

SD-WAN

INTERVIEW QUESTIONS GUIDE

One Step Closer Towards Your Dream Job...

Q1. WHAT IS THE CISCO SD-WAN SOLUTION?

Traditional Wide Area Networks (WAN) was designed using MPLS for connectivity where majority of branch office traffic flows within an enterprise's intranet boundary. However New Cloud Applications (SaaS) like Microsoft Office 365 and Salesforce.com, and Public Cloud Services (IaaS) like Amazon Web Services (AWS) and Azure are changing traffic patterns.

Today, majority Enterprise Traffic flows to Public Clouds and the Internet. This change creates new requirements for security, application performance, cloud connectivity, WAN Management, and operations. Cisco SD-WAN offers a new way to manage and operate WAN Infrastructure. Cisco SD-WAN is a cloud based solution that delivers a secure, flexible, and rich services architecture.

Q2. WHAT ARE THE KEY BENEFITS OFFERED BY CISCO SD-WAN?

Better User Experience – Deploy applications in minutes on any platform and with consistent user experience.

Greater Agility – Simplify deployment and operation of your WAN and get faster performance using less bandwidth. Deploy your WAN over any type of connections like MPLS, Internet, or 4G LTE.

Threat Centric Security – Securely connect users to applications in minutes and protect your data from WAN Edge to the Cloud..

Q3. WHICH PROBLEM CAN A CISCO SD-WAN OVERCOME?

Here are some problems which Cisco SD-WAN Solutions can offer:

- Establish a transport independent WAN for high diversity and low cost
- Meet Service Level Agreements (SLAs) for business critical and real time applications
- Provide End-to-End Segmentation for protecting critical enterprise compute resources
- Extend seamlessly into Public Cloud provide Optimal User Experience for SaaS and IaaS Applications

Q4. WHICH SECTORS AND INDUSTRIES HAVE DEPLOYED THE CISCO SD-WAN SOLUTIONS?

Cisco has one of the most widely deployed enterprise-grade SD-WAN solutions within the industry. Large deployments have been made in sectors like retail, healthcare, financial services, energy, and many more. The solution is deployed across Fortune 2000 enterprises with thousands of production sites in major industries including healthcare, manufacturing, retails, energy, oil and gas, insurance, finance, government, logistics, and distribution as some examples.

Q5. HOW DO YOU MANAGE AND OPERATE CISCO SD-WAN?

Cisco SD-WAN is a centrally managed, orchestrated, and operated solution with a cloud-hosted Cisco vManage GUI management and provisioning platform, vSmart controller, and vBond orchestration layer at the heart of the solution.

vSmart Controllers are the centralized brain of the solution that implements policies and connectivity between SD-WAN branches.

vBond Orchestrator facilitate the initial bring-up by performing authentication and authorization of all elements into the network.

Cisco vManage manages the entire solution. Cisco's GUI based centralized management and provisioning platform for day 0, day 1, and day n+ for the entire Cisco SD-WAN infrastructure.

Q6. WHAT ARE vSMART CONTROLLERS?

vSmart Controllers are the centralized brain of the solution that implements policies and connectivity between SD-WAN branches. The centralized policy engine in Cisco vSmart Controllers provides policy constructs to manipulate routing information, access control, segmentation, extranets, and service chaining.

Q7. WHAT ARE vBOND ORCHETRATORS?

The **vBond Orchestrator** facilitates the initial bring-up by performing authentication and authorization of all elements into the network. Cisco vBond Orchestrator also provides the information on how each of the components connects to other components. Cisco vBond Orchestrator plays an important role in facilitating Cisco SD-WAN devices that sit behind the Network Address Translation (NAT) to communicate with the network.

Q8. WHAT IS THE CISCO vMANAGE?

Cisco vManage maintains the entire solution. Cisco's GUI based centralized management and provisioning platform for day 0, day 1, and day n+ for the entire Cisco SD-WAN infrastructure. You can login to the Cisco vManage dashboard to centrally manage the WAN. Cisco vManage provides the ability to manage all aspects of the WAN—from provisioning, monitoring, and upgrading routers to application visibility and troubleshooting the WAN.

Q9. HOW IS CISCO SD-WAN DEPLOYED AT BRANCH OFFICES AND DATA CENTER NETWORK OR REGIONAL HUBS?

Branch office and regional data center hub sites can be deployed and connected using either virtual or physical secure routers. Enterprise customers and service providers can gain rich services like WAN optimization and firewall or basic WAN connectivity for physical or virtual platforms across the branch, WAN, or cloud as follows:

Physical

- Branch-Cisco vEdge Series Routers
- Branch-Cisco 1000 Series Integrated Services Routers (ISR)
- Branch-Cisco 4000 Series ISR
- Branch/Regional Hub/Data Center -Cisco ASR 1000 Series Aggregation Services Routers (ASR)

Virtual

- SD-Branch-Cisco
- 5000 Series Enterprise Network Compute System (ENCS) and Integrated Services Virtual Router (ISRv)
- Network-Hub/colocation/data center -Cisco Cloud Services Platform 5000 and Cloud Services Router 1000V (CSR1000V)

Public Cloud (IaaS)

- Amazon Web Services
- Microsoft Azure
- Google Cloud Platform

Q10. IS THE CISCO SD-WAN SOLUTION SECURED?

Cisco is bridging networking and security together like no other vendor. With Cisco SD-WAN we provide highly effective and scalable security that is easy to manage, deploy, and maintain, empowering businesses to adopt the latest cloud services with confidence.

Cisco SD-WAN is built based on the zero-trust model and multilayer security encrypts all data for protection from the WAN edge to the cloud. All of the Cisco SD-WAN components mutually authenticate each other and all of the edge devices are authorized before they are allowed into the network. Every packet across data plane, control plane, and management plane that flows through the network is encrypted using Secure Socket Layer (SSL) and IP Security (IPsec) technologies. The Cisco SD-WAN Solution has differentiated integrated capabilities to build a large-scale IPsec network across tens of thousands of branches.



INDIA'S MOST TRUSTED NETWORKING TRAINING COMPANY

Q11. DOES CISCO SD-WAN SOLUTION SUPPORT NETWORK SEGMENTATION AND WHAT ARE THE BENEFITS?

Yes, the Cisco SD-WAN solution supports network segmentation. Segmentation provides secure logical isolation on the SD-WAN network, where each segment is defined as a separate VPN and controlled centrally by access-control policies.

Some of the Benefits of Segmentation:

- Increased security-Isolate your network from outside attackers and create secure separation within multiple application segments.
- Acquisitions can be integrated on the parent network and yet kept separate. Policies control what applications the acquired company can access.
- Guest Wi-Fi can be maintained on a separate, low-priority segment and offloaded onto the Internet at the closest exit points.
- Business partners can each be defined on a separate segment, or on a collective business-partner network segment.
- Policies control the access of business partners to data-center applications.

Q12. WHAT ARE THE SD-WAN SECURITY CAPABILITIES AND WHICH PLATFORMS SUPPORT SD-WAN SECURITY?

Cisco SD-WAN Security capabilities include an application-aware enterprise firewall, intrusion prevention, DNS Layer Enforcement (Cisco Umbrella), and URL filtering. Cisco SD-WAN reduces complexity by having a single management interface (vManage) for both the network and security.

Platform	Enterprise firewall	Enterprise firewall application-awareness	Intrusion prevention system	URL filtering	DNS web layer security (Umbrella)
Cisco vEdge 100, 1000, 2000, and 5000 series	Yes	DPI using Qosmos	X	X	Yes
Cisco CSR	Yes	Yes	Yes	Yes	Yes
Cisco ISRV/ ENCS 5000 series	Yes	Yes	Yes	Yes	Yes
Cisco 4451, 4351, 4331, 4321, and 4221ISRs	Yes	Yes	Yes	Yes	Yes
Cisco 1111X-8PISR	Yes	Yes	Yes	Yes	Yes
Cisco 1111-4P, 1111-8P, 1116-4P, and 1117-4PISRs	Yes	Yes	X	X	Yes
Cisco ASR 1001-HX, 1002-HX, 1001-x, and 1002-X=	Yes	Yes	X	X	Yes

Q13. CAN THE CISCO SD-WAN SOLUTION PROVIDE OPTIMIZATION FOR IaaS and SaaS PLATFORMS LIKE AWS, MICROSOFT AZURE & OFFICE 365, GOOGLE, CISCO WEBEX, SALESFORCE.com, ETC?

The Cisco SD-WAN fabric connects users at the branch to applications in the cloud in a seamless, secure, and reliable fashion. Cisco delivers this comprehensive capability for Infrastructure-as-a-Service and Software-as-a-Service (IaaS/SaaS) applications with Cisco Cloud OnRamp, and is currently available with vEdge series platform SD-WAN solutions.

With Cloud OnRamp, the SD-WAN fabric continuously measures the performance of a designated application through all permissible paths from a branch (i.e. MPLS, Internet, and 4G LTE). The Cisco SD-WAN fabric automatically makes real-time decisions to choose the best-performing path between the end users at a remote branch and the cloud application. Enterprises and service providers have the flexibility to deploy this capability in multiple ways and according to business needs and security requirements.

Q14. HOW IS THE CISCO SD-WAN SOLUTION ORDERED?

Cisco SD-WAN software is included with each vEdge series routing device and platform and can be enabled on some 1000 and 4000 ISRs, ASR 1000Series Routers, the Cisco ISRv on the ENCS 5000, and with the Cisco CSR 1000V on the Cloud Services Platform 5000 Series with the latest Cisco IOS-XE software.

Each device requires a Subscription License (three or five years) for Cisco SD-WAN software. The License Fee is charged per branch device. The License Fee is dependent on service bandwidth and feature content, with a single set of software licenses that includes security and access to ongoing innovation and the latest threat Intelligence License Bundles include:

Cisco DNA Essentials – Includes Basic Connectivity, Security, and Application Visibility

Cisco DNA Advantage – Includes everything in Cisco DNA Essentials plus Flexible Connectivity, Advanced Security, and Enhanced Application Visibility

Cisco DNA Premier (replaces Cisco ONE Cisco DNA Advantage) – Includes everything in Cisco DNA Essentials and Cisco DNA Advantage plus Advanced Application Policy and Experience using Analytics and Assurance, and WAN Optimization.

The subscription price of SD-WAN software includes cloud-hosted vManage, vSmart and vBond devices, 24-hour daily Cisco SD-WAN support, next-day hardware replacement for Cisco SD-WAN platforms, software upgrades on all components, and the cost of hosting Cisco SD-WAN controllers in the Cisco SD-WAN cloud.

Q15. DOES CISCO SD-WAN SOLUTION SUPPORT MULTI-TENANCY?

Yes, a Service Provider can manage multiple customers, called tenants, from vManage that is running in multitenant mode. All tenants share a single vBond orchestrator. All tenants share the service provider's domain name, with each tenant having a subdomain name to identify the tenant.

For example, the service provider, fruit.com, might manage the tenants, mango (mango.fruit.com) and plum (plum.fruit.com). For each tenant, you configure one or more vSmart controllers and vEdge routers the same way that you configure these devices on a single-tenant vManage Network Management System (NMS).

Q16. IS CISCO'S SD-WAN SOLUTION PROGRAMMABLE AND DOES IT SUPPORT APIs?

Yes, the Cisco SD-WAN Solution is Open and Programmable and with open APIs, Cisco SD-WAN provides service providers and partners the opportunity to create new and unique services, including operational and business support systems.

With Cisco SD-WAN you can access the available Representational State Transfer (REST) APIs, create API calls, obtain device and interface information using code, pass parameters and write applications, and work on innovative solutions.

As part of the SD-WAN developer resources availability and learning content, there are two additional resources that are a great value added service for developers:

- DevNet Ecosystem Exchange
- DevNet Code Exchange

Q17. WHAT IS DEV-NET ECOSYSTEM EXCHANGE?

DevNet Ecosystem Exchange makes it easy to find and share an application or solution built for Cisco platforms. Business leaders and developers alike can use this online portal to discover partner solutions that span all Cisco platforms and products. Currently, this central repository for developers contains over 1300 solutions.

Q18. WHAT IS DEV-NET CODE EXCHANGE?

DevNet Code Exchange gives developers a place to access and share software to quickly build next-generation applications and workflow integrations. It offers a curated list of sample code, adaptors, tools, and SDKs available on GitHub and written by Cisco and the DevNet Community. Code Exchange spans Cisco's entire portfolio and is organized according to Cisco platform and product areas.

Q19. WHAT ARE THE HIGHLIGHTS OF CISCO ROUTING SOFTWARE SUBSCRIPTION?

