

Innovations in Ethernet Encryption (802.1AE - MACsec) for Securing High Speed (1-100GE) WAN Deployments

Authors

Craig Hill
Distinguished Systems Engineer
U.S. Federal Area

Stephen Orr
Distinguished Systems Engineer
U.S. Public Sector

Introduction

Over the course of the past decade, customer demand for increasing Wide Area Network (WAN) bandwidth has been driving the networking industry to continually innovate in order to increase WAN transport speeds. Thus, we have witnessed the evolution from Asynchronous Transport Mode (ATM) to Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) and, more recently, innovations in Ethernet and optical. Ethernet and optical have now emerged as the de facto standards and we have seen

speeds grow from 10-Gb, 40-Gb, and now to 100-Gb speeds with no end of growth in sight.

Demand for increased bandwidth continues, driven by cloud services, mobile devices, and massive increases in video traffic. With the shift to cloud and mobile services, the need for ever-faster WAN transport speeds continues in order to handle the traffic created by locating applications and data off-premises.

While link speeds and demand for bandwidth continue to increase, the innovation of encryption technologies for securing these high-speed links, specifically for the service providers, cloud providers, large enterprises and governments, has failed to keep up. Furthermore, customers want to simplify their network operations and reduce the amount of protocol layers and complexity they are implementing in these high-speed networks, including the recent interest to hide network layer information in transit (IP addresses and protocol port numbers).

This document provides an in-depth look into:

- How Cisco is addressing this dilemma of link speed bandwidth outpacing the encryption technologies currently available
- Encryption innovations led by Cisco, including a detailed introduction to WAN Media Access Control Security (MACsec)
- How Cisco is giving the 10-year old 802.1AE MACsec standard a technology “face lift” and innovating to meet the new customer demands for high-speed WAN encryption (1G – 100G+) for WAN data center interconnect, branch back-haul, and Metro Ethernet
- Detailed use cases and analyses from the perspective of enterprise customers as well as service providers offering transport services (Metro Ethernet, IP/Multiprotocol Label Switching [MPLS], as well as cloud service providers)
- A comparison of MACsec and IPsec, but also how each technology complements the overall Cisco® encryption solution portfolio and, in some cases, can be combined

The Growing Interest in High Speed Encryption


For many years, IP Security (IPsec) was synonymous with encryption in the WAN, specifically over the Internet. It has been the dominant encryption solution for customers back-hauling business traffic from remote and branch office locations, as well as being the encryption choice of most Virtual Private Network (VPN) clients. IPsec is an encryption solution operating at the IP layer of the Open Systems Interconnection (OSI) model and is flexible in that it can operate over any IP transport including private and public (Internet) transport. Many large-scale IPsec deployments are currently in operation across enterprise and government networks today.

IPsec has proven to be extremely flexible, transport agnostic, and capable of scaling to thousands of end devices. It is, however, proving to be more challenging from an overall throughput perspective for newer applications and cloud providers. Several shifts in new applications and the explosion of cloud are changing designs, including:

- Increasing bandwidth demands over the WAN for branch offices, application deliveries, video content distribution, and data center intraconnections.
- Fewer applications are run locally in branch locations, and thus driving the need for higher speed transport.
- Highly resilient cloud computing architectures driving high-speed data center replication across geographically dispersed locations.
- Traffic pattern changes to a more any-to-any model, dictated by trends such as cloud, machine-to-machine (M2M) communications, and the Internet of Things (IoT) and Internet of Everything (IoE).
- Encryption landscape that is changing in the U.S. government (Commercial Solutions for Classified CSfC, transport security) that is driving the need for high-speed layered encryption solution offerings.

As noted previously, cloud computing and new applications continue to emerge that are changing the traffic patterns of routed networks, as well as outpacing the encryption rates traditional IPsec can support. As shown in **Figure 1**, using IPsec as an example, the encryption performance capabilities are no longer aligned with links speeds as the links move to 40/100G and beyond. For example, some of the higher performing IPsec engines in routers today target approximately 75-Gbps IPsec performance, unidirectional flow, at 1400-byte packet sizes. As the applications require bi-directional flow patterns, that number gets cut in half to approximately 37 Gbps. Then, introduce Internet Mix (IMIX) traffic patterns or smaller packet sizes based on the application being encrypted, and the overall IPsec performance drops further.


Figure 1. Link Speeds Outpacing IP Encryption Example with Cisco ESP-200



Furthermore, if the deployment requires that all traffic leaving/traversing the router must be encrypted, the overall throughput of the router is now restricted to the performance of the IPsec engine which, in most cases, can be a fraction of the router’s aggregate forwarding capabilities. This is a huge factor from an economics perspective of cost per bit through the router and MACsec changes the encryption cost per bit through routing elements. For deployments requiring encryption and the capability of leveraging an Ethernet transport (public or private), MACsec offers a simplified, line-rate, per port encryption option for secure next-generation deployments.

As shown in **Figure 2**, MACsec, as the name implies, is MAC layer or link layer encryption and offers encryption equal to that of the Ethernet port rates (1/10/40/100Gbps) bidirectionally regardless of the packet size, executing the encryption function in the physical layer (PHY) of the Ethernet port. Unlike IPsec, which is typically performed on a centralized application-specific integrated circuit (ASIC) optimized for accelerating encryption, MACsec is enabled on a per-port basis with no performance impact.

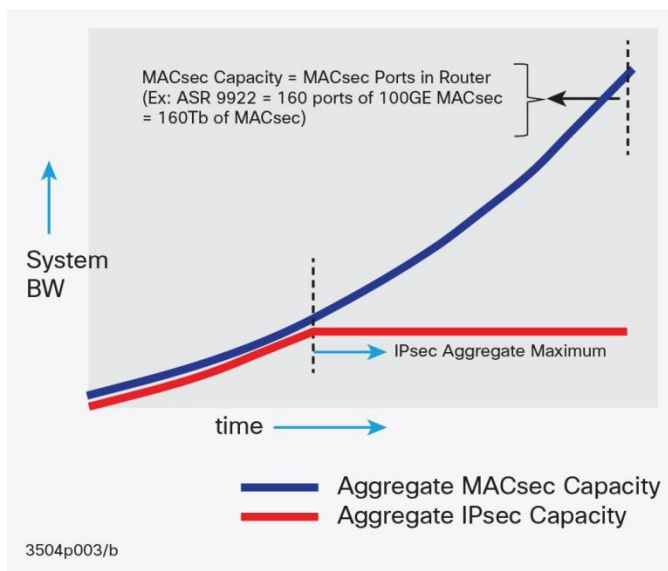
Figure 2. Link Speeds Aligning with Encryption Using MACsec (Example Using ASR 9000 100G MACsec)



For a router capable of forwarding terabits of traffic, IPsec encryption will be the bottleneck and limiting factor of maximum throughput of the device. For example, if a router has multiterabit forwarding capabilities, and ten 100-GE ports require encryption at line-rate, the MACsec solution offers 1 Tbps of AES-256 encryption on each port, regardless of the packet size, so the overall encryption throughput utilizing MACsec can leverage the full forwarding capability of the router, while also offering encryption of each bit on the Ethernet wire.

As shown in **Figure 3**, the encryption capacity, as it relates to the entire chassis of the routing platform, is exponential as it relates to MACsec versus IPsec. For IPsec-based systems, the encryption engine is maximized based on the encryption off-load mechanism.¹ For MACsec, encryption grows by MACsec port capabilities. Using the Cisco ASR 9000 series Aggregated Services Router system as an example, a Cisco ASR 9922 has 20 usable slots. Adding 20 line cards that support 8-port 100 GE MACsec each, this system can support an aggregate of 160 Terabits of AES-256/GCM MACsec encryption within a single chassis.

