VERSA NETWORKS

WHITE PAPER

# Reducing Header Overhead with Tunnel-less SD-WAN

# Table of Contents

## Introduction

Traditional IP packet forwarding adds a 20-byte header to the ethernet frame to pass traffic from one segment to the next, and only considers the destination IP address in the forwarding decision. In contrast, in making forwarding decisions a modern SD-WAN has the ability to distinguish between and take into account different applications, different transport types, and different path performance. These capabilities are what allow for the creation of SD-WAN policies to choose the most appropriate path for a given application, secure the application flow if required, and apply network performance enhancements (i.e., FEC or packet replication) as needed. Two different application flows, from a single source IP address to a single destination IP address, can have differentiated forwarding based upon the SD-WAN policy.

While the transmission overhead that SD-WAN tunnel encapsulation and encryption adds to IP packets to accomplish the above is small, it may be noticeable on networks that are bandwidth constrained or on ones that utilize external encryption methods. In the scenario of severely bandwidth-constrained networks, every byte transmitted is a precious resource. This document explores how so-called "tunnel-less SD-WAN" can alleviate SD-WAN overhead-related issues in such scenarios.

## Tunneled SD-WAN overhead explained

A traditional SD-WAN receives IP packets (see Figure 1) at the user network interface (normally, referred to as the LAN interface or the ingress interface), identifies the application flow associated with those packets, evaluates application flow for the appropriate policies (e.g., security, forwarding, encryption, etc.), and then forwards the packets. Since an SD-WAN forwards packets based upon application flow, the network must be able to distinguish between two IP packets with the same source and destination IP addresses.


*Figure 1 - IP packet header view*

To do this, an SD-WAN encapsulates the packet with a header that encodes this differentiation into the IP packet. And, as is often the case, encrypts the IP packets for transport across unsecure transports (see Figure 2).


*Figure 2 - SD-WAN IP packet header view*

This is known as "tunneled SD-WAN." The tunneled SD-WAN encapsulation and encryption increases header overhead from the 20-bytes added by the IP layer. This additional overhead can be as little as 50 bytes, but can be as many as 200 bytes depending on vendor implementation and depending on the scenario. The impact to the application flow is minimal if the size of the

maximum transmission unit (MTU) for the application flow is large. However, if the size of the payload is small, then this is a significant increase in the packet overhead.

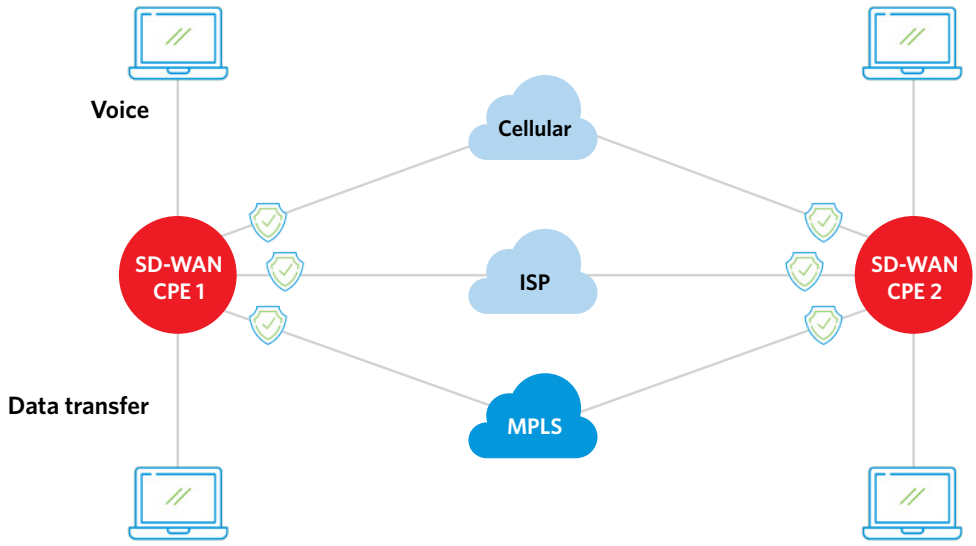


*Figure 3 - Voice and data transfer diagram*

Consider two different application flows as depicted in Figure 3. One is a data transfer flow with an MTU size of 1450 payload. The other is a voice flow with a MTU size of 60 bytes payload. Addition of the IP header to the payloads makes the data transfer flow a 1470-byte packet size and the voice flow an 80-byte packet size. This is a two-percent increase for the data transfer flow, but a 33 percent increase for the voice flow.