As part of the intent-based networking journey, Cisco is strengthening its SD-WAN software portfolio with the following new capabilities:

- Advanced Security features that deliver the right security in the right place at the right time
- Software-as-a-Service (SaaS) optimization features that make Microsoft Office 365 run faster
- Unified Access Security and Multifactor Authentication

Q20. WHAT ARE THE BENEFITS OF SD-WAN AND ROUTING SUBSCRIPTION OFFER?

The new licensing offers bring these customer benefits:

Latest Innovations through Simple Subscription Tiers:

Simplicity in purchasing and using via Cisco DNA Essentials, Advantage, or Premier Software Suites

Management Flexibility

Choice of Cloud or On-premises Management

Availability Across the Routing Stack:

Across the Cisco ASR 1000 Aggregation Series Service Routers, Cisco 1000 Integration Series Routers (ISR 1000) and Cisco 4000 Integration Series Routers (ISR 4000), Cisco Cloud Services Routers 1000 Series (CSR 1000V), Cisco 5000 Series Enterprise Compute System (ENCS 5000), and Cisco vEdge Routers

Software License Portability:

Between generations of hardware and between product families (for example, vEdge to ISR, across ISR platforms); ability to renew and scale what you want, when you want.



Baldev Singh
CCIE SECURITY #37094



Saurabh Yadav
Triple CCIE (R&S, Sec, SP) #46962



Sudhanshu Bhat
CCIE VOICE #41212



Surendra Singh
CCIE R&S #60346

**GET TRAINED BY
CERTIFIED AND
WORLD CLASS TRAINERS**

Q21. WHY SHOULD ONE OPT FOR SDN?

Application Experience

- Predictable SLA on all critical enterprise applications
- Application aware policies with real time enforcement around network problems
- Multiple hybrid active-active links for all scenarios

Best in Class Integrated Security

- Zero-trust foundation with authentication and encryption
- Segmentation to isolate and protect critical assets with cloud, partner networks, guest wireless etc.
- Enterprise firewall, IPS, AMP, DNS-layer enforcement, URL filtering, A/V and SSL decryption proxy integrated into SD-WAN

Cloud Optimized

- Seamlessly extend the WAN to multiple public clouds
- Real-time optimized performance for Office365, Salesforce and other major SaaS applications
- Optimized workflows for AWS and Azure

Q21. WHY SHOULD ONE OPT FOR SDN? (continued...)

Operational Simplification

- Single management dashboard for configuration and management of WAN, cloud and security Template-based, Zero
- touch provisioning for all locations
- Full automation with RESTful Integration into existing tools

Rich Analytics

- Granular Visibility of applications and infrastructure enables rapid failure correlation and mitigation
- Sophisticated forecasting and what-if analysis for effective resource planning
- Insightful recommendations for policy changes based on traffic patterns

Q22. WHAT IS THE LATEST SOFTWARE RELEASE VERSION FOR CISCO XE SD-WAN IMAGE SUPPORTED ON CISCO 1000 & 4000 SERIES ISRs, ASR 1000 SERIES & 5000 SERIES ENCS PLATFORM?

the latest software release version for the Cisco IOS XE SD-WAN image supported on the Cisco 1000 and 4000 Series ISRs, ASR 1000 Series, and 5000 Series ENCS platforms is:

Cisco IOS XE SD-WAN Software Release 16.11.1.

Q23. HOW DOES A CUSTOMER CHOOSE AN OPTIMAL ROUTING OR SD-WAN SOLUTION?

The customer can consider these five steps while choosing an Optimal Routing or SD-WAN Solution:

Step 1: Identify license tier

Step 2: Select the bandwidth

Step 3: Pick the license term

Step 4: Choose on-premises or cloud managed

Step 5: Determine platform for future scale

Q24. WHAT IS VIPTELA SD-WAN?

Viptela was formed by the triumvirate of Alcatel Lucent, Cisco, and Juniper Networks network architects working on Software Defined Networking (SDN) at the WAN level. The resulting solution Viptela Secure Extensible Network (SEN) tackled architecture transformation.

Q25. HOW DID VIPTELA SD-WAN HELP IN ARCHITECTURAL TRANSFORMATION?

Transport Independence:

Viptela SD-WAN disaggregates the service from the physical network, building an overlay on top of whatever forms of connectivity an organization has. This enables transport independence, not tied to any particular form of service.

Security At Routing Scale:

Viptela SD-WAN provides security in the form of encryption and device authentication. The founders applied their expertise in routing protocols to develop a solution that provides encryption and security from any-to-any perspective. The Viptela router can connect all entities and automatically route traffic between those as if they were on one seamless VPN connection.

Network-Wide Segmentation:

Because Viptela technology enables the overlay, the company can segment the network on an end-to-end basis. The Viptela SD-WAN allows an enterprise to build multiple logical topologies any way they want, and each of these different segments of the network can have different encryption schemes.

Q25. HOW DID VIPTELA SD-WAN HELP IN ARCHITECTURAL TRANSFORMATION? (continued...)

Enforce Policy and Business Logic Centrally:

Each network location enforces the policies of a specific location, but all of the locations are influenced by the centralized controller. If necessary, an organization can have multiple controllers to meet resiliency requirements.

Insert Layer 4-7 Services on Demand:

Viptela SD-WAN enables Layer 4–7 network services to be advertised, enabling organizations to spin up any third-party service on the network and connect it to the Viptela overlay. Then anyone wanting to use those services sets a centralized policy to direct traffic to that particular location.

Q26. WHAT ARE THE COMPONENTS OF VIPTELA SD-WAN?

Here are the Components of Viptela SD-WAN:

vSmart Controller

Central management of routing, policy, security, segmentation, and authentication of devices

vManage

A centralized dashboard for configuration and management

vEdge Routers

Full-featured IP routers that perform standard functions such as BGP, OSPF, ACLs, QoS, and various routing policies in addition to the overlay communication

vBond Orchestrator

Initial authentication and authorization of all elements into the network; provides the information on how each of the components connects to other components.

Q27. WHAT IS DOMAIN ID?

A domain is a logical grouping of vEdge routers and vSmart controllers that demarcates the span of control for the vSmart controllers. Each domain is identified by a unique integer, called the domain ID. Currently, you can configure only one domain in a Viptela Overlay Network.

Within a domain, vEdge routers can connect only with the vSmart controllers in their own domain. The vBond orchestrator is aware of which vSmart controllers are in which domain, so that when new vEdge routers come up, the vBond Orchestrator can point those routers to the vSmart controllers in the proper domain. However, the vBond orchestrator is never a member of a domain.

Within a domain there is full synchronization of routing information among the vSmart controllers and vEdge routers, and there is scope for route aggregation and summarization. An organization can divide up its network into domains to serve desired business purposes. For example, domains can correspond to a large geographic area or to data centers so that each data center and the branches for which it is responsible are contained within a single domain.

Q28. WHAT ARE OMP ROUTES?

Prefixes that establish reachability between end points that use the OMP-orchestrated transport network. OMP routes can represent services in a central data center, services at a branch office, or collections of hosts and other end points in any location of the overlay network. OMP routes require and resolve into TLOCs for functional forwarding. In comparison with BGP, an OMP route is the equivalent of a prefix carried in any of the BGP AFI/SAFI fields.

Q29. WHAT ARE THE TYPES OF OMP ROUTES?

OMP Routes

Prefixes that establish reachability between end points that use the OMP-orchestrated transport network. OMP routes can represent services in a central data center, services at a branch office, or collections of hosts and other end points in any location of the overlay network. OMP routes require and resolve into TLOCs for functional forwarding. In comparison with BGP, an OMP route is the equivalent of a prefix carried in any of the BGP AFI/SAFI fields.

Transport Locations (TLOCs)

Identifiers that tie an OMP route to a physical location. The TLOC is the only entity of the OMP routing domain that is visible to the underlying network, and it must be reachable via routing in the underlying network. A TLOC can be directly reachable via an entry in the routing table of the physical network, or it must be represented by a prefix residing on the outside of a NAT device and must be included in the routing table. In comparison with BGP, the TLOC acts as the next hop for OMP routes.

Service Routes

Identifiers that tie an OMP route to a service in the network, specifying the location of the service in the network. Services include firewalls, Intrusion Detection Systems (IDPs), and load balancers.

Q30. WHAT IS SITE-ID?

A site is a particular physical location within the Viptela Overlay Network, such as a branch office, a data center, or a campus. Each site is identified by a unique integer, called a Site-ID. Each Viptela device at a site is identified by the same site ID. So within a data center, all the vSmart controllers and any vEdge routers are configured with the same site ID. A branch office or local site typically has a single vEdge router, but if a second one is present for redundancy, both routers are configured with the same site ID.



WORLD CLASS INFRASTRUCTURE

Q31. WHAT IS SYSTEM IP ADDRESS?

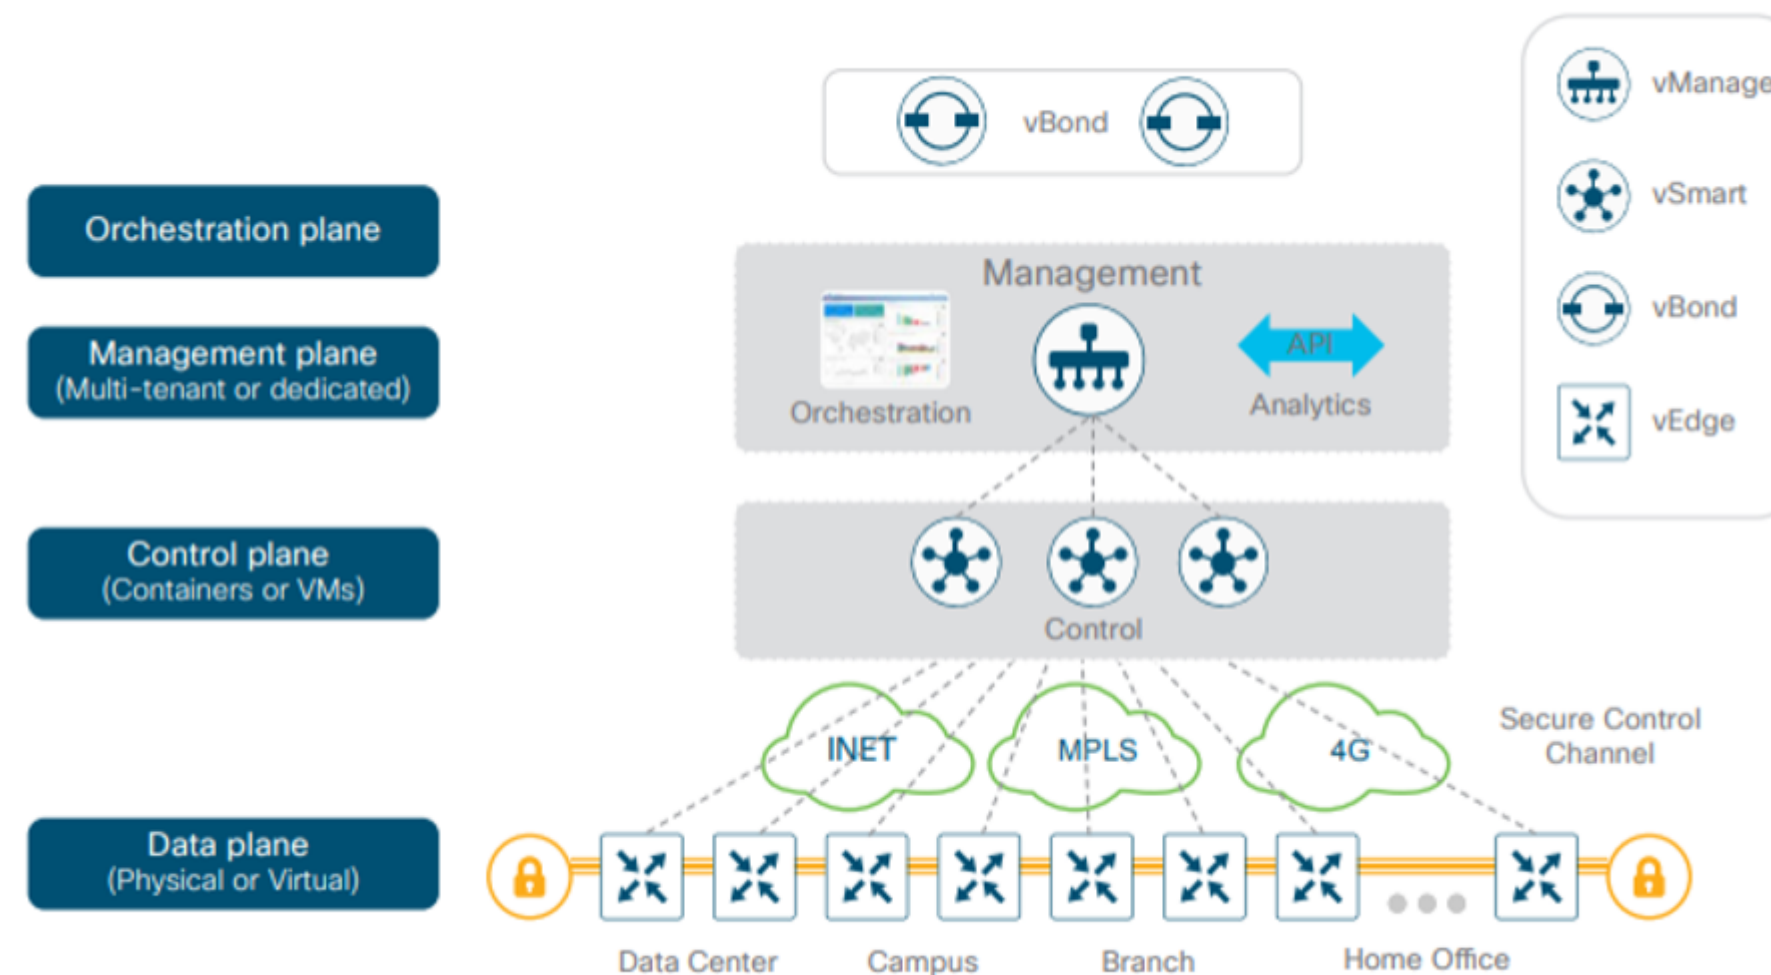
Each vEdge router and vSmart Controller is assigned a system IP address, which identifies the physical system independently of any interface addresses. This address is similar to the Router ID on a regular router. The system IP address provides permanent network overlay addresses for vEdge routers and vSmart controllers, and allows the physical interfaces to be renumbered as needed without affecting the reachability of the Viptela device. You write the system IP address as you would an IPv4 address, in decimal four-part dotted notation.