Figure 3. Aggregate MACsec vs. IPsec Encryption System Capacity



Note: While MACsec offers a new set of high-speed encryption capabilities, IPsec is now, and will remain, a vital element to network designs, offering an extremely agile design option when IP (public or private) is the transport available. MACsec offers network designers another option when Ethernet can be leveraged as the end-to-end WAN/Metro transport and high-speed encryption is vital to the overall business requirement.

Overview of MACsec

In 2006, 802.1AE was standardized by the IEEE 802.1 working group. 802.1AE-2006 defines Media Access Control Security, or MACsec, which enables devices on point-to-point or shared Ethernet networks to provide confidentiality, integrity, and authenticity for user data. MACsec supports and facilitates:

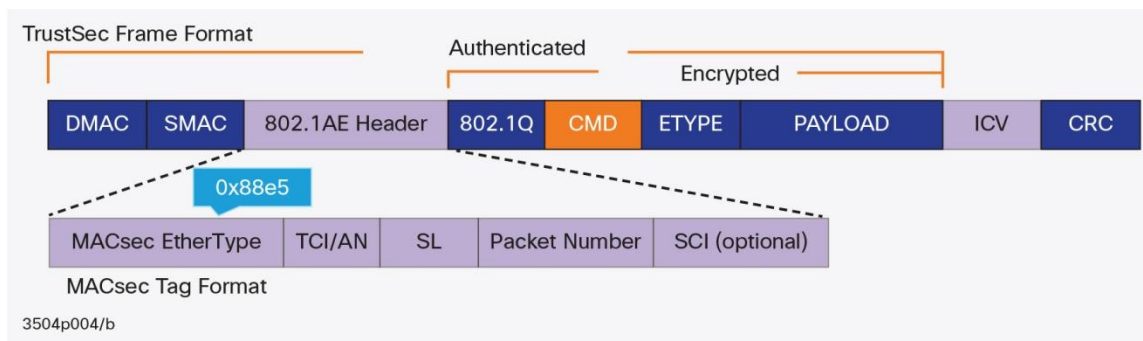
- Maintenance of correct network connectivity and services
- Isolation of denial of service attacks
- Localization of any source of network communication to the LAN of origin

¹ Depending on the routing system, this can be done either with a specific line card or onboard ASIC internal to the platform.

- The construction of public networks offering service to unrelated or possibly mutually suspicious customers, using shared LAN infrastructures
- Secure communication between organizations, using a LAN for transmission
- Incremental and nondisruptive deployment, protecting the most vulnerable network components²

While MACsec is based on the standard Ethernet frame format, an additional 16-byte MACsec Security Tag or SecTAG was included, as well as a 16-byte Integrity Check Value (ICV) at the end of the frame. Walking through the MACsec frame format depicted in **Figure 4**, there are no changes to the destination and source MAC address.

Figure 4. MACsec Header Format



However, the new SecTAG field, which is 16 octets long, is as follows:

- MACsec EtherType: The first two octets and the value are set to 0x88e5 and designate that the following frame is a MACsec frame.
- TCI/AN: The third octet is the TAG Control Information (TCI)/Association Number field. The TCI designates the MACsec version number, if confidentiality or integrity are used alone.
- SL: The fourth octet is **short length**, which is set to the length of the encrypted data.
- PN: Octets 5 through 8 are the packet number and are used for replay protection and the construction of the initialization vector (along with the secure channel identifier [SCI]).
- SCI: Octets 9 through 16 are the secure channel identifier. Each connectivity association (CA) is a virtual port and each virtual port is designated a secure channel identifier that is the concatenation of the MAC address of the physical interface and a 16-bit port ID.

The MACsec frame can be both encrypted and authenticated to provide privacy and integrity. MACsec utilizes the Galois/Counter Mode Advanced Encryption Standard (AES-GCM) for authenticated encryption and Galois Message Authentication Code (GMAC) if only authentication, but not encryption is required.


The current MACsec standard encrypts all fields after the SecTAG which obfuscates fields such as MPLS labels, 802.1P and 802.1Q from the original Ethernet frame. Subsequently, any intermediary network device that may require those tags are not able to see them as the Ethernet frame traverses the underlying transport between encrypted stations.³

² IEEE Standard for Local and Metropolitan Area Networks – Media Access Control (MAC) Security, IEEE 802.1AE-2006.

³ MACsec header offset capabilities is not defined in the standard and is open to vendor implementations.

Figure 5 is an example of the MACsec encryption process applied per link. As the Ethernet frame enters the PHY layer encapsulation process, the MACsec encryption is applied on egress of the router. Decryption of MACsec frames are performed on ingress to the router interface.

Figure 5. MACsec Using Per-Hop Encryption

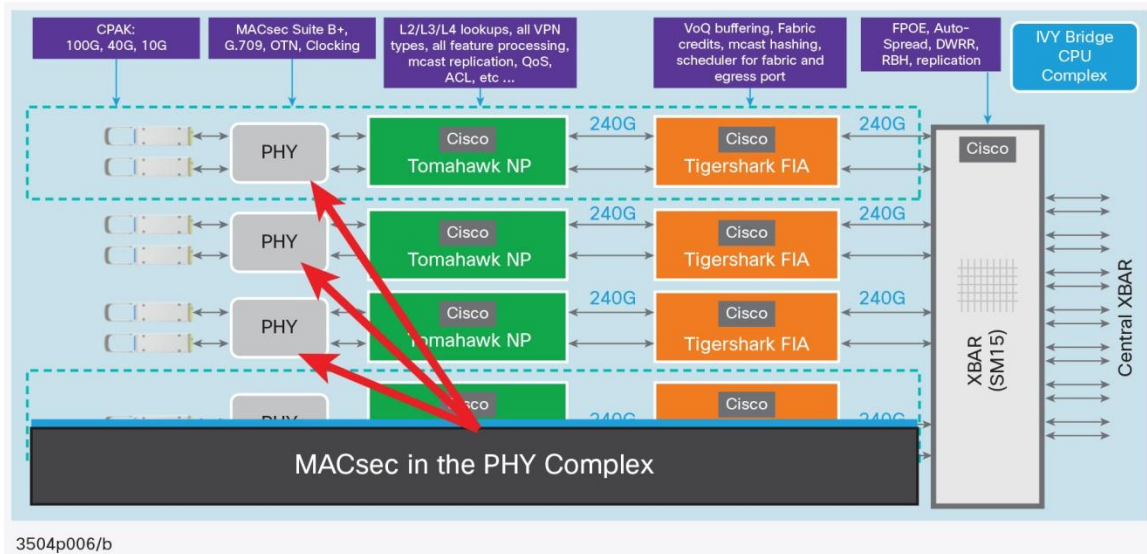


Leveraging the hop-by-hop, or per-link, nature of the MACsec decryption/encryption process on ingress/egress in the frame forwarding procedure offers several advantages over end-to-end encryption technologies like IPsec or Transport Layer Security (TLS)/Datagram Transport Layer Security (DTLS):

- Layer 2-Layer 7 services can be performed on the Ethernet frame or IP Packet since it is “in the clear” prior to being encrypted on egress
- No impact on Ethernet frame markings 802.1p for Quality of Service (QoS), 802.1Q tag, Q-in-Q tags
- Policy routing, QoS, filtering
- Low latency/low overhead
- Support for Jumbo frames
- No need for complex policy statements to define “interesting” traffic to encrypt

As depicted in **Figure 6**, the Cisco ASR 9000 100 GE line card, on ingress the MACsec frame is decrypted in the PHY prior to performing all ingress functions (MPLS label imposition, queuing, scheduling, access control lists [ACLs], etc.). On egress, the process is reversed such that Layer 2–Layer 7 services are performed prior to MACsec encryption of the frame, which is done on the PHY.

