The Role of Dynamic IPsec Tunnels in Modern SD-WAN Networks

Dynamic IPSec tunnels are often cited as a must-have requirement in SD-WAN networks. But the actual need is much narrower, given the performance of today's networks. Is the cost worth the benefit?

Dynamic IP Security (IPsec) tunnels were invented by Cisco and first deployed in Cisco's Dynamic Multipoint VPN (DMVPN). The requirement was driven by the need to achieve quality of experience for voice communication between VoIP endpoints. Dynamic tunnel features were first deployed in 2004, when internet and service provider WAN networks were at lower levels of performance. As network speeds have increased and IP access and exchange points have proliferated closer to customer sites, latencies have been reduced significantly. The original need for dynamic tunnels is not as acute now as it was 15 years ago.

Software-Defined WAN (SD-WAN) has become the de facto standard for deployment of WAN aggregation in modern enterprise networks. It owes its popularity to its simplified approach to delivering important enterprise use cases as well as its awareness of the performance of multiple WAN network paths and performance-aware routing. Like any engineering consideration, dynamic tunnels come at the price of complexity and load added to both the network and the SD-WAN control system. The decision to deploy dynamic tunnels requires careful consideration of the need and an analysis of whether they will actually improve network performance given the additional cost of the load and complexity they introduce.

Contents

Rationale for dynamic tunnels

When Cisco introduced dynamic IPsec tunnels in DMVPN in the mid 2000s, it was in response to real network latency impairments in early customer VoIP implementations. The issue involved the latency budget. Service provider network latencies across major geographic distances were eating up a significant proportion of the maximum delay tolerable for good-quality voice conversations.

There are four criteria to consider when assessing performance along a network path and how it will impact voice performance:

- Latency: Delay of packet delivery from source to destination
- Jitter: Variations in the delay of packet delivery
- Packet loss: Too much traffic in the network causes the network to drop packets
- Burstiness of jitter and packet loss: Loss and discards tend to occur in bursts

Dynamic tunnels are primarily aimed at addressing latency and jitter. The premise is that packets will traverse a shorter physical distance as well as through fewer network elements, each of which can contribute to delay and jitter. It is not at all clear that a direct site-to-site tunnel will improve jitter, as that depends primarily on congestion levels and the number of network elements on the route between the sites.

Dynamic tunnels do not effectively address loss, as the tunnel may or may not bypass congested nodes, depending on the actual elements between the two sites. In fact, the bursty nature of both jitter and loss within the network can often result in thrashing between dynamic, direct tunnels and static ones.

The performance benefits of dynamic tunnels are highly dependent on the service provider routing topology within the region and between regions in your branch aggregation network. For example, while your network provider may offer full-mesh connectivity at the routing layer, does it provide the physical topology that would actually shorten the path? Network providers often use regional aggregation hubs to connect and provide service for multiple tenants. Based on the topology of these hub sites and connectivity between them, the actual path length may not differ significantly enough from your own regional aggregation hub topology. In this case it is often better to use static, always-on tunnel topologies provided in your SD-WAN to build your own virtual network segments over which you can exercise full SD-WAN policy control, visibility, and topological control. It will also be easier to debug, since you are using only static topologies.

Delay, loss, and jitter in modern networks

Since the first introduction of dynamic IPsec tunnels in DMVPN, network providers have continued to invest in increased performance, capacity, and reliability in their networks. They have been continuously improving by reducing delay and jitter – impairments that dynamic tunnels can protect against – as well as loss. You already get the benefits of these investments.

Let's take a look at the target and actual performance numbers for two major U.S. service providers across global theaters. To do this we need to slice the problem a couple of ways and look at performance within a theater and performance between theaters, as the latter captures overseas routes.

Table 1. 12-month average latencies across theaters (August 2018 to August 2019)

	AT&T	Verizon
Regional Averages for Latency (ms)	**Last 12 Months Average**	**Last 12 Months Average**
US	30.90	32.09
Europe	14.61	11.13
Asia Pacific	57.75	86.52
Cala	101.62	95.88

Table 1 shows average round-trip latencies within the specified region. Callers usually notice round-trip voice delays of 250 ms or more. ITU-T G.114 recommends a maximum one-way latency of 150 ms (300 ms round-trip). Since this includes the entire voice path, your own network should have transit latencies of considerably less than 150 ms.