Q32. WHAT IS TRANSPORT LOCATION (TLOC)?


A TLOC, or transport location, identifies the physical interface where a vEdge router connects to the WAN Transport Network or to a NAT gateway. A TLOC is identified by a number of properties, the primary of which is an IP address–color pair, which can be written as the tuple {IP-address, color}. In this tuple, IP address is the System IP address and color is a fixed text string that identifies a VPN or traffic flow within a VPN. OMP advertised TLOCs using TLOC routes.

Q33. EXPLAIN THE CISCO SD-WAN SOLUTION ARCHITECTURE?

- The Cisco SD-WAN solution is comprised of separate orchestration, management, control, and data planes
- The orchestration plane assists in the automatic on boarding of the SD-WAN routers into the SD-WAN overlay
- The management plane is responsible for central configuration and monitoring
- The control plane builds and maintains the network topology and makes decisions on where traffic flows
- The data plane is responsible for forwarding packets based on decisions from the control plane



Q34. EXPLAIN THE SD-WAN TOPOLOGY?



This sample topology depicts two sites and two public Internet transports. The SD-WAN controllers, the two vSmart controllers, and the vBond orchestrator, along with the vManage management GUI that reside on the Internet, are reachable through either transport.

At each site, vEdge routers are used to directly connect to the available transports. Color is used to identify an individual WAN transport; different WAN transports are assigned different colors, such as MPLS, private1, biz-internet, metro-Ethernet, LTE, etc. The topology uses a color called biz-internet for one of the Internet transports and a color called public-internet for the other.

The vEdge routers form a Datagram Transport Layer Security (DTLS) or Transport Layer Security (TLS) control connection to the vSmart controllers and connect to both of the vSmart controllers over each transport. The vEdge routers securely connect to vEdge routers with IPsec tunnels at other sites over each transport. The Bidirectional Forwarding Detection (BFD) protocol is enabled by default and will run over each of these tunnels, detecting loss, latency, jitter, and path failures.

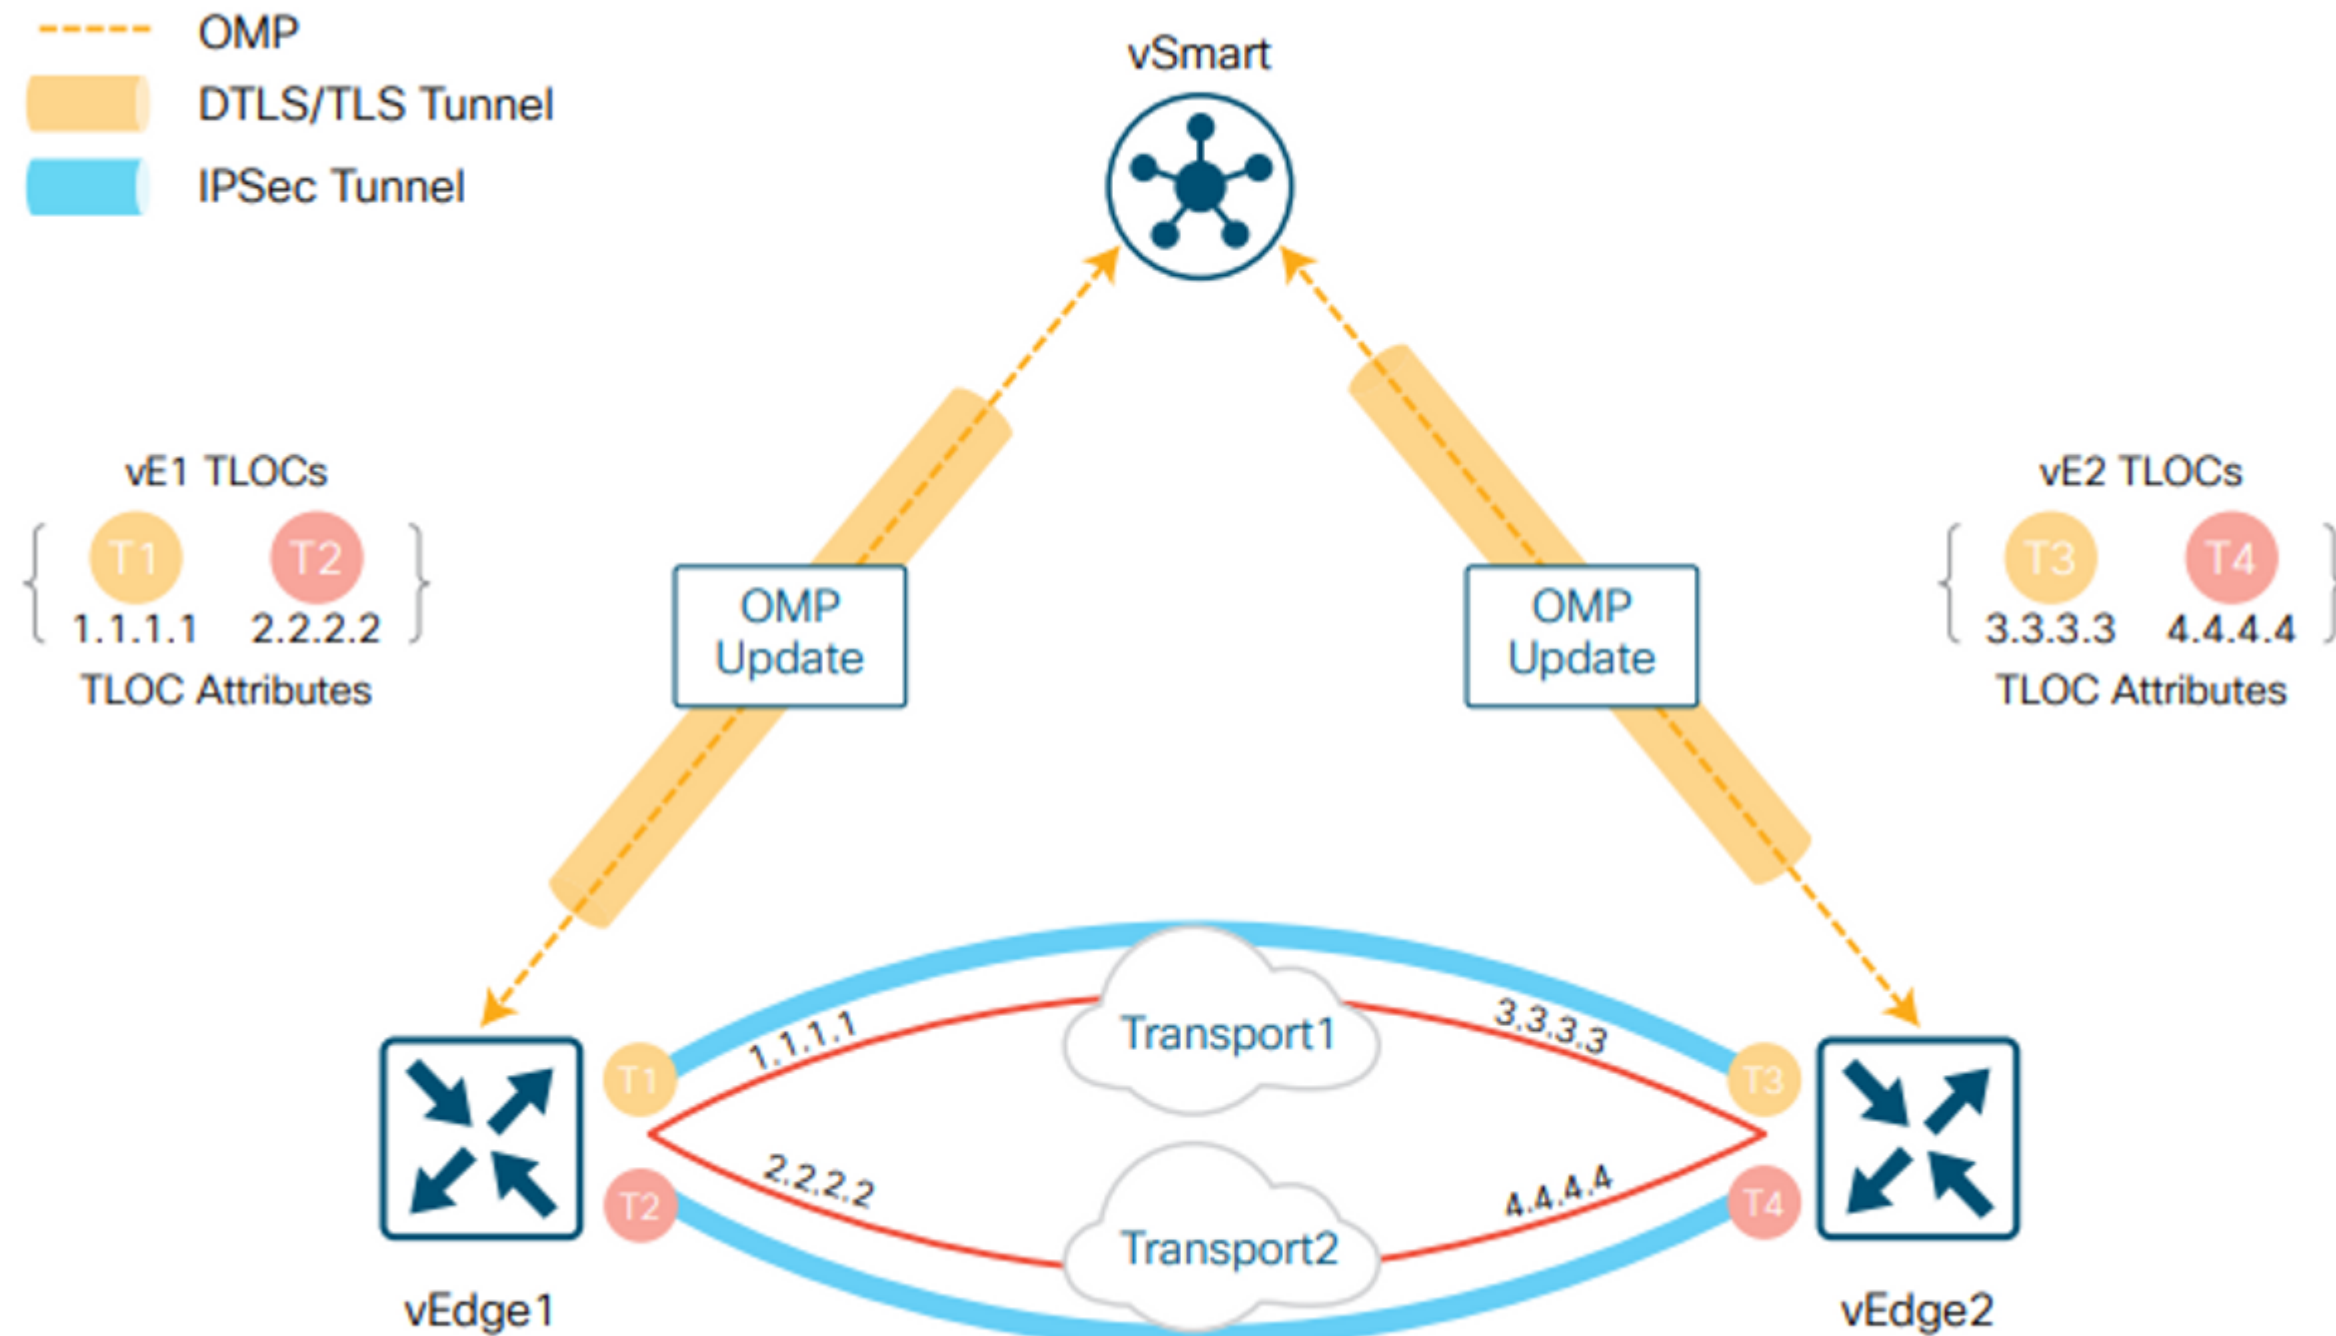
Q35. WHAT IS A COLOR?