Figure 6. MACsec in the PHY – ASR 9000 “Tomahawk” Line Card Example



There have been two amendments to the original 802.1AE specification since its release in 2006:

- Amendment 1: 802.1AEbn-2011—This amendment adds GCM-AES-256 as an optional cipher in addition to the mandatory GCM-AES-128.⁴
- Amendment 2: 802.1AEbw-2013—This amendment extended packet numbering adds two additional cipher suites GCM-AES-XPB-128 and GCM-AES-XPB-256. As link speed increased to 10/40/100 Gig—the original 32-bit packet number (PN) field was not adequate enough to handle the higher speed interfaces and could cause a new security association to occur every 5 minutes for a 10G interface. Subsequently, the PN field was increased to 64 bits providing for 264 Ethernet frames to be transmitted before a new security association is required.⁵

The MACsec Key Agreement (MKA)

MKA Overview and Terminology

The MACsec Key Agreement (MKA) is included as part of the IEEE 802.1XREV-2010 Port-Based Network Access Control Standard. The purpose of MKA is to provide a method for discovering MACsec peers and negotiating the security keys needed to secure the link. There are three ways defined within the 802.1 standard for the generation of keying material for use with MKA:

- Pre-shared Keys (PSK)
- The master session key which is a product of a successful Extensible Authentication Protocol (EAP) authentication
- Key distributed from an MKA key server

Note: A working knowledge of 802.1X and the Extensible Authentication Protocol (EAP) is assumed and not included as part of this document.

⁴ IEEE Standard for Local and Metropolitan Area Networks—Media Access Control [MAC] Security Amendment 1, IEEE 802.1AEbn-2011

⁵ IEEE Standard for Local and Metropolitan Area Networks—Media Access Control [MAC] Security Amendment 2, Extended Packet Numbering, IEEE 802.1AEbw-2013

First, we'll take a moment to describe the basic nomenclature of the components used as part of MKA, the key hierarchy, and distribution process. **Table 1** lists the MKA terminology and their definitions.

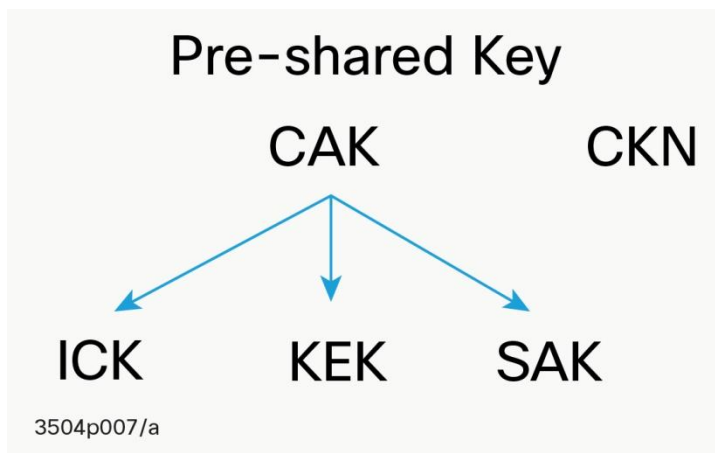
Table 1. MKA Terminology

Acronym	Definition
MKA	MACsec Key Agreement: Defined in IEEE 802.1XREV-2010 is a key agreement protocol for discovering MACsec peers and negotiating keys.
EAP	Extensible Authentication Protocol: Defined by IETF RFC 3748, a flexible framework for authentication for Wired and Wireless Local Area Networks is encapsulated via 802.1X.
CA	Secure Connectivity Association: A security relationship between MACsec-capable devices on a LAN or WAN.
MSK	Master Session Key: Generated during EAP exchange. Supplicant and authentication server use the MSK to generate the CAK.
CAK	Connectivity Association Key: Is either a manually entered Pre-shared Key, derived from the MSK if an EAP method is used or a key delivered from a MKA Key Server. The CAK is a long-lived master key used to generate all other keys used for MACsec.
CKN	Connectivity Association Key Name: Identifies the CAK.
ICK	Integrity Check Key (ICK): Used to prove an authorized peer sent the message.
KEK	Key Encrypting Key: Used to protect the MACsec keys (SAK).
SAK	Secure Association Key: Derived from the CAK and is the key used by the network device ports to encrypt traffic for a given session.
KS	Key Server: Responsible for selecting and advertising a cipher suite and generating the SAK.

MKA Key Hierarchy

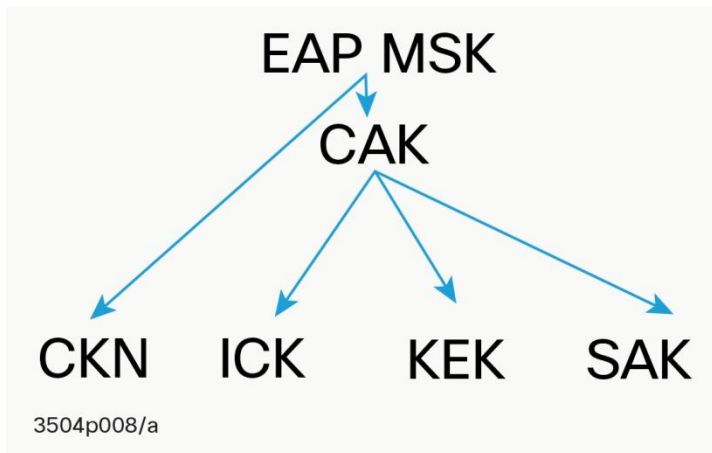
Two methods can be utilized to derive the MACsec Encryption Keys: manual pre-shared keys or 802.1X/EAP. As shown in **Figure 7**, when pre-shared keys are used the pre-shared key (PSK) is equal to the connectivity association key (CAK) and the connectivity association key name (CKN) must be manually entered and is stored in the device's configuration. The CAK is then used to generate the rest of the MACsec encryption keys (ICK, KEK, and SAK).

Figure 7. Pre-shared Key Derivation



As depicted in **Figure 8**, when 802.1X/EAP is utilized the master session key (MSK) is generated as a by-product of the EAP Authentication process. The CAK is then derived from the MSK. Unlike the pre-shared key method where the connectivity association key name is manually entered, the CAK is also derived from the MSK. As with the pre-shared key method, the remainder of the MACsec keys derived from the CAK.

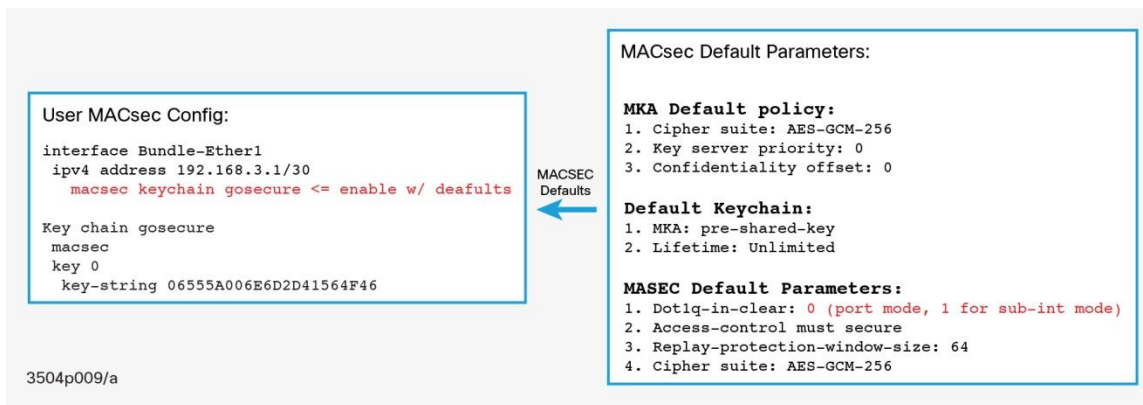
Figure 8. 802.1X/EAP Key Derivation



MACsec Policies – Applying MACsec to an Interface

MACsec policy choices are configured on a per-interface or subinterface basis and designate whether the link is encrypted or not. **Figure 9** shows a minimal configuration on an interface utilizing the default parameters.

