Omada SD-WAN White Paper

Contents

1 Overview

Omada software-defined wide area network (SD-WAN) is a one-click WAN networking solution built on software-defined networking (SDN) and VPN technologies, designed for prosumers and small- and medium-sized businesses (SMBs). It enables one-click networking across multiple sites, supports site-to-site data communications, and ensures secure data transmission.

1.1 Introduction

Omada SD-WAN applies SDN to wide-area networking as a VPN-based solution. In the control plane, sites exchange tunnel information over Secure Sockets Layer (SSL) to establish SD-WAN tunnels. Omada SD-WAN also advertises private routes between sites via Open Shortest Path First (OSPF). In the data plane, it uses generic routing encapsulation (GRE) to forward packets and IPsec tunnels to ensure secure data transmission, providing secure and reliable connectivity for enterprise networks, data centers, and other environments distributed across a wide geographic area.

1.2 Background

In real-world networks, multi-site connectivity is typically implemented by provisioning MPLS (Multiprotocol Label Switching) private lines or by building VPN links. Private lines are expensive, have long setup periods, and costs grow linearly—or even exponentially—as more branches are added. In multi-site deployments, VPN configuration is often complex and time-consuming, typically requiring devices to have public IP addresses.

To address these challenges, SD-WAN emerged. By using SDN, organizations can centrally manage devices across all sites and configure site devices with one-click provisioning, enabling automated VPN configuration and connectivity across sites. In addition, peer-to-peer (P2P) hole punching reduces dependence on public IP addresses to some extent, allowing SMBs to build VPN-based networks without requiring public IP addresses.

1.3 Purpose

This white paper aims to:

1. Explain the principles and advantages of Omada SD-WAN in detail.
2. Provide configuration guidance and recommended deployment strategies for Omada SD-WAN.

3. Describe typical Omada SD-WAN use cases and solution scenarios.

By reading this white paper, you will gain a comprehensive understanding of how SD-WAN works and where it can be applied.

2 Key Technologies

Omada SD-WAN combines SDN, VPN (GRE over IPsec), OSPF, NAT traversal, and P2P hole punching. It enables reliable connectivity and secure communications between sites with simple cloud-based configuration.

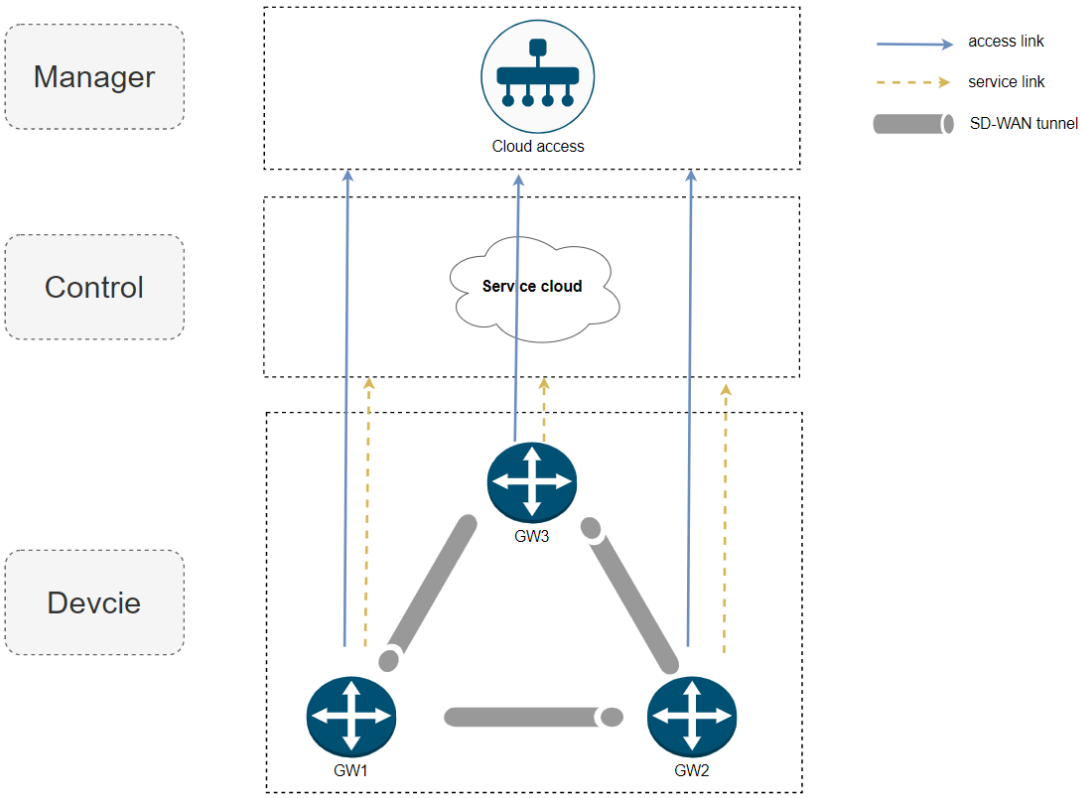

Figure 1: SD-WAN Network Model

The Omada SD-WAN network model includes:

1. **Gateway:** The edge device in the user network.
2. **Cloud Access:** The cloud service that manages gateways and orchestrates SD-WAN networking.
3. **Service Cloud:** The cloud service that relays Interactive Connectivity Establishment (ICE) information between gateways, including virtual IP, private IP/port, public IP/port, NAT type, etc.
4. **Group ID:** The group identifier. Each SD-WAN group has its own group ID, and only devices in the same group can join the SD-WAN network.
5. **Device ID:** The unique identifier of an SD-WAN-capable device within a site. The

service cloud uses this identifier for authentication.

6. **SD-WAN Tunnel:** The data path between two SD-WAN gateways. Sites exchange packets in SD-WAN tunnels to establish site-to-site connectivity.
7. **Access Link:** The connection between a device and cloud access. Cloud access uses this link to control and manage the device.
8. **Service Link:** The connection between a device and the service cloud. This link carries ICE information.

The end-to-end workflow is as follows:

1. **Orchestration by cloud access**

   Cloud access orchestrates SD-WAN networking by selecting the participating devices and their WAN and LAN networks. It then sends control messages (e.g., virtual IP, WAN, LAN networks, and the peer device ID) to each participating device over the cloud access link.
2. **Local interface setup and NAT discovery**

   After receiving the message, each device creates a GRE interface, configures the virtual IP, binds the interface to the corresponding WAN, and probes the Session Traversal Utilities for NAT (STUN) server to discover the public IP/port and NAT type.
3. **ICE exchange**

   Each device sends its local ICE information (e.g., public IP/port, NAT type, and virtual IP) to its peer device through the service cloud and requests the peer's ICE information.
4. **P2P hole punching**

   After both devices receive the peer's ICE information, they perform P2P hole punching to the peer's public IP/port and establish a UDP connection.
5. **IPsec negotiation**

   After the UDP connection is established, the devices negotiate an IPsec tunnel over that connection.
6. **OSPF enablement and route exchange**

   Each device enables OSPF on the GRE interface, advertises the corresponding LAN networks, and learns routes to the peer LAN network.
7. **Data forwarding**

   After the routes are learned, the devices forward traffic through the established IPsec tunnel.

The following chapter describes the key technologies in more detail: SDN, GRE over

IPsec, OSPF, NAT traversal, and P2P hole punching.

2.1 SDN

2.1.1 Basic Concepts

SDN is a network architecture decoupling the control plane from the forwarding plane. The core idea is to extract the distributed, closed control logic from traditional network devices and manage it through a centralized (or logically centralized) software system, while the devices focus on efficient and reliable packet forwarding.

2.1.2 Benefits

Key benefits include:

- **Global visibility and centralized control**
  A centralized (or logically centralized) control plane gives the network a holistic view and enables consistent decision-making across the entire topology, which helps optimize resource allocation in complex, multi-path environments.
- **Enhanced flexibility and programmability**
  Network behavior is defined and adjusted through software interfaces and policy abstractions rather than device-by-device static configuration. This helps networks respond faster to changing business requirements and supports automation, dynamic adjustments, and centralized operations.
- **Reduced coupling and complexity**
  By separating control from forwarding, the control logic can evolve independently while the forwarding devices remain stable, improving scalability and long-term maintainability.

2.1.3 Market Demand

From a market perspective, SDN developed in direct response to the scalability, flexibility, and operational limitations of traditional networks.

- **Growing scale and complexity**
  Multi-branch, multi–data center, and multi-cloud environments make it difficult to efficiently manage traditional networks that rely on manual configuration and distributed control. The market urgently needs a network architecture that can provide a centralized perspective, unified control, and automation capabilities to reduce operating costs and improve network controllability.
- **Higher expectations for agility**

The widespread adoption of cloud applications, mobile offices, real-time services, and cross-regional access requires networks to rapidly adapt and change on demand. SDN enables the shift from a static infrastructure to a dynamic resource platform through software control and policy abstraction.

- **Closer alignment with business needs**

  Networks are no longer simply passively carrying service traffic; instead, they need to provide differentiated support based on service type, quality requirements, and security requirements. This shift demands that network control capabilities be closer to the service layer. SDN is a crucial foundation for achieving deep collaboration between networks and services.

- **Broader adoption in new models**

  In emerging network models, such as SD-WAN, cloud networks, and Network as a Service (NaaS), the market demand for SDN is further amplified. The centralized control, programmable interfaces, and abstraction capabilities provided by SDN have become an important prerequisite for supporting the implementation of these new network forms, rather than optional enhancements.

2.2 GRE over IPsec

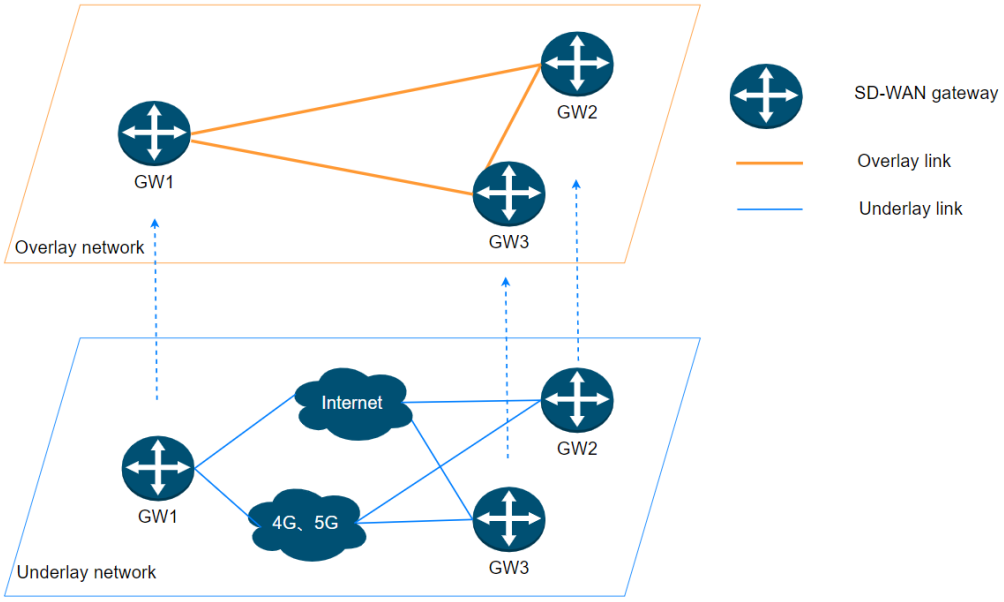

Figure 2: Layered Network Model

As shown in Figure 2, devices rely on physical links such as fiber, 4G, and 5G for reachability in the underlay network. The overlay network is a logical network built on top of the underlay and centrally controlled by the SD-WAN system. It carries business traffic, enforces policies, and supports traffic steering and link optimization. GRE over IPsec

provides a logical tunneling mechanism that assumes minimal underlay capabilities while ensuring secure data transmission.

2.2.1 Basic Concepts

GRE over IPsec is a layered tunneling approach that combines GRE's generic encapsulation with IPsec's security services to build a communication channel over untrusted networks, providing both flexible overlay transport and secure connectivity.

2.2.2 Benefits

- **Clear layering and separation of responsibilities**
  This is a key advantage of GRE over IPsec. Connectivity and security are handled by different mechanisms, which helps avoid overloading a single tunnel technology and improves architectural flexibility and scalability.
- **Layered transport/encapsulation**
  GRE can transparently encapsulate a variety of Layer 3 traffic types, allowing internal addressing, routing, and protocols to be carried without being constrained by the underlay.
- **Secure**
  IPsec provides standardized encryption and integrity protection for the GRE-encapsulated traffic, supporting secure communications over the public internet.

2.2.3 Market Demand

GRE over IPsec is widely adopted because enterprises need both secure connectivity and flexible overlay transport. As networks evolve from single campuses to multi-branch, multi–data centers, and multi-cloud environments, leased-line connectivity becomes less attractive due to cost and limited agility.

In this context, enterprises need a general-purpose tunneling solution that can run over the internet and multiple transport types while still providing data protection and isolation. GRE over IPsec preserves the overlay network's ability to abstract topology and carry traffic consistently across sites while ensuring security.

