SDWAN Edge Discovery

https://datatracker.ietf.org/doc/draft-dunbar-idr-sdwan-edge-discovery/

Linda Dunbar
Sue Hares
Robert Raszuk
Kausik Majumdar
Nov 2020

Intention of the draft

- Purpose:
  - Tag the client routes with Extended Communities for tunnel-encaps + Color
  - Using BGP UPDATE to advertise properties of hybrid underlays of SDWAN
    - Properties: Security Associations, port in SDWAN
- Based on:
  - *draft-ietf-idr-tunnel-encaps-20*

Basic Schemes

- **UPDATE U1: Client routes UPDATE:**
  - NLRI: Prefix: 10.1.1.1
  - Nexthop set to the address of the Edge node, e.g. C-CPE2 (for recursive lookup)
  - Encapsulation Extended Community
    (Tunnel-Type=SDWAN-Hybrid or other existing tunnel types);
  - Color Extended Community;

- **UPDATE U2: Detailed tunnel attributes of the underlay networks:**
  - Tunnel-egress-endpoint Sub-TLV
  - Extended Port Sub-TLV (Includes the ISP property Sub-TLV (Optional))
  - Geo Location Sub-TLV
  - IPSec Sub-TLVs or
  - IPsec SA IDs (for preconfigured IPSec SSs)