Figure 9. MACsec Policies



One important configuration parameter for the MACsec policy is whether the interface participates in MKA/MACsec. The available MACsec policies are Must Secure, Should Secure, and Should Not Secure.

- **Should-Not-Secure:** The switch does not perform MKA. If another network device sends MKA protocol frames, they are ignored. The network device sends and receives unencrypted traffic only.
- **Should-Secure (default):** The switch attempts MKA. If MKA succeeds, the switch sends and receives encrypted traffic only. If MKA times out or fails, the network device permits unencrypted traffic.
- **Must-Secure:** The network device attempts MKA. If MKA succeeds, only encrypted traffic is sent or received. If MKA times out or fails, the connection is treated as an authorization failure by terminating the session and retry authentication after a quiet period.

What Is WAN MACsec?

As was mentioned earlier, MACsec has been in existence since 2006, so what is changing? For one, Ethernet has evolved beyond a private LAN transport to include a variety of WAN transport options and offerings as well. Ethernet is widely offered by service providers as their primary transport offering to customers, either as a point-to-point circuit replacement of older T1/T3, SONET/SDH, or back-haul services to private IP VPN services, Metro Ethernet offerings, to name a few.

The WAN MACsec offering is standards based but offers additional capabilities not found in earlier MACsec capabilities. More specifically, MACsec can be leveraged by enterprise customers over public carrier Ethernet offerings, allowing customers to adapt to the public carrier Ethernet service offering and capabilities (or restrictions).

New enhancements for WAN MACsec include:

1. 802.1Q Tag in the Clear
2. Standard IEEE 802.1X-rev MACsec Key Agreement
3. Integrated MACsec authentication adaptability over public Carrier Ethernet transport

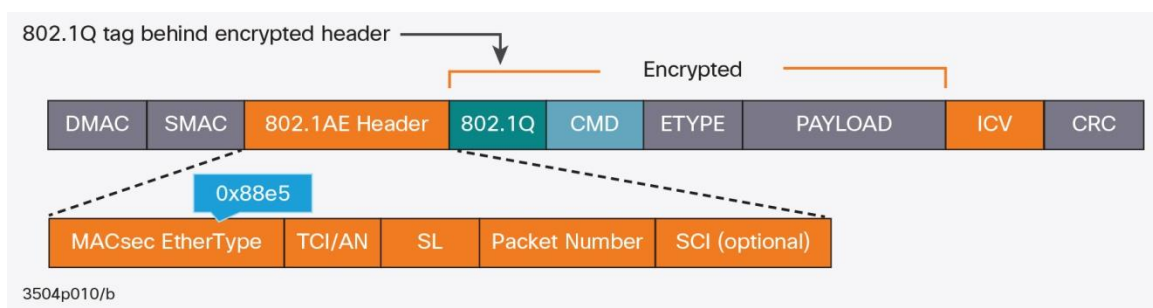
1. 802.1Q Tag in the Clear

This enhancement offers the ability to expose the 802.1Q tag outside the encrypted MACsec header. Exposing this field offers a multitude of design options with MACsec, and in some cases of public Carrier Ethernet transport providers, is necessary for leveraging certain transport services (see use case section).

While offering high-speed encryption, the multipoint use case exposes limitations and impracticalities in recent MACsec-offered solutions. Why: Because earlier MACsec solutions did not offer the ability to expose the 802.1Q tag in the header, requiring a physical Ethernet connection on the central site, per branch. This was not a realistic design due to complexity of cabling, cost of each port, and “box” real estate required in the router to terminate this 1-to-1 remote site to physical-port requirement.

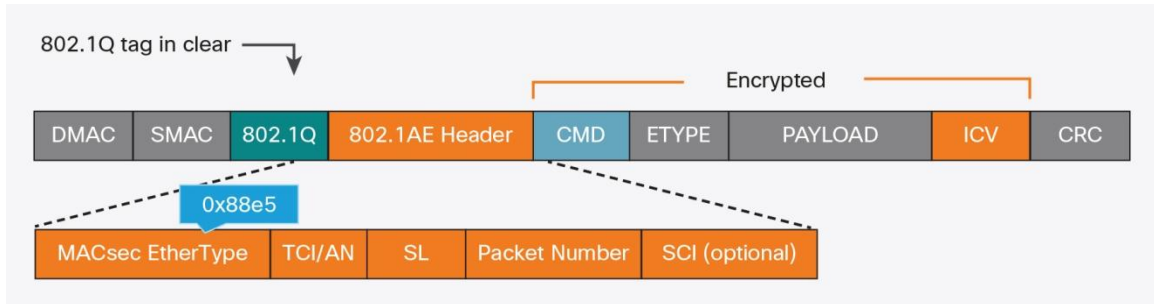
As described earlier, and as shown in **Figure 10**, the original MACsec header format encoded the 802.1Q tag as part of the encrypted payload, thus hiding it from the public Ethernet transport, which also limited the topologies and network design options that could be leveraged when transporting Ethernet frames over public or private Ethernet.

Figure 10. WAN MACsec – 802.1Q Tag Encrypted



With 802.1Q tag in the clear, as shown in **Figure 11**, the 802.1Q tag is encoded outside the 802.1AE encryption header, exposing the tag to the private and public Ethernet transport, which opens up a multitude of design options when using WAN MACsec.

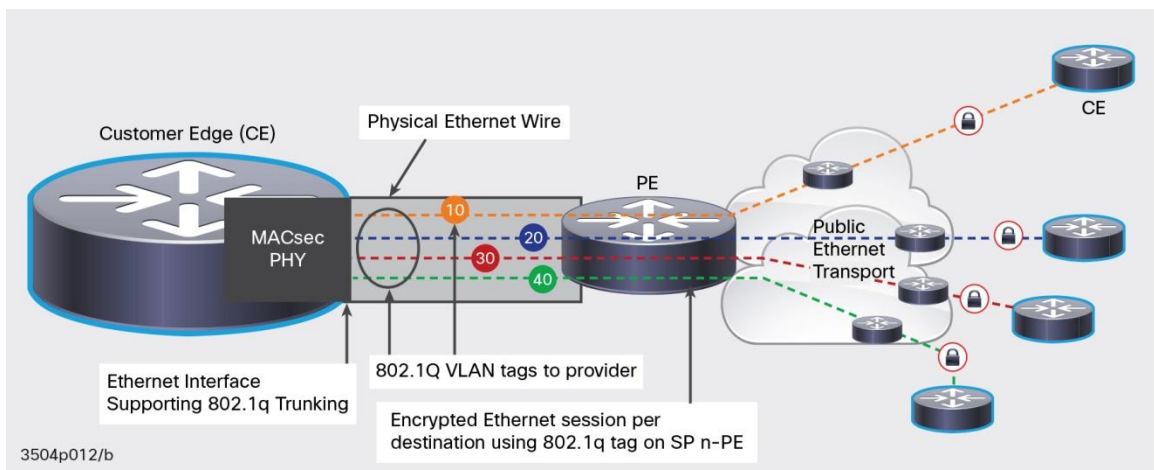
Figure 11. WAN MACsec – 802.1Q Tag in the Clear Example



One of the primary use cases designers are looking to leverage with this new “tag in the clear” capability is the ability to build hub/spoke networks with WAN MACsec over public Ethernet Virtual Private Line (E-LINE) services.