2.3 OSPF

2.3.1 Basic Concepts

OSPF is an interior gateway protocol (IGP) that uses a link-state algorithm to maintain topology information within an autonomous system and compute the best paths to destinations. By flooding link-state information, OSPF enables all participating nodes to converge on a consistent view of the topology, which helps produce stable and predictable

forwarding behavior.

OSPF's primary role is to describe topology and reachability, rather than to perform traffic scheduling or security functions directly. In modern architectures, OSPF often runs over tunnels or virtual networks to provide a reliable routing foundation for higher-level policy and control systems.

As an open and broadly supported protocol, OSPF remains one of the most commonly used dynamic routing protocols in enterprise and service-provider networks.

2.3.2 Benefits

Key benefits include:

- **Fast convergence and stability**
  As a link-state protocol, OSPF can quickly propagate topology changes and independently recompute paths on each node, which is well-suited for environments that require high network continuity.
- **Scalability through hierarchy**
  OSPF areas help limit the scope of routing updates and reduce control-plane overhead in large networks, resulting in clearer and more structured topology management.
- **Standards-based interoperability**
  As an open protocol, OSPF works well in multi-vendor and multi-device environments, making it easy to deploy uniformly in complex network architectures and serving as a common routing layer between different network technologies.

2.3.3 Market Demand

From a market perspective, OSPF's long-term viability stems from the continuous demand from enterprise and industry networks for stable and universal internal routing capabilities. As networks grow in scale and complexity, static routing gradually reveals its limitations in maintainability and scalability, necessitating dynamic routing.

At the same time, networks are evolving toward multi-branch, multi-region, and virtualized deployments, where topology changes are more frequent. This increases the need for routing that adapts quickly with minimal manual intervention—an area where OSPF aligns well.

In hybrid network environments with SD-WAN networking and internet-connected data centers, OSPF is still widely used to distribute baseline topology and routing information, providing stable input for centralized control, traffic steering, and policy decisions.

2.4 NAT Traversal and P2P Hole Punching

2.4.1 Basic Concepts

With the widespread adoption of CGNAT, multi-layer NAT, mobile-network NAT, and multi-WAN deployments, direct reachability between endpoints is often not available by default. Solutions that rely on public-IP direct access have limited success in practice. Pure relay-based approaches improve reachability, but they can increase latency, consume more cloud bandwidth, and raise long-term operating costs. As a result, NAT traversal has become a critical capability in remote access solutions. P2P hole punching is often used as a practical measure of real-world usability because it directly affects connection success rates, performance, and platform resource consumption.

NAT traversal and P2P hole punching refer to techniques that help two endpoints—each behind NAT—establish an end-to-end communication path by exchanging network parameters, identifying NAT characteristics, and attempting hole punching. The goal is to avoid long-term reliance on relays and to create a viable transport path for subsequent encrypted tunnel establishment.

NAT traversal requires cloud coordination and typically involves:

1. Each side discovering its current network information, including the publicly visible address, port mapping, and NAT type.
2. A cloud service coordinating and exchanging the connection parameters between both sides.
3. The endpoints choosing an appropriate hole-punching strategy based on the NAT environments and beginning the hole-punching process.

2.4.2 Benefits

Key benefits include:

- Improved end-to-end connectivity in NAT environments.
- Reduced ongoing dependence on relay forwarding.
- Shorter data paths, which can reduce latency and improve transfer efficiency.
- Lower long-term relay bandwidth and forwarding costs for the platform.

2.4.3 Market Demand

Environments without public IP addresses, with complex NAT, or with mobile access are now common in remote connectivity scenarios. Without effective NAT traversal and P2P hole punching, solutions must rely heavily on relays, which can significantly impact user experience, cost, and scalability. For home networks, micro and small businesses, remote offices, and remote operations, strong NAT adaptability through P2P hole punching has

become a foundational market requirement.

3 Key Technology Principles

3.1 SDN

3.1.1 Protocol Introduction

SDN is not a single protocol but a network architecture decoupling the control plane and the forwarding plane. Its core objective is to extract network control logic from the underlying devices and manage it centrally or logically through software while maintaining efficient and stable data forwarding. Actual interaction is primarily achieved through TLS, HTTPS, and HTTP/2 protocols.

3.1.2 Principle Analysis

The basic principle of SDN can be summarized as follows: the execution capability of "how to forward data" is retained in the network devices, while the decision-making logic of "why to forward data this way" is moved to a centralized or logically centralized control system. Therefore, SDN does not replace traditional protocols, but rather builds upon them, achieving collaborative optimization across devices, links, and regions through centralized control logic.

3.1.3 Framework Analysis

In the Omada solution, cloud access manages various devices via SDN, acting as the brain of the SD-WAN network.

3.1.4 Application Scenario Analysis

The software-defined control capabilities provided by SDN offer fundamental support for centralized orchestration, cross-region management, and automated operation and maintenance.

3.2 GRE over IPsec

3.2.1 Protocol Introduction

GRE is a general-purpose tunneling protocol that operates above the network layer. It encapsulates and forwards various network layer protocols, enabling logical connectivity across network environments. The core design goal of GRE is to provide a lightweight and universal encapsulation mechanism that allows original packets to be transparently transmitted through intermediate networks without altering their internal structure. RFC

1701 defines the GRE packet format, as shown in Figure 3.

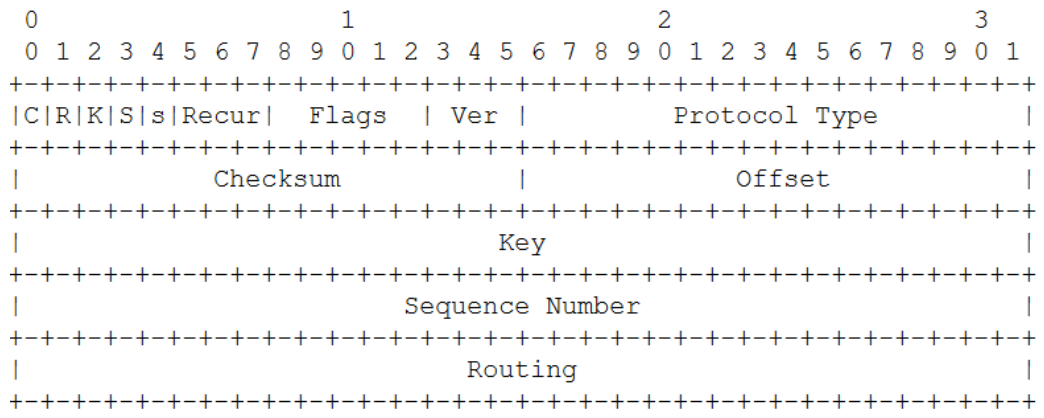

Figure 1: GRE Packet Format

The fields and their meanings are shown in Table 1:

Field	Length	Description
C	1 bit	Checksum verification bit. If this bit is 1, it indicates that a checksum field is inserted into the GRE header. If it is 0, the GRE header does not contain a checksum field.
R	1 bit	Route verification bit. If this bit is 1, it indicates that an offset field and a source route field are inserted into the GRE header. If either the checksum verification bit or the route bit is 1, or both bits are 1, it indicates that a checksum field and an offset field are inserted into the GRE header.
K	1 bit	Key verification bit. If this bit is 1, it indicates that a key field is inserted into the GRE header. If it is 0, the GRE header does not contain a key field.
S	1 bit	Sequence number verification bit. If this bit is 1, it indicates that a sequence number field is inserted into the GRE header. Otherwise, it indicates that the GRE header does not have a sequence number field.
s	1 bit	Strict source routing bit. RFC1071 recommends setting this field to 1 only when the routing information field contains strict source routing

		information.
Recur	3 bits	Indicates the number of encapsulation layers a GRE message is encapsulated in. This field prevents messages from being encapsulated an unlimited number of times. It increments by 1 after each GRE encapsulation operation. If the number of encapsulation layers is greater than 3, the message is discarded.
Flags	5 bits	Reserved field. It must be set to 0.
Ver	3 bits	Version field. It must be set to 0. Version 1 is used in PPTP (RFC2637).
Protocol Type	16 bits	Passenger protocol type. Its values and meanings are the same as "ETHER TYPES" in RFC1700.
Checksum	16 bits	Checksum field for the GRE header and its payload.
Offset	16 bits	Offset field. Represents the byte offset from the start byte in the active routing information field to the active source route entity (SRE). This field is only carried when the route identifier bit or checksum verification bit is set to 1, and is only valid when the route identifier bit is set to 1.
Key	32 bits	Keyword field. Used by the tunnel receiver to verify received messages.
Sequence Number	32 bits	Message sequence number field. A 32-bit unsigned integer inserted by the encapsulation node.
Routing	32 bits	Route information field. It is an optional field and appears only when the route identifier bit is set to 1. The routing information field contains a series of SREs.

Table 1: GRE Fields

From a protocol perspective, GRE does not focus on data confidentiality, integrity, or authentication. Its main function is to re-encapsulate source packets and achieve cross-network forwarding through an outer IP header. GRE is often considered a pure tunneling and interconnection technology, rather than a secure tunneling technology.

Therefore, GRE over IPSec combines the general encapsulation capabilities of GRE

with the data security capabilities of IPsec to achieve a tunnel system that simultaneously satisfies network connectivity, logical abstraction, and data security. In this combined mode, GRE is used to perform general encapsulation of internal network packets, allowing various service traffic to be transparently carried in the overlay network; IPsec is used to encrypt and protect the integrity of the traffic encapsulated by GRE, giving data anti-eavesdropping and anti-tampering capabilities when transmitted in the underlay network. Through a clear hierarchical division of labor, the two logically form an "encapsulation first, encryption later" tunnel structure.

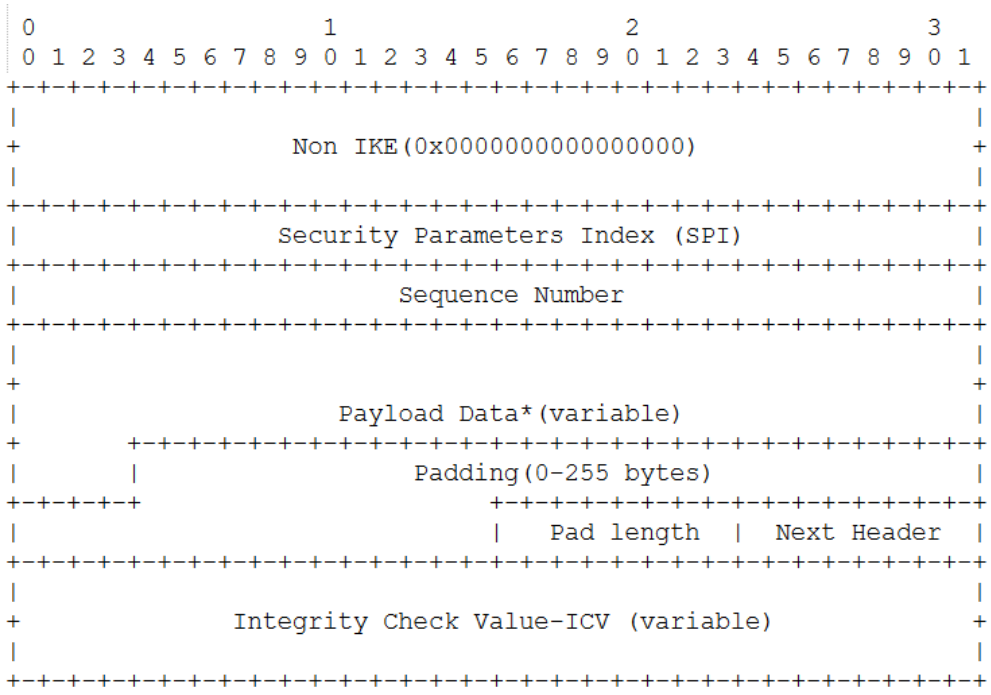

Figure 2: IPsec Packet Format

GRE over IPsec embodies a typical SD-WAN layered design concept:

- GRE is responsible for the construction and hosting of the overlay
- IPsec is responsible for data security protection on the underlay

This combination avoids a single protocol taking on too many roles, making the capabilities of each layer clearer and more replaceable.

3.2.2 Principle Analysis

The core principle of GRE over IPsec is to build a controllable and secure logical communication channel over an untrusted network through a separated design of encapsulation and encryption.

In the forwarding path, the original service packet first enters the GRE tunnel for encapsulation, forming a complete overlay payload. After GRE encapsulation, the packet retains the original service semantics and internal addressing structure, with only an

additional outer logical encapsulation header used for overlay forwarding. This stage achieves network abstraction and the construction of connectivity relationships.

Next, the entire GRE-encapsulated packet is processed by IPsec. Through encryption and integrity-check mechanisms, IPsec protects the data. IPsec does not inspect the specific traffic types carried inside GRE. Instead, it treats them as transparent payload and encrypts them for transmission, thereby creating a protected data channel over the underlay network.