On vEdge routers, the color attribute helps to identify an individual WAN transport tunnel. You cannot use the same color twice on a single vEdge router. Colors by themselves have significance. The colors metro-ethernet, mpls, and private1, private2, private3, private4, private5, and private6 are considered private colors. They are intended to be used for private networks or in places where you will have no NAT addressing of the transport IP endpoints, as the expectation is that there is no NAT between two endpoints of the same color.

When a vEdge router uses a private color, it will attempt to build IPSec tunnels to other vEdge routers using the native, private, underlay IP. The public colors are 3g, biz, internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, public-internet, red, and silver. With public colors, vEdge routers will try to build tunnels to the post-NAT IP address (if there is NAT involved).

If you are using a private color and need NAT to communicate to another private color, the carrier setting in the configuration dictates whether you use the private or public IP address. Using this setting, two private colors will establish a session when one or both are using NAT.

Q36. DRAW TLOC ROUTES AND OMP



Q37. WHAT IS A VIRTUAL PRIVATE NETWORK (VPN)?

In the SD-WAN overlay, Virtual Private Networks (VPNs) provide segmentation, much like Virtual Routing and Forwarding instances (VRFs) that many are already familiar with. Each VPN is isolated from one another and each have their own forwarding table. An interface or subinterface is explicitly configured under a single VPN and cannot be part of more than one VPN. Labels are used in OMP route attributes and in the packet encapsulation, which identifies the VPN a packet belongs to.

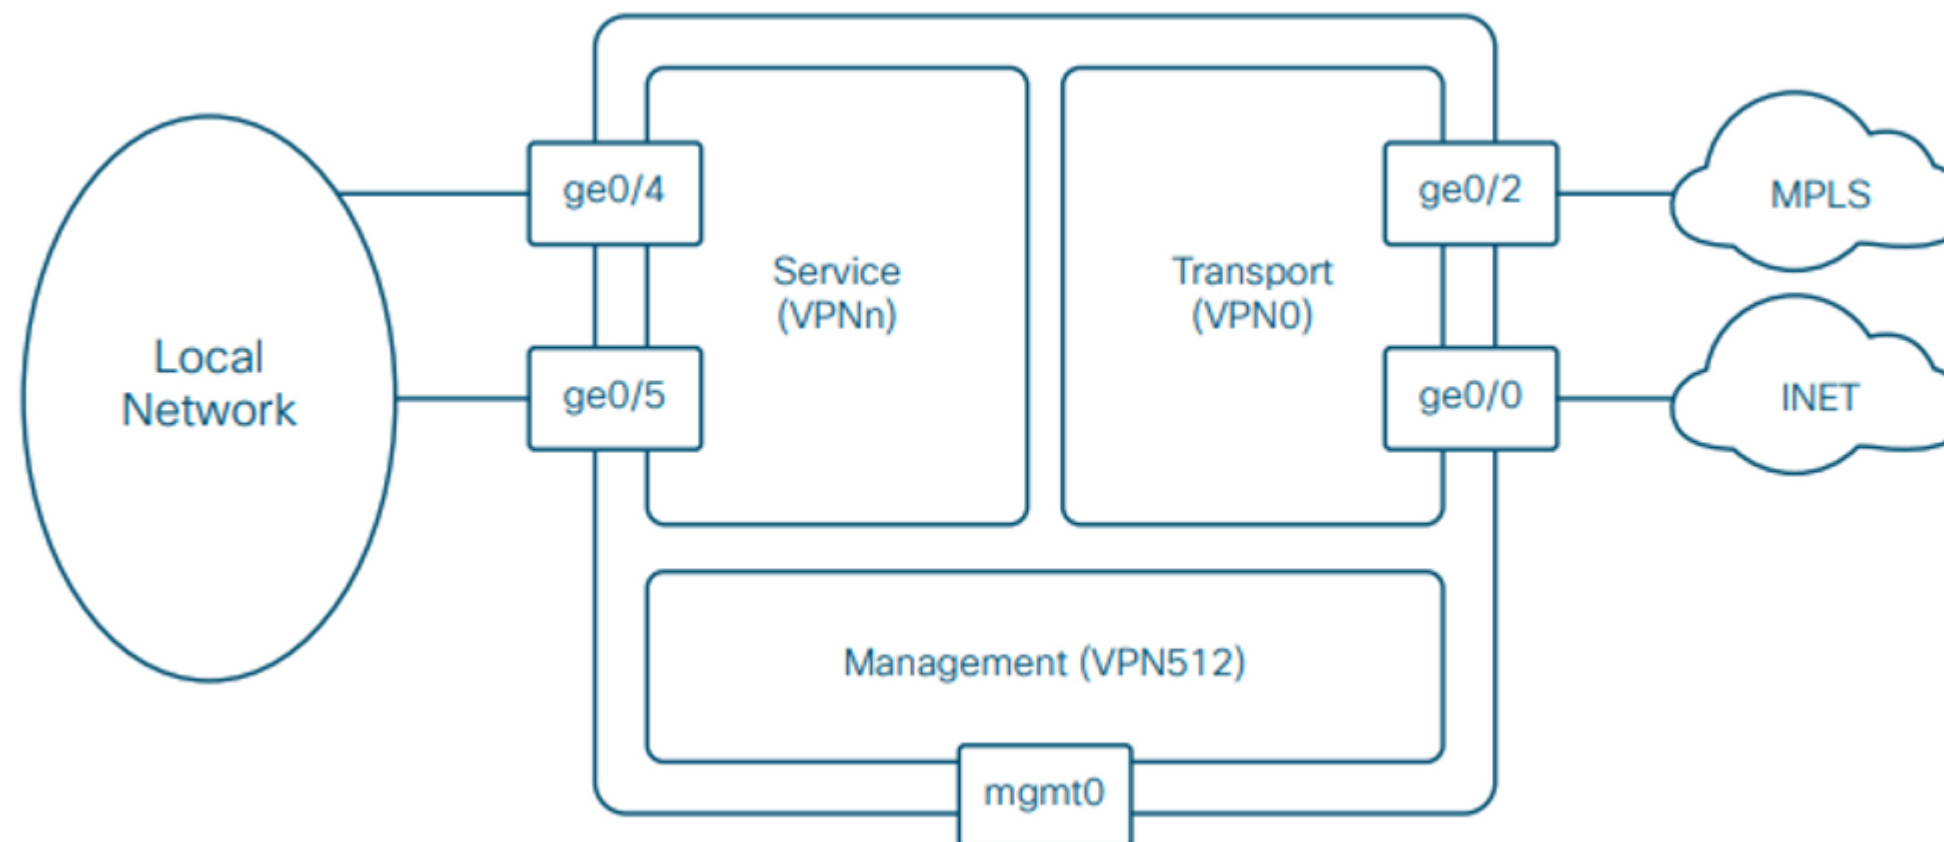
The VPN number is a four-byte integer with a value from 0 to 65530. There are two VPNs present by default in the vEdge devices and controllers, VPN 0 and VPN 512.

- VPN 0 is the transport VPN. It contains the interfaces that connect to the WAN transports. Secure DTLS/ TLS connections to the vSmart or between vSmart and vBond controllers are initiated from this VPN. Static or default routes or a dynamic routing protocol needs to be configured inside this VPN in order to get appropriate next-hop information so the control plane can be established and IPsec tunnels can connect to remote sites.
- VPN 512 is the management VPN. It carries the out-of-band management traffic to and from the Cisco SD-WAN devices. This VPN is not carried across the overlay network.

Q38. WHAT IS A SERVICE-SIDE VPN?

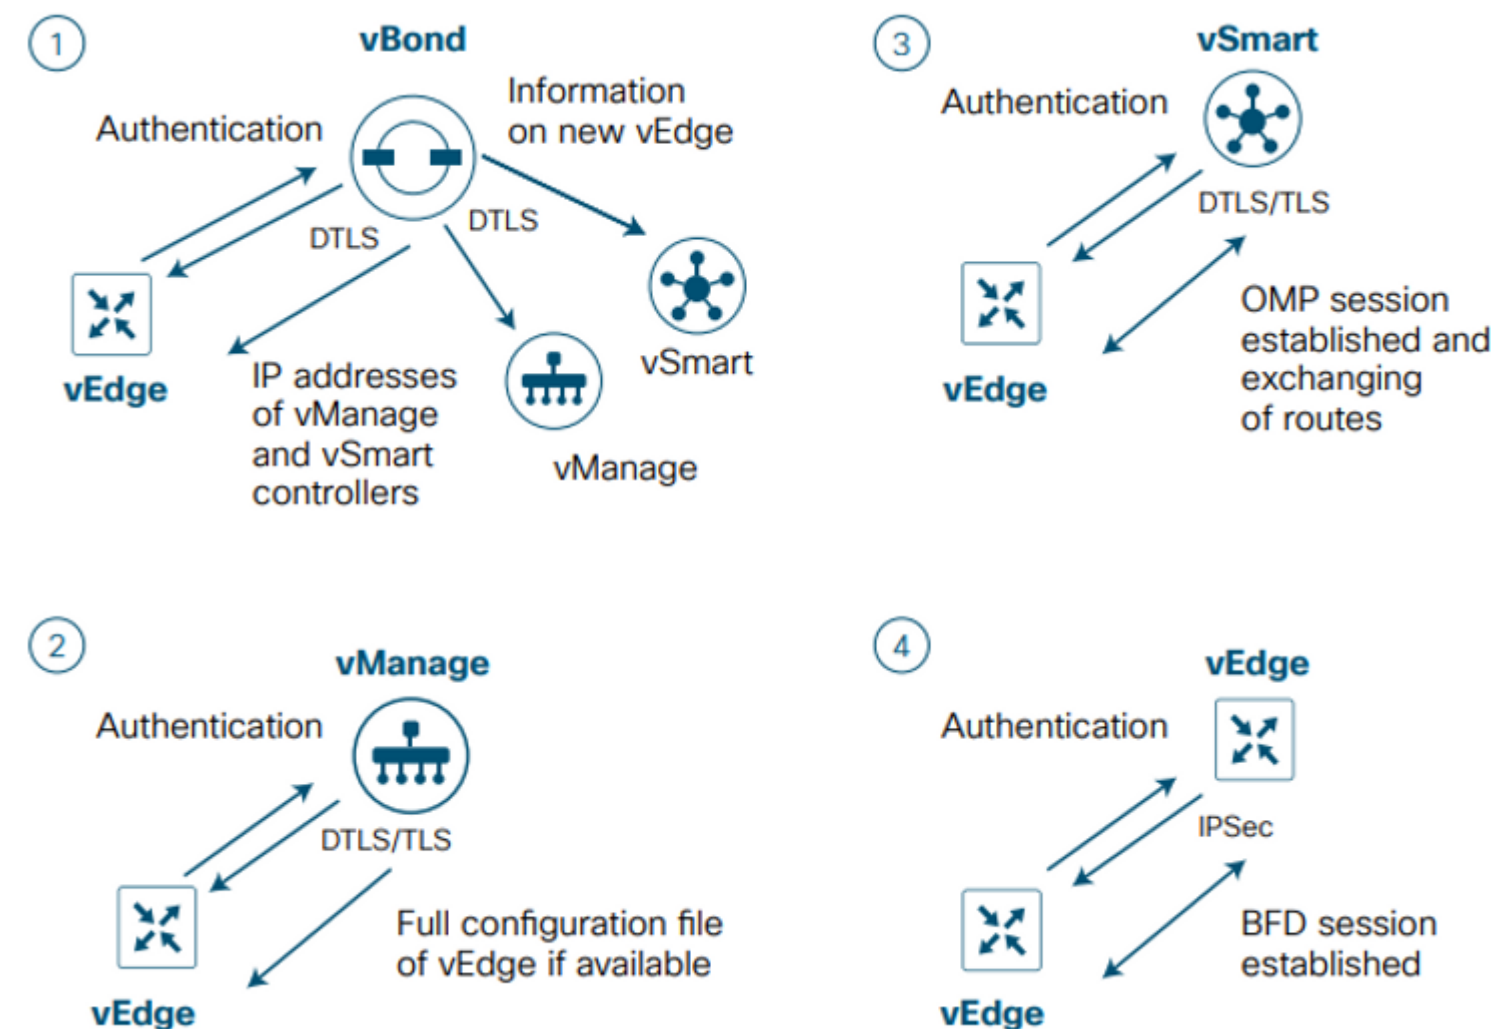
In addition to default VPNs that are already defined, one or more service-side VPNs are needed to be created that will contain interfaces that will connect to the local-site network and carry user-data traffic. These VPNs can be enabled for features such as OSPF or BGP, Virtual Router Redundancy Protocol (VRRP), QoS, traffic shaping, or policing. User traffic can be directed over the IPsec tunnels to other sites by redistributing OMP routes received from the vSmart controllers at the site into the service-side VPN routing protocol. In turn, routes from the local site can be advertised to other sites by advertising the service VPN routes into the OMP routing protocol, which will be sent to the vSmart controllers and redistributed to the other vEdge routers.