DETAILS – AT HIGH LEVEL

UPDATE #1: Client Routes IP SAFI =1 NLRI

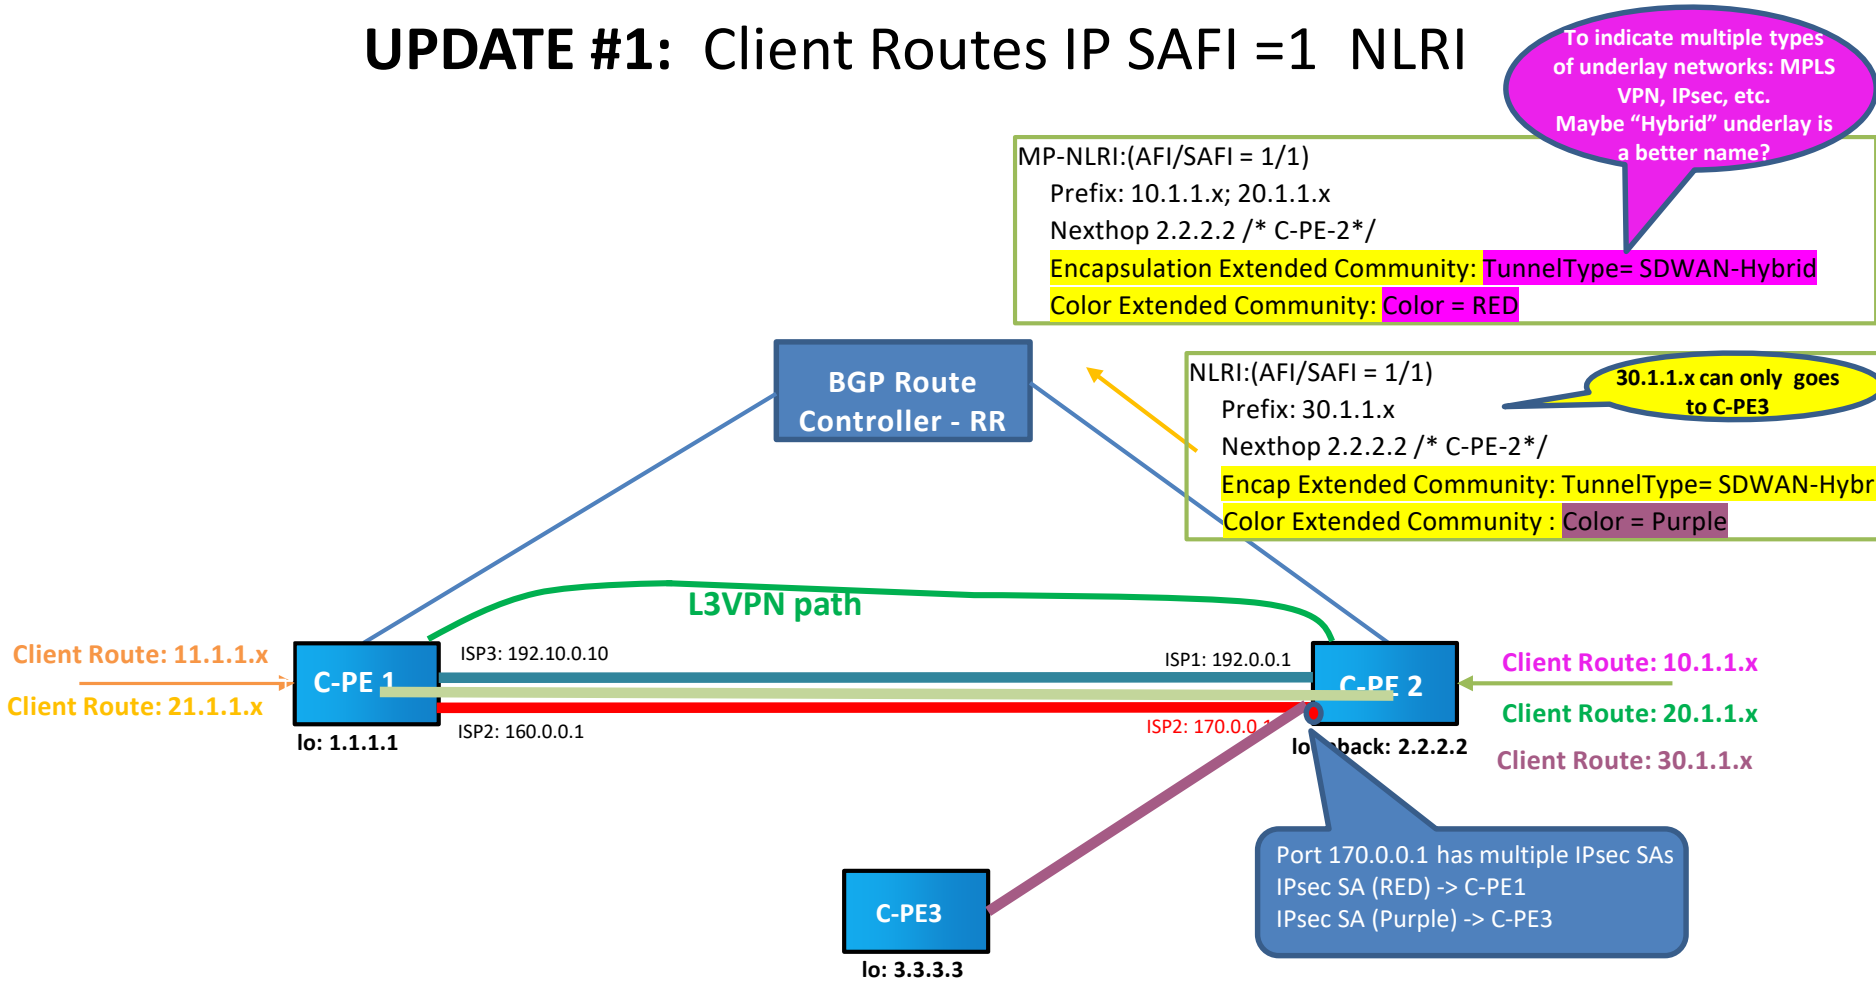

BGP UPDATE #2: Detailed Tunnel attributes: “Color” encoded in SDWAN NLRI