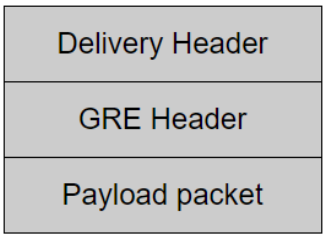

Figure 5: GRE Packet Structure

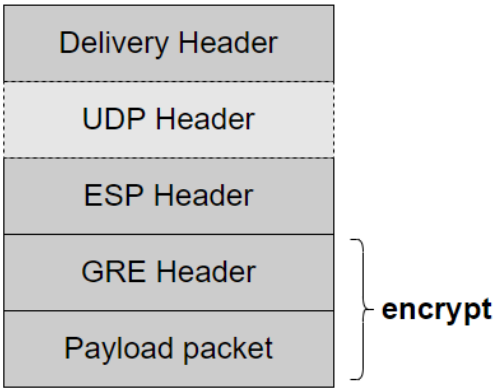

Figure 3: GRE over IPsec Packet Structure

3.2.3 Framework Analysis

From an SD-WAN architecture perspective, GRE over IPsec can be regarded as a typical implementation of a layered tunnel framework. Its overall structure can be abstracted into three cooperating layers:

1. **Overlay Transport Layer (GRE)**

This layer uses GRE tunnels to build logical virtual links for SD-WAN. With GRE, internal network addresses, routing relationships, and service traffic are encapsulated in a unified manner, allowing the control plane to schedule and manage traffic without relying on the underlay network's capabilities.

2. **Data Security Layer (IPsec)**

This layer provides encryption and integrity protection for the overlay data carried by GRE. It ensures service traffic has the necessary security properties when transmitted across an untrusted network environment while hiding internal communication details.

3. **Underlay Transport Layer (Base IP Network)**

The underlay network is responsible only for forwarding the encrypted IP packets

between tunnel endpoints. It does not interpret the traffic carried inside the tunnel, nor does it participate in control policies or the overlay topology.

3.2.4 Application Scenario Analysis

In typical SD-WAN deployments, GRE over IPsec is widely used for secure interconnection across the public internet, multiple links, and heterogeneous network environments. For interconnection requirements between enterprise branches and headquarters, or between data centers and cloud environments, this combined tunnel can both provide a stable and unified overlay network view and meet the data security requirements for communication over the public internet.

In multi-link SD-WAN scenarios, GRE tunnels can be carried over different underlay links, while IPsec provides consistent data protection for each tunnel. The SD-WAN system can dynamically schedule traffic based on tunnel quality to achieve load balancing, service steering, and fast switchover, without being affected by differences in underlying security mechanisms.

Moreover, in environments that require both complex service transport and security/compliance, GRE over IPsec combines logical encapsulation and security protection to enable enterprises to quickly build a WAN with unified control and encryption protection without modifying existing service systems. This characteristic makes it one of the commonly used and mature tunnel implementations in SD-WAN solutions.

3.3 OSPF Technology

With the GRE over IPsec tunnel, SD-WAN builds a stable transmission channel over untrusted networks that provides both logical encapsulation and data security. However, a reliable tunnel underlay is not sufficient to form a complete WAN system. In real-world deployments with multiple nodes and multiple paths, a mechanism is still required to perceive the network topology, maintain reachability, and quickly adjust paths when links change.

In this context, dynamic routing protocols become a key component in the GRE over IPsec tunnel system. Among them, OSPF, as a mature and standardized link-state routing protocol, is often used in SD-WAN encapsulated networks to carry overlay topology information and provide basic reachability control.

3.3.1 Protocol Introduction

OSPF (Open Shortest Path First) is a link-state interior gateway routing protocol. It is mainly used to advertise network reachability information within an autonomous system and to select optimal forwarding paths through a consistent topology calculation process.

Functionally, OSPF's core responsibility is to maintain a consistent view of the network topology, rather than to perform data forwarding or provide security protection.

In SD-WAN scenarios, OSPF often runs over established GRE or GRE over IPsec tunnels and is used as a control protocol in the overlay network. Its role is no longer limited to path computation in traditional physical networks; it is now also used to describe reachability between logical tunnels, connectivity of edge nodes, and area-based topology structures.

OSPF packets are carried directly in IP packets with protocol number 89. There are five packet types: Hello, Database Description (DD), Link State Request (LSR), Link State Update (LSU), and Link State Acknowledgment (LSAck). These five packet types share the same 24-byte header format, as shown below.

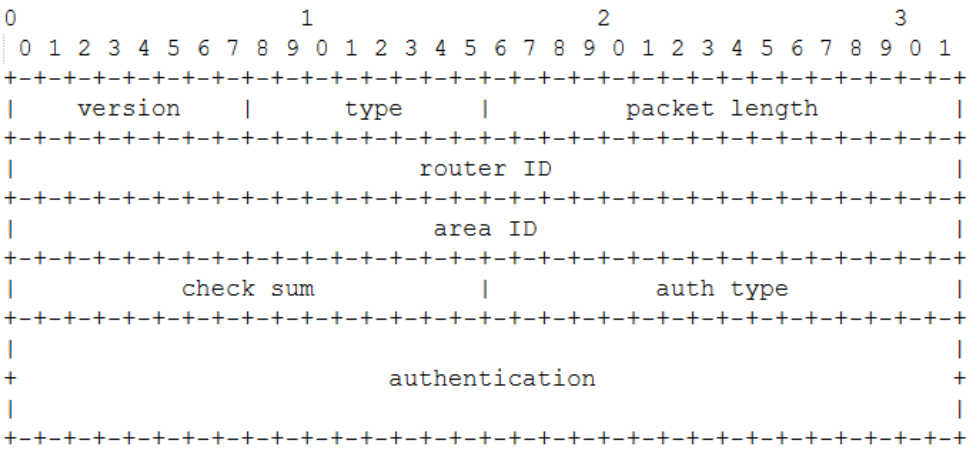

Figure 4: OSPF Protocol

The fields and their meanings are shown in Table 2:

Field	Length	Description
Version	1 byte	OSPF version. For OSPFv2, the value is 2.
Type	1 byte	OSPF packet type: 1 = Hello; 2 = DD; 3 = LSR; 4 = LSU; 5 = LSAck.
Packet length	2 bytes	Total length of the OSPF packet (including the header), in

		bytes.
Router ID	4 bytes	Router ID of the sending router.
Area ID	4 bytes	Area to which the packet belongs.
Checksum	2 bytes	Checksum of the entire packet excluding the authentication field.
Auth Type	2 bytes	Authentication type: 0 = None; 1 = Simple password; 2 = Cryptographic authentication.
Authentication	8 bytes	Authentication field, dependent on Auth Type: Undefined when Auth Type = 0; Password when Auth Type = 1; Includes Key ID, authentication data length, and sequence number when Auth Type = 2. This field contains only the authentication data length, not the authentication data itself.

Table 2: OSPF Fields and Meanings

3.3.2 Principle Analysis

OSPF operates based on link-state advertisements (LSAs) and global topology computation. Each node advertises its local connectivity, contributing to a complete logical topology view. Based on this shared topology database, each node computes the best path to each destination network.

In an SD-WAN architecture combined with GRE over IPsec, the operating environment of OSPF changes significantly:

- The “links” it perceives are not physical interfaces, but logical tunnels.
- The “adjacencies” it maintains reflect connectivity between tunnel endpoints.
- The paths it computes represent logical forwarding entries within the overlay network.

Therefore, in this scenario, OSPF functions more like an overlay topology synchronization and convergence mechanism. When a tunnel state changes, OSPF can

detect it quickly at the logical layer and adjust path selection, providing continuous reachability for upper-layer services.

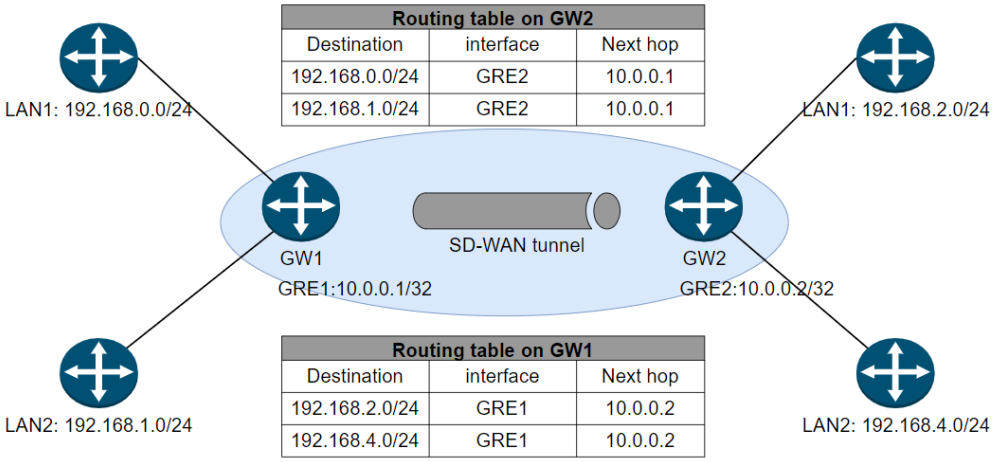

Figure 8: OSPF Principle

As shown in the figure above, GW1 and GW2 first use OSPF to advertise their local private subnets to each other over the SD-WAN tunnel. Then they learn routes to the peer private subnets and install them into their respective routing tables, thereby enabling private network connectivity between GW1 and GW2.

3.3.3 Framework Analysis

In the SD-WAN system, OSPF typically sits between the tunnel system and policy control, providing basic routing and topology synchronization. Logically, its role can be summarized as a three-layer collaboration:

1. **Lower layer: Tunnel and transport system (GRE over IPsec)**

   Provides stable, encrypted logical links, enabling secure point-to-point communication between nodes.

2. **Middle layer: Routing and topology synchronization (OSPF)**

   Maintains node relationships and reachability information in the overlay network over the tunnels, ensuring a consistent network-wide view of topology changes.

3. **Upper layer: SD-WAN policy decision and control**

   Based on the reachability provided by OSPF, dynamically steers traffic according to application identification, performance metrics, and policy rules.

Within this framework, OSPF does not perform intelligent steering; instead, it provides a stable, generic, and standardized topology foundation for the SD-WAN control system. This decouples complex service policies from the underlying routing computation logic, improving overall maintainability and scalability.

3.3.4 Application Scenario Analysis

In practical SD-WAN deployments, OSPF is mainly used in scenarios with high requirements for topology consistency and basic reachability. For example, in enterprise networks with multiple branches and multiple regions, OSPF can be used to maintain node relationships in the overlay network and reduce management complexity caused by static configuration.

After GRE over IPsec tunnels are established, OSPF can automatically detect connectivity changes of each tunnel and recompute paths at the logical layer, ensuring continuous forwarding of service traffic. This capability is especially important in scenarios where links change frequently or branch nodes join dynamically.

For interconnection with traditional networks or data center environments, OSPF—as a mature routing protocol—helps integrate the overlay network with existing routing systems, supporting a smooth SD-WAN adoption into the existing network architecture.

Overall, OSPF in SD-WAN primarily serves as a provider of basic routing and topology convergence rather than the ultimate traffic decision-maker. Its value lies in providing a stable and predictable structural foundation for the overlay network.

3.4 NAT Traversal and P2P Hole Punching

In SD-WAN networks, branch sites often use private IP addresses to conserve public IP resources. Only after NAT translates private IP addresses into public IP addresses can users at the site access other sites. When packets sent by a gateway pass through a NAT device, the IP address changes. If the post-NAT public IP address cannot be obtained, an SD-WAN tunnel cannot be established between gateways.

NAT traversal and P2P hole-punching techniques are designed to address this problem. Their goal is to establish an end-to-end path as much as possible in complex network environments, providing a base channel for subsequent secure communication.

3.4.1 Protocol Introduction

STUN (Session Traversal Utilities for NAT) is a lightweight network protocol used to address P2P communication issues caused by NAT devices. The evolution of the related RFCs is as follows:

RFC 3489 (obsolete) → RFC 5389 (current standard) → RFC 7350 (update)

Although RFC 5389 improves and extends the protocol, the core mechanisms defined in RFC 3489—such as the Binding Request/Response model, NAT type detection, and message integrity protection—remain the foundation of current NAT traversal techniques. In practice, STUN is often used together with technologies such as TURN and ICE to form a

complete NAT traversal solution.

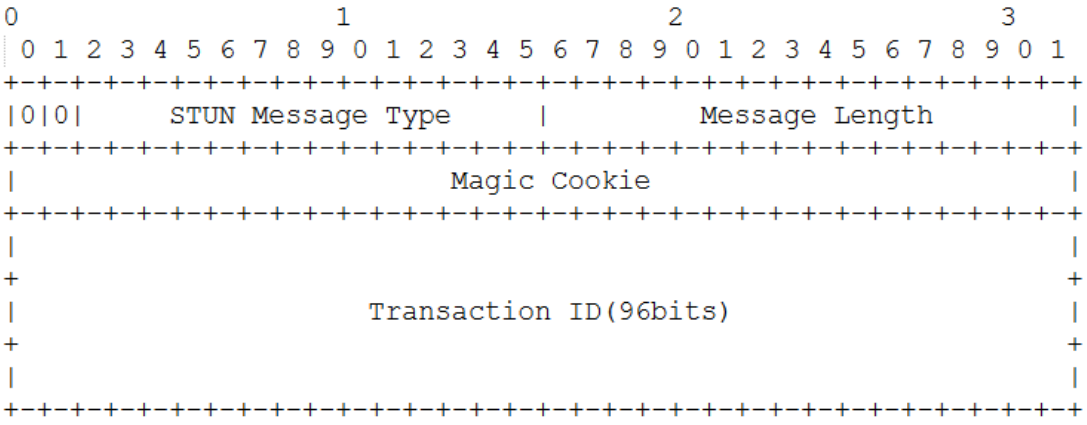

Figure 5: STUN Header

- The first 2 bits are 0, used to distinguish STUN from other protocols.
- STUN Message Type: Defines the message class and method.