As you can see from the table, regional round-trip latencies are well below the limits for good-quality VoIP, and in fact there is significant headroom left in the 250- to 300-ms budget after network latency is taken into account. What we can take from this is that a simple hub-and-spoke topology will more than suffice for voice and other delay-sensitive traffic.

But it also tells us that when we have tighter requirements for real-time applications, and those applications are peer-to-peer in nature, we need to dive further into the topology of our service provider network to ensure that a direct site-to-site tunnel will actually be shorter topologically and geographically such that it will save significant delay budget, with margin. Deploying site-to-site tunnels blindly, without this analysis, may result in periodic instability in your network, which can be much more problematic in terms of user perception and satisfaction.

But what about interregion transit? Transit between theaters incurs transoceanic distances as well as more network elements and possible interprovider exchanges. Even so, examination of Table 2 shows that network latencies are still below the tolerable round-trip limits for VoIP performance, leaving reasonable delay margin in the budget.

Table 2. 12-month average latencies across regions (August 2018 to August 2019)

AT&T	
Inter Region Latency (ms)	**Last 12 Months Average**
San Francisco to Tokyo	93.09
New York to London	70.99
Singapore to Los Angeles	167.31
New York to Rio De Janeiro	108.56

Verizon	
Inter Region Latency (ms)	**Last 12 Months Average**
Trans Pacific	105.21
Trans Atlantic	70.69
EMEA to Asia Pacific	123.79
Argentina to US	117.74

It's interesting to note that there are only so many transcontinental fiber paths that networks can follow between overseas destinations. All SD-WAN tunnels, dynamic or static, will have to traverse the same transoceanic links and interprovider hops to get between major geographies. Establishment of dynamic IPsec tunnels just won't help these paths that much.

What this tells us is that for real-time traffic over global networks, a high-speed core network is required as well as the ability to engineer VPN paths across that core for minimum delay. Static tunnels are the right approach because of the ability to control the path, optimize traffic distribution (in the case of multiple paths), and determine the path when things go wrong, which will simplify operations for failover, repair, and debugging.

So what does this imply? You should carefully consider the hard requirements you have for deploying dynamic tunnels in your network. Are you following engineering rules prescribed in an earlier WAN deployment that may no longer be relevant? As engineers, our responsibility is to ensure that we evaluate deployment of any network feature against need, cost, and whether it achieves a desired result. These should all be measurable, because any feature remedy applied must provably make our network run better or cheaper. There are costs for turning on features that are not needed or provide minimal benefits, in the form of operational cost, processing load, and time spent troubleshooting.