While these numbers are significant on the voice flow, these are default values for IP traffic and the application flows are required to have these header values. But if SD-WAN added an additional 100 bytes of encapsulation and encryption to each of these flows, then the data flow would be a 1590-byte packet size with an increase of 8 percent header overhead (6 percent increase due to SD-WAN), and this packet would likely need to be fragmented as the maximum transmission unit size is most likely 1500 over the internet. One packet would go as a 1500-byte packet (20 bytes of IP header, 100 bytes of SD-WAN encryption and encapsulation, 20 bytes of original IP header, and 1360 bytes of User Data) and the second packet would be 230 bytes (20 bytes of IP header, 100 bytes of SD-WAN header, 20 bytes of original IP header and the 90 bytes of payload that would not fit in the first IP packet). While the voice flow would not require fragmentation of the original payload, the relative effect on the header overhead is greater. The original 60 bytes of payload results in an SD-WAN 200-byte IP packet (20 bytes of IP header, 100 bytes of SD-WAN encryption and encapsulation, 20 bytes of the original IP header and 60 bytes of user data.) This results in a 233 percent increase in the size of the packet as compared to the payload or 150 percent increase over a traditional IP packet. This means that 2.5 IP packets of voice flow would be able to be transmitted via a traditional IP network vs 1 IP packet of voice flow for an SD-WAN network.

## How much SD-WAN overhead optimization is possible

Does SD-WAN need all this overhead to distinguish between IP packets and application flows? The answer is that in some SD-WAN solutions, the SD-WAN overhead can be optimized in a

way that reduces the SD-WAN overhead. This optimized SD-WAN is often referred to as (so-called) "tunnel-less" SD-WAN, a term that is a bit of a misnomer. The traffic does still require that some additional header information be added to the IP packets to distinguish among application flows, but it is the case that the SD-WAN overhead can be reduced by 50% or more. This depends on many different factors, which range from the version of IP being transmitted (IPv4 or IPv6), the version of IP being utilized by the SD-WAN overlay (IPv4 or IPv6), and the proprietary implementation for this optimization.

Assuming a 50 percent reduction in the SD-WAN header overhead (from 100 bytes to 50 bytes), we can then extrapolate that the flow would need to transmit a 150-byte packet instead of the 200 bytes previously required. This is a 25% savings in the total packet size and a reduction of 35% in the total header size. Depending on implementation, the header savings could approach 50%.

## Complex design with an external encryption mechanism

SD-WANs are used widely today across networks to address the needs of a variety of deployments. Some of these scenarios can be simple and some can be more complex in topologies or in technical requirements. Consider Figure 4, which illustrates the use of SD-WAN to categorize applications (top secret vs non-mission activity) utilizing external encryption devices (HAIPE) to optimally traffic steer application flows across disparate transport networks (Satellite GEO and MEO and MPLS).