After the 20-byte STUN header, there are zero or more attributes. Each attribute has a variable length, but the total length is always a multiple of 4 bytes. The attribute structure is shown below.

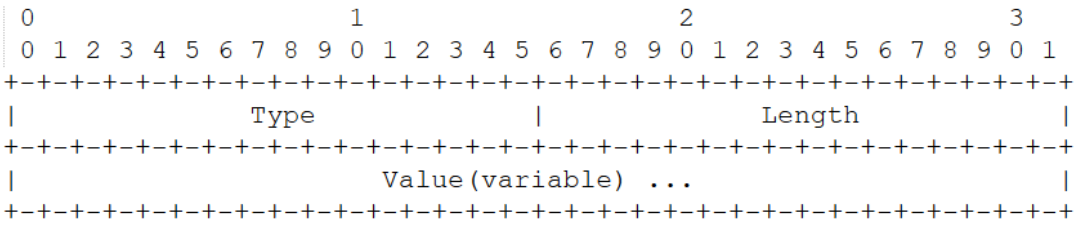

Figure 6: STUN Attributes

- Type: 16-bit attribute type
- Length: Attribute length (length of Value) in bytes; must be a multiple of 4 bytes
- Value: Attribute data in network byte order

3.4.2 Principle Analysis

The P2P hole-punching mechanism depends on the NAT types and specific NAT policies on both sides, as described below:

1. NAT Types

NAT types refer to the four types of NAT: Full Cone NAT, IP Restricted Cone NAT, Port Restricted Cone NAT, and Symmetric NAT. In P2P and VoIP connections, different NAT types significantly affect the establishment of external connections.

(1) Full Cone NAT

Full Cone NAT is the most permissive type: it allows any external host to communicate with an internal host as long as the internal host has initiated communication.

a) When an internal host initiates a connection, the NAT maps the internal IP/port to an external IP/port.
b) The mapping is static: all traffic from the same internal IP/port maps to the same external IP/port.
c) Any external host that knows the mapped external IP/port can send traffic to the internal host via that mapping.

(2) IP-Restricted Cone NAT

IP-Restricted Cone NAT adds restrictions on external access: only external hosts that have previously communicated with the internal host can continue to reach it.

a) Outbound connections are mapped to an external IP/port.
b) The NAT allows inbound traffic only from an external IP address that the internal host has previously sent to.
c) If an external host's IP address has not previously communicated with the internal host, packets from that IP are dropped.

(3) Port-Restricted Cone NAT

Port-Restricted Cone NAT further restricts access by requiring the external host to use the same source port used in the previous communication.

a) Outbound connections are mapped to an external IP/port.
b) The NAT allows inbound traffic only from the same external IP address and port that the internal host previously sent to.
c) If the external host uses a different port, the NAT rejects the traffic.

(4) Symmetric NAT

Symmetric NAT is the most restrictive type. Each connection from an internal host creates a different mapping. Even for the same internal IP/port, different external destinations result in different external mappings.

a) Each outbound connection request creates a unique external IP/port mapping.
b) Only the specific external host associated with that mapping can send traffic back through it.
c) Requests from the same internal host to different external destinations generate different external mappings.

2. P2P Hole-Punching Process

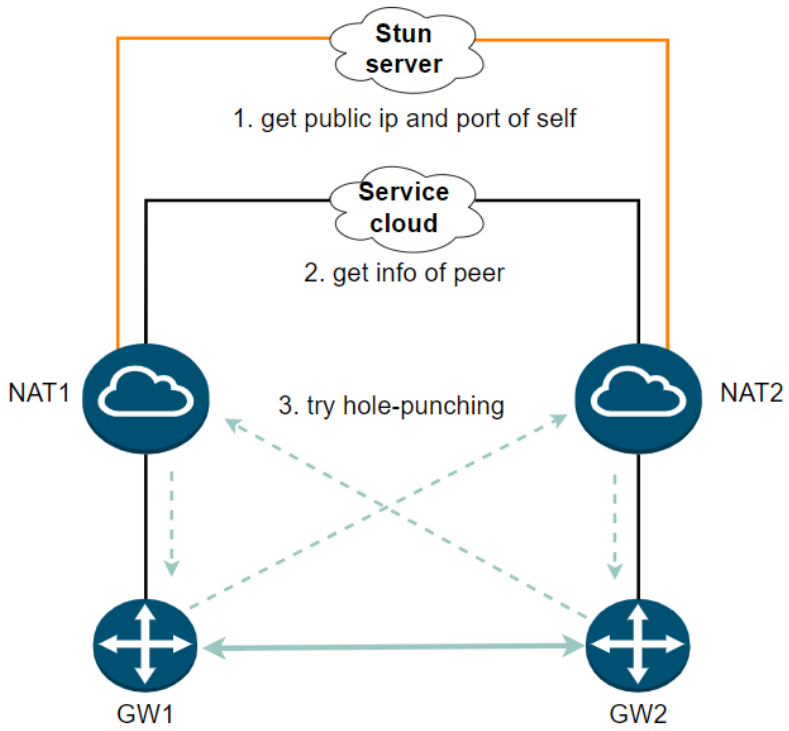

Figure 7: P2P Hole-Punching Process

1. **Obtain public-facing information**

The client and the device each use a STUN-like mechanism to obtain their current public-facing address and port.

2. **Exchange connection parameters**

Both sides exchange this public-facing information and other required parameters via the cloud.

3. **Enter the hole-punch preparation state**

The client and the device prepare to initiate peer-to-peer communication attempts based on the peer information.

4. **Send UDP packets to the peer**

Under coordination, both sides send UDP packets to each other's public-mapped addresses almost simultaneously.

5. **Form a reachable path**

If the returning packets satisfy the current NAT conditions, a reachable underlying path is established.

6. **Establish the subsequent encrypted tunnel**

After the underlying path is established, an encrypted tunnel is built based on it.

Additionally, when NAT types are known, different strategies can be applied in advance for different type combinations to reduce redundant hole-punch attempts and improve success rates. Under the above types, if the peer is behind Symmetric NAT, direct traversal

is generally difficult. However, in practice, techniques such as port prediction can achieve a relatively high probability of success.

3.4.3 Connectivity Boundaries and Reachability Enhancement Strategies

As indicated by the analysis above, whether P2P hole punching succeeds depends not only on whether the client and device have completed public-facing address discovery and connection-parameter exchange, but also on whether the network environments on both sides allow a stable return path. NAT mapping behavior, inbound traffic filtering rules, UDP mapping lifetime, carrier network policies, and changes in endpoint network state can all affect the final result.

Therefore, in the current SD-WAN version, multiple reachability enhancement strategies are introduced during the P2P connection process to improve direct-connect success rates and recovery capability.

1. Multi-round hole punching with coordinated connection setup

Under cloud coordination, the client and device perform multiple rounds of UDP hole-punching attempts. Continuous probing exchanges, NAT mapping refresh, and connection-state synchronization increase the probability of forming a usable communication path. This mechanism improves adaptability in complex NAT environments, but its effectiveness still depends on actual network policies and NAT behavior.

2. Multi-candidate path probing and dynamic retries

In situations with complex NAT mappings, unstable link states, or public-side reachability changes, the system can combine probing across multiple candidate paths, dynamic retries, and connection-state feedback to increase the probability of discovering a working communication path. This mechanism enhances path discovery in complex networks but does not guarantee direct-connect success under all NAT types or network policies.

3. Keepalive and mapping maintenance

After a P2P path is established, the device can send periodic keepalive packets to maintain NAT mappings, reducing the probability of disconnection caused by UDP mapping timeouts. Keepalive behavior should be balanced against network load, endpoint power consumption, and link stability to achieve an appropriate trade-off between reliability and overhead.

4. Path re-probing and reconnection

When the system detects connection errors, unreachable links, network switching, or connection state changes, it can re-trigger public-facing address discovery, parameter

exchange, and the P2P hole-punching process. This mechanism helps adapt to dynamic environments such as mobile network switching, multi-WAN egress changes, and NAT mapping changes, improving recovery capability.

These mechanisms can improve P2P direct-connect success rates and connection stability, but they cannot fully eliminate the risk of direct-connect failure in complex NAT, CGNAT, or strict firewall scenarios. Therefore, the NAT traversal capability in the current version should be positioned as a reachability enhancement capability with a "P2P direct-connect first" approach. For scenarios where direct connectivity still cannot be established, future versions will further improve overall connection availability through a relay mechanism.

3.4.4 Relay Capability Planning in Future Versions

To further improve connection availability in complex network environments, future SD-WAN versions plan to introduce a relay mechanism as a complement to P2P direct connectivity.

With this mechanism, the system will still prioritize P2P direct connections. When the network environment does not support end-to-end direct connectivity, the client and device can establish an indirect communication path through a relay node, which forwards traffic between them. Since relay does not depend on direct reachability between the client and device, it can cover a wider range of complex network environments.

A P2P direct path typically provides a shorter path, lower latency, and reduced platform-side bandwidth consumption. A relay path improves availability but introduces an additional forwarding hop, which may result in higher latency, lower throughput, and increased cloud resource consumption. Therefore, future connection policies will prioritize P2P direct connectivity, using relay as an availability fallback, balancing connection success rate, performance, and platform resource usage.

3.4.5 Framework Analysis

Architecturally, NAT traversal and P2P hole punching rely on the coordination among the client, the device, and the cloud. The client and device each obtain information about their respective network environments and perform the actual hole-punching attempts. The cloud, meanwhile, exchanges the parameters required for connection setup and performs coordination. The relationship in the system can be summarized as:

1. The cloud is responsible for parameter exchange and coordination.
2. The client and device are responsible for collecting endpoint-side network information and performing peer-to-peer hole punching.

This module sits between cloud-coordinated access and encrypted tunnel establishment, serving as the "reachability bridge" of the overall access system.

3.4.6 Application Scenario Analysis

NAT traversal and P2P hole punching are especially suitable for remote-access scenarios where both sides are behind private networks, including home broadband, mobile networks, CGNAT, multi-level NAT, and multi-WAN environments. Without effective NAT traversal, remote access solutions often have to rely on relay forwarding, which imposes clear limitations on latency, bandwidth cost, and scalability.

Therefore, NAT traversal capability directly affects the practicality of remote access solutions in real-world network environments and is a key foundation for evaluating system reachability and efficiency.

4 Application Scenarios and Solutions

Under the accelerated push toward enterprise digital transformation, distributed business architectures have become the norm, and networking demands for multi-branch, cross-region, and hybrid-cloud environments continue to surge. When faced with challenges such as rapid growth in distributed sites, increasingly complex heterogeneous network environments, and more dynamic business traffic, traditional WAN architectures expose key pain points:

- Long deployment cycles,
- High O&M costs,
- Insufficient reliability,
- Blurred security boundaries.

This chapter explains the core architectural mechanisms and scenario-based deployment strategies of our SD-WAN solution, providing technical support for enterprises to build a resilient, intelligent, and secure next-generation WAN.

4.1 Omada SD-WAN Solutions

The Omada SD-WAN solution can be categorized by device and management approaches into the ER Series standalone gateway solution and the Fusion Series cloud-managed solution. The two solutions differ in applicable scenarios, management modes, network topologies, and recommended node scale. Users can choose based on network size, deployment cost, management method, and business reliability requirements.

4.1.1 ER Series Standalone Gateway SD-WAN Solution

The ER Series solution applies to standalone gateway devices such as the ER8411, ER7206 v2, and ER605 v2. This solution relies on a Local Omada Controller or a Cloud-Based Controller for unified centralized management and is suitable for multi-site networking scenarios that require network stability, deployment flexibility, and cost control.

Under this solution, SD-WAN supports a Hub-and-Spoke topology. In this mode, ER Series devices allow users to configure whether to establish direct connections between Spokes. With direct Spoke-to-Spoke connectivity enabled, traffic can bypass the Hub and does not need to be relayed through it.

- ***Typical Network Topology***

  Hub -Spoke Topology

The Hub-Spoke topology is suitable for typical scenarios where the headquarters interconnects with multiple branches. The central node acts as the Hub, carrying primary business traffic and inter-branch connectivity traffic. Each branch node acts as a Spoke and establishes a connection to the Hub through an SD-WAN tunnel.

Users can select a high-performance gateway as the Hub—such as ER8411—based on the central node's concurrent connection capacity, throughput, business load, and overall network scale. For branch nodes, where the business load is typically lower, users can flexibly choose more cost-effective devices such as ER605 v2, thereby reducing deployment cost while maintaining overall network stability.

This mode is well-suited for building a highly stable and cost-effective multi-branch SD-WAN network. The network scale can reach up to 600 nodes.