Figure 12 depicts a hub/spoke design leveraging WAN MACsec 802.1Q tag in the clear to support remote site connectivity over a public ELINE service. In this example, the hub site router is leveraging a Layer 3 (IP) sub-interface per 802.1Q virtual local area network (VLAN) that is associated with each remote site branch, through coordination with the carrier Ethernet provider, so that each remote branch site leverages the exact 802.1Q tag to E-LINE circuit association. The result is a highly flexible MACsec hub/spoke design that eliminates the older solutions that required a physical interface “per remote site” on the hub router.

Figure 12. MACsec Tag in the Clear for a Hub/Spoke Design



Typical deployments for hub/spoke MACsec solutions are most typically used by enterprise, commercial, and federal customers that have encryption requirements that exceed what IPsec can offer and/or are looking to eliminate the complexity that IPsec requires in certain designs.

It should be noted that while Cisco WAN MACsec solution can leverage tag in the clear for virtual segmentation of connections, these tags can also leverage the 802.1p bits carried in that tag, for QoS service offerings. Without this capability, the QoS offerings will be much more coarse and typically very limited.

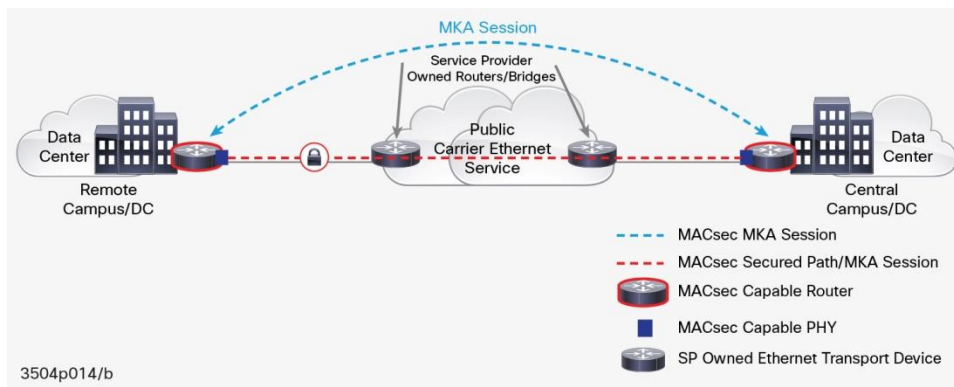
2. Standard IEEE 802.1X-rev MACsec Key Agreement

As detailed earlier in this document, MKA is a standard keying agreement method so, while not new, supporting MKA in WAN MACsec provides support for multivendor interoperability when required or needed for standards-based certification processes (for example, Federal Information Processing Standard FIPS) and Common Criteria).

3. Integrated MACsec Authentication Adaptability over Public Carrier Ethernet Transport

The ability to provide the operator with the capabilities to adapt MACsec key ringing functions over any public Ethernet carrier is viewed by some as a vital capability on the Customer Edge (CE) router. In other words, the CE must comply with the standards the service provider leverages in its Carrier Ethernet offering. This directly applies to carriers that are leveraging a MAC address lookup function in the provider backbone bridges to forward frames through their transport network. While this may seem irrelevant, it has been proven to wreak havoc on the MKA “keying process” when deploying MACsec over the Carrier Ethernet transport (**Figure 13**).

Figure 13. WAN MACsec Site-to-Site Example and Components



To elaborate further, Metro Ethernet Forum (MEF) standards dictate a specific set of well known MAC addresses deemed as “for me” frames to the carrier Ethernet forwarders—meaning transit Carrier Ethernet switches consume the frames containing these MAC addresses into their control plane for processing. Current MACsec and MKA implementations leverage an EAP over LAN (EAPoL) packet for MKA key negotiation and these EAPoL MAC addresses fall under the MEF “well known” MAC addresses for consumption. This means that customers deploying MACsec over a public Carrier Ethernet transport that operate this Ethernet service with Carrier Ethernet switches that consume these EAPoL frames, cannot leverage MACsec across these providers.

To mitigate this problem, Cisco introduced the ability for the operator deploying WAN MACsec to change the EAPoL destination address⁶ and/or EtherType to an address that is defined in the provider’s bridge as “uninteresting.”

⁶ EAPoL destination address change link – http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/macsec/configuration/xs-3s/macsec-xe-3s-book.html#concept_EBD97E2A36A142B2957F33C0EB8E1DFA

Example:

The “eapol destination-address” command allows the operator to change the destination MAC address of an EAPoL packet that is transmitted on an interface towards the service provider Ethernet transport, as shown in **Figure 14** (CLI commands from IOX-XE on the Cisco 1001-X).

Figure 14. EAP over LAN (EAPoL) Configuration Example

<pre>... interface GigabitEthernet0/0/4 macsec dot1q-in-clear 1* macsec replay-protection-window-size 100 eapol destination-address broadcast</pre>	Leverage “broadcast” address as the destination EAPoL address. Provider switch will forward as standard “broadcast” Ethernet frame.
---	---

As shown above in Figure 14, this configuration capability overcomes the limitation imposed by the Carrier Ethernet provider and allows the consumer of the Ethernet transport to leverage MACsec over any public Carrier Ethernet network.

Carrier Ethernet Transport Description and WAN MACsec Use Case Overview

Carrier Ethernet Service Description

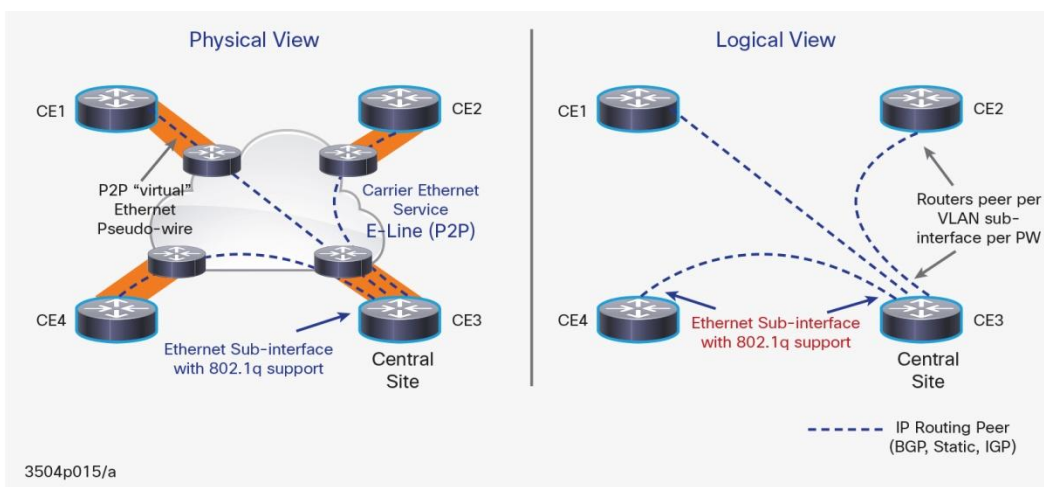
Like any network deployment over a public transport, it is vital that the designer understands the capabilities of that transport and the implications it can have on the behavior of the applications, routing design, and service offerings traversing it. Public Ethernet transport is no different and, in some views, introduces more complexities. While the Metro Ethernet Forum (MEF) categorizes Ethernet transport options, the focus in this document will be on E-LINE and E-LAN.

E-LINE Service Description

An E-LINE carrier Ethernet service offering would be analogous to any of the point-to-point transport options available, including T1/T3, OC-3/12/48/192, as well as channelized offerings.

As the name implies, E-LINE is a point-to-point service transport, limiting the devices connected over that specific service link to two, creating a peer between each. As shown in **Figure 15**, each router has a dedicated Ethernet link between each site, creating an Ethernet wire service.