The following figure demonstrates VPNs on a vEdge router. The interfaces, ge0/2 and ge0/0, are part of the transport VPN; ge0/4 and ge0/5 are part of the service VPN, which is attached to the local network at the site; and the mgmt0 port is part of VPN512. Note that while physical interfaces are displayed in the diagram, the interfaces in the transport and service VPNs could be sub-interfaces instead.



Q39. HOW CAN YOU BRING THE vEDGE INTO THE OVERLAY?

In order to join the overlay network, a vEdge router needs to establish a secure connection to the vManage so that it can receive a full configuration, and it needs to establish a secure connection with the vSmart controller so that it can participate in the overlay network. The discovery of the vManage and vSmart happens automatically and is accomplished by first establishing a secure connection to the vBond orchestrator.



Q39. HOW CAN YOU BRING THE vEDGE INTO THE OVERLAY? (CONTINUED...)

Step 1: Through a minimal bootstrap configuration or through the Zero-Touch Provisioning (ZTP) process, the vEdge router will first attempt to authenticate with the vBond orchestrator through an encrypted DTLS connection. Once authenticated, the vBond orchestrator sends the vEdge router the IP addresses of the vManage Network Management System (NMS) and the vSmart controllers. The vBond orchestrator also informs the vSmart controllers and vManage of the new vEdge router wanting to join the domain.

Step 2: The vEdge router begins establishing secure DTLS or TLS sessions with the vManage and the vSmart controllers and tears down the session with the vBond orchestrator. Once the vEdge router authenticates with the vManage NMS, the vManage will push the full configuration to the vEdge router if available.


Step 3: The vEdge router attempts to establish DTLS/TLS connections to the vSmart controllers over each transport link. When it authenticates to a vSmart controller, it will establish an OMP session and then learn the routes, including prefixes, TLOCs, and service routes, encryption keys, and policies.

Step 4: The vEdge router will attempt to establish an IPSec tunnel to TLOCs over each transport. A TLOC on a private transport color attempts to connect to TLOCs on both public and private colors, and a TLOC on a public color tries to connect to other TLOCs on public colors by default. The restrict keyword on the tunnel will only build tunnels between TLOCs of the same color. BFD will then run over these established connections.

Q40. HOW CAN YOU BOOTSTRAP THE vEDGE ROUTER?

With the bootstrap configuration method, the idea is to configure the minimum network connectivity and the minimum identifying information along with the vBond orchestrator IP address or hostname. The vEdge router will attempt to connect to the vBond orchestrator and discover the other network controllers from there. In order for you to bring up the vEdge router successfully, there are a few things that need to be present on the vEdge:

- Configure an IP address and gateway address on an interface connected to the network, or alternatively, configure Dynamic Host Configuration Protocol (DHCP) in order to obtain an IP address and gateway address dynamically. The vEdge should be able to reach the vBond through the network.
- Configure the vBond IP address or host name. If you configure a host name, the vEdge router needs to be able to reach a DNS server in order to resolve it. You do this by configuring a DNS server address under VPN 0.
- Configure the organization name, system IP address, and site ID. Optionally, configure the host name.

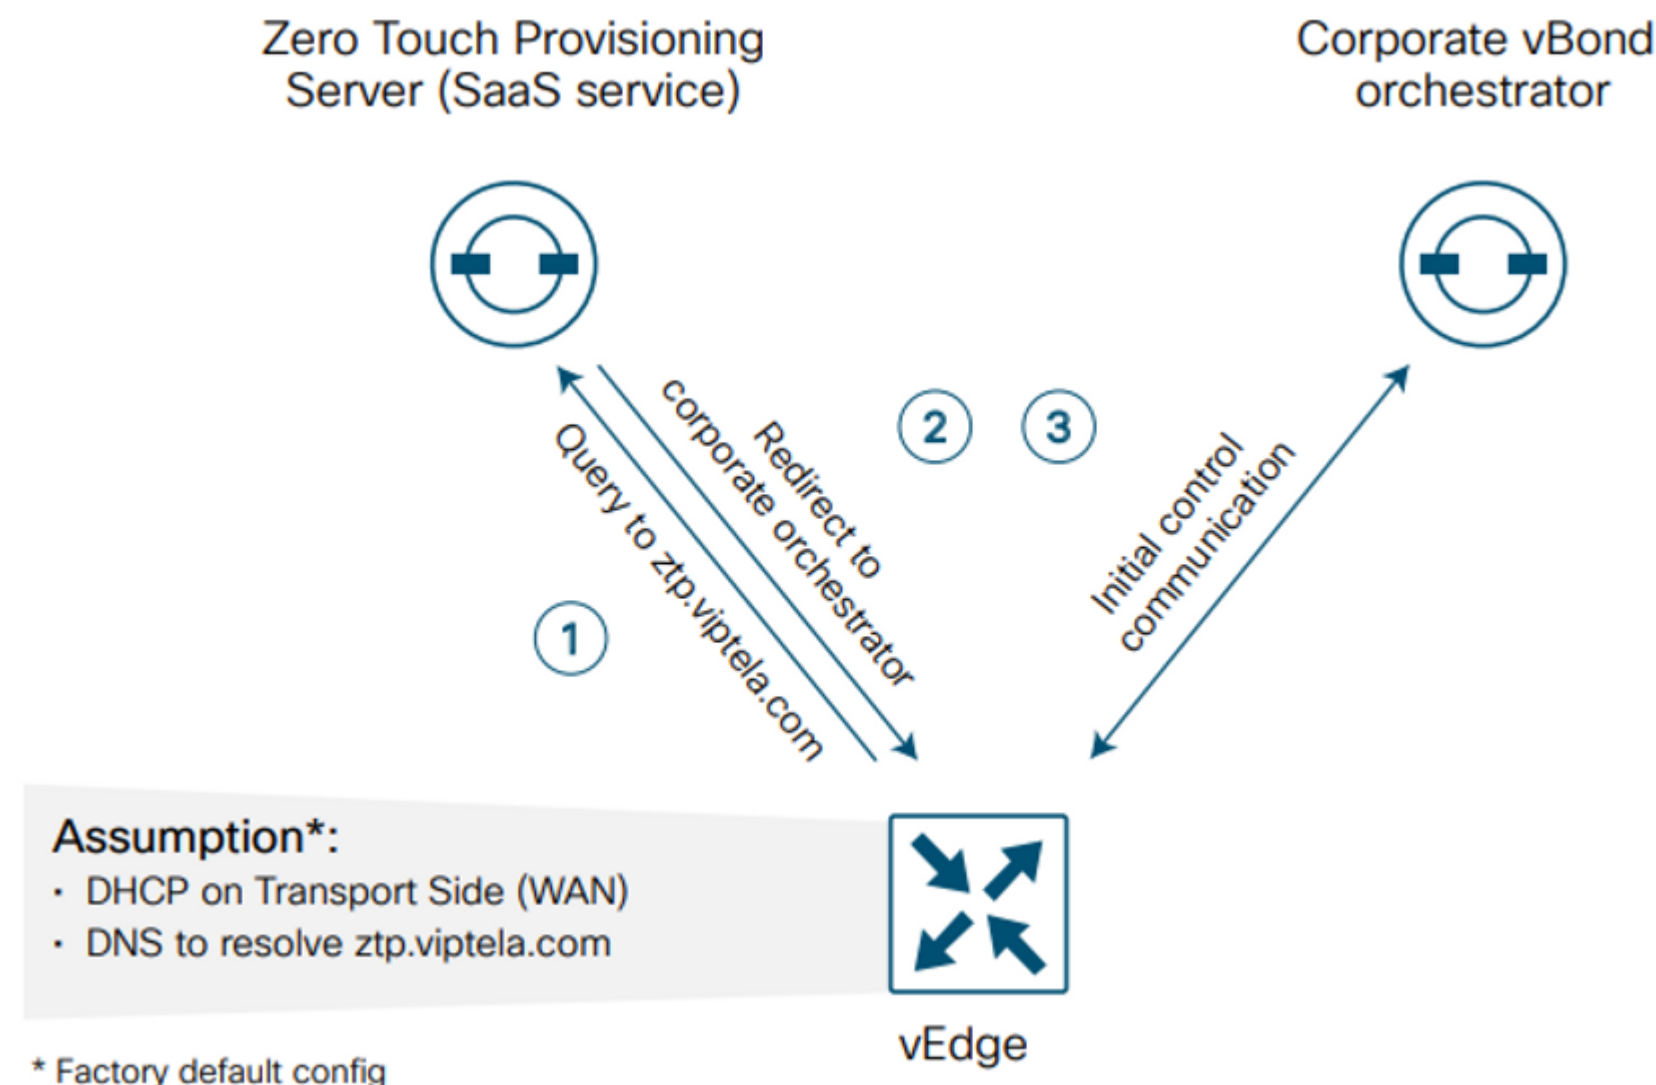
 <p>Rasika Joshi COMPUCOM</p> 	 <p>Nakul Sonare ZENSAR TECHNOLOGIES</p> 	 <p>Minal Ghubde ZENSAR TECHNOLOGIES</p> 	 <p>Ketan Panpate COMPUCOM</p> 
 <p>Mayank Chauhan COMPUCOM</p> 	 <p>Prakash Datta TRIOS</p> 	 <p>Vaibhav Bindu TATA COMMUNICATIONS LIMITED</p> 	 <p>Mangesh Dhongade COMPUCOM</p> 
 <p>Amol Sankpal SECURVIEW</p> 	 <p>Piysuh Shrungare ZENSAR TECHNOLOGIES</p> 	 <p>Khyati Sawant SOPHOS</p> 	 <p>Abhijeet Kamble COMPUCOM</p> 
 <p>Pankaj Sakore NETSCOUT</p> 	 <p>Akshita Mankad COMPUCOM</p> 	 <p>Amey Malotkar AGC NETWORKS LIMITED</p> 	 <p>Arpit Kansal CRYSTALVOXX</p> 
 <p>Pravin Valkunde SECURVIEW</p> 	 <p>Rohit Naidu TATA COMMUNICATIONS LIMITED</p> 	 <p>Hussain Sagwadiya ZENSAR TECHNOLOGIES</p> 	 <p>Pravin Kawale SECURVIEW</p> 

SHAPING CAREER

EMPOWERING FUTURE

Q41. WHAT IS A ZERO TOUCH PROVISIONING (ZTP) PROCESS?

ZTP is an automatic provisioning procedure which starts when the vEdge router is powered up for the first time. The vEdge will attempt to connect to a ZTP server with the hostname `ztp.viptela.com`, where it will get its vBond orchestrator information. Once the vBond orchestrator information is obtained, it can then subsequently make connections to the vManage and vSmart controllers in order to get its full configuration and join the overlay network.



Q42. WHAT ARE THE REQUIREMENTS FOR ZTP?

There are a few requirements for ZTP as follows:

- With the hardware vEdge appliances, only certain ports are pre-configured by default to be a DHCP client interface and can be used for ZTP. The following table outlines the ports that must be plugged into the network for ZTP to work.
- The Gateway Router for the vEdge router in the network should have reachability to public DNS servers and be able to reach ztp.viptela.com.
- In vManage, there must be a device configuration template for the vEdge router attached to the vEdge device. The system IP address and site ID need to be included in this device template in order for the process to work. The ZTP process will not succeed without this.

Q43. WHAT ARE CONTROLLERS CONNECTIONS?