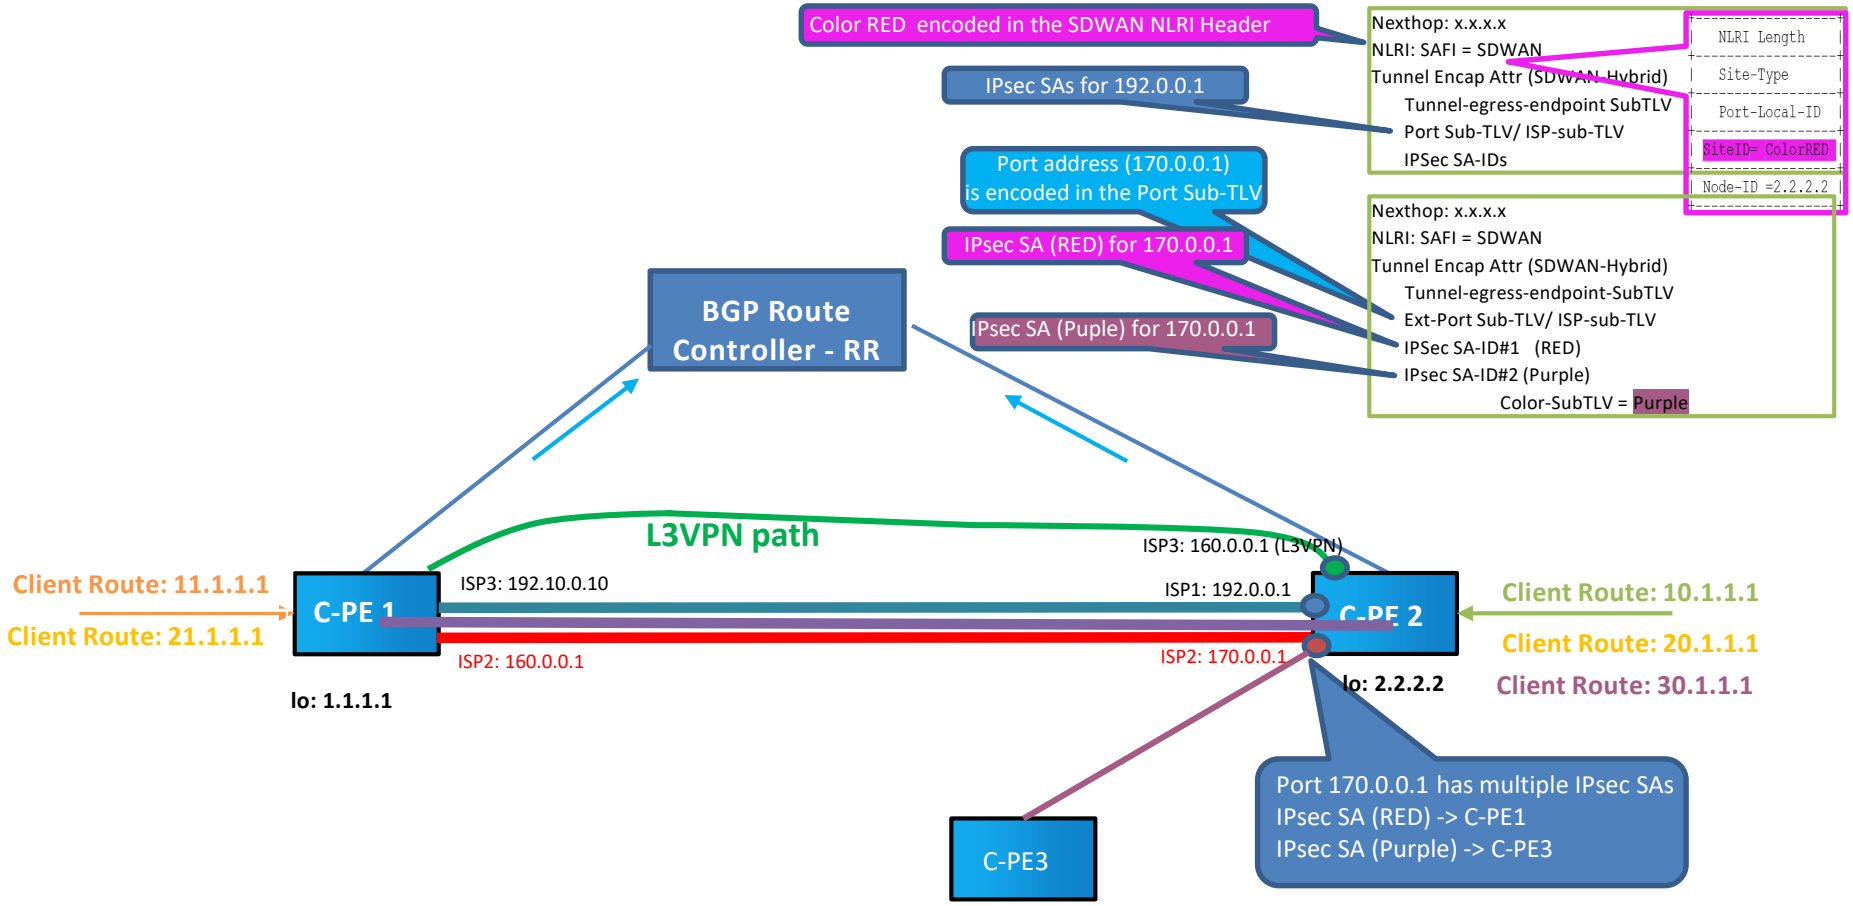

When IPsec SAs are pre-configured, IPsec SA-IDs can be included in the SDWAN-Hybrid Tunnel Path Attribute