Figure 15. Router Peering Model over MEF E-LINE Service



From a logical perspective, the impact on the routing design designates each Ethernet wire as a point-to-point link, designating a /30 address, and the routing protocol sees these as point-to-point peers as well. At the central site in Figure 11, prior to support for 802.1Q tags in the clear, each Ethernet wire was connected as a physical Ethernet link. Having 802.1Q tags in the clear completely changes this design (details of this use are below). Some of the keys for E-LINE services include:

- Deterministic peering and bandwidth per virtual Ethernet wire (physical or logical)
- Deterministic QoS, as a Service Level Agreement (SLA) can be applied per Ethernet virtual connection (EVC), quantifying the amount of BW per logical connection
- Deterministic traffic shaping per physical/logical connection, eliminating the overrun of a particular site (based on how the logical connections are sized).

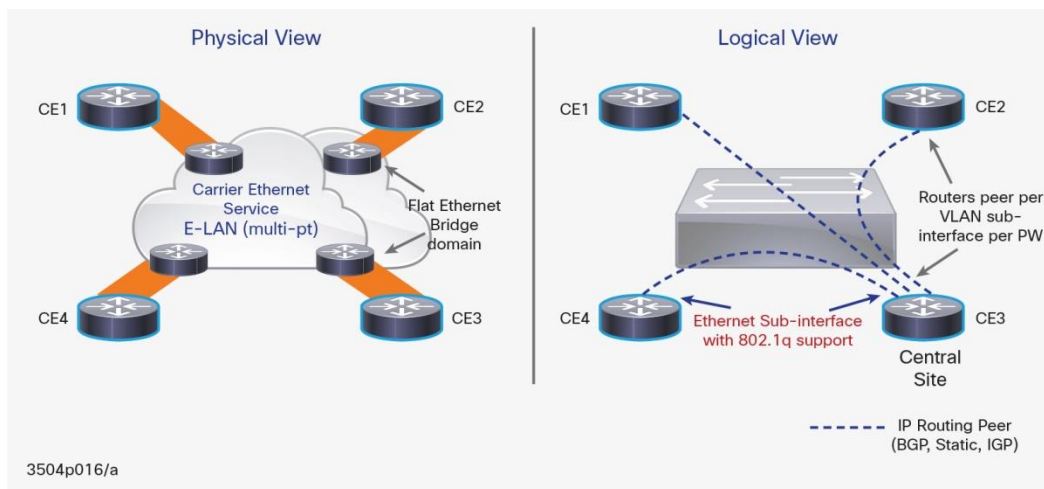
While E-LINE services can be more costly, they do offer more deterministic and prescriptive connections and QoS policies (specifically shaping per-EVC connection).

E-LAN Overview

Unlike E-LINE, which easily maps to point-to-point services, E-LAN offers a point-to-multipoint delivery option that in some ways simplifies configuration, but also poses scaling challenges to a routing network that deserves serious design considerations, as will be covered below.

As can be seen in **Figure 16**, E-LAN emulates more of an “Ethernet Switch” in its forwarding paradigm versus E-LINE, which emulates a point-to-point Ethernet wire. The E-LAN service offers both positives and challenges. From the positive view, IP addressing is simple as each interface connected into the bridge domain (think of this as a flat Ethernet switch or a single VLAN) and each device has “any to any” communications with each CE. The challenges posed with E-LAN services are multi-fold in that while a single bridge domain is simple, it offers enough rope for the designer to hang him/her-self. How so?

Figure 16. Router Peering Model over MEF E-LAN Service



In Figure 16, consider each CE has a single physical attachment into the bridge domain. Each router creates a routing adjacency (consider an Interior Gateway Protocol [IGP] in a broadcast segment) with every other router on the network, so applying N-1 in Figure 16, each router establishes three routing adjacencies. While four routers are minimal, consider a design with 100 routers attached to the same bridge domain, again, applying N-1 routing adjacencies. Each PE now contains 99 IGP routing adjacencies, so routing design best practices should apply when using an E-LAN service or any flat broadcast domain public transport.

Some have compared E-LAN transport to the old ATM days of LANE, which is a fair and accurate analogy. In addition to N-1 challenges, the QoS model in this configuration is indeterministic. In other words, each CE does not have the view of which CE is sending to another CE. In Figure 12, consider CE3 above front-ends the data center in the example and 90 percent of the traffic enters through it from CE1, CE2, and CE4. There is no way for CE1, CE2, and CE4 to know at which rate traffic is being sent, and if the access speed of each CE is 1G, CE3 could be overrun by traffic. To summarize, although E-LAN simplifies the connection model of transport options, and is also typically a less expensive transport, it can pose challenges that need to be recognized in the overall design, including:

- The N-1 maintenance of routing protocol adjacencies
- QoS can have blind spots with a difficulty in rate limiting per CE as the bandwidth is more of a share model or ingress rate limiting function
- The inability to efficiently prune routers from a given multicast tree exists

MACsec, however, is transparent to the details of the transport network. There are other challenges seen with MACsec that will be highlighted and covered in the use cases in this document.


Examples of Use Cases for WAN MACsec

Secure High-Speed Data Center and Cloud Interconnection

The requirements for securing the connection of multiple data center and cloud links typically target a smaller amount of links, while requiring the highest level of bandwidth requirements. Sometimes called data center interconnect (DCI), WAN MACsec is an ideal encryption solution for interconnecting data center or multilocation cloud environments.

As shown in **Figure 17**, WAN MACsec is transparent to large packet sizes and/or maximum transmission units (MTU) in that any variation of small, large, or IMIX traffic patterns do not impact the performance of the encryption process of the flow. Meanwhile, WAN MACsec can offer a broad scale of Ethernet rates, from 1 Gbps and 10 Gbps, to well beyond N-x 100 Gbps links (leveraging link bundle technologies). WAN MACsec in these environments can leverage any form of dark fiber, Metro Ethernet service provider service, or any Ethernet transport service over a WAN that offers an Ethernet service delivery point hand-off. Needless to say, DCI is one of the primary use cases that can leverage WAN MACsec for encryption for typical high-speed bandwidth requirements.

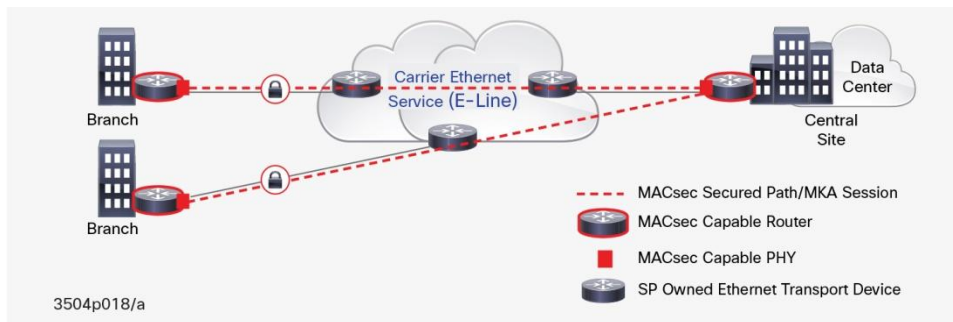
Figure 17. WAN MACsec Point to Point E-LINE Example



Secure High-Speed Branch Router Backhaul

The backhauling of remote branch sites for government, enterprise, or commercial organizations is critical to any business. In the consumer space, this can be important for remote stores and point-of-sale kiosks, and in the enterprise and government space, it is crucial for remote agencies and offices (**Figure 18**).

Figure 18. WAN MACsec Point to Multipoint E-LINE Example



WAN MACsec offers a secure high-speed alternative for remote branch backhaul that is challenged with IPsec performance limitations and its impact from variable packet size flows. While the use cases existed in the past for leveraging MACsec in this hub and spoke networks, the implementation was not available due to the lack of 802.1Q tag in the clear, as each remote site would consume a physical link per branch router, making the option not deployable.