The secure sessions between the vEdge routers and the controllers (and between controllers), by default are DTLS, which is User Datagram Protocol (UDP)-based. The default base source port is 12346. The vEdge may use port hopping where the devices try different source ports when trying to establish connections to each other in case the connection attempt on the first port fails. The vEdge will increment the port by 20 and try ports 12366, 12386, 12406, and 12426 before returning back to 12346. Port hopping is configured by default on a vEdge router, but you can disable it globally or on a per-tunnel-interface basis. It is recommended to run porthopping at the branches, but disable this feature in the controllers, and data center, regional hub, or a place where aggregate traffic exists. Control connections on vManage with multiple vCPUs will have a different base port for each vCPU core.

For vEdge routers that sit behind the same NAT device and share a public IP address, you do not want each vEdge to attempt to connect to the same controller using the same public IP and port

umber. In this case, you can configure an offset to the base port number of 12346, so the port attempts will be unique among the vEdge routers. A port offset of 1 will cause the vEdge to use the base port of 12347, and then port-hop with ports 12367, 12387, 12407, and 12427. Port offsets need to be explicitly configured, and by default, the port offset is 0.

Q43. WHAT ARE CONTROLLERS CONNECTIONS? (continued...)

Alternatively, you can use TLS to connect to the vManage and vSmart controllers, which is TCP-based instead of UDP-based. vBond controller connections will always use DTLS, however. TCP ports will originate on the vEdge from a random port number destined to the base port of 23456, and control connections with multiple vCPUs will have a different base port for each vCPU core, similar to the DTLS case.

IPSec tunnels and BFD from a vEdge router to another vEdge router use UDP with similar ports as defined by DTLS.

Ensure that any firewalls in the network allow communication between vEdge routers and to controllers. Ensure that they are configured to allow return traffic as well. The following table is a summary of the ports used, assuming the controllers are configured to not use port-hopping.

Q44. WHAT ARE CONFIGURATION TEMPLATES?

Configurations and policies apply to vEdge routers and vSmart controllers which enable traffic to flow between the data center and the branch or between branches. An administrator can enable configurations and policies through the Command-Line Interface (CLI) using console or Secure Shell (SSH) on the vEdge device, or remotely through the vManage GUI.

To configure a vEdge device or controller on the network using the vManage GUI, an administrator applies a device template to a vEdge router or multiple vEdge routers. These templates can be CLI-based or featurebased. While you can create CLI-based templates, we recommend feature-based templates because they are modular, more scalable, and less error-prone. Each device template is made up of several feature templates that describe the interface configurations, tunnel configurations, and local routing behavior.

Templates are extremely flexible, and there are a number of approaches to putting templates together. You can choose to have more variables inside your template, which will result in less feature templates, or you can have less variables but more feature templates. For example, you can choose to enable NAT as a variable or a global value. You can create one interface feature template and choose to enable or disable NAT through a variable, or you can create two different feature templates, one with NAT disabled and one with NAT enabled, and choose the most appropriate feature template to use, depending on the device template. In any case, you should add a detailed description of each feature and device template in detail in the GUI and create very descriptive variable names so that it is very clear what each template and variable is.

Q44. WHAT ARE CONFIGURATION TEMPLATES? (continued...)

When designing configuration templates, it is helpful to think about how operations may interact with the templates on a day-to-day basis. It might be useful to use variables for interface names so that interfaces can be moved for troubleshooting purposes, without having to create new feature templates to accomplish it. It also might be helpful to create variables for states of interfaces and routing protocols for troubleshooting reasons, such as allowing the disabling of an interface or a BGP neighbor by just changing a variable.

Q45. WHAT IS A DEVICE TEMPLATE?

Device templates are specific to only one vEdge model type, but you may need to create multiple device templates of the same model type due to their location and function in the network. Each device template references a series of feature templates which makes up the entire configuration of the device. A device template configuration cannot be shared between vEdge models, but a feature template can span across several model types and be used by different device templates.

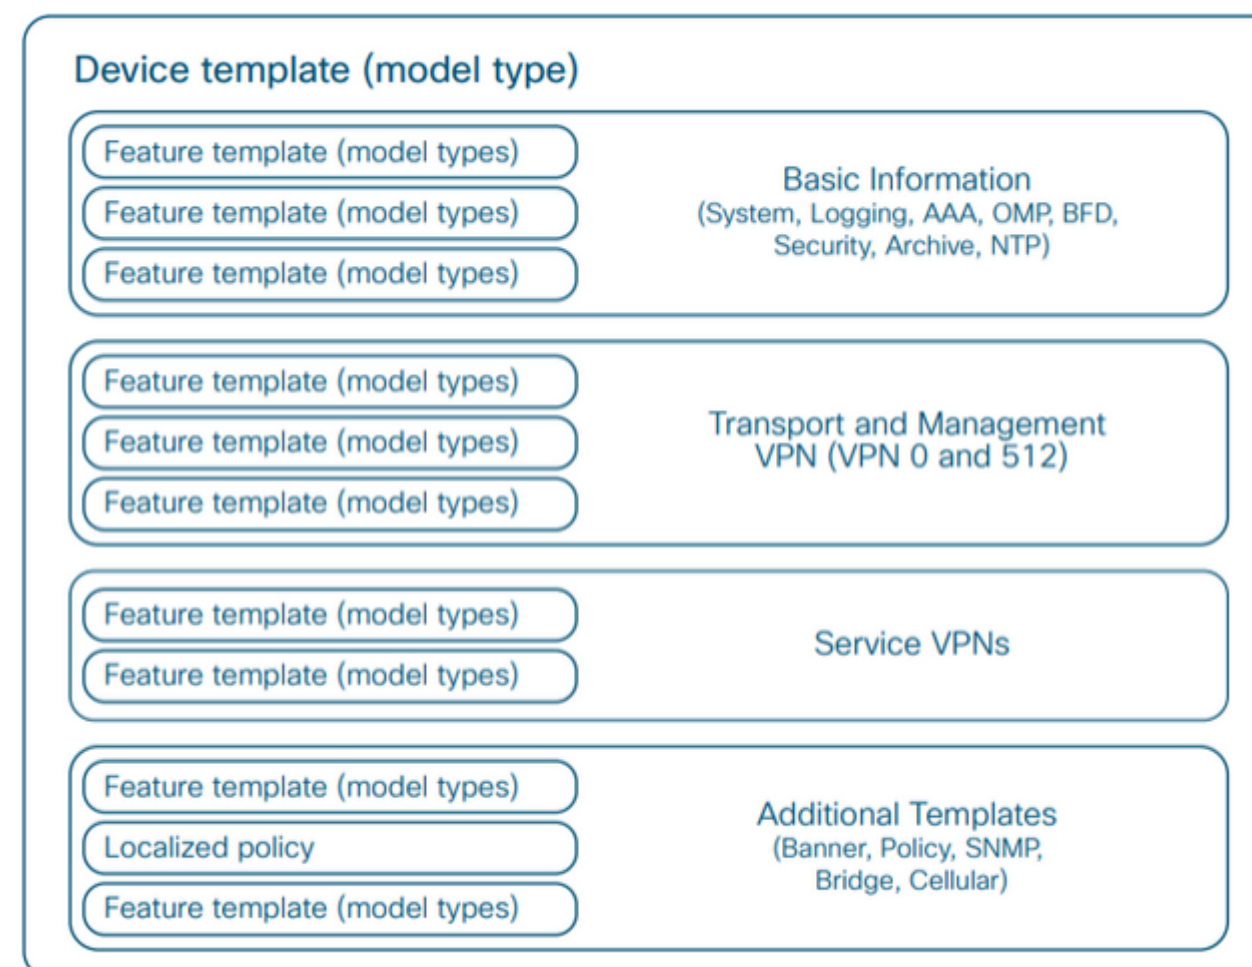
Q46. MENTION DEVICE TEMPLATE COMPONENTS.

Basic Information - This section includes system, logging, AAA, OMP, BFD, security, archive, and NTP feature templates.

Transport and Management VPN - This section includes the templates used to configure VPN 0 and VPN 512, which includes BGP, OSPF, VPN interface, VPN interface cellular, VPN interface GRE, and VPN interface PPP feature templates.

Service VPN - This section includes the templates used to configure the service VPNs, which contains the BGP, IGMP, Multicast, OSPF, PIM, VPN interface, VPN interface bridge, VPN interface GRE, VPN interface IPsec, VPN interface Natpool, and DHCP server feature templates.

Additional Templates - This section includes banner, Simple Network Management Protocol (SNMP), bridge, localized policy, and cellular feature templates.



Q47. WHAT ARE CONFIGURING PARAMETERS?

An administrator uses vManage to configure device and feature templates, specifying variables where needed since templates can apply to multiple vEdge devices that have unique settings.

When configuring values of parameters inside of feature templates, there is often a drop-down box that gives you three different types of values:

Global - When you specify a global value, you specify the desired value, either by typing the value into a text box, selecting a choice from a radio button, or selecting a value from a drop-down box. Whatever value you select will be applied to all devices the feature template is applied to.

Device-specific - When you specify a device-specific value, you will create a variable name. The value for this variable will be defined when the device template is applied.

Default - When you specify a default value, a default value will be applied to all devices the feature template is applied to. If there is a specific value, it will appear in a textbox in grey scale.

In the illustration below, Timezone is shown as a global, device-specific, or default value. A variable name is entered when specifying the device-specific value.

Q47. WHAT ARE CONFIGURING PARAMETERS? (continued...)