Figure 1. Traffic patterns in modern networks

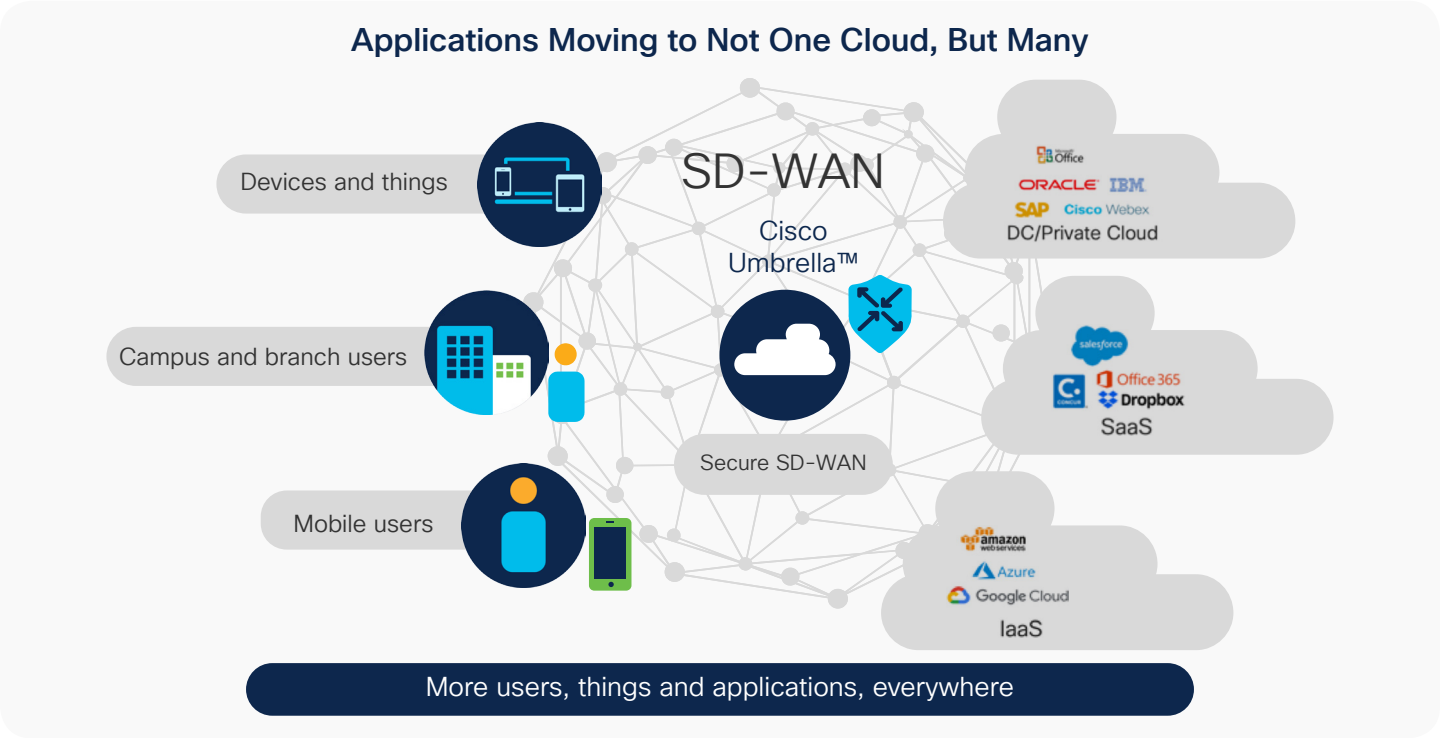

The past few years have seen major changes in the way applications are deployed and accessed in the network as well as in the variety of devices that we see connected to our WAN networks. The rise of cloud service providers and Software as a Service (SaaS) has altered traffic patterns as well. In response SD-WAN has changed the deployment paradigm with simplified deployment, policy control, integrated cloud connectivity, and SDN principles applied to the WAN.

What these changes mean is that even more traffic than in the past is northbound-southbound. As SaaS has become more and more prevalent, applications that were once peer-to-peer distributed have moved to the cloud. For example, unified communications are now centralized from a control plane perspective, Session Initiation Protocol (SIP) trunking is prevalent, and more real-time traffic is traveling to a central hub site before reaching its destination. SaaS application traffic also travels to an edge Point of Presence (PoP) before reaching the centralized application servers. Cloud networking has increased the centralization of traffic flows even more than in the past, forcing more hub-and-spoke traffic patterns. This further diminishes the effectiveness of direct, site-to-site tunnels, whether static or dynamic.

Costs of dynamic, site-to-site tunnels

Any implementation of dynamic tunnels comes with costs that should be weighed against the benefits before you decide to deploy them in your network. Some of the costs to consider are as follows:

- **Load and scale:** In SD-WAN, tunnel performance is measured by active probing. These probes consist of packets that can be used to measure delay, jitter, and loss. The accuracy of the measurement and the ability to react faster to degradation are dependent on the frequency of the probing. Increased accuracy and reaction time will increase load on the SD-WAN device CPU as well as overhead on the tunnel. Depending on the number and frequency of dynamic tunnels being created, this load can be highly dynamic. If load exceeds the capacity of the SD-WAN device, this can manifest itself in failures that may not point to an obvious cause. Significant time and debugging effort may be required to determine the root cause of such failures.