WAN MACsec completely changed this by introducing the 802.1Q tag in the clear capability. With this capability on the central site router (shown in Figure 11), network designers, for example, can now leverage this ability to apply a logical Layer 3 subinterface per remote site that is a subrate of bandwidth from say the 10 GE PHY, offering subrate capabilities from the physical bandwidth. The subinterface can support E-LINE or E-LAN services and can also support hierarchical traffic shaping to align with the prescribed subrate interface.

While IPsec with DMVPN (and Cisco Intelligent WAN [IWAN]) is the dominant solution for remote branch office backhaul, WAN MACsec alternatives target a smaller number of remote branches with high-speed encryption. With new innovations in WAN MACsec with 802.1Q tag in the clear, the flexibility of the hardware solutions is available to leverage an alternative MACsec solution.

The word of caution is necessary when referring to target smaller number of remote branches as relates to the security association (SA) scale of the PHY for MACsec. Each physical Ethernet interface supporting MACsec has an SA key limitation as described by the vendor. For example, in the case of the Cisco ASR 1001-X, each 10GE PHY scales to 64 SAs. Allocating for hitless key rollover, 64 is cut in half, and thus targeting a maximum of 32 branch sites per interface. For the ASR 9000 100 GE interface, the limit is 256 SAs, so the maximum SA is hardware specific and is an important design element for WAN MACsec.

Note: A hub site router can leverage multiple 10GE interfaces, so the limit is not per router, rather per interface.

Secure IP/MPLS and Metro Ethernet Backbone Networks

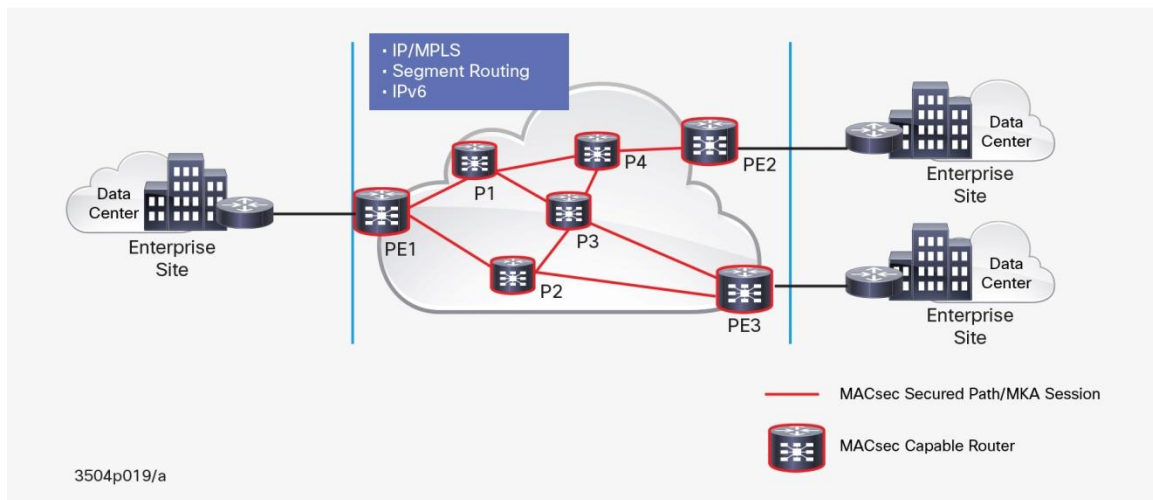
Given the fact that MPLS offers no native encryption solutions, the predominant encrypted MPLS solutions typically leverage the concept of running MPLS over generic route encapsulation (GRE) tunnels or User Datagram Protocol (UDP) as explained in RFC 4023 (MPLS over IP) in combination with IPsec to encrypt the IP tunnel encapsulation. While there are several highly scalable methods for simplifying MPLS over GRE with IPsec, these offerings can limit several key capabilities that MPLS offers, specifically using MPLS traffic engineering, as most MPLS over GRE solutions are overlays and suffer from performance limitations due to the requirement for using IPsec.

WAN MACsec eliminates all of these challenges found in MPLS over GRE + IPsec as the MPLS labels are transparent to MACsec in the Ethernet transport, offering native line-rate encryption of the Ethernet frame + MPLS label between P and PE router in the core.

Referring to **Figure 19**, MACsec is a per-hop encryption function between each of the PE and P routers, but so are the MPLS label push/pop functions. For example, PE 1, P1, P4, and PE 2 are supporting MPLS forwarding and MACsec on each link or interface. As an MPLS packet traverses the label switch path (LSP) from PE 1 to PE 2 through P1 and P4, the functions at each node includes:

- Ingress into the Interface: De-encrypt MACsec header from Ethernet frame, pop the MPLS label, perform MAC-layer rewrite.
- Egress out of the Interface: Push MPLS label onto new Ethernet frame, encrypt Ethernet frame via MACsec function.

Figure 19. Example: Securing MPLS Backbone with MACsec



The per-hop encryption of MACsec offers high-speed per link encryption, while offering complete transparency to the functions of an IP/MPLS architecture (MPLS services, traffic engineering, relevant protocols, etc.). The per-hop nature of MACsec allows encryption at 100GE+ transport speeds, while preserving necessary per-hop functions of MPLS, such as label push/pop, traffic engineering, MPLS operations, administration, and maintenance (OAM) functions, as well as net flow and any future analytics enhancements.

Note: MACsec is being leveraged as the encryption recommendation for newly offered segment routing⁷ capabilities, for all of the reasons listed previously, specifically offering 100-Gbps encryption while remaining transparent to the segment routing control and data plane functions and service requirements needed per hop.

Secure PE-CE Links for Managed Private IP VPN Transport

Although this use case option requires coordination with the service provider, adding WAN MACsec to elements of service provider transport offerings can enhance a multitude of capabilities and enhancements to the services. Consider a service provider that offers an IP MPLS VPN transport. A typical connection model to the IP VPN service includes a backhaul link from the Customer Edge (CE) device to the Provider Edge (PE) device. In most cases, these links are unsecure and rely on the customer CE to encryption end to end.

In the case of an service provider (SP)-managed service, SPs could leverage WAN MACsec to offer customers a secure encrypted PE to CE backhaul link to the provider cloud (e.g. PE device). The SP could extend this security service end to end if the SP expanded MACsec into their IP/MPLS MPLS backbone, as shown above in Figure 19. In a managed CPE offering, assuming the service met agency/customer security mandates, this secure transport offering (driven by MACsec) would eliminate the need for the CE to deploy IPsec overlay solutions (DMVPN or GET VPN). This type of deployment would greatly reduce the complexity for the end customers, not having to deploy a secure IP VPN overlay, while reducing operating expenses (OpEx) on the SP side and expanding the SP service catalog to their end customers.

