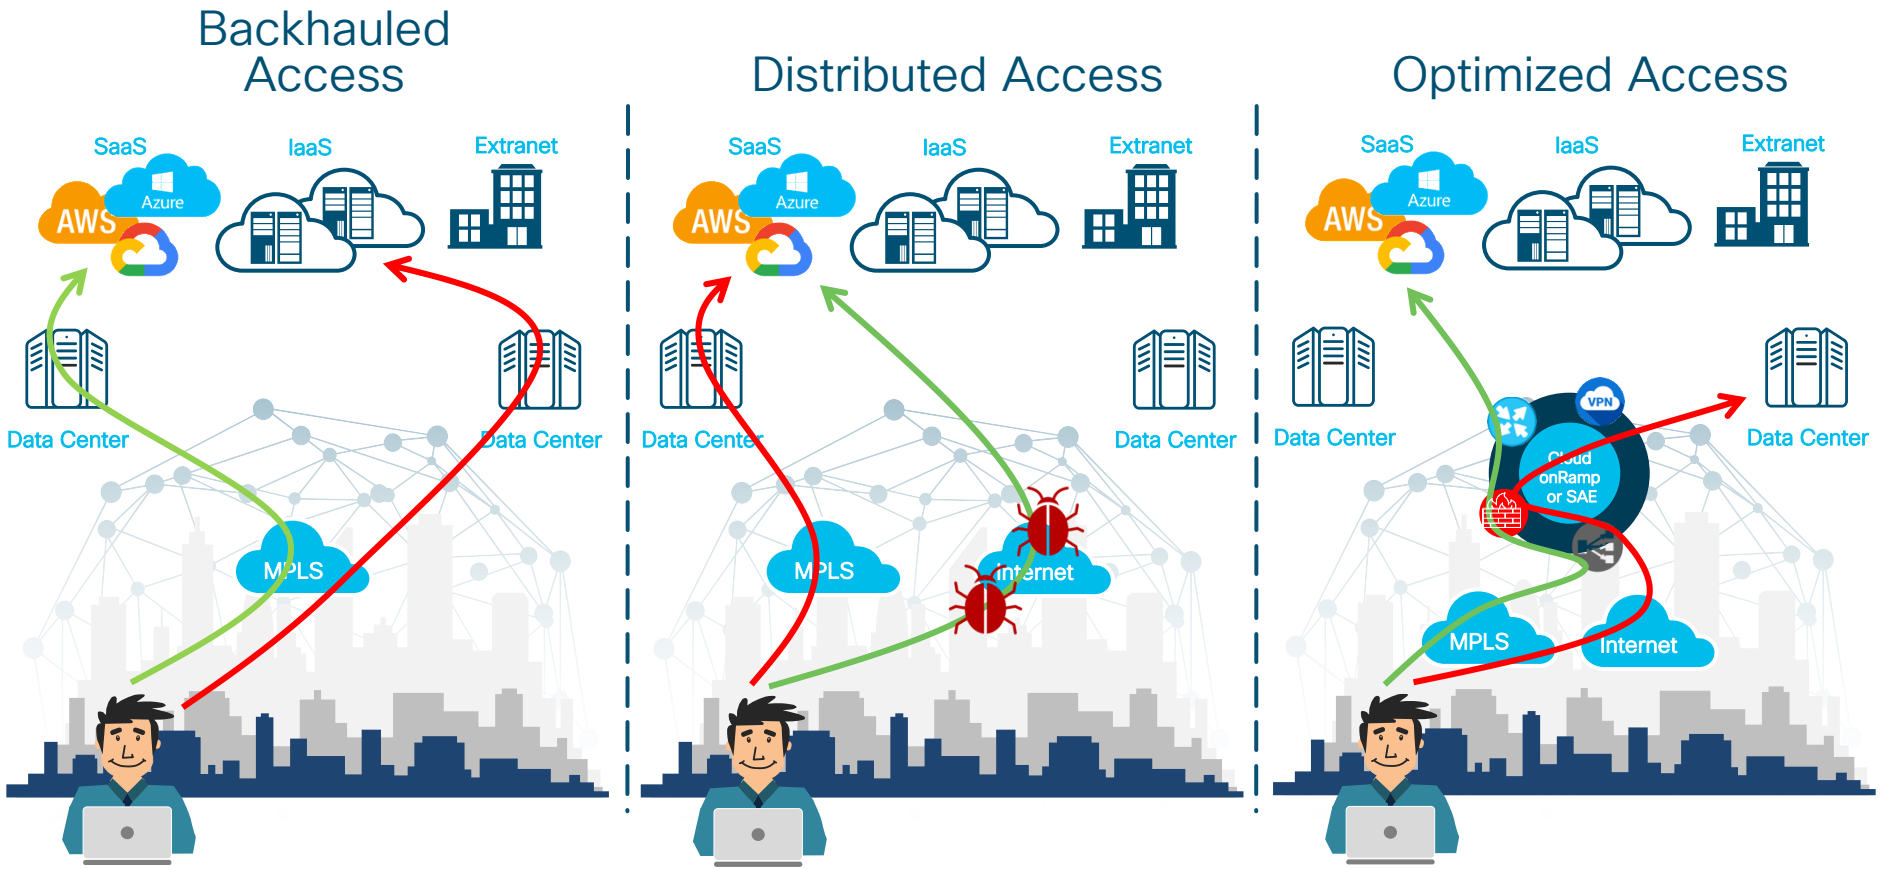

# Modern Hierarchical Global WAN Design

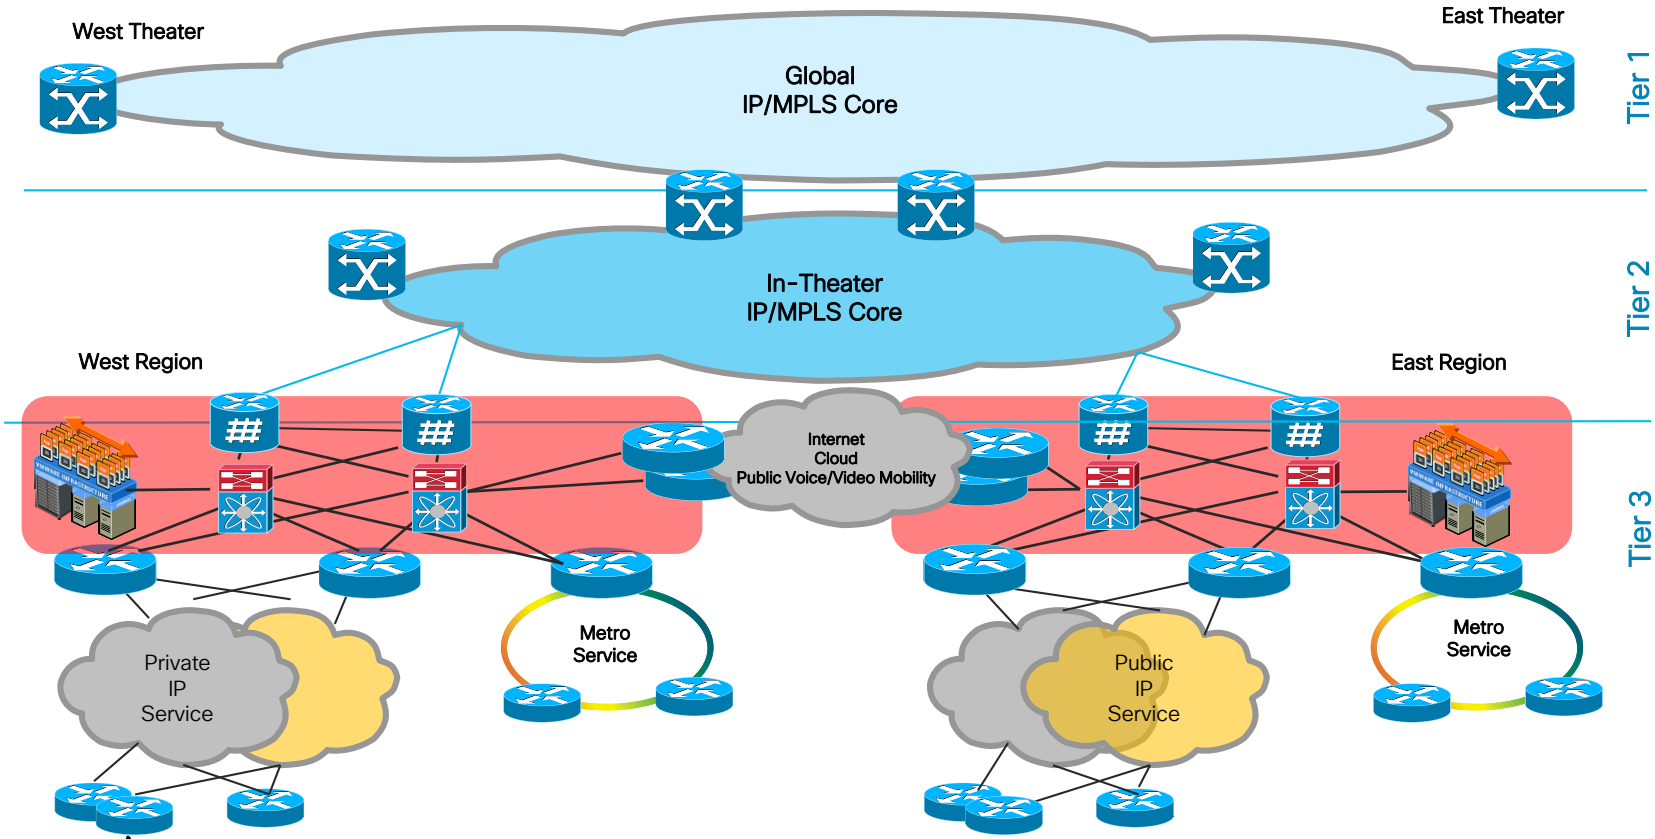