Basic Configuration GPS Tracker Advanced

Overlay ID ☒ 1

Timezone ☒ UTC

Hostname

Location

☒ Global

☐ Device Specific > Enter Key

☒ Default system_timezone

☒ America/New_York

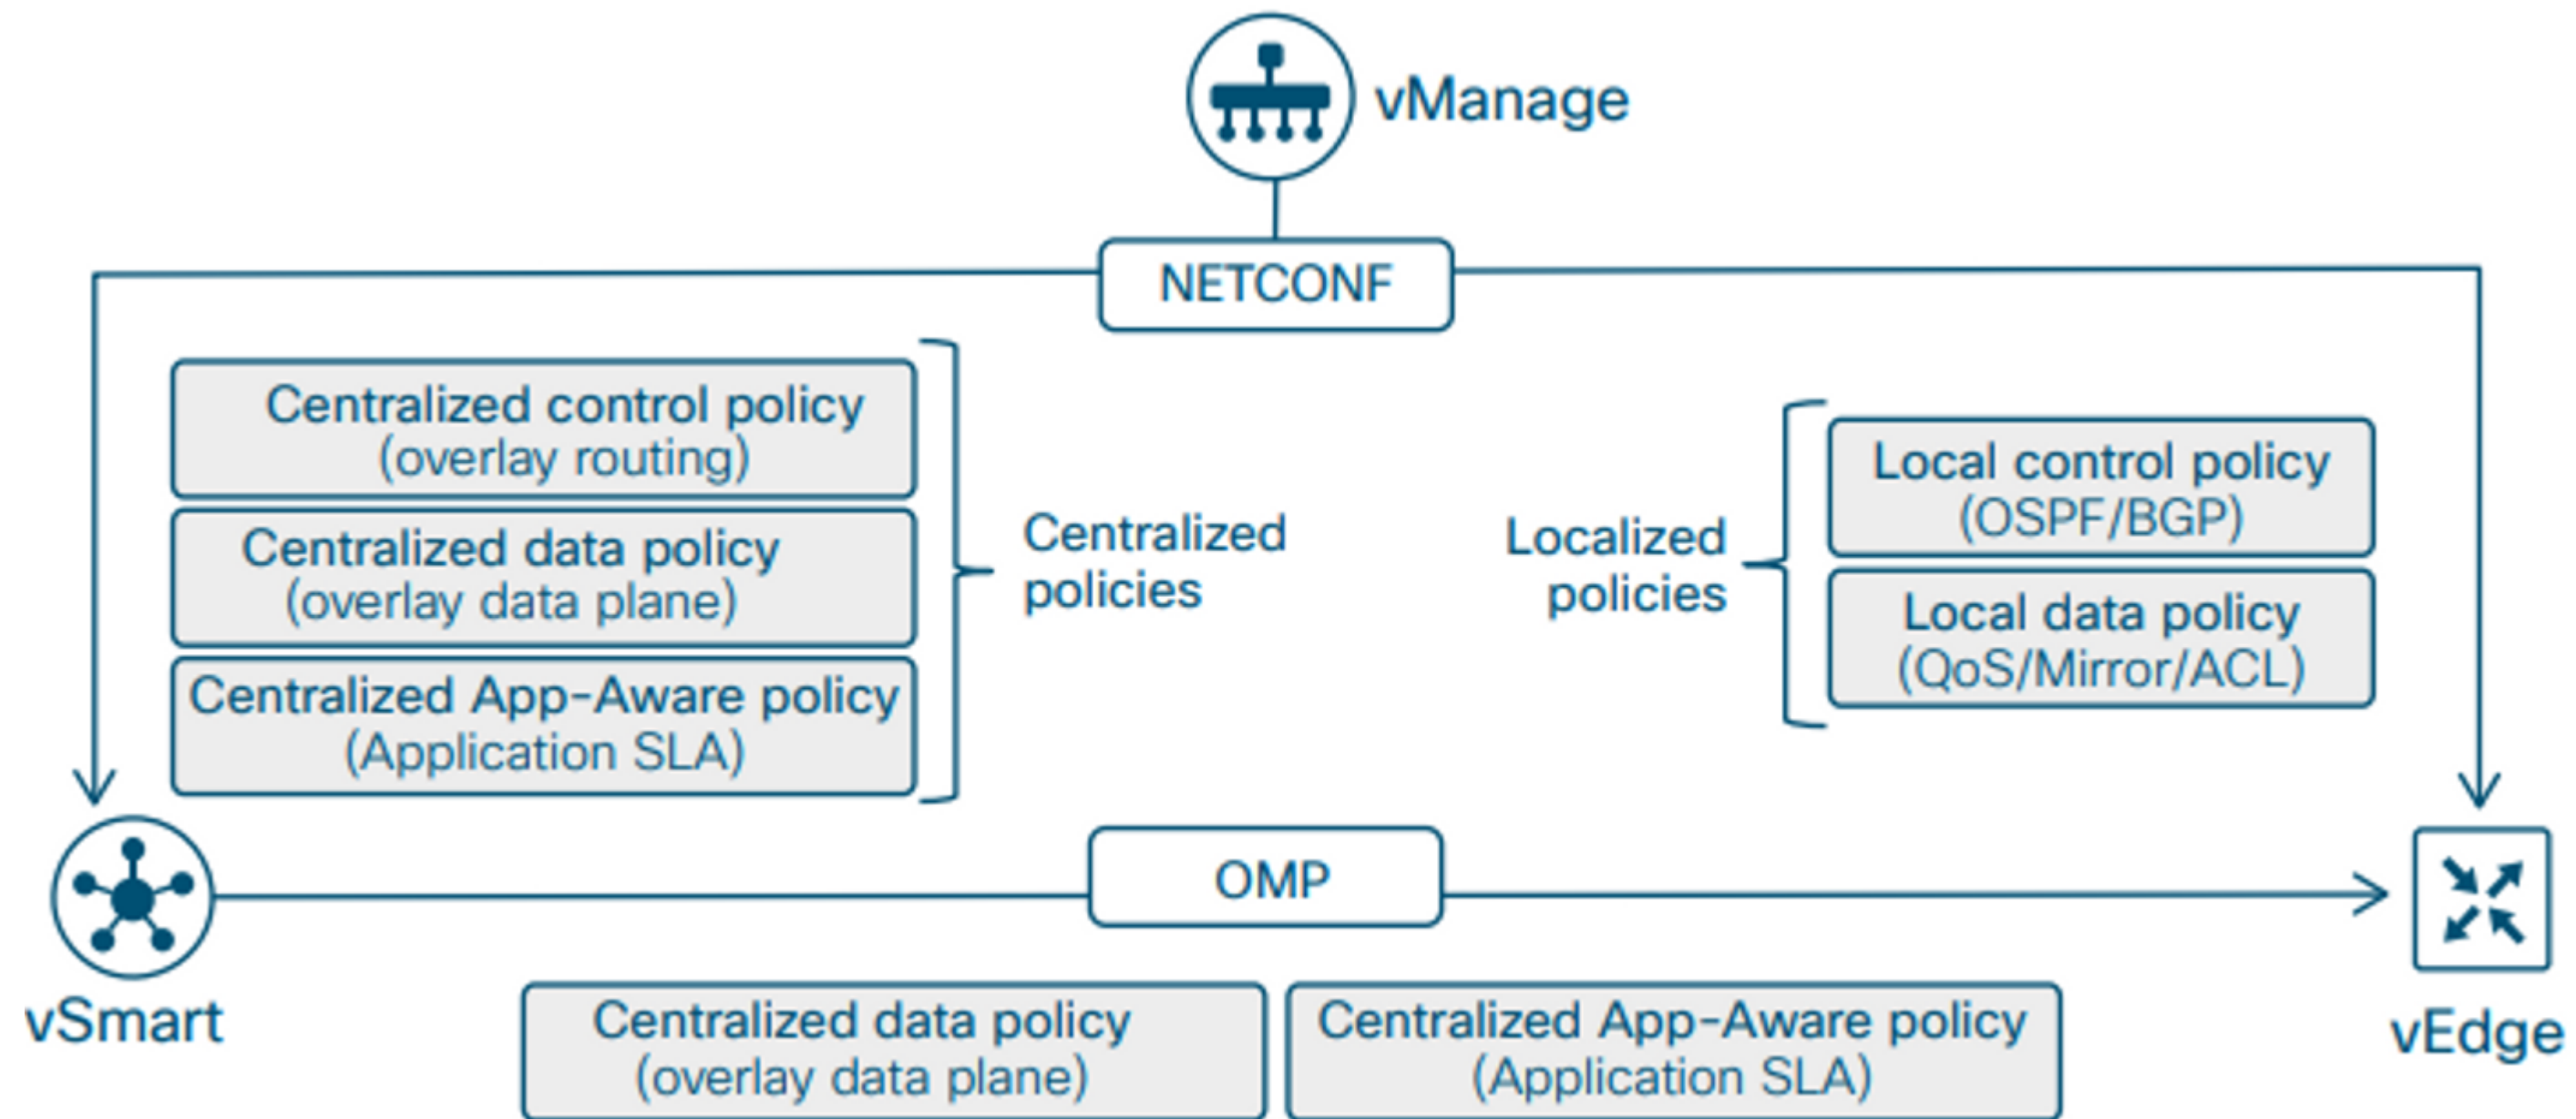
☐ [system_timezone]

☒ UTC

Q48. HOW TO DEPLOY DEVICE TEMPLATE?

Once feature templates are configured, the device template configuration is completed by referencing the desired feature template in each configuration category (system, AAA, BFD, VPN, VPN interface, etc.). Once a device template is configured, it can be attached to a specific vEdge device. Once attached, you will be required to fill in the values for any variables in the template for each vEdge the template will apply to before the configuration can be deployed. You can enter values through the vManage GUI directly, or by filling out a .csv file that can be uploaded. The .csv file method allows you to deploy a large number of vEdge routers quickly and more easily. vManage will then modify the configuration of the targeted vEdge devices in the database and then push out the entire configuration to the intended vEdge routers on the network. When making an update to a feature or device template, the application will happen immediately if there are devices attached to those templates. If the configuration gets pushed out and if there is an error, such as an incorrect value format or a reference to a loop-back interface that doesn't exist, the template configuration rolls back to its previous state before the edit.

Q49. HOW DO CENTRALIZED AND LOCALIZED POLICIES LOOK?



Q50. HOW WOULD YOU CONFIGURE LOCALIZED POLICY?

There are three steps for applying localized policy:

Step 1: In the vManage GUI, create the localized policy under Configuration>Policies and select the Localized Policy tab. Before Release 18.2, the policy is added as a CLI policy. Starting in Release 18.2, a policy configuration wizard was created to assist with policy creation.

Step 2: In the device template, under the Additional Templates section next to Policy, reference the name of the localized policy.

Step 3: Reference any policy components, like route policies and prefix lists, inside the feature templates.

When you are creating a device template and referencing a feature template that already has a route policy or prefix list or another localized policy component configured in it, you must have a policy name referenced in the device template before you can create or update the device template. If a device is already attached to an existing device template, you must first attach a localized policy to the device template before referencing any localized policy elements within the feature templates that are associated with that device template.

You can only apply one localized policy to a vEdge device. Within this policy, you will create both control and data policies components; prefix-lists, route-policies, as-path lists, community-lists, QoS class-maps, qosmap policies, mirror and policing policies, rewrite-rule policies, and access lists will all be included in this one localized policy.

RECENT CCIE FROM I-MEDITA



HELPING STUDENTS BECOME CERTIFIED

Q51. HOW WOULD YOU CONFIGURE CENTRALIZED POLICY?

When configuring centralized policy in the vManage GUI, there are three main components:

Lists

Lists are used to group related items so you can reference them as a group. They are used when applying policy or used in matching or actions within the policy definitions. You can create lists for applications, color, data prefixes, policers, prefixes, sites, SLA classes, TLOCs, and VPNs. Data prefixes are used in data policies to define data prefixes, and prefixes are used in control policies to match on route prefixes.

.

Policy Definition

The policy definitions control the aspects of control and forwarding. Within the policy definition is where you create policy rules, specifying a series of match-action pairs which are examined in sequential order. There are several types of policy definitions: app-route policy, cflowd-template, controlpolicy, data-policy, and a vpn-membership policy.

.

Policy Application

The policy is applied to a site list

Q52. WHAT ARE THE TYPES OF POLICY DEFINATION?

There are several different types of policy definitions:

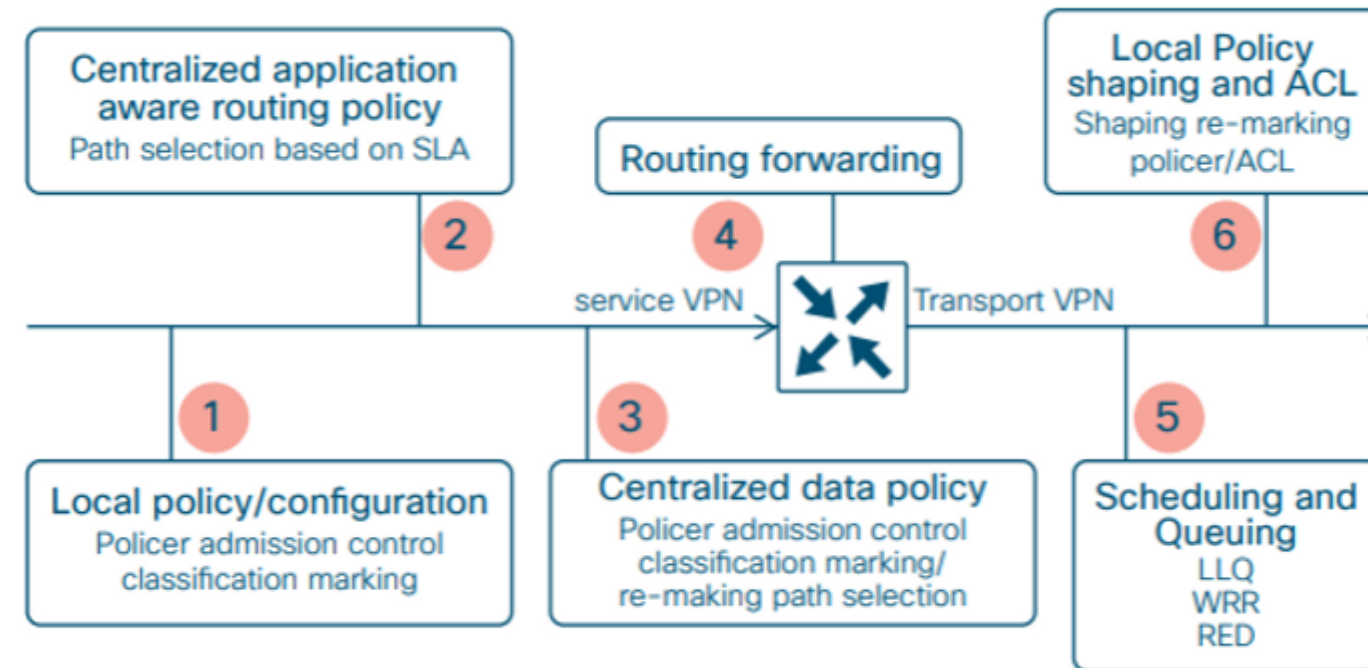
- **App-route Policy** - Allows you to create an application-aware routing policy which tracks path characteristics such as loss, latency, and jitter. Traffic is put into different SLA categories (loss, delay, and jitter), and traffic is directed to different paths depending on the abilities to meet the SLA categories
- **Cflowd Template** - Allows you to enable cflowd, which sends sampled network data flows to collectors
- **Control Policy** - Operates on the control plane traffic and influences the routing paths in the network
- **Data Policy** - Influences the flow of data traffic based on the fields in the IP packet header
- **VPN Membership Policy** - Can restrict participation in VPNs on vEdge routers and the population of their route tables.

Control policy examines the routes and TLOC attributes in the routing information and modifies attributes that match the policy. This policy is unidirectional and can be applied to a site list in an inbound or outbound direction. The direction is from the perspective of the vSmart controller. A policy applied to a site list in the inbound direction means that policy would affect routes coming from the sites on the site list and actions would be applied on the receive side of the vSmart controller. A policy applied to a site list in the outbound direction means the policy would affect routes going to the sites on the site list and actions would be applied to the sending side of the vSmart controller.