- **Efficacy:** In SD-WAN, traffic is mapped to tunnels based on the best-performing tunnel path. This mapping cannot be done with dynamic tunnels, as the path cannot be measured before the tunnel is established. Thus, your SD-WAN control system can only "assume" that the direct path will be faster. If it turns out to be wrong, the traffic will thrash back to the static, preestablished path.
- **Path continuity:** Any implementation of dynamic tunnels must ensure connectivity through some other path during the tunnel setup period. Otherwise, packets may be dropped while waiting for the tunnel to be established. Since this happens at the beginning of the flow, it is often control packets that get dropped, interfering with session setup and creating drops or miscalls for which it is very hard to find the root cause.
- **Security:** Creation of direct, site-to-site tunnels may have implications for security policy. It may require deployment of firewall scanning at each site, and this should be considered part of the overall cost of deployment.
- **Quality of Service (QoS):** Different sites have different packet processing capacity. Shaping must be applied on the tunnel ingress in order to ensure that higher-capacity sites do not flood lower-capacity ones. Therefore, any dynamic site-to-site tunnel implementation should be able to dynamically determine the target site interface bandwidth in order to configure the source shape rate toward the target site. Static configuration of these bandwidths is cumbersome and error prone, with potentially large static lists at numerous sites.
- **Troubleshooting:** Since dynamic tunnels are ephemeral, debugging associated problems is difficult. Reproduction can be time-consuming. Good syslog and debug messages (such as those produced in Cisco IOS® Software) are essential. The thoroughness of such messages can be developed only with time and experience; for example, the assurance capabilities of DMVPN have been augmented significantly over its more than 15 years of deployment.

Dynamic tunnels in large-scale routing environments

When Cisco created dynamic tunnels in DMVPN, its original implementation distributed a full set of underlay and overlay routes to all sites. That way each site would have a direct underlay next-hop route to every other site such that when the tunnel endpoint route was resolved, it resolved the WAN IP of the destination site router itself. A tunnel could then be set up directly. It soon became apparent that in large-scale environments, the number of routes began to scale unmanageably for smaller site routers with limited processing scale and memory. Any dynamic tunnel implementation must manage this route scale to ensure that smaller site routers are not overwhelmed by route table size or convergence processing.

DMVPN handled the issue via Next-Hop Resolution Protocol signaling. The next-hop information was advertised by the hub to the routers selectively, sharing direct routes only after determining the existence of traffic flow destined between two peers on the same DMVPN network. There was no packet loss during this advertisement and dynamic tunnel setup because the traffic was already flowing between peers via through preexisting hub site routes.

Integrating scalable dynamic tunnels into SD-WAN

Cisco® SD-WAN further simplifies and builds upon what DMVPN did. With Overlay Management Protocol (OMP), dynamic tunnel establishment can be integrated as part of OMP updates in a way that allows for fast, more scalable tunnel establishment. Because of experience gained from a wide range of DMVPN deployments over many years, Cisco engineers have learned a lot about the network dynamics that can impact dynamic tunnel scalability, and many of these learnings have driven the development of the new IPsec architecture found in Cisco SD-WAN.

With OMP, the control plane is now separated from the data plane, and routes can be advertised to all sites or just a subset of sites. These OMP routes can be tagged with a number of different attributes. One of these route attributes can be "Dynamic." When you want to enable dynamic tunnels, you will use the route policy workflow in vManage to tag the OMP Transport Location (TLOC) routes (WAN IPs) registered by any site or subset of sites as Dynamic.

You can then filter the TLOCs such that vSmart will advertise only a hub TLOC and a selected set of TLOCs marked as Dynamic. For TLOCs marked as Dynamic, the VPN tunnel will be established on demand, rather than statically. You can control the topology and sites where dynamic tunnels are allowed via vManage policies, minimizing the set of routes that must be carried at each site.

A word about loss

Your SD-WAN also provides specialized features to mitigate the effects of loss. These include the ability to turn on forward error correction to reconstruct errored packets at the destination or to create duplicate, redundant traffic streams that can be sent on diverse paths through the WAN. Both of these features consume significant additional bandwidth on your WAN, which comes at a real cost to your organization in terms of both dollars and efficiency.

Most demonstrations of these features insert loss artificially to show the benefit of the feature in operation. But before you decide to deploy these features, take a look at the recent statistics of service provider loss in real networks shown in Table 3.