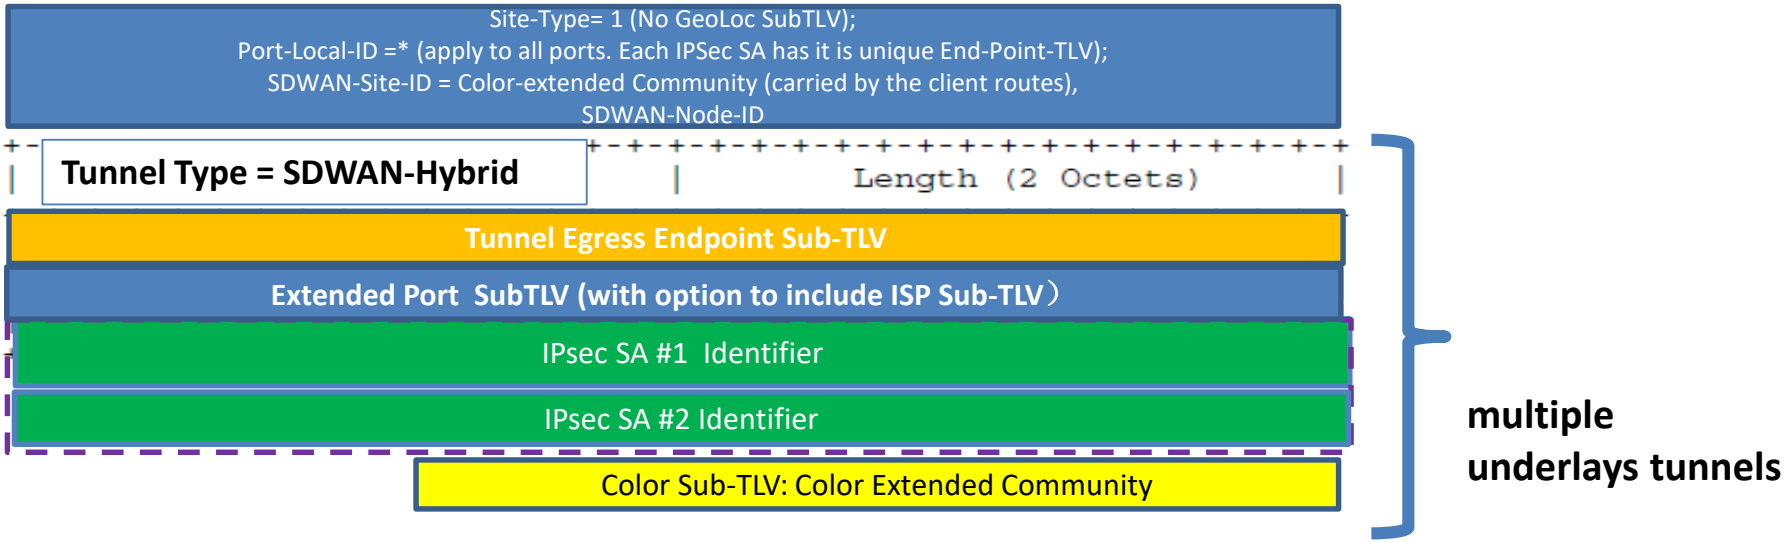

NLRI SAFI = SDWAN: Tunnel Encap Attribute: with option to include IPsec SA sub-TLVs