- ***Prerequisites***

For ER Series SD-WAN deployment, all participating ER gateways must be managed by the same Local Omada Controller or Cloud-Based Controller and belong to the same organization. Only after unified management requirements are met can SD-WAN topology configuration and policy provisioning be completed on the controller side. The ER gateway acting as the Hub, as well as any Spoke devices that require direct connections, must have a public IP address.

- ***Network Setup Entry***

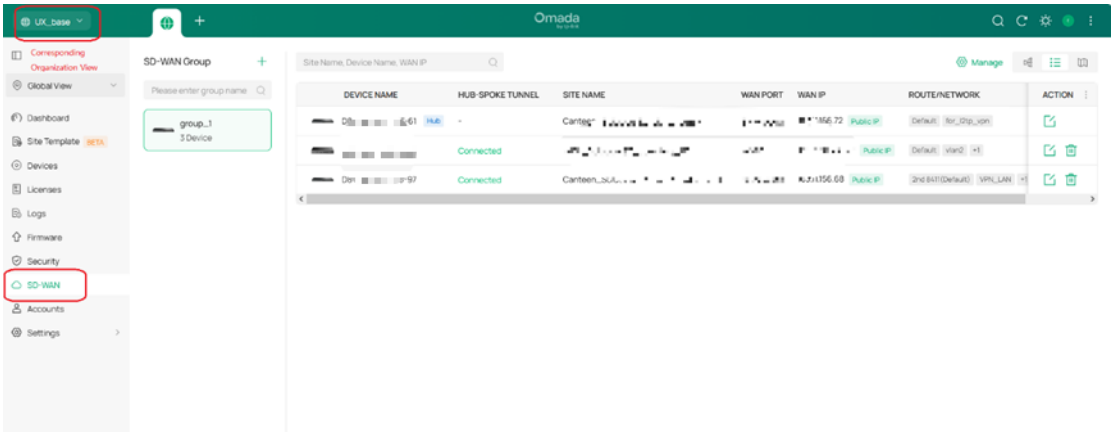

Figure 12: ER Series SD-WAN Solution Network Setup Entry

4.1.2 Fusion Series SD-WAN Solution

The Fusion Series solution targets cloud-managed deployment scenarios and currently supports SD-WAN networking only when devices are under cloud management.

- ***Typical Network Topology***

  Full-Mesh Topology

  The Fusion Series currently supports a **Full-Mesh** topology. On the SD-WAN feature page in the Cloud Portal, users can select Fusion devices that have already been onboarded to the cloud and complete SD-WAN interconnection configuration among multiple nodes.

  This solution has a relatively simplified deployment process and is suitable for scenarios where users want to quickly enable multi-site interconnection via the cloud and reduce the deployment and maintenance costs of an on-premises controller. No device is required to have a public IP address.

  At present, the recommended node count for Fusion Series Full-Mesh deployments is fewer than 20 nodes.

- ***Prerequisites***

  Users must deploy one Fusion device at each site and complete cloud onboarding using a TP-Link ID (devices only need to be bound to the same TP-Link ID; the onboarding method is not limited to a specific approach), so that the devices can access the Cloud Portal for unified management.

- ***Network Setup Entry***

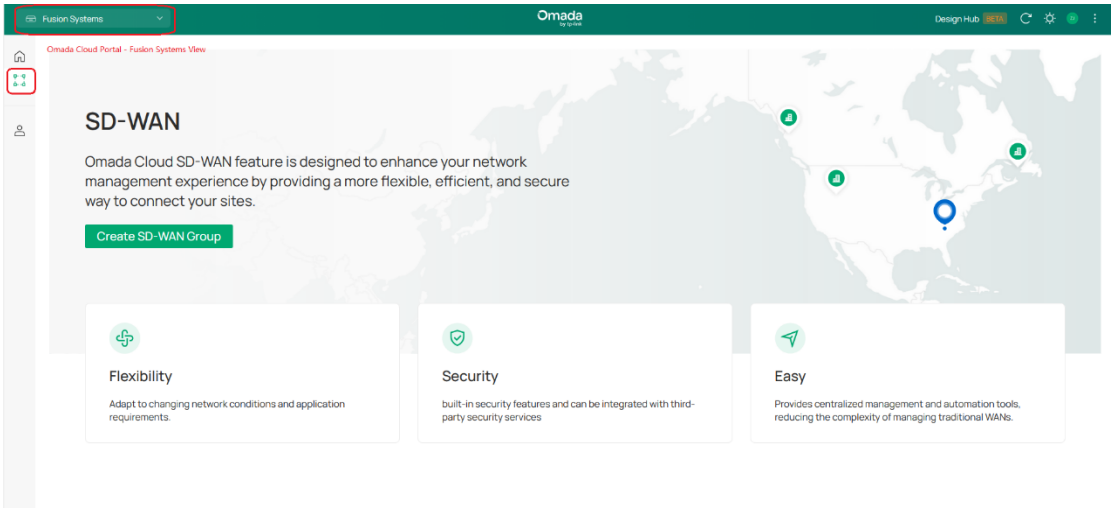

Figure 13: Fusion Series SD-WAN Solution Network Setup Entry

4.2 Typical Solution Scenarios

For scenarios with a relatively large number of nodes, where a headquarters site must interconnect with a large number of branches and both device performance and deployment cost must be balanced, the ER Series Hub-Spoke solution is recommended as the first choice. By pairing a high-performance Hub gateway with cost-effective Spoke gateways, this solution achieves a balance between stability and cost efficiency.

For scenarios with a smaller node scale, where rapid cloud deployment is preferred and direct interconnection among sites is required, the Fusion Series Full-Mesh solution can be selected. This solution relies on the Cloud Portal for unified configuration and is suitable for lightweight, cloud-managed multi-site SD-WAN deployments.

Below are two representative example scenarios.

4.2.1 Small-to-Medium ePOS Retail Chains

In franchise-based ePOS retail chain scenarios, multi-store owners typically deploy one Fusion device at each store, allowing it to serve as both the controller and gateway. Each store site can be quickly onboarded via Bluetooth and bound to the same TP-Link ID, enabling cloud access and unified management.

After devices are onboarded, store owners can log in to the Cloud Portal, select the relevant store sites on the SD-WAN feature page, and quickly complete SD-WAN networking configuration among branches. This deployment method requires no additional on-premises controller and does not depend on a large data center or a headquarters Hub node to participate in networking, effectively reducing deployment complexity and O&M costs.

This type of scenario primarily focuses on interconnecting multiple stores, with a

controllable network scale—typically fewer than 20 nodes. Therefore, the Fusion SD-WAN solution is a better fit for lightweight, multi-branch, and cloud-managed retail chain scenarios.

4.2.1.1 Network Topology

Figure 14 shows a Full-Mesh SD-WAN topology for small-to-medium chain retail scenarios: each store deploys one Fusion device and is managed centrally through cloud management (via a cloud-based controller). Based on SD-WAN, full-mesh interconnection tunnels are established among stores to enable direct branch-to-branch communication without relying on a headquarters Hub node for traffic relay.

This topology is suitable for multi-store retail scenarios with typically fewer than 20 nodes and without a large data center or headquarters hub. It reduces topology planning and on-site O&M complexity while enabling fast, secure, and lightweight branch interconnection.

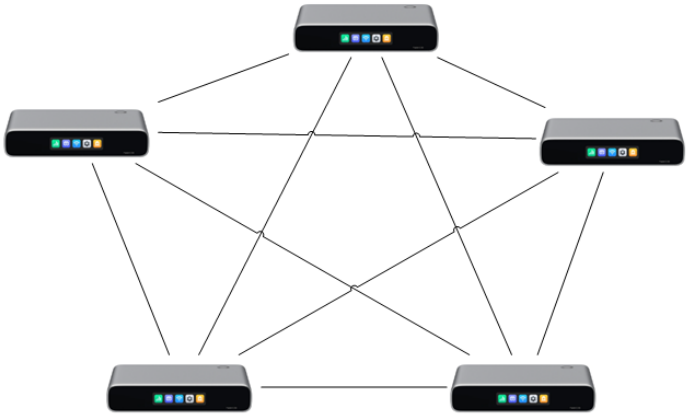

Figure 14: Full-Mesh SD-WAN Topology for Small-to-Medium Chain Retail

4.2.1.2 Pain Points and Requirements

In small-to-medium ePOS retail chains, individual multi-store owners, and franchise scenarios, the number of stores often grows in a distributed manner, and sites are geographically dispersed, while users' overall network O&M capability and professional networking experience are relatively limited. Therefore, during multi-branch WAN deployment and daily operations, the following typical pain points are encountered.

1. Complex multi-branch deployment and high on-site configuration costs

In WAN expansion and deployment scenarios with dozens or even hundreds of branch nodes—such as retail stores, franchise stores, and remote offices—traditional routing and switching devices typically rely heavily on on-site network engineers for initial configuration and commissioning.

Such devices often require per-device configuration via CLI or local web UI for basic

network parameters, routing policies, VPN parameters, and security policies. This approach leads to long deployment cycles and low implementation efficiency, and it demands strong networking expertise from on-site personnel.

When facing a large number of repetitive configuration tasks, manual configuration is also prone to input errors or parameter inconsistencies, causing network issues such as subnet conflicts, incorrect routing configuration, VPN tunnel establishment failures, or inter-site service inaccessibility—thereby increasing subsequent troubleshooting and O&M costs.

2. Complex Hub-Spoke planning with high skill requirements

In classic Hub-Spoke star WAN topologies, Spoke branch nodes typically interconnect and forward traffic through a headquarters Hub/central node. While this mode suits traditional headquarters-to-branch scenarios, real deployments require comprehensive planning that considers business traffic, link quality, device performance, and site scale.

For example, users must select an appropriate Hub node and assess the Hub device's throughput capacity, concurrent tunnel capacity, and business load. They must also decide whether traffic between different Spokes should be relayed through the Hub or whether direct Spoke-to-Spoke connectivity should be used to improve access efficiency.

For individual multi-store owners, franchisees, or small IT teams with limited network O&M capabilities, topology planning, central node selection, load evaluation, and path design have relatively high learning and implementation barriers, which can negatively affect deployment efficiency and network stability.

3. Multi-layer NAT is common, and traditional VPN success rates are limited

These users generally do not have public IP addresses. In real-world WAN environments, many edge branch stores are deployed behind carrier-grade NAT, dynamic public addressing, or enterprise multi-layer NAT. Site devices often cannot obtain fixed, globally routable public IP addresses.

Under these network conditions, traditional IPsec VPN solutions heavily depend on public addressing, port mapping, and reachability between both ends. When both VPN endpoints are behind NAT and lack public IPs or fixed port mappings, devices may fail to complete IKE negotiation, key exchange, and tunnel establishment, resulting in failed inter-site VPN networking or reduced stability.

Therefore, for small-to-medium multi-store retail scenarios, an SD-WAN solution must provide stronger NAT traversal capability, automated networking, and cloud-based unified orchestration to reduce reliance on public IP addresses, on-site configuration, and professional network personnel.

4.2.1.3 Fusion Series SD-WAN Solution

The Fusion Series SD-WAN solution is designed for multi-branch scenarios such as small-to-medium ePOS retail chains, individual multi-store owners, and franchisees. Through centralized cloud management, Full-Mesh topology, automated configuration provisioning, and secure tunnel establishment, it reduces traditional WAN deployment dependence on professional network staff, fixed public IP addresses, and on-site manual configuration—helping users quickly achieve secure interconnection among multiple stores.

1. Centralized cloud management to reduce multi-branch deployment complexity

Fusion devices integrate both controller and gateway capabilities, eliminating the need for an additional on-premises controller at store sites. After power-on and Bluetooth-based quick onboarding, devices can be bound to the same TP-Link ID and connected to the Cloud Portal for unified management by a cloud-based controller.

During SD-WAN deployment, branch routing configuration, VPN tunnel parameters, site interconnection relationships, and network policies are centrally orchestrated in the cloud and provisioned uniformly. Users do not need to log in to each device for CLI configuration. This transforms the traditional on-site, engineer-dependent deployment model into a remote, centralized, visual unified configuration and management approach—shortening service rollout time and reducing the risk of manual configuration errors.