# Modern Hierarchical Global WAN Design

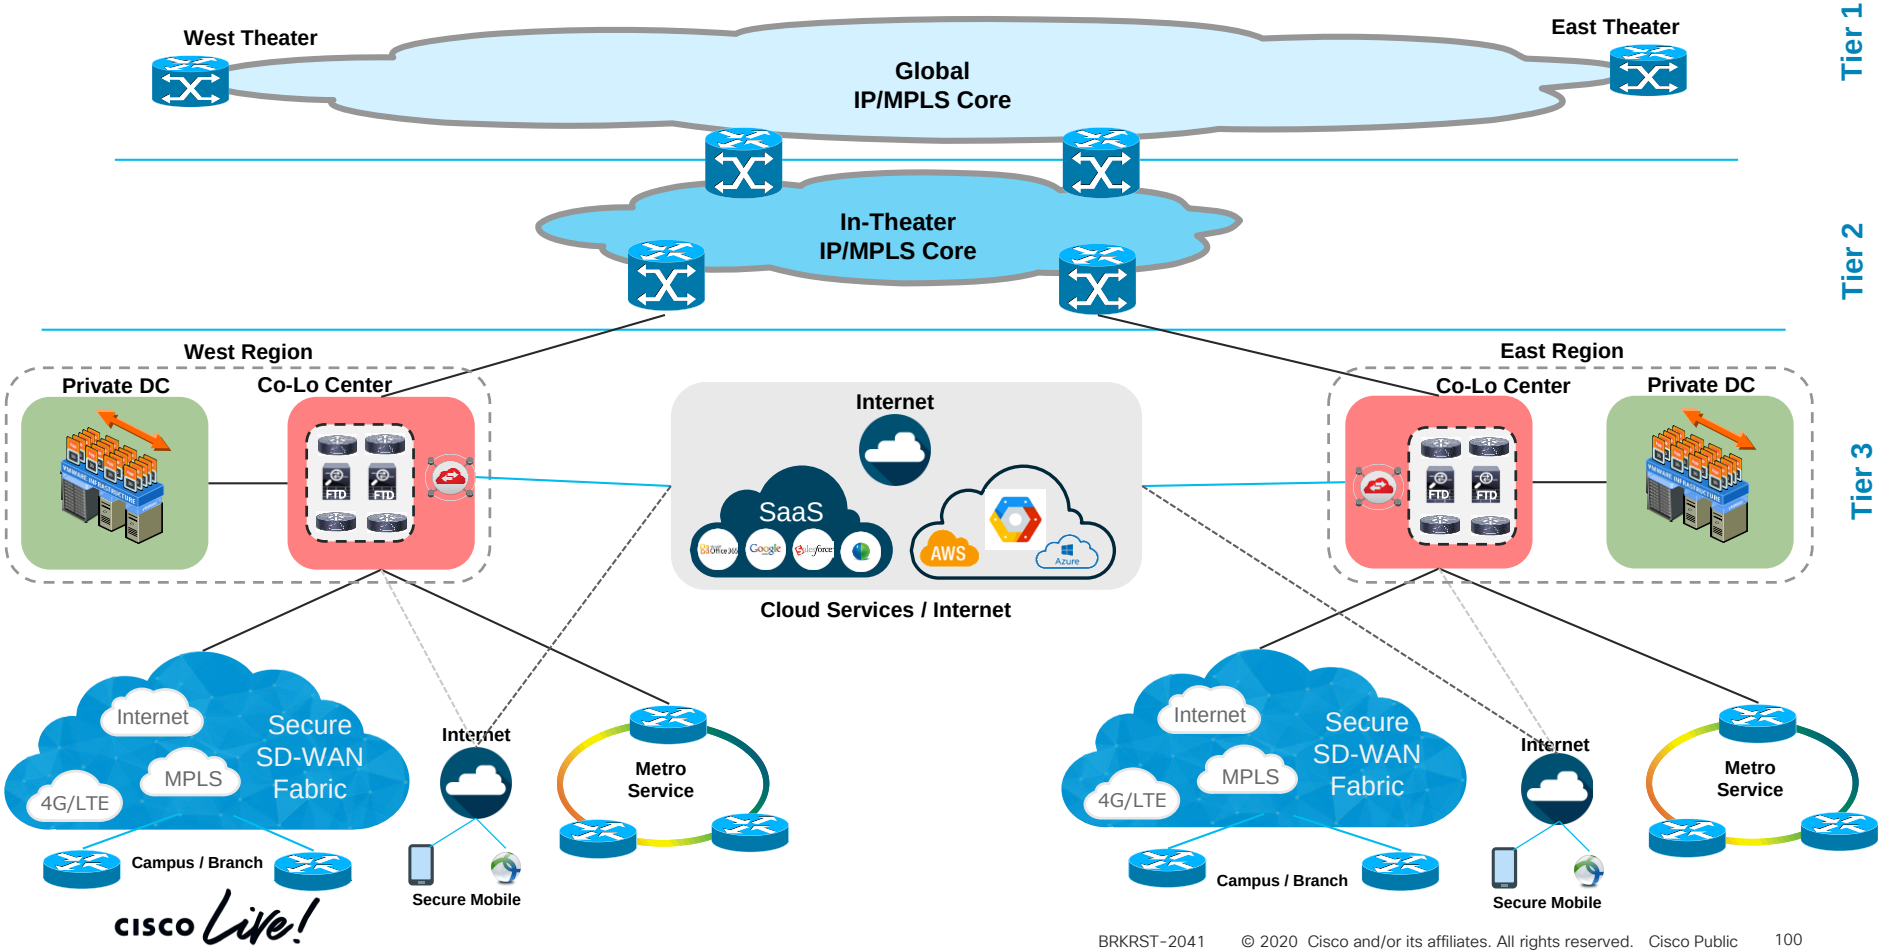

# WAN Architectures and Design Principles
## Key Takeaways