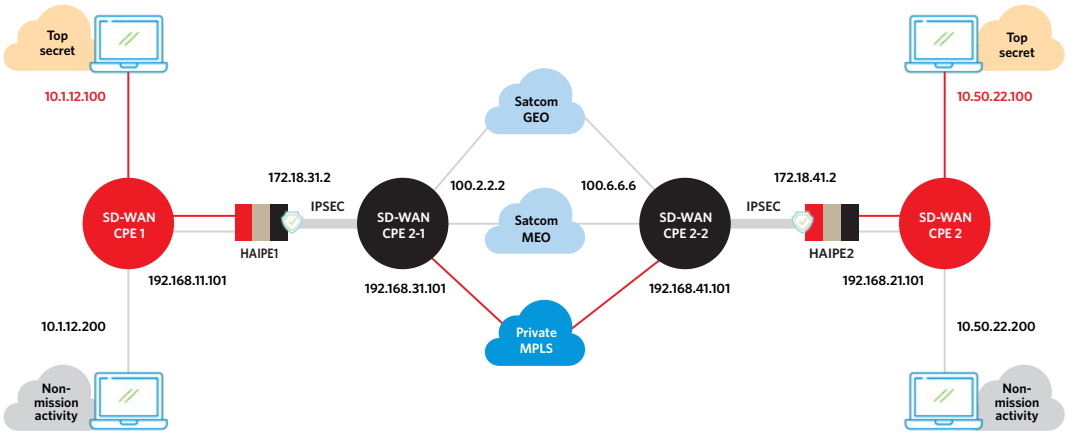


*Figure 4- Multiple SD-WAN and external encryption*

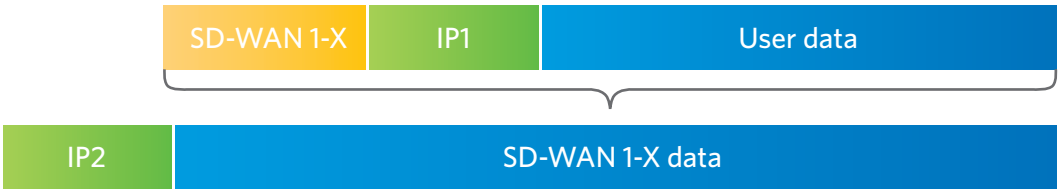


*Figure 5 - Complex SD-WAN IP packet header view*

As shown in Figure 5, this complex design requires that the first set of SD-WAN devices would encapsulate the IP packets with the SD-WAN headers to distinguish the top-secret applications from the non-mission applications.

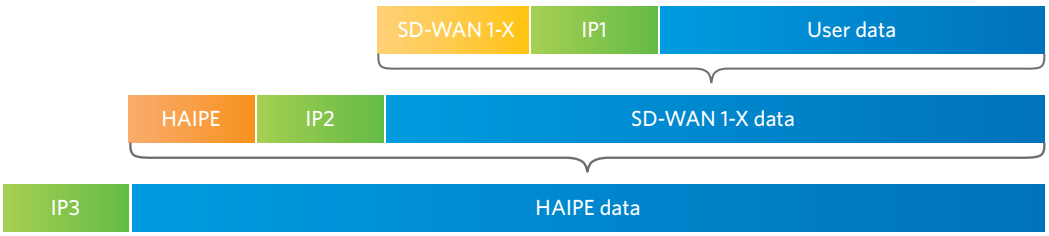


*Figure 6 - Complex HAIPE IP packet header view*

As shown in Figure 6, the external encryption devices would encrypt those flows keeping only the DSCP markings from the SD-WAN packets. The inner SD-WAN, which determines which transport network to utilize, would then add the SD-WAN header to distinguish the different SD-WAN flows and steer them to the appropriate transport networks.