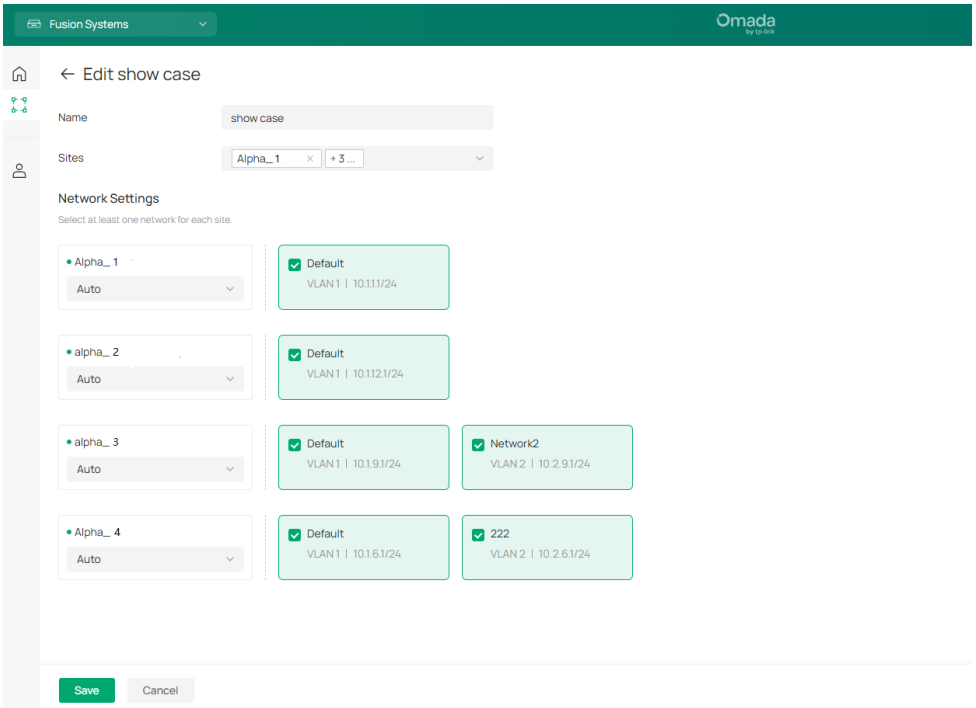

Figure 15: Cloud Portal Page for Unified Configuration of Fusion Full-Mesh SD-WAN

2. Automatic Full-Mesh networking to simplify topology planning and improve branch-

to-branch efficiency

Fusion SD-WAN supports a Full-Mesh topology, allowing direct interconnection tunnels between stores without relying on a headquarters Hub node for centralized relay.

In individual multi-store owner and franchise retail scenarios, there is typically no large data center or headquarters Hub, and the number of stores is usually fewer than 20. Full-Mesh mode avoids complex tasks such as manually planning a Hub node, evaluating central-node load, or deciding between Spoke direct connectivity and Hub relay paths—lowering the barrier to topology design.

At the same time, inter-branch traffic can be forwarded directly, reducing detours typical of Hub-Spoke deployments and improving lateral access efficiency between stores.

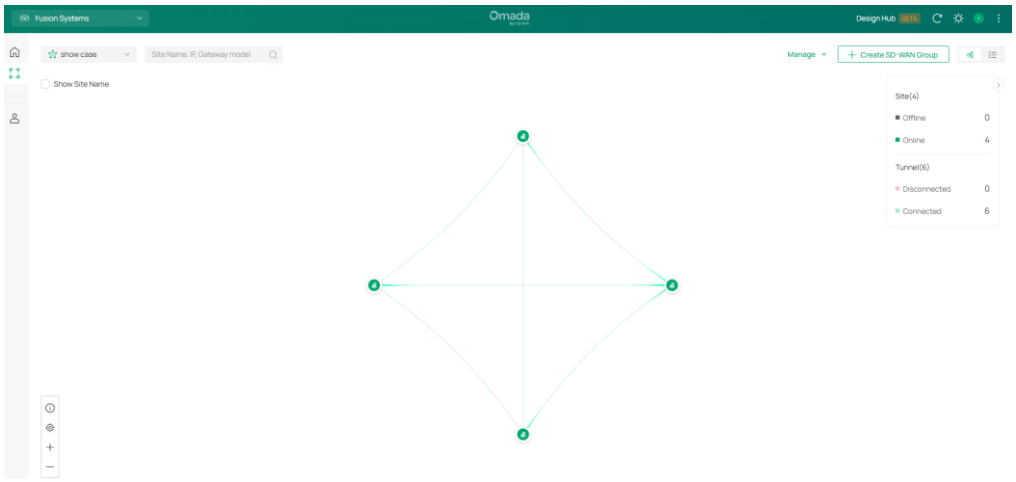

Figure 16: SD-WAN Full-Mesh Topology Diagram

3. Cloud-based signaling coordination and secure tunnel establishment to improve success rates in complex NAT environments

In many retail stores and remote branch scenarios, devices are often behind carrier-grade NAT, dynamic public addressing, or enterprise multi-layer NAT. When neither end has a fixed public IP address, traditional IPsec VPN tunnels can be difficult to establish due to peer reachability issues or IKE negotiation failures.

Fusion SD-WAN leverages a cloud-based controller for centralized signaling exchange. Branch devices maintain secure connections to the cloud through encrypted control channels and automatically obtain global topology information, routing configuration, and tunnel parameters. The controller assists with site discovery, topology synchronization, and tunnel orchestration, reducing reliance on fixed public IPs, port mapping, and manual VPN parameter configuration.

With this mechanism, Fusion SD-WAN improves networking success rates and deployment flexibility in complex NAT environments, providing more stable and low-barrier

secure interconnection for multi-store scenarios.

The Fusion Series also supports features such as underlay link redundancy and LAN conflict detection. Refer to the next section for more details.

4.2.2 Medium-Sized Logistics Enterprises

For medium-sized logistics enterprises that have an enterprise data center or a typical Hub-node requirement, such enterprises have branches nationwide. The number of nodes typically exceeds 100 and branch-site traffic load is relatively low, while headquarters/Hub traffic load is higher.

In this scenario, ER6 Series gateways can be used to meet branch access requirements, while ER8 Series gateways handle the high throughput demand at the Hub node. All nodes are centrally managed via an Omada Controller at the headquarters or a Cloud-Based Controller, enabling centralized policy provisioning and network orchestration.

This deployment mode allows the ER Series standalone gateway SD-WAN solution to ensure network stability while delivering better cost efficiency, reasonable load distribution, and efficient interconnection of large-scale branch nodes—meeting the enterprise's nationwide multi-branch networking requirements.

4.2.2.1 Network Topology

A Hub-Spoke SD-WAN topology for a medium-sized logistics enterprise typically uses the enterprise headquarters or data center as the Hub node, with logistics branches, outlets, or offices nationwide connecting as Spoke nodes.

The Hub node can deploy a high-performance ER8 Series gateway to carry centralized business access, inter-branch connectivity, and unified egress traffic. For Spoke sites with frequent and heavy traffic, direct Spoke-to-Spoke connections can be configured to reduce Hub load. Since branch sites generally have lower business load, they can deploy cost-effective ER6 Series gateways.

All ER gateways are centrally managed by the HQ Omada Controller or a Cloud-Based Controller, and SD-WAN establishes secure interconnection tunnels between the Hub and each Spoke—enabling centralized networking, unified operations, and cost-effective expansion for 100+ branch nodes.

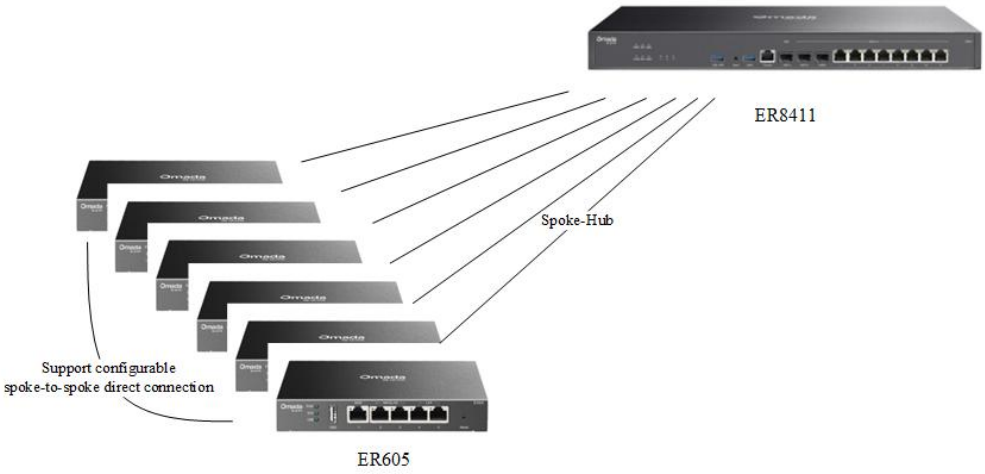

Figure 17: Hub-and-Spoke SD-WAN Topology for a Medium-Sized Logistics Enterprise

4.2.2.2 Pain Points and Requirements

For mid-sized multi-site enterprises—especially medium-sized logistics enterprises with 100+ branches—SD-WAN networking must not only support large-scale site access but also ensure long-term stable, secure, and manageable WAN interconnection among sites. In such scenarios, the enterprise typically has a headquarters or data center as a Hub node, and branches nationwide connect as Spokes, resulting in relatively high architecture complexity. Therefore, the following key pain points are commonly encountered during deployment and operations.

1. Overlay stability highly depends on underlay reliability

SD-WAN enables inter-site connectivity via overlay tunnels, but overlay stability fundamentally depends on the availability of the underlying physical underlay links. In enterprise networks with 100+ sites, if a site deploys only a single WAN uplink or the gateway does not support underlay failover, then the corresponding site's SD-WAN overlay tunnel may fail to establish or remain up when the WAN link experiences a physical disconnection, dial-up failure, carrier outage, or public address change.

In Hub-Spoke deployments, the Hub often carries centralized access and inter-site traffic forwarding for many branches, so its link stability has an especially significant impact on the entire network. If the Hub or key Spokes lack effective link redundancy and automatic failover, a physical-layer single point of failure can easily occur, leading to overlay tunnel interruption and affecting inter-site connectivity and access stability.

2. Complex path selection and link scheduling in multi-WAN scenarios

Medium-sized enterprise branches often deploy multiple WAN links—such as primary broadband, backup broadband, LTE/4G/5G backup links, or uplinks from different carriers—

to improve reliability. In SD-WAN deployments, WAN links differ in role, IP type, interface type, and availability. Selecting the most appropriate underlay link to carry overlay tunnels becomes a key consideration for administrators.

Without automated link scheduling, administrators must manually determine priority among the Primary WAN, Backup WAN, and Load Balanced WAN, while also considering differences in reachability between public and private IP addressing, stability differences across interface types (e.g., SFP/RJ45/LTE/USB/etc.), and failover strategies after link failures. For 100+ sites, per-site manual planning and maintenance are both labor-intensive and prone to inconsistent policies, which can lead to tunnel establishment failures, suboptimal link selection, or delayed failover.

Some enterprise workloads may even have special security, compliance, or billing requirements, requiring SD-WAN overlay tunnels to be pinned to specific physical WAN interfaces to prevent automatic switching. Such scenarios require the system to support both automated scheduling and deterministic path management.

3. High risk of LAN subnet conflicts across sites, impacting inter-site routing

During business expansion, rapid branch rollout, organizational changes, M&A, or cross-department network integration, LAN-side addressing plans across sites often lack a unified design, leading to overlapping private IPv4 address spaces. For example, multiple branches may reuse the default subnet 192.168.1.0/24, or different regions may have built networks independently early on using the same office subnet.

In traditional network solutions, once LAN subnets overlap across sites, inter-site route advertisement and service access are directly affected, potentially causing unreachable routes, abnormal traffic forwarding, or ambiguous destination access. To fully resolve conflicts, it is often necessary to re-plan IP addressing across the entire network and adjust many endpoints, servers, gateways, DHCP scopes, and static routes. This process is complex, high-risk, and likely to cause prolonged business interruption.

For medium-sized enterprises with 100+ sites, LAN conflicts not only increase upfront planning complexity but also significantly raise O&M risk during later expansion, M&A onboarding, and network consolidation. Therefore, an SD-WAN solution must provide automatic subnet detection, conflict risk alerts, and conflict route blocking, helping administrators identify issues during configuration and preventing conflicting settings from entering production networks.

4.2.2.3 ER Series Standalone Gateway SD-WAN Solution

To address pain points faced by medium-sized multi-site enterprises (100+ branches)—including link stability, WAN redundancy, overlay dependency, and LAN subnet

conflicts—the ER Series standalone gateway SD-WAN solution delivers highly reliable and controllable multi-branch networking through hardware performance, centralized management, and intelligent link scheduling.

1. **Multi-WAN redundancy and intelligent scheduling**
    - **Automatic link selection (Auto mode):** Supports multiple WAN interfaces and automatically selects the optimal underlay WAN link to establish overlay tunnels using a dynamic priority algorithm that considers Primary/Backup roles, public/private IP, physical interface type, etc.
    - **Failover:** When the primary WAN fails, the system automatically recalculates and switches to the backup WAN to keep overlay tunnels available.
    - **Deterministic path (Manual mode):** Supports pinning key business links to specific WAN interfaces to ensure traffic strictly uses designated WAN uplinks for security, billing, or compliance needs.

    These mechanisms mitigate the risks of single-WAN uplinks, strong underlay dependency, and easy overlay interruption, ensuring reliable connectivity between the Hub and Spokes.
2. **Centralized control and remote configuration**
    - **Unified management:** All ER Series gateways can be managed centrally via the HQ Omada Controller or a Cloud-Based Controller.
    - **Remote SD-WAN configuration:** The controller centrally provisions routing, tunnel parameters, and network policies, reducing per-device manual configuration.
    - **Rapid rollout:** New branches can quickly join the network and automatically obtain global topology and configuration profiles, enabling zero-touch deployment.

    This addresses issues such as complex large-scale branch deployment, heavy manual configuration workload, and error-prone operations, while ensuring consistent configuration baselines across the network.