Table 3. Recent intercity U.S. network loss measurements

AT & T	Network Loss																								
U.S. Network Loss* in %																									
City Pairs	Atl																								
Austin	0	Aus																							
Cambridge	0	0	Cam																						
Chicago	0	0	0	Chi																					
Cleveland	0	0	0	0	Cle																				
Dallas	0	0	0	0	0	Dal																			
Denver	0	0	0	0	0	0	Den																		
Detroit	0	0	0	0	0	0	0	Det																	
Houston	0	0	0	0	0	0	0	0	Hou																
Indianapolis	-	-	-	0	0	-	-	-	-	Ind															
Kansas City	0	0	0	0	0	0	0	0	0	-	Kan														
Los Angeles	0	0	0	0	0	0	0	0	0	-	0	LA													
Madison	-	-	-	0	-	-	-	-	-	-	0	0	Mad												
Nashville	0	0	0	0	0	0	0	0	0	-	0	0	-	Nas											
New Orleans	0	0	0	0	0	0	0	0	0	-	0	0	-	0	NO										
New York	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	NY									
Orlando	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	Orl								
Philadelphia	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	Pa							
Phoenix	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	0	Phx						
San Antonio	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	0	0	SA					
San Diego	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	0	0	0	SD				
San francisco	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	0	0	0	0	SF			
St. Louis	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	StL		
Seattle	0	0	0	0	0	0	0	0	0	-	0	0	-	0	0	0	0	0	0	0	0	0	0	Sea	
Washington	-	-	-	-	-	-	0	-	-	-	-	0	-	-	0	-	0	0	0	0	0	0	0	0	Was
Current Overall Average: 0.00%																									

As you can see, there is no measurable loss rate in this network. There may be bursts or anomalies, but they are not statistically significant, and at any rate, SD-WAN performance measurements have the ability to react and route around any short-term loss incidents.

It is further interesting to note that even on transcontinental paths, service providers are able to achieve four-nines reliability of packet delivery. Loss is very much a thing of the past in modern networks – one less thing for you to worry about.

Dynamic tunnels: What about security and scale?

Another key factor when considering an SD-WAN architecture for your network is how IPsec secure tunnels are established. The method used can have a big impact on your network scale and may be noticed only at the most critical times – such as when a large companywide videocast is starting or during an SD-WAN appliance failure at an aggregation site. Events of these types can drive very high load on the CPU of SD-WAN data plane devices as they struggle to keep up with a rush of tunnel-establishment requests from remote devices. Some devices may not be able to keep up or may lose routing connectivity if the load on the control plane CPU becomes too high, starving other processes.

These types of problems come from the peer-to-peer nature of IPsec and Internet Key Exchange (IKE) tunnel establishment and the fact that in most networks, the far endpoint of the tunnel is a large-scale aggregation device. This aggregation device often has too many peers to keep up with if they all want to establish tunnels at the same time. In addition, IKE protocol negotiation takes time to establish a secure channel for key exchange, and it does this every time key exchange is required.