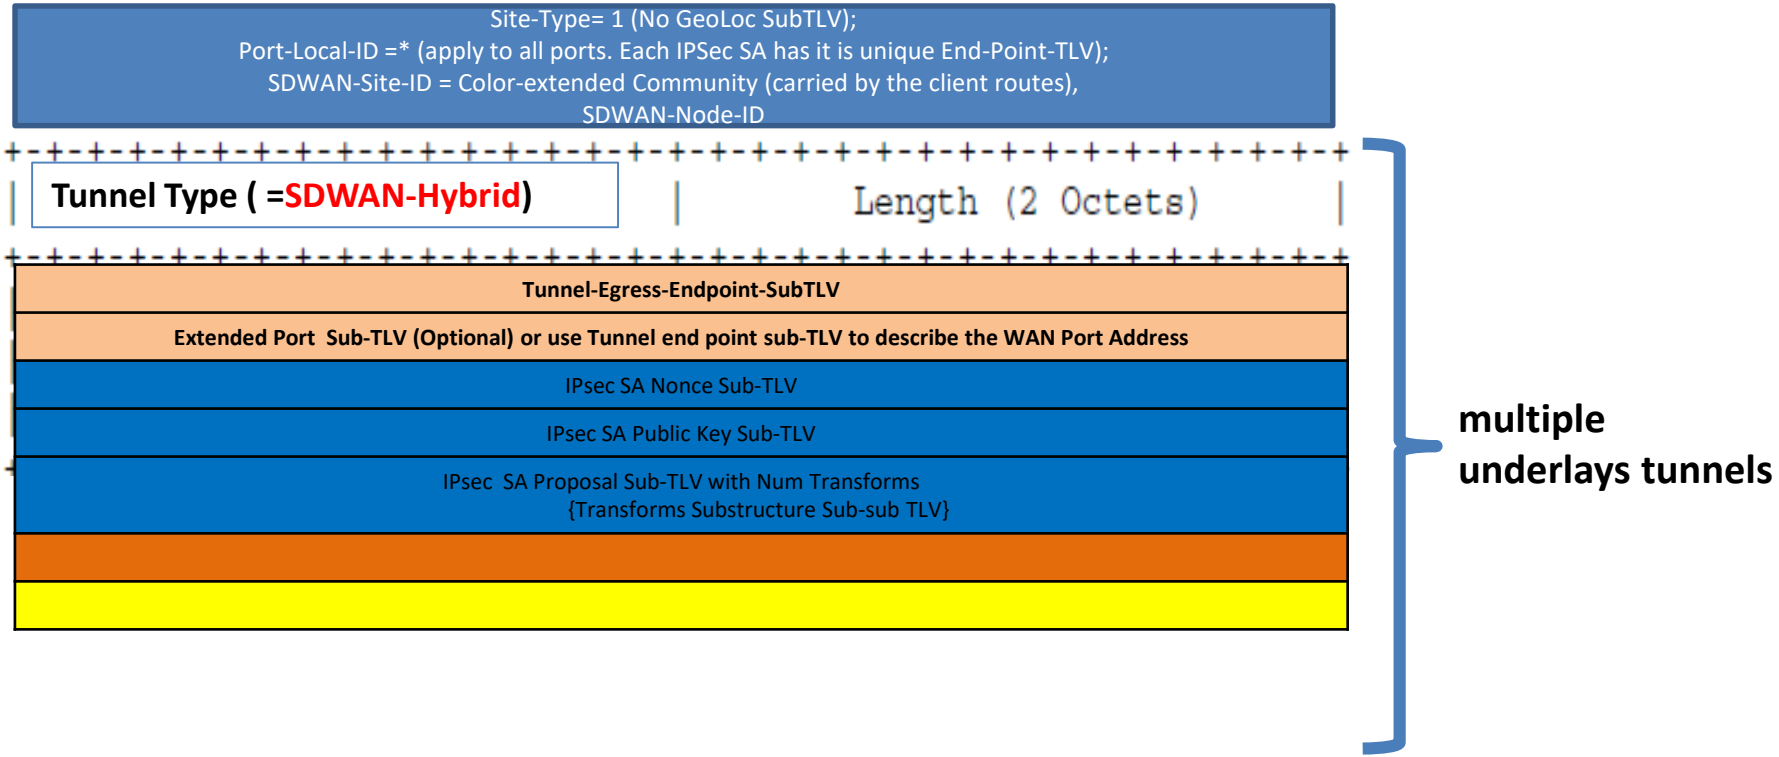

SDWAN SAFI (=74) NLRI Encoding Format for Underlay Network Properties Advertisement