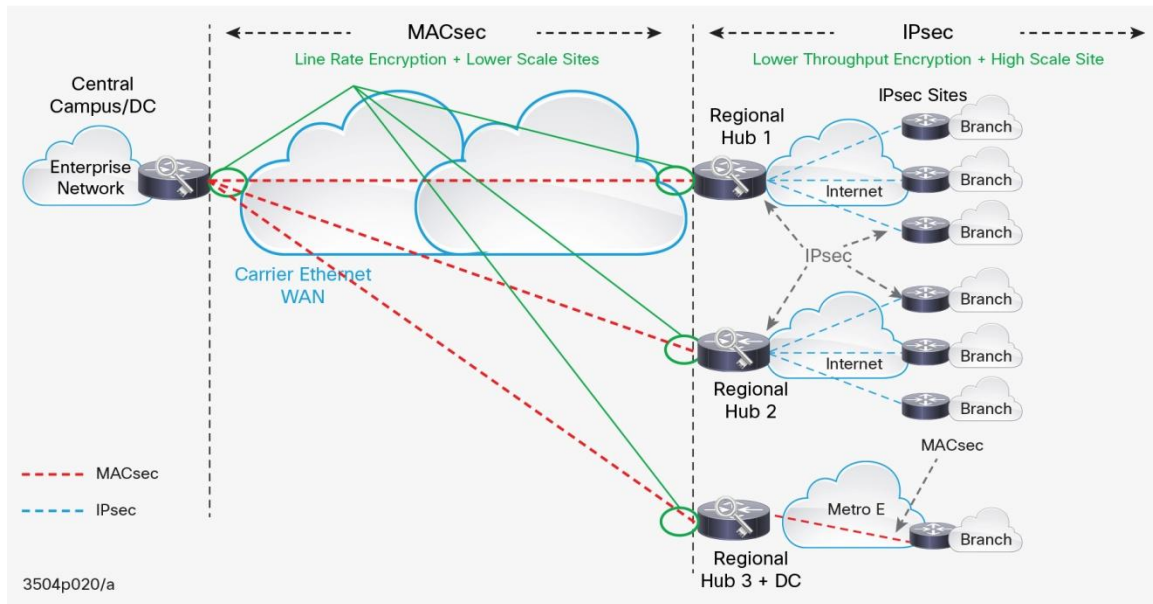
Hybrid Design Using WAN MACsec with IPsec

In large public sector and enterprise WAN designs, it is common to see deployments that require multiple tiers, and in some cases, leverage a mix of public and private transport. For example, the aggregation sites may be connected via IP/MPLS within the agency's private MPLS network, while the back-haul of smaller, more remote branch locations, leverage inexpensive IP transport services. This hybrid multitier design approach is ideal for mixing encryption technologies as well, using IPsec for smaller, lower-speed connections, and MACsec from the data center and in the core or aggregation IP/MPLS backbone, where the link speeds are much higher.

Consider the hybrid design example in **Figure 20**. In a typical 2-tier design, the option would be to leverage WAN MACsec in the IP/MPLS backbone (regional hubs and DC edge routers) where links speeds could target 10-100 Gbps. The branch locations, typically requiring lower speed links but higher volume of locations, can leverage IPsec with DMVPN or Cisco IWAN, to take advantage of the higher scale site termination IPsec and DMVPN offers.

⁷ IETF - Segment Routing Architecture: <https://tools.ietf.org/html/draft-ietf-spring-segment-routing-09>

Figure 20. Typical Two-Tier WAN Design



This hybrid encryption design approach leverages the strengths of each encryption technology, with IPsec targeting higher scale SAs with lower encryption throughput, and MACsec optimizing the solution through extremely high-speed, lower-scale SAs and transparency for MPLS labels, Segment Routing, without the need for MPLS over GRE tunnels.

Comparing MACsec to IPsec

While this document highlights the clear advantages WAN MACsec offers network designers needing higher speed encryption solutions, it is important to not position encryption solutions against one another, specifically MACsec and IPsec. Rather, understand the strengths and limitations in each, and attempt to properly position as the requirements in the design dictate the use of each technology, specifically in relation to:

- Transport availability and feature offerings
- Performance requirements of the solutions and/or application traversing the WAN
- Scale of the design and requirements (number of spokes, connected endpoints, aggregate encryption peers, etc.)

A very important guideline to remember when evaluating the various options beyond IPsec (because IPsec can run over any offered transport, optimal or not) is that the underlying transport dictates the available encryption options that can be leveraged.

Figure 21 provides a basic set of guidelines to use when evaluating MACsec and IPsec encryption options. Start with the orange boxes, which completely eliminate that aspect from the technology, and then continue to evaluate. Blue assumes support, with limitations.

Figure 21. Ethernet and IP Encryption Positioning Matrix

Design Component	MACsec	IPsec
Topology - Point to Point	E	E
Topology - Multipoint Capable (P2MP, MP2MP)	E	E
Transport Service Support		
Ethernet (P2P, Point to Multipoint)	E	E
IP (MPLS VPN, broadband, Internet)	NS	E
Optical/Lambda/Dark fiber	E	S : speed limited
Logical Link Segmentation (802.1Q/sub-int capable)	E	E
Leverage legacy transport (T1/E1/T3/E3, SONET/SDH)	NS	E
Encryption Performance		
Encryption Line rate per the PHY interface (1/10/40/100G)	E	NS
Encryption process NOT dependent on physical interface	NS	E
Encryption not <u>limited</u> by packet size, MTU, PPS of engine	E	S : Impacts perf
Scale		
Hub Site Scale (Hub/Spoke Topology)	S : PHY SA Scale	E : (1000+ sites)
Simplicity of Configuration		
Transparent to IPv4/v6, MPLS, IGP/BGP, IP Multicast	E	S : needs GRE

■ E = Excellent
 ■ S = Supported (with Limitations)
 ■ NS = Not Supported

To summarize some of the findings from the matrix in Figure 21:

- MACsec supports line-rate encryption performance (100 Gbps+), regardless of the MTU and packet size
- MACsec is transparent to upper layer protocols (IPv4/v6, MPLS labels)
- IPsec is extremely flexible from an underlying transport perspective (completely agnostic)
- IPsec supports massive scale (DMVPN moving beyond 4000 connections) from an SA termination perspective
- MACsec support will be dictated by the hardware's Ethernet PHY capabilities
- In some cases, either solution will work and experience and desire from a designers perspective will dictate the choice

As is evident from the matrix and highlights, there are multiple decision factors that need to be accounted for that can relate to business requirements, transport available in a country or region, governing certifications in an organization, as well as other criteria (operations comfort, expertise, future direction, etc.). As an evaluator and designer, it is important to carefully understand short and long-term requirements and to understand the applications requirements as this is the most important element as it relates to end-user experience.

Note: When evaluating the various encryption solutions that exist in network deployments today, it is important to understand that the evaluation should not be that one technology is superior to the other. The key focus in the evaluation process should be which technology best meets the business objectives of the end user and application experience, and/or simplifies operations, along with offering the most cost-effective solution. These technologies are never a one size fits all, so it is up to the network designer to understand the holistic view of the network so the proposed solution aligns with the business objectives and services the network aims to offer. This is a key component when evaluating encryption technologies or any network solution and transport overall.

Summary

As the demand for increased bandwidth is driven by the expansion of cloud services, mobile devices, and massive increases in video traffic, the requirement for encryption rates to align at speeds beyond 100 Gbps at any packet size, is vital. While IPsec continues to be the predominant network encryption solution, MACsec is the choice for next-generation high-speed encryption in federal and enterprise, cloud, and service provider transport networks. In addition, WAN MACsec offers a high-speed encryption solution integrated with the MAC PHY (versus an external encryption device), as well as enhancements needed for operators to deploy seamlessly over any public carrier Ethernet, optical, or DWDM/OTN transport network, while leveraging open standard key agreement solutions.

As network architects and designers evaluate encryption solutions moving forward, WAN MACsec offers encryption rates never before seen at speed 40 and 100 gigabits and beyond. While WAN MACsec is that de facto high-speed solution moving forward, it should not be thought of as a replacement for IPsec, but rather another set of tools in the encryption tool bag moving forward, and in some cases, deployed in combination with IPsec in larger scale deployments.

References

- Cisco Application Engineered Routing (for example, segment routing): <http://www.cisco.com/c/en/us/solutions/service-provider/application-engineered-routing/index.html>
- Cisco Visual Networking Index: <http://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf>
- MPLS VPN over Dynamic Multipoint VPN (DMVPN): <http://www.cisco.com/c/en/us/products/security/dynamic-multipoint-vpn-dmvpn/index.html>
- MPLS VPN over Multipoint Generic Route Encapsulation (mGRE): http://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/network-virtualization-solutions/white_paper_c11-726689.pdf
- NSA Suite B Cryptography: http://www.nsa.gov/ia/programs/suiteb_cryptography/




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)