Q53. WHAT IS THE ORDER OF OPERATIONS?

Following is the order of operations on a packet as it traverses from service VPN to transport VPN on a vEdge router:

1. Local policy/configuration - includes QoS classification, policer, and marking
2. Centralized application-aware routing policy
3. Centralized data policy - includes QoS classification, policer, marking, and path selection
4. Routing/forwarding
5. Scheduling and queueing
6. Local policy shaping and ACL - includes shaping, re-marking, and policer



Packet flow through the vEdge router (from service interface to WAN-Transport interface)

From the ordering, it's possible for a centralized data policy to overwrite the actions of a local data policy configuration, and it's also possible for a centralized data policy to influence the path selection that is different than what was chosen as part of the application-aware routing policy. Keep this information in mind as you define the policies for the network.

Q54. WHAT IS THE TRAFFIC SYMMETRY OF DPI?

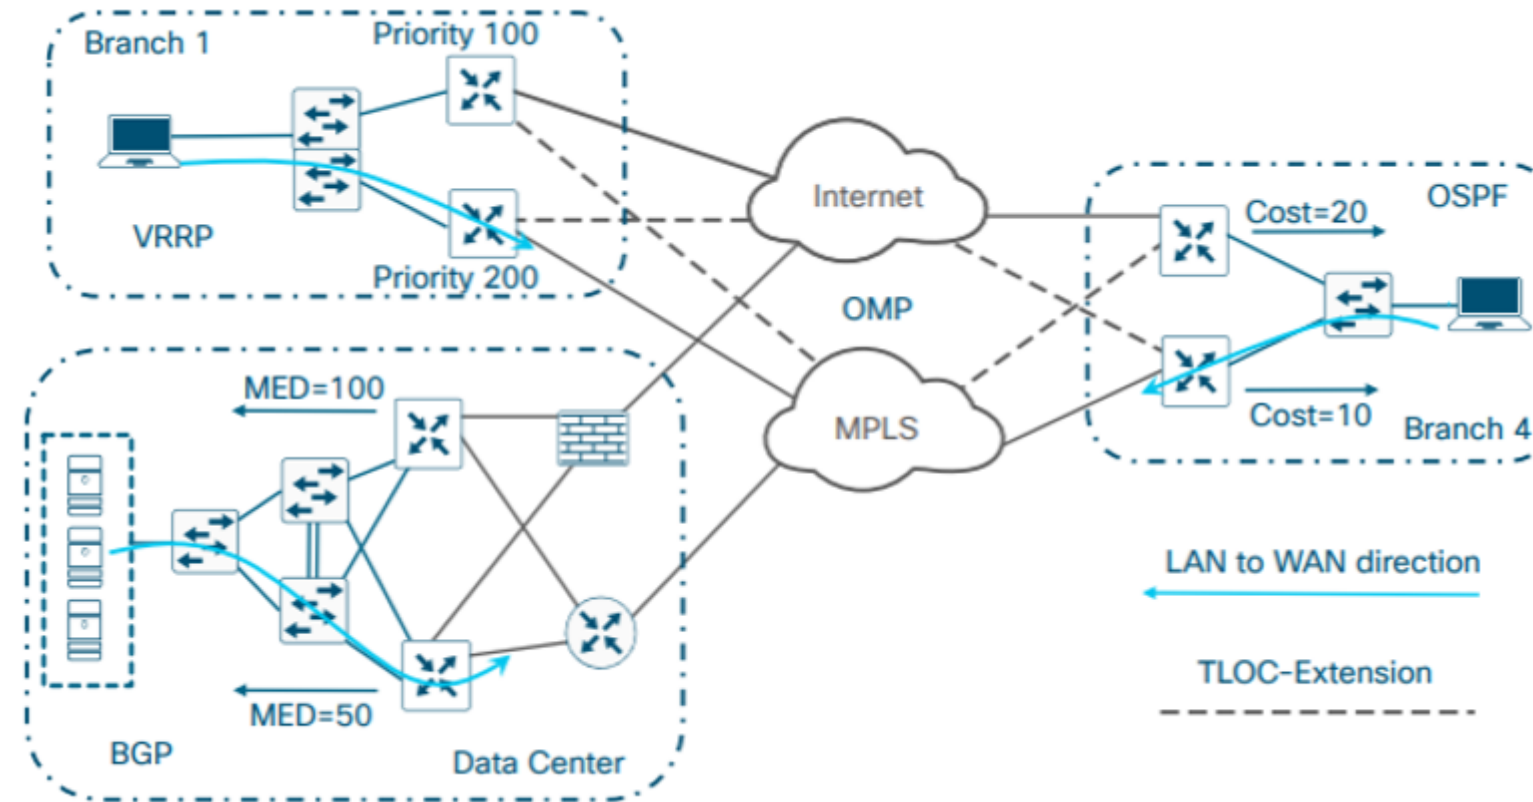
Application-aware routing uses Deep Packet Inspection (DPI) for matching on applications within the policy. In order for DPI on a vEdge router to be able to classify most application traffic, it is important that the vEdge router sees network traffic in both directions. In dual-vEdge sites without any policy enabled, equal cost paths exist over each transport and to each vEdge router, and network traffic is hashed depending on fields in the IP header. Traffic is unlikely to always be forwarded to the same vEdge router in both the LAN-to-WAN direction and the WAN-to-LAN direction. To maintain symmetric traffic, it is recommended to set up routing so that traffic prefers one vEdge over another at dual-vEdge router.

To ensure symmetry, traffic needs to prefer one router in both directions, from the LAN to the WAN and from the WAN to the LAN. There are different ways to accomplish this.

To influence traffic in the LAN to WAN Direction:

- For VRRP, use VRRP priority to prefer one vEdge router over the other.
- For OSPF, use the cost metric, configured either on the interface of the neighbouring switch itself or through a route policy on the vEdge router that modifies the metric of routes redistributed from OMP to OSPF.
- For BGP, use a route policy and set AS path prepend or Multi-Exit Discriminator (MED) on routes redistributed from OMP to BGP.

Q54. WHAT IS THE TRAFFIC SYMMETRY OF DPI?



To influence traffic in the WAN-to-LAN direction over the overlay, you can influence an OMP attribute or set the IPSec Tunnel Preference. When BGP or OSPF is redistributed into OMP, the MED setting for BGP and the cost for OSPF is automatically translated into the OMP origin metric, which is used in the decision making for picking the best route.

Some common methods to Influence Traffic for the WAN-to-LAN direction:

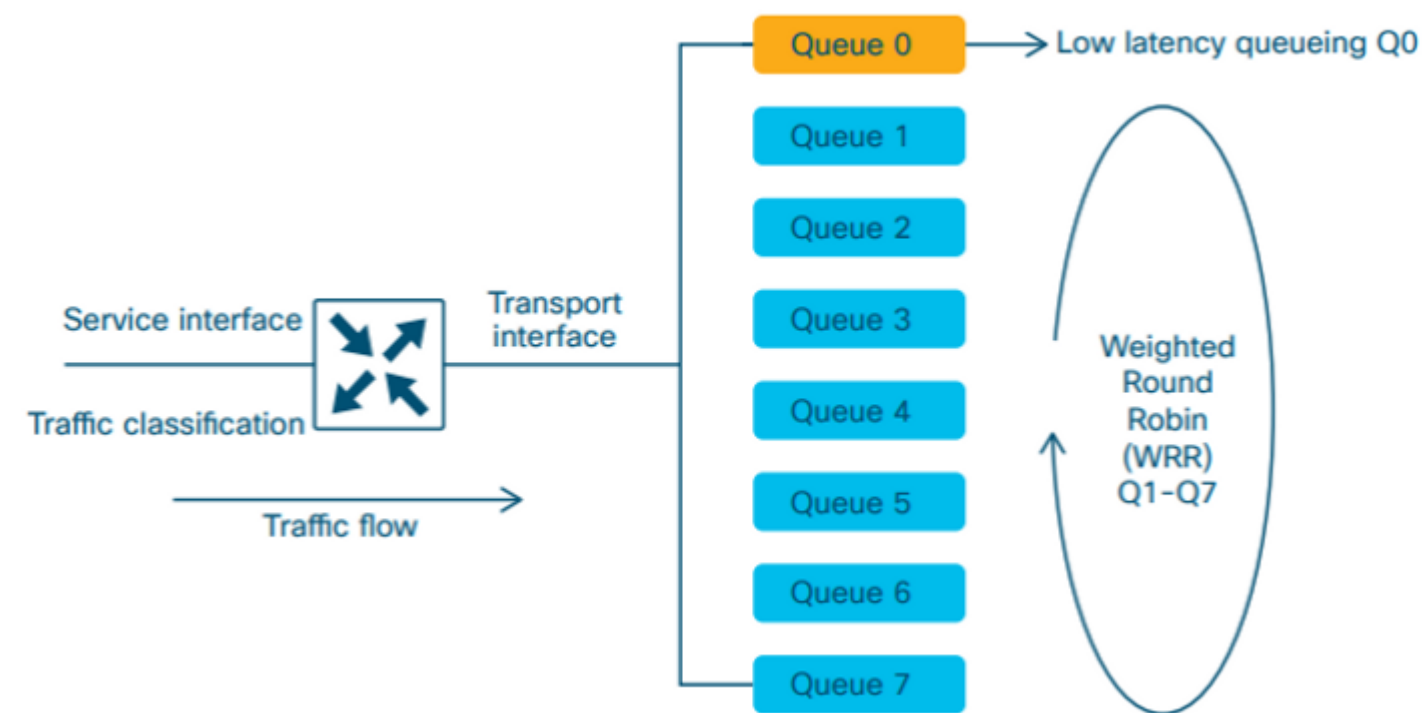
- For BGP, use a route-policy and set MED (metric) on routes inbound from the LAN BGP neighbors
- For OSPF, use vEdge Router interface cost to set the metric on routes coming into the LAN Interface
- For any vEdge Router, use IPSec tunnel preference to influence which is the preferred vEdge through the WAN overlay

Q55. WHAT IS QUALITY OF SERVICE?

QoS is frequently deployed on the WAN transport, as bandwidth is less there, compared to the LAN side. QoS can only be applied to physical interfaces and not subinterfaces.

Each vEdge router physical interface has eight queues, labeled 0-7. Queue 0 uses Low Latency Queueing (LLQ), while queues 1-7 use Weighted Round Robin (WRR) for scheduling. Queue 0 (using LLQ) is a strict-priority traffic queue, meaning delay-sensitive traffic that is assigned to this queue is transmitted before packets in other queues. In addition, tail-drop is the only congestion avoidance algorithm used for this queue, and the queue is strictly policed.

By default, control traffic and BFD traffic (both marked as DSCP 48 decimal) use queue 0, while user traffic uses queue 2.



Q55. WHAT IS QUALITY OF SERVICE? (continued...)

There are a few steps to take when configuring QoS on vEdge router hardware:

- Map QoS forwarding classes to output queues. There are eight queues, 0 through 7. Queue 0 is reserved for control traffic and low-latency queueing (LLQ) traffic. If you assign traffic queue 0, it must be configured for LLQ scheduling. This is configured through localized policy.
- Configure the QoS scheduler for each forwarding class. This assigns the scheduling method (LLQ or WRR), bandwidth percentage, buffer percentage, and drop algorithm (Random Early Detection [RED] or tail drop) to each forwarding class. The bandwidth and buffer percentages must add up to 100 percent. This is configured through localized policy.
- Create a QoS map, where all of the QoS schedulers are grouped. This is configured through localized policy.
- For vEdge 5000 routers (and vEdge cloud routers), use the cloud-qos command to enable QoS scheduling and shaping for the transport-side tunnel interfaces. Use the cloud-qos-service-side command to enable QoS scheduling and shaping for the service-side interfaces. This is configured within the localized policy.
- Create a re-write policy (optional). This policy changes the DSCP value in the tunnel header and allows you to map to a smaller number of DSCP values that the service provider might support in the provider cloud. This is configured through localized policy.

Q55. WHAT IS QUALITY OF SERVICE? (continued...)

- Define an access list to match traffic and assign to forwarding classes. This can be configured either through centralized policy as a traffic data policy or through localized policy.
- Apply the access list to an interface. It is typically applied as an inbound list on the service-side interface. This is configured either through centralized policy by specifying direction when the policy is applied (typically, the from-service option is used), or this can be configured by specifying the access list created in the localized policy in the VPN Ethernet template.
- Apply the QoS map and, optionally, the re-write policy, to an egress interface. This is configured through the VPN Interface Ethernet template of the desired interfaces in VPN 0.

Looking for Networking Training?

Join our CCNA, CCNP, CCIE, F5, Checkpoint, Palo Alto & Fortinet Certification Courses

[Click here to Sign Up for a Free Demo Session](#)

REGISTER FOR FREE DEMO