MP-REACH-NLRI Path Attribute
AFI = 1; SAFI = SDWAN
<Site-Type, Port-Local-ID, SDWAN-Color, SDWAN-Node-ID>
Geo-location Sub-TLV (Optional)
Tunnel Encaps Attribute (23)
Tunnel Type: SDWAN-Hybrid (to indicate hybrid underlay tunnels)
Tunnel-egress-endpoint Sub-TLV
Extended Port Sub-TLV (with or without NAT, Optional, including the ISP-subTLVs)
VxLAN, GRE, etc. for Secure VPN underlays

IPsec SA Nonce Sub-TLV
IPsec SA Public Key Sub-TLV
IPsec SA Transform Sub-TLV
{Transforms Substructure Sub-sub TLV}

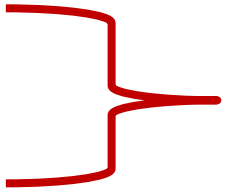

List of IPsec-SA Attributes.
Or
IPSec-SA IDs

NLRI: SDWAN SAFI = 74

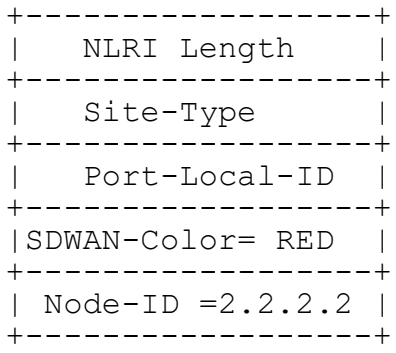

- Site Type: 1 octet value. The SDWAN Site Type defines the different types of Site IDs to be used in the deployment. The draft defines the following types:

  Site-Type = 1: For simple deployment, such as all edge nodes under one SDWAN management system, a simple identifier is enough for the SDWAN management to map the site to its precise geolocation.

  Site-Type = 2: to indicate that the value in the site-ID is locally significant, therefore, need a Geo-Loc Sub-TLV to fully describe the accurate location of the node. This is for large SDWAN heterogeneous deployment where Site IDs has to be described by proper Geo-location of the Edge Nodes [LISP-GEOLoc].

- Port local ID: SDWAN edge node Port identifier, which can be locally significant. The detailed properties about the network connected to the port are further encoded in the Tunnel Path Attribute. If the SDWAN NLRI applies to multiple ports, this field is NULL.

- SDWAN-Color: used to identify a common property shared by a set of SDWAN edge nodes or a group of WAN ports, such as the Color encoded in the Client Routes' "Color Extended Community" or property of a specific geographic location shared by a group of prefixes.


We want
your feedback!