Figure 2. Scalable key distribution in Cisco SD-WAN

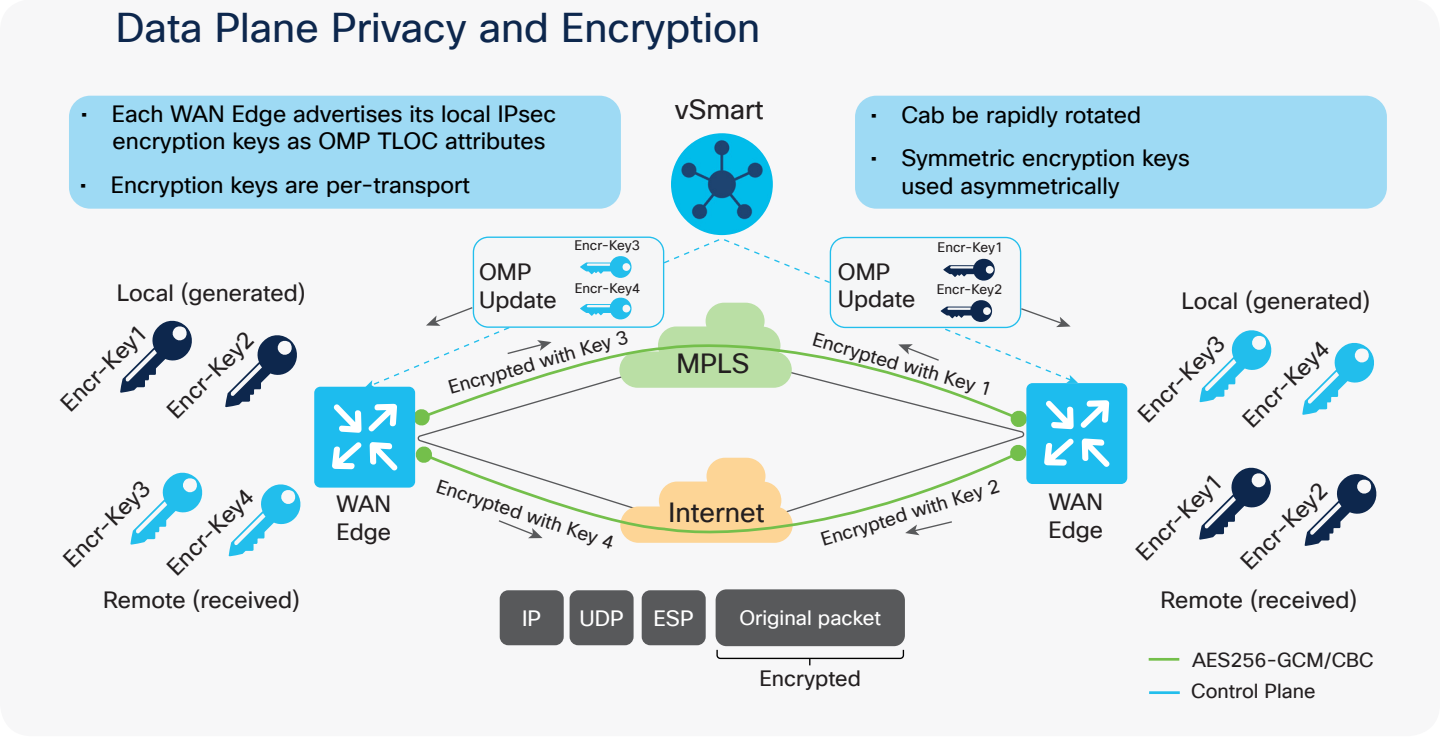

In Cisco SD-WAN, this problem is eliminated by creating a highly secure, preestablished control channel among all SD-WAN devices. The channel is authenticated by Public Key Infrastructure (PKI) certificate exchange and encrypted with the latest AES265-GCM ciphers. It's persistent and thus available for securely exchanging all kinds of information, including IPsec keys for secure tunnel establishment between SD-WAN peers. Figure 2 illustrates how

OMP advertises IPsec keys securely over this preestablished channel. This means that keys can be exchanged and prestaged at SD-WAN peer devices and are available immediately when needed. This prestaging eliminates the scale issues with peer-to-peer IKE negotiation described earlier.

Authentication

It is important to understand how your VPN tunnels are authenticated. Some SD-WAN solutions do not really automate VPN establishment in a fundamental way. They only automate the distribution of Preshared Keys (PSKs) among the VPN peers. They still have the weakness of PSK-based implementations, including management of the PSKs and appropriate selection and rotation of these keys. PKI methods are a much more secure way of establishing authenticity.

Alternatives

How should you approach a decision about whether or not to deploy dynamic site-to-site tunnels? First, consider your service provider or ISP network performance. Most service providers track latency, jitter, and loss, and those statistics are available to you. SD-WAN assurance features allow you to track this performance in real time.

If you find network performance acceptable with enough margin, hub-and-spoke topologies work just fine for voice and many other real-time applications. In addition, hub-and-spoke networks require the minimum tunnel scale at remote sites, create less load on network devices, and are easier to control from a policy perspective as well as troubleshooting. Start here first. Many Cisco SD-WAN customers are running hub-and-spoke topologies with excellent delay performance. Consider regionalizing your network hubs more aggressively if needed. Finally, if optimized paths are the only option, engineer them using static tunnels, only between sites where they are needed.

Conclusion

Cisco DMVPN introduced dynamic tunnels almost 15 years ago, in a time when private and public network performance was very different than it is today.

Dynamic tunnels can be a useful tool in some network scenarios, but they should be used only for applications that need them, and their efficacy should be tested to ensure that the promised performance is achieved. If you do decide that they are needed, dig deeper into the implementation to ensure that the overall product architecture is sound from a security point of view: Are tunnels properly authenticated with PKI certificates, and can the data plane devices support a high rate of tunnel reestablishment?

Cisco believes that dynamic tunnels are a useful tool, but in a narrower application space than in the past. While such support is nice to have, you should consider whether the underlying SD-WAN implementation is sound before making a decision. Cisco SD-WAN with OMP has been architected to address the scale and security concerns related to dynamic tunnels that we've talked about here. It is something that we plan to support in a future release of Cisco SD-WAN.

C11-743108-00 12/19