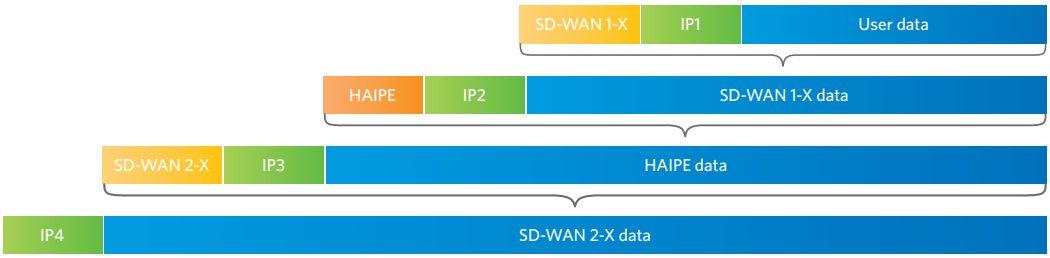


*Figure 7 - Complex SD-WAN level 2 IP packet header view*

But this now implies the application of SD-WAN headers and one encryption header to the IP packet (see Figure 7). If we assume that each header is adding 100 bytes of header, then that is a total of 300 bytes of header added to each packet. This reduces the maximum payload from 1480 bytes per packet to a maximum of 1120 bytes. While this increases the number of packets needed to be sent for the data transfer flow, the real impact is on the voice flow, where the 80-byte IP packet is now a 440-byte packet (20 bytes of outer IP header, 100 bytes of Level 2 SD-WAN encryption and encapsulation, 20 bytes of inner #3 IP header, 100 bytes of HAIPE encryption and encapsulation, 20 bytes of Inner #2 IP header, 100 bytes of SD-WAN Level 1 encryption and encapsulation, 20 bytes of original IP header, and 60 bytes of user data). This results in a 633 percent increase as compared to the 60-byte payload and a 450 percent increase when compared to the 80-byte IP packet for the voice flow. If the SD-WAN layers could reduce their overhead by 50% (from 100 bytes to 50 bytes each), then the resulting IP packet would be 340 bytes which is a 23 percent savings for the entire packet and a 32 percent reduction in the header overhead of the packet. If encryption were turned off in the SD-WAN header, since the design includes an external encryption mechanism, then the SD-WAN headers could be reduced even further and could result in 50% or more header overhead reduction.

## Why SD-WAN routing is worth it

Some would ask why even consider SD-WAN in these use cases, if the SD-WAN header adds such a significant additional amount of data needing to be transmitted. Consider Figure 4, where the top-secret data was delay sensitive, but the non-mission activity was not delay sensitive.

In this case, our data transfer flow might be considered the non-mission activity, while the voice flow would be the top-secret data. By using SD-WAN, the initial SD-WAN device can

distinguish between the top-secret application and the non-mission activity. By inserting the label, after the external encryption device encrypts either flow, the internal SD-WAN can utilize the DSCP to differentiate the two flows and send the non-mission activity over the higher latency satellite links and the top-secret flow over the better delay path of the MPLS. If these devices were traditional routers, the packets would have followed the same path for both top-secret and non-mission activity, because each had a source IP within the same network and were sending packets to a destination IP within the same network.

## Transitioning from tunneled to tunnel-less SD-WAN

So how does (so-called) tunnel-less SD-WAN differ from tunneled SD-WAN? As already discussed, tunnel-less SD-WAN reduces the SD-WAN overhead that needs to be added to the IP packet to determine the application flow. This is done by SD-WAN starting out as a tunneled SD-WAN solution. The first packet of the application flow traverses the SD-WAN network via tunnels. But each SD-WAN device assigns meta data to the packet. Once the first packet has been sent and the response has returned, and all the SD-WAN devices have assigned their local application flow meta data, the flow can transition from tunneled mode to tunnel-less mode. There are optimizations that happen through the use of this meta data which result in a significant reduction in the SD-WAN encryption and encapsulation headers. While the greatest reduction in SD-WAN headers occurs when encryption is not included, this option is only relevant for designs that include an external encryption mechanism. Otherwise, lack of encryption on unsecure transports would pose a risk to data security.

## Conclusion

Tunnel-less SD-WAN has advantages over the tunneled SD-WAN implementations where bandwidth is at a premium. Even just a savings of 50 bytes per packet could mean many more packets can be transmitted over the same connection. In scenarios where an external encryption mechanism is deployed, tunnel-less SD-WAN can provide a savings in the header overhead, while allowing for transport of the packets over the most appropriate connection.

## About Versa Networks

Versa Networks, the leader in single-vendor Unified SASE platforms, delivers AI/ML-powered SSE and SD-WAN solutions. The platform provides networking and security with true multitenancy, and sophisticated analytics via the cloud, on-premises, or as a blended combination of both to meet SASE requirements for small to extremely large enterprises and service providers.

Thousands of customers globally with hundreds of thousands of sites and millions of users trust Versa with their mission critical networks and security. Versa Networks is privately held and funded by Sequoia Capital, Mayfield, Artis Ventures, Verizon Ventures, Comcast Ventures, BlackRock Inc., Liberty Global Ventures, Princeville Capital, RPS Ventures and Triangle Peak Partners. For more information, visit https://www.versa-networks.com or follow Versa Networks on X (Twitter) @versanetworks.

Versa Networks, Inc, 2550 Great America Way, Suite 350, Santa Clara, CA 95054
+1 408.385.7660 | info@versa-networks.com | www.versa-networks.com