DETAILS – AT FIELDS IN SUBTLV

Tunnel Path Attributes and Sub-TLVs inside the SDWAN NLRI

Extended Port (NAT) Sub-TLV

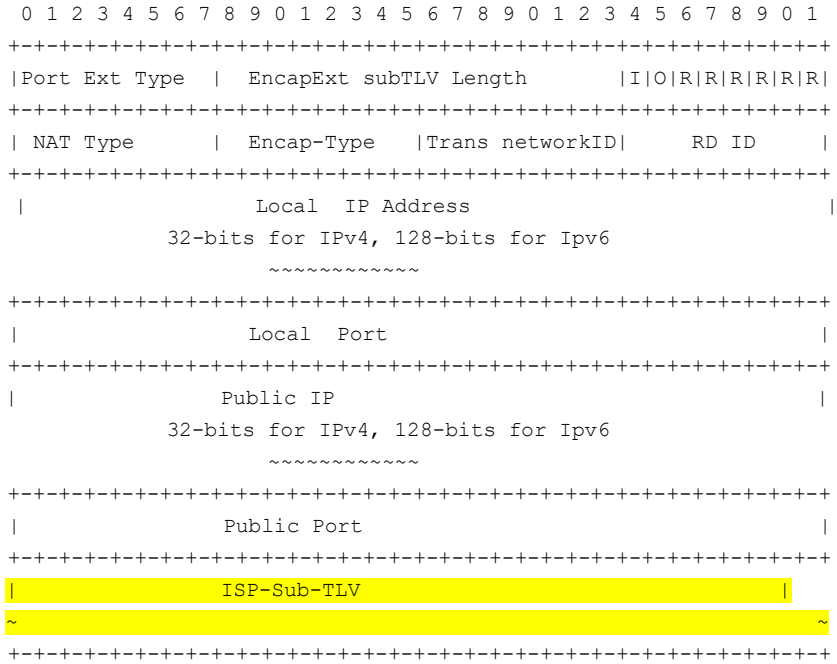

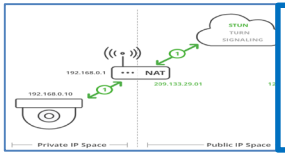

Edge node get NAT properties via STUN requests/responds.
Peers may not be able to access the STUN server.

- Flags:

-I bit (CPE port address or Inner address scheme)
  If set to 0, indicate the inner (private) address is IPv4.
  If set to 1, it indicates the inner address is IPv6.
- O bit (Outer address scheme):
  If set to 0, indicate the public (outer) address is IPv4.
  If set to 1, it indicates the public (outer) address is IPv6.
- R bits: reserved for future use. Must be set to 0 now.

- NAT Type：
  without NAT; 1:1 static NAT; Full Cone; Restricted Cone; Port Restricted Cone; Symmetric; or Unknown (i.e. no response from the STUN server).
- Encap Type：
  the supported encapsulation types for the port facing public network, such as IPsec+GRE, IPsec+VxLAN, IPsec without GRE, GRE (when packets don't need encryption)
- Transport Network ID：
  Central Controller assign a global unique ID to each transport network；
  RD ID：Routing Domain ID， Need to be global unique.
- Local IP：The local (or private) IP address of the port; If NAT is not used, this field is set to NULL.
- Local Port：used by Remote SDWAN edge node for establishing IPsec to this specific port. If NAT is not used, this field is set to NULL.
- Public IP：The IP address after the NAT.
- Public Port：The Port after the NAT.

ISP of the Underlay Network Sub-TLV

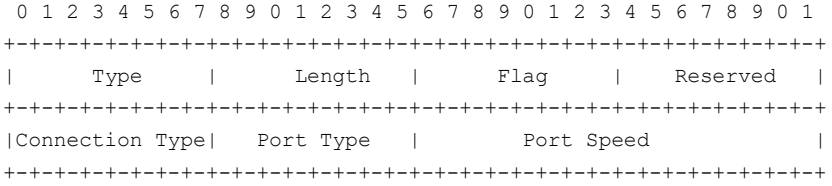

- Type: To be assigned by IANA
- Length: 6 bytes.
- Flag: a 1 octet value.
- Reserved: 1 octet of reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.
- Connection Type: There are two different types of WAN Connectivity. They are listed below as:
  - Wired – 1
  - Wireless – 2
  - LTE – 3
  - 5G – 4
- Port Type: There are different types of ports. They are listed below as:
  - Ethernet – 1
  - Fiber Cable – 2
  - Coax Cable – 3
  - Cellular – 4
- Port Speed: The port seed is defined as 2 octet value. The values are defined as Gigabit speed.

Encoding
for
IPsec Properties

Two Types of IPsec SA attributes (only use one) Sub-Sub-TLV

- Full Set: with multiple Sub-TLVs for full suite of IPsec SA attributes
  - Nonce Sub-TLV
  - Public Key Sub-TLV
  - Proposal Sub-TLV: to indicate the number of Transform subTLVs to be included
    - Transforms Substructure Sub-TLV
- Simple Set: Simple Deployment with limited number of parameters
  - One Sub-TLV to represent Public Key, Nonce, ReKey, Transform

Nonce Sub-TLV, Public Key Sub-TLV

- Nonce Sub-TLV:

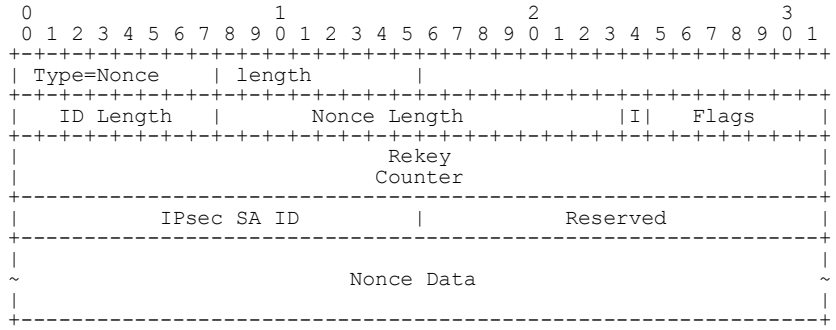

IPsec SA ID - The 2 bytes IPSec SA ID could 0 or non-zero values. It is cross referenced by client route's IPSec Tunnel Encap IPSec-SA-ID or IPSec-SA-Group Sub-TLV in Section 5 of the Draft. When there are multiple IPsec SAs terminated at one address, such as WAN port address or the node address, they are differentiated by the different IPsec SA IDs.

- Public Sub-TLV:

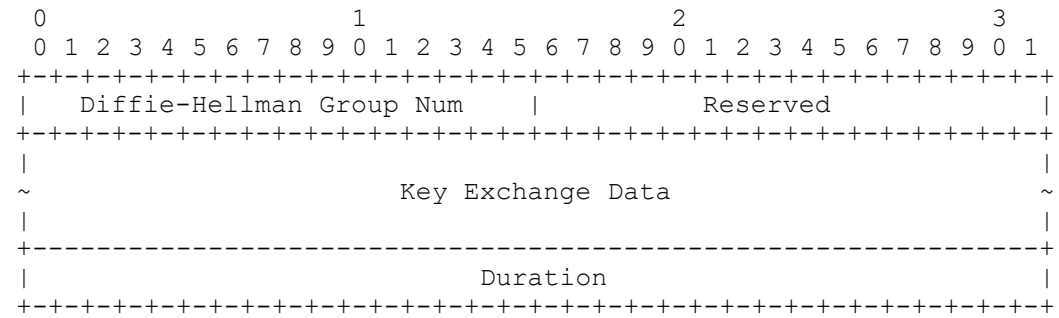

Simplified IPsec SA attributes advertisement

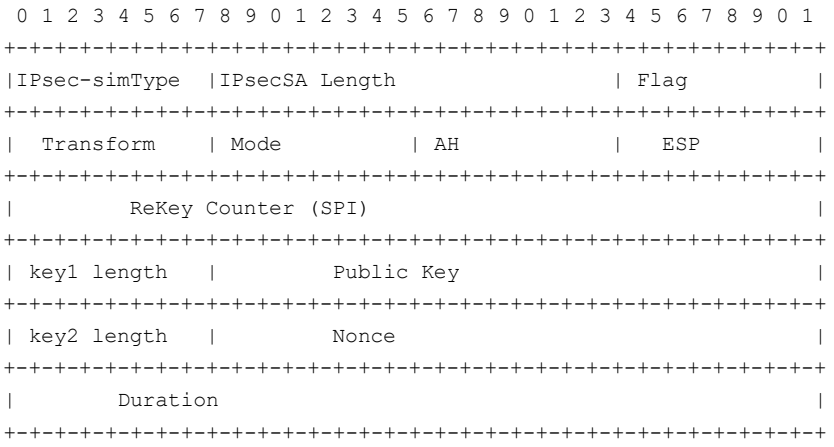

- IPsec-simType: to be assigned by IANA.
- Flags: for future usage.
- Transform (1 Byte): the value can be AH, ESP, or AH+ESP.
- Mode (1 byte): Indicate Tunnel Mode or Transport mode
- AH (1 byte): AH authentication algorithms supported, which can be md5 | sha1 | sha2-256 | sha2-384 | sha2-512 | sm3.
- ESP (1 byte): ESP authentication algorithms supported, which can be md5 | sha1 | sha2-256 | sha2-384 | sha2-512 | sm3.
  - Each SDWAN edge node can have multiple authentication algorithms; send to its peers to negotiate the strongest one. Default algorithm is AES-256.
  - When node supports multiple authentication algorithms, the “Transform Sub-TLV” described by [SECURE-EVPN] can be used to describe the additional algorithms supported by the node.
- Rekey Counter (Security Parameter Index)
- Public Key: IPsec public key
- Nonce: IPsec Nonce
- Duration: SA life span.