- The goal is for a simple, modular, hierarchical, structured design
- Business, technical, and physical requirements and constraints must all be considered
- Desired WAN availability and services have design implications
- Evolving technology is driving new WAN designs
- Leveraging Internet, Cloud, and CoLo now fundamental

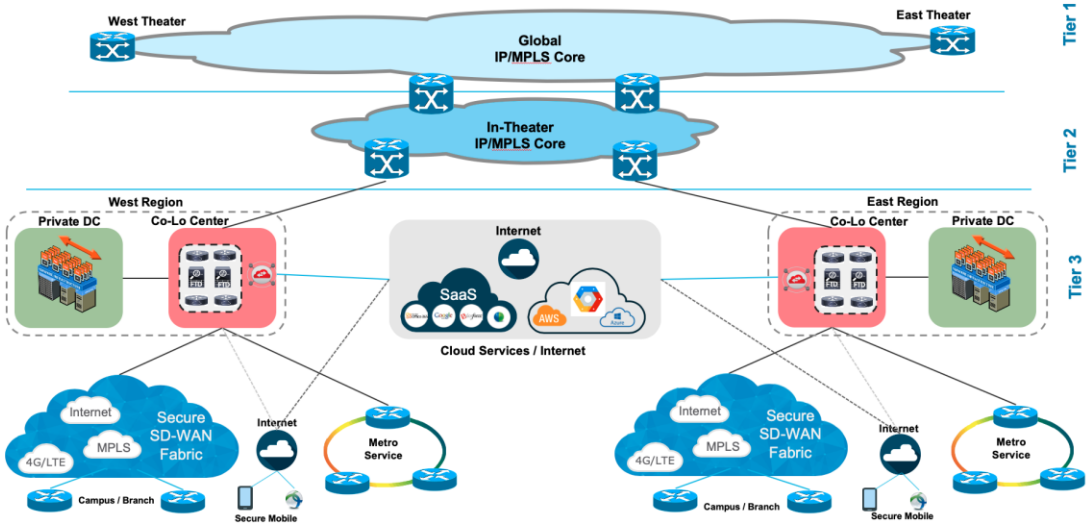

One final time, the Main Message:

*Foundational Design is key to WAN Architecture*

TECCRS-2014
SD-WAN Technical Deep Dive — 8 Hours

TECRST – 2191
SD-WAN design, deploy and best practices — 4 Hours

TECCRS-3006
ENFV Deep Dive and Hands on Lab — 8 Hours

**Cisco SD-WAN**

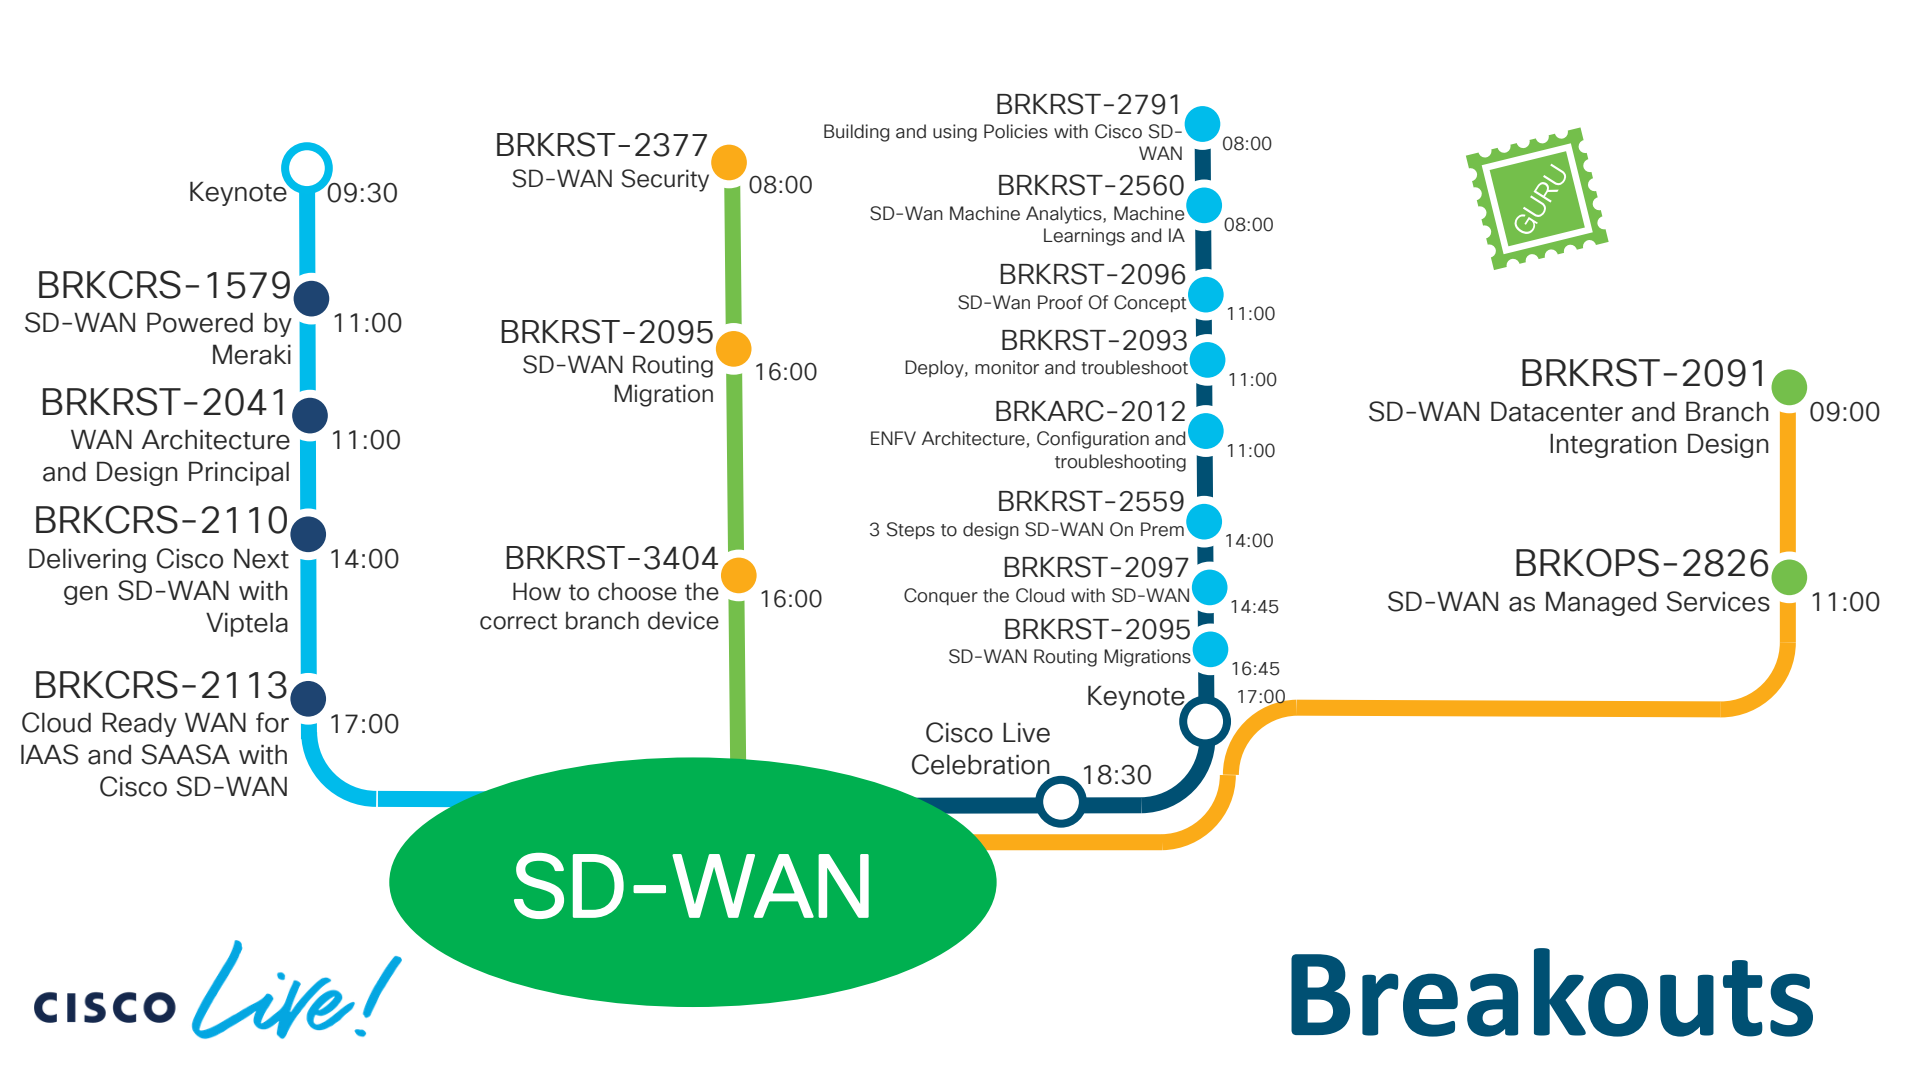
Keynote 09:30
BRKCRS-1579
SD-WAN Powered by Meraki
11:00
BRKRST-2041
WAN Architecture and Design Principal
11:00
BRKCRS-2110
Delivering Cisco Next gen SD-WAN with Viptela
14:00
BRKCRS-2113
Cloud Ready WAN for IAAS and SAASA with Cisco SD-WAN
17:00
BRKRST-2377
SD-WAN Security
08:00
BRKRST-2095
SD-WAN Routing Migration
16:00
BRKRST-3404
How to choose the correct branch device
16:00
BRKRST-2791
Building and using Policies with Cisco SD-WAN
08:00
BRKRST-2560
SD-Wan Machine Analytics, Machine Learnings and IA
08:00
BRKRST-2096
SD-Wan Proof Of Concept
11:00
BRKRST-2093
Deploy, monitor and troubleshoot
11:00
BRKARC-2012
ENFV Architecture, Configuration and troubleshooting
11:00
BRKRST-2559
3 Steps to design SD-WAN On Prem
14:00
BRKRST-2097
Conquer the Cloud with SD-WAN
14:45
BRKRST-2095
SD-WAN Routing Migrations
16:45
Keynote
17:00
Cisco Live Celebration
18:30
GURU
BRKRST-2091
SD-WAN Datacenter and Branch Integration Design
09:00
BRKOPS-2826
SD-WAN as Managed Services
11:00
SD-WAN
cisco Live!
Breakouts

# Call to Action

As you leave ask yourself these three questions:

- Is it a simple design?
- What are the critical business requirements?
- Are you leveraging the available technology?

# Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

cisco Live!

# Continue your education

Demos in the Cisco campus

Walk-in self-paced labs

Meet the engineer 1:1 meetings

Related sessions

# Thank you

cisco Live!

You make **possible**