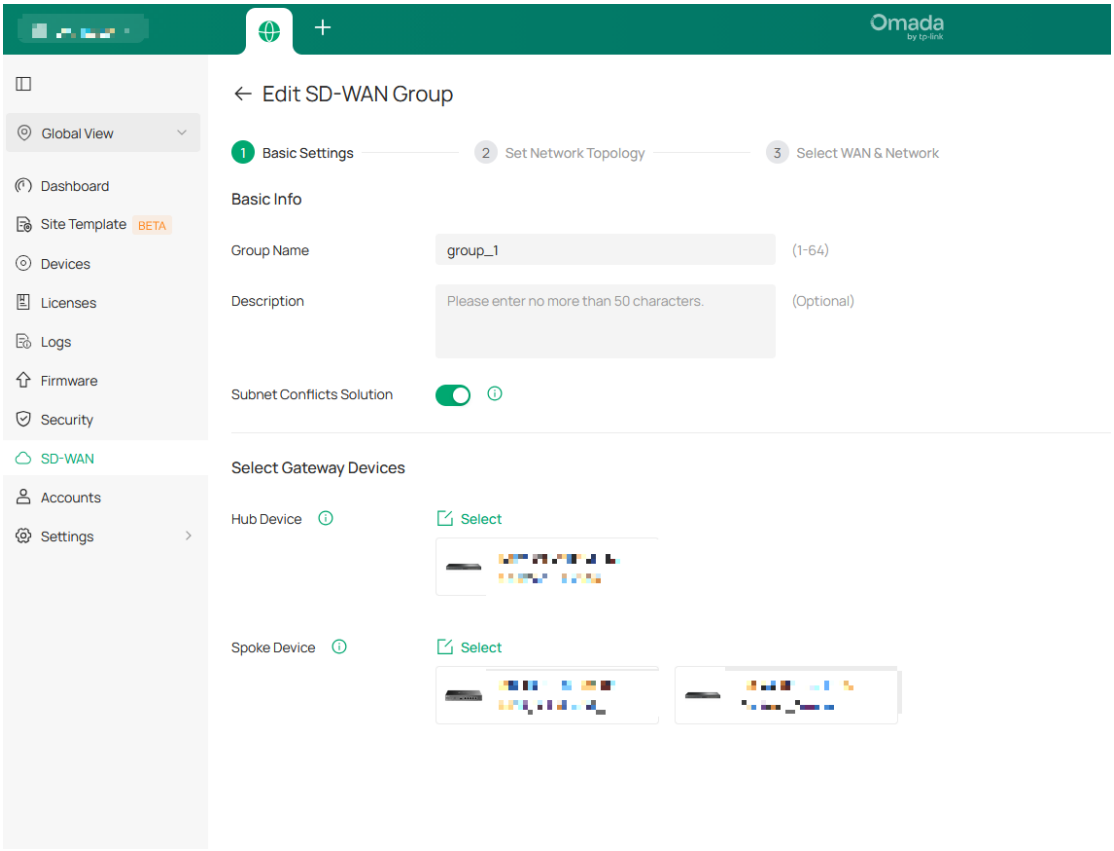

Figure 18: Unified Configuration Page for ER Hub-and-Spoke SD-WAN in an Organization

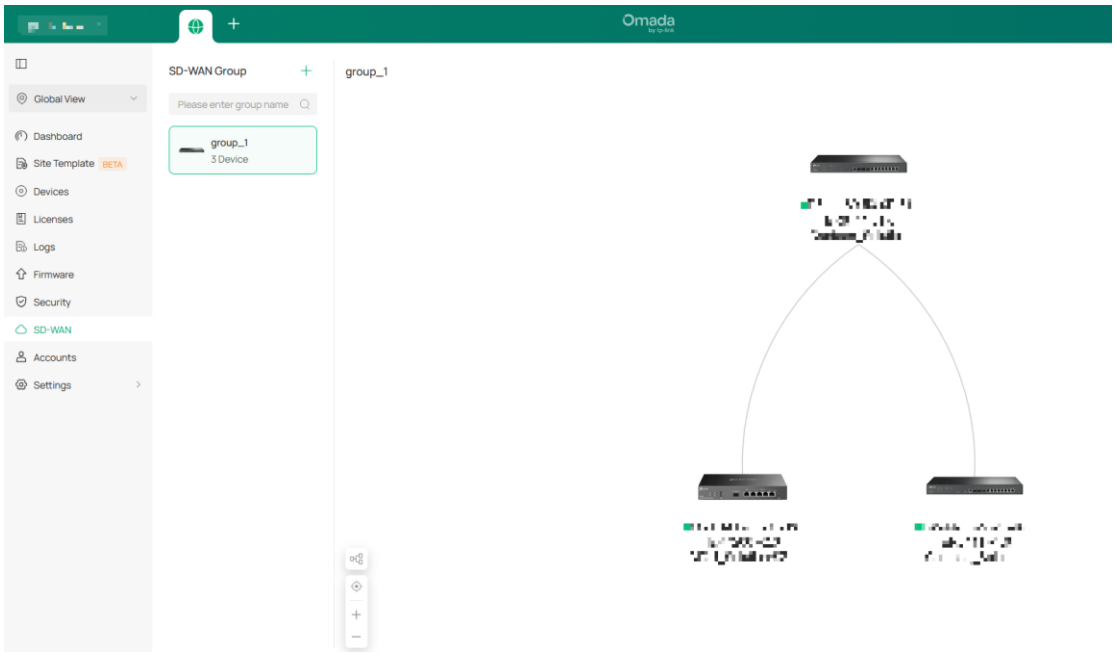

Figure 19: SD-WAN Hub-and-Spoke Topology Diagram

3. Overlay tunnel self-healing and stability assurance

- Through real-time interaction with the controller, ER Series gateways can

automatically detect WAN status and rebuild overlay tunnels, ensuring inter-site communication is not interrupted by single-link failures.

- In Hub-Spoke deployments, a high-throughput Hub node using an ER8 Series high-performance gateway can carry centralized traffic, while branches use lower-load ER6 Series gateways—achieving a balance between cost and performance.

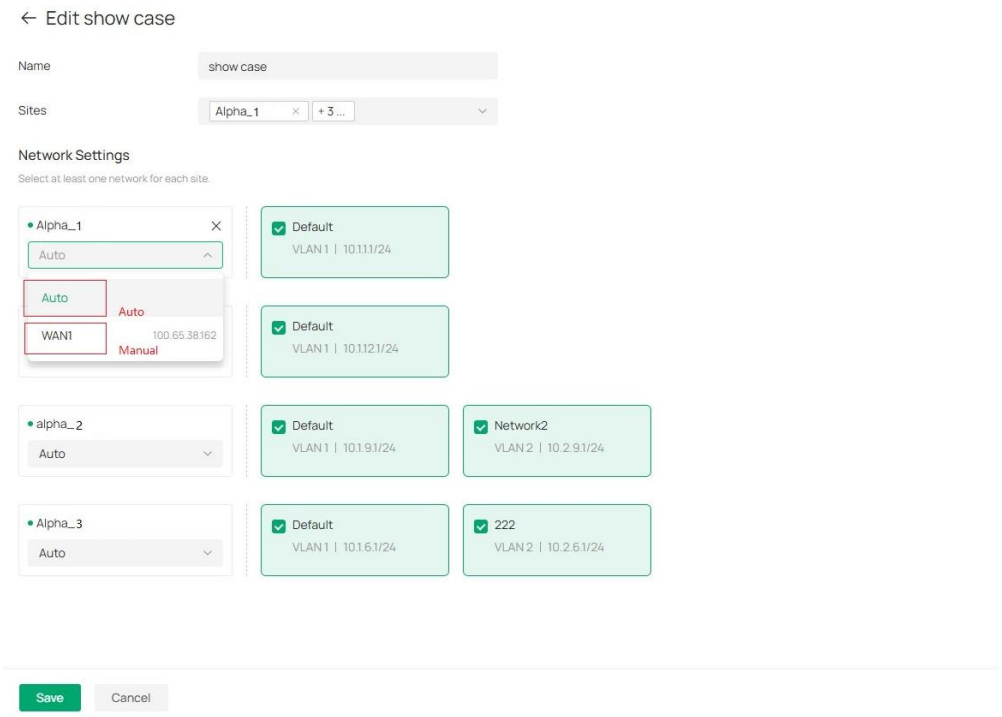

Figure 20: Underlay Link Redundancy Configuration

4. **Subnet conflict detection and security policy**
    - During site onboarding or route advertisements, the controller automatically collects LAN subnet information and compares it against the global topology database. If overlapping or conflicting subnets are detected, the system provides visual alerts and blocks conflicting route configurations.
    - Prevents inter-site communication failures caused by LAN subnet conflicts and ensures business continuity.

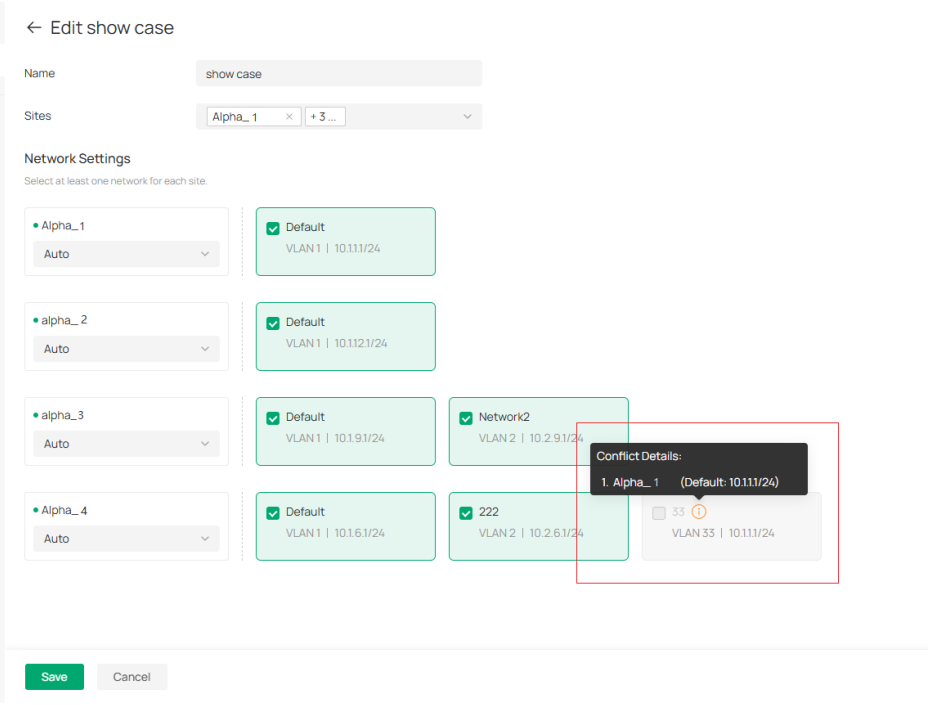

Figure 21: Subnet Conflict Detection Alerts

The ER Series also supports all features mentioned in the Fusion Series SD-WAN solution, such as NAT traversal.

4.3 Operation and Maintenance Troubleshooting

This section provides a brief SD-WAN troubleshooting procedure.

When SD-WAN networking is abnormal—such as tunnels not being established or inter-site services being unreachable—troubleshoot in the following order:

Device status → Configuration delivery → Underlay link → NAT/P2P → Tunnel status → Route learning → Service connectivity

1. Confirm device online status

Confirm whether the gateways participating in the SD-WAN network are online, and check whether their connections to Cloud Access / Service Cloud are functioning properly. If a gateway is not properly connected to cloud services, the controller may be unable to push SD-WAN configuration or synchronize SD-WAN status.

Confirm that the relevant devices have joined the same SD-WAN Group. Otherwise, SD-WAN connections between sites may fail to establish because the devices are not in the same SD-WAN group.

2. Verify SD-WAN configuration has been correctly delivered

Confirm that the selected sites, gateways, WAN ports, and LAN subnets are correct, ensuring the controller configuration matches the actual network topology.

Check for LAN subnet conflicts. If different sites use identical or overlapping LAN

subnets, route advertisements may be abnormal or inter-site access may fail. Additionally, confirm that there are no misconfigured ACLs, firewall rules, or access control policies to avoid cases where SD-WAN tunnels are established, but business traffic is blocked by policy.

3. Check underlay link status

Check whether the gateway WAN interface is up, and confirm that it can access the internet or cloud services normally.

In multi-WAN scenarios, confirm that the WAN link currently selected by SD-WAN matches expectations. If Auto mode is used, verify that the system-selected underlay WAN link is available. If Manual mode is used, focus on whether the manually bound WAN interface is online, whether dial-up has succeeded, and whether it has internet connectivity.

If the underlay link is unstable or unreachable, overlay tunnels typically cannot be established or maintained.

4. Check NAT / P2P hole-punching status

Confirm whether the device has successfully obtained the public-mapped address, port information, and NAT type, and verify that both ends have successfully exchanged ICE candidates / STUN probing information.

If both devices are behind symmetric NAT, P2P direct connectivity is typically difficult to achieve. In such cases, a relay path may be required, or the network environment may need adjustment. Pay close attention to NAT type, port mapping changes, and the impact of carrier network restrictions on tunnel establishment.

5. Check tunnel establishment status

Confirm whether IPsec negotiation has succeeded, including the IKE phase, key exchange, encryption suite matching, and peer identity authentication.

Then check whether the GRE interface has been created and is up. If IPsec is established but GRE is not up, further verify the tunnel interface, peer address, encapsulation parameters, and underlying link reachability.

If tunnels repeatedly flap, first investigate WAN link instability, NAT mapping changes, IPsec negotiation errors, or abnormal control-channel status.

6. Check route learning and the forwarding table

Confirm whether the OSPF neighbor relationship has been established and whether the local device is correctly advertising its local LAN subnets.

Confirm whether the local device has learned the peer LAN routes and whether the next hop of those routes points to the SD-WAN tunnel interface. If the tunnel status is normal but services are unreachable, focus on whether routes are missing, route

preferences are abnormal, or whether incorrect default routes or policy-based routing are affecting forwarding paths.

7. Validate service connectivity

For service validation, follow a progression from near to far and from lower layers to the application layer:

- First, ping the peer gateway or the GRE virtual IP to confirm basic connectivity across the tunnel.
- Then, ping the peer LAN gateway to confirm inter-site Layer 3 reachability.
- Finally, test access to specific application servers, service ports, or applications.

If ping works but service access fails, focus on ACLs, firewall/security policies, port-allow rules, and whether the application server itself is listening/healthy.

8. Review logs and alerts

If the issue still cannot be located using the steps above, further review controller alerts, gateway system logs, VPN negotiation logs, WAN interface state-change records, and routing neighbor-change records.

Use logs to determine at which stage the issue occurs—for example, configuration not delivered, WAN unreachable, NAT hole-punching failure, IPsec negotiation failure, GRE not up, OSPF neighbor anomalies, or business traffic blocked by policy—so the troubleshooting scope can be narrowed further.

5 Representative Models

Fusion Series:

- Fusion 2.5G Gateway

Built to simplify deployment, reduce costs, and streamline network management for installers, MSPs, and SMBs.

With license-free cloud management, five 2.5G ports, Bluetooth setup, a built-in controller, touchscreen diagnostics, and secure remote access through Omada LightLink VPN, Fusion 2.5G makes business networking faster and easier.

ER Series:

- ER8411

A high-performance Omada 10G VPN gateway that provides two 10G SFP+ ports, up to ten WAN ports, hardware-accelerated VPN, and centralized management through Omada SDN. Suitable for enterprise multi-WAN aggregation and high-bandwidth internet egress.

- ER7206

An Omada Gigabit VPN gateway that supports multi-WAN via SFP/RJ45, USB LTE backup, centralized cloud management, and multiple VPN types. Suitable for SMB multi-line access and secure interconnection.

- ER605

An entry-level Omada Gigabit VPN gateway that supports up to three WAN ports, Omada SDN management, SD-WAN Controller mode, and common VPN options. Suitable for small offices or branch networks.

6 Future Outlook

As digital transformation accelerates, enterprise WAN requirements are shifting toward more complex heterogeneous links, larger branch footprints, and stricter security isolation. Building on the existing capabilities—efficient unified cloud deployment, Full-Mesh secure connectivity, priority-based automatic link scheduling, and cross-site LAN subnet conflict detection—we will continue to evolve our SD-WAN solution.

Going forward, we will further expand SD-WAN across high-availability and disaster-recovery architectures, automated orchestration, and built-in network security. We plan to deliver the following key capabilities.

6.1 ER/Fusion Hybrid Networking

Today, Omada SD-WAN supports two deployment models based on device type: an ER-series standalone gateway model and a Fusion-series cloud-managed model. Next, we will extend hybrid networking between the ER and Fusion series, so gateways of different types can work together in the same SD-WAN system.

In a typical deployment, ER-series gateways are used as headquarters, data center, or core Hub nodes to handle higher throughput, larger tunnel scale, and more complex multi-WAN access. Fusion-series gateways are used for stores, lightweight branches, or temporary sites to take advantage of fast turn-up, cloud management, and all-in-one deployment. With ER/Fusion hybrid networking, you can select the appropriate device for each site based on size, cost, and performance requirements without using a single device

type across the project.

6.2 Enhanced P2P/Relay Reachability Without a Public IP

In real-world deployments, branch sites are often behind multi-layer NAT, CGNAT, mobile networks, or dynamic public IP addresses. Traditional VPNs rely heavily on public IPs, port mappings, and mutual reachability. Today, SD-WAN improves direct-connect success rates in no-public-IP scenarios through NAT traversal and P2P hole punching. Next, we will add relay capabilities to create a complete reachability model: P2P first with relay as a fallback.

Future connection logic will evolve as follows:

1. Try P2P first.

Use STUN/ICE-like mechanisms to learn each side's public-reachable address, port mapping, and NAT type, and perform multiple hole-punching attempts coordinated by the cloud. If the network allows it, establish an end-to-end direct path between sites to reduce latency and cloud bandwidth usage.

2. Failover to relay when direct connectivity fails.

If both sides are behind symmetric NAT, strict firewalls, ISP restrictions, or frequently changing port mappings, a stable P2P path might not be possible. In this case, automatically select a relay node that forwards inter-site traffic to ensure the tunnel comes up and applications remain reachable.

3. Support dynamic path switching.

After the relay path is established, continue periodic P2P re-probing. If network conditions change and a P2P path becomes available again, switch back from relay to a direct path to balance availability and performance.

The goal is not to guarantee direct connectivity in every scenario, but to maximize connection success rates under complex internet conditions, keeping essential services available through relay when P2P is unavailable.

6.3 High Availability and Load Balancing

To further improve link reliability in typical Hub-and-Spoke WAN topologies, the system will build on existing underlay failover and introduce enhanced multi-WAN support and

intelligent primary/backup scheduling.

- **Active-Active and dynamic routing:** In future networking deployments, both Spoke and Hub nodes will support two uplinks—a Primary WAN and a Secondary VPN WAN—to participate in the SD-WAN network. With Active-Active (hot-standby) enabled by default, the two WAN interfaces will establish four (2×2) tunnels simultaneously. The system will then use dynamic routing to calculate the optimal tunnel and perform intelligent path selection.
- **Multi-dimensional failure detection and failover:** The new disaster-recovery architecture will correlate internal tunnel status with external physical link status. The device will monitor WAN interface state (Up/Down) in real time and combine it with link SLA metrics to detect reachability and quality. The system can quickly withdraw routes for failed tunnels and switch traffic to the standby link, and perform a smooth failback when the primary link recovers to maximize business continuity for critical services.

6.4 Automated Provisioning in Batches

In projects with dozens or even hundreds of branch sites, manually configuring gateways, WAN, LAN, VPN, and routing policies device by device significantly increases delivery cost and often leads to configuration inconsistencies. Going forward, Omada SD-WAN will further enhance ZTP and Template Configuration, so new sites can be created from manual setup to standardized batch configuration.

We plan to add the following capabilities:

1. Zero-touch provisioning (ZTP).

After a branch device is powered on and connected to the network, it will automatically connect to the cloud and complete device claiming, site identification, and baseline configuration downloading. On-site staff will only need to cable and power on the device—no need to understand complex VPN, routing, or WAN policies.

2. Site templates.

The controller will provide standardized site templates, such as templates for retail stores, logistics sites, office branches, and temporary sites. Templates can predefine WAN policies, LAN subnets, VLANs, DHCP, SD-WAN Group, route advertisement scope, and access control policies.

3. Batch parameter configuration.

When sites differ only in IP subnets, site names, WAN types, or egress policies, use variables to generate configurations in batches and avoid copy-and-paste errors.

4. Configuration baselines and inconsistency detection.

The system will check all sites against a configuration baseline, identify devices that deviate from the template, and support one-click remediation, batch rollback, or difference confirmation to ensure consistency at scale.

This direction can significantly reduce the deployment barrier for large SD-WAN projects and make the solution a better fit for MSPs, SIs, chain stores, and multi-branch enterprises that require batch rollout.

6.5 Link Health and Intelligent Operations

As the number of sites grows, SD-WAN operation challenges shift from how to configure to how to continuously confirm that the network is healthy. Going forward, the system will strengthen link health monitoring, anomaly detection, and fault identification to allow administrators to rapidly detect issues, assess impact, and receive actionable recommendations.

We plan to enhance the following capabilities:

1. Link SLA monitoring.

The system will continuously monitor latency, packet loss, jitter, bandwidth utilization, tunnel re-establishment counts, and uptime stability across WAN links and SD-WAN tunnels. These metrics can be used to generate site-, link-, and tunnel-level health scores.

2. Abnormal trend detection.

The system will detect risks early when a link shows periodic packet loss, rising latency, frequent reconnects, or NAT mapping changes, instead of waiting until services are fully down to raise an alarm.

3. Layered troubleshooting guidance.

The controller will generate guided troubleshooting steps in this order to help users quickly pinpoint the failure stage: device online status → cloud connectivity → WAN reachability → NAT/P2P status → IPsec/GRE tunnels → OSPF routing → ACL/firewall policies → service connectivity.

4. Operation visualization.

The topology view will clearly display site status, tunnel status, link quality, current paths, failed links, and alarm severity, so users do not have to jump across multiple pages to troubleshoot.

With link health and intelligent operations, SD-WAN can evolve from a configuration tool into an operation and maintenance platform that is observable, diagnosable, and decision-assistive.

6.6 Intelligent Routing

Currently, SD-WAN supports path selection based on link priority and availability. Next, we will evolve toward intelligent routing, so the system can dynamically choose the best forwarding path based on application type, link quality, and policy requirements.

Future intelligent routing will focus on:

1. Application-based routing.

The system will identify different traffic types, such as POS payments, video surveillance, file transfers, remote desktops, voice/video meetings, and general office traffic, and select paths based on each application's latency, loss, and bandwidth requirements.

2. Link-quality-based dynamic path selection.

When the primary link is still up but shows high latency, high packet loss, or abnormal jitter, the system will move critical traffic to a higher-quality backup link instead of relying only on interface Up/Down status.

3. Policy-based deterministic routing.

For traffic with security or compliance requirements, such as payments, finance, and core business systems, administrators will be able to pin traffic to a specific egress, tunnel, or Hub to avoid the path uncertainty introduced by automatic steering.

4. Multi-path coordination.

In Full-Mesh, Hub-and-Spoke, and future ER/Fusion hybrid networking deployments, the system will evaluate direct paths, Hub transit paths, relay paths, and multi-WAN underlay paths, and automatically balance performance, reliability, and cost.

The goal is to ensure SD-WAN not only connects, but also routes correctly, fails over quickly, and stays stable based on service needs. This will further improve the Omada SD-WAN application in multi-branch offices, chain retail, logistics scheduling, and remote operations.

7 Appendix

7.1 Glossary

Abbreviation	Full Term	Category	Description
CBC	Cloud-Based Controller	Architecture / Core	The centralized management core of SD-WAN. Delivers global policies, adopts

			and manages devices, schedules links, monitors status, and provides orchestration control.
CPE	Customer Premises Equipment	Architecture / Edge	Customer-side equipment. Refers to an SD-WAN gateway deployed at an enterprise branch or HQ edge. It provides local traffic access, overlay tunnel encapsulation, path selection, and local policy enforcement.
Overlay	Overlay Network	Network Layer	A virtual logical network built on top of the physical network. Uses tunneling and encryption to abstract underlay differences and carry application traffic, security policies, and QoS.
Underlay	Underlay Network	Network Layer	The physical transport network infrastructure (e.g., MPLS, internet, and 5G). Provides basic IP packet forwarding and the transport for the overlay.
Full-Mesh	Full-Mesh Topology	Topology	A full-mesh model that enables direct data connectivity between branch nodes, significantly reducing lateral latency and eliminating the HQ gateway bottleneck.
Hub-and-Spoke	Hub-and-Spoke Topology	Topology	A star topology. All traffic between the branch nodes (Spokes) is forwarded through the HQ (Hub) as a transit node.
Failover	Failover	Reliability	Redundant switchover. When the primary link fails or degrades, the system quickly moves traffic to a heterogeneous backup link (e.g., 5G) or a standby tunnel to keep services running.
Active-	Active-Active	Link Mode	A hot-standby mode. Multiple WAN

Active	Mode		interfaces establish tunnels simultaneously to provide load balancing and extremely high link redundancy.
UDP Hole Punching	UDP Hole Punching	NAT Traversal	A technique that allows nodes behind NAT to establish an end-to-end tunnel without a public IP by coordinating outbound UDP traffic and creating NAT mappings.
ACL	Access Control List	Security Policy	A 5-way (i.e., source IP, destination IP, source port, destination port, and protocol) filtering policy that defines security boundaries and isolates sensitive services.
IPsec	Internet Protocol Security	Security / Tunneling Protocol	A suite of protocols that uses encryption and authentication to protect the confidentiality and integrity of data transmitted through overlay tunnels.
GRE	Generic Routing Encapsulation	Tunneling Protocol	A technology that encapsulates multiple protocols in an IP tunnel and is commonly used to build overlay links in SD-WAN.
STUN	Session Traversal Utilities for NAT	NAT Traversal	A protocol used to detect NAT type and discover public-mapped addresses. It is a key building block for UDP hole punching.