SDWAN Edge Discovery

draft-dunbar-idr-sdwan-edge-discovery-03

Linda Dunbar
Sue Hares
Robert Raszuk
Kausik Majumdar
August 2021

SDWAN DRaft

- Use case: BESS/RTGWG drafts
- New Tunnel type (SDWAN-Hybrid)
  - Client routes: link to hybrid tunnel with IP-Sec (U1)
  - Tunnel next hop (SAFI-DSWAN Hybrid) – pass information regarding tunnel (IP-SEC SA, end point, Encaps)
- Purpose of Discussion: Answer questions

SDWAN topology (from IETF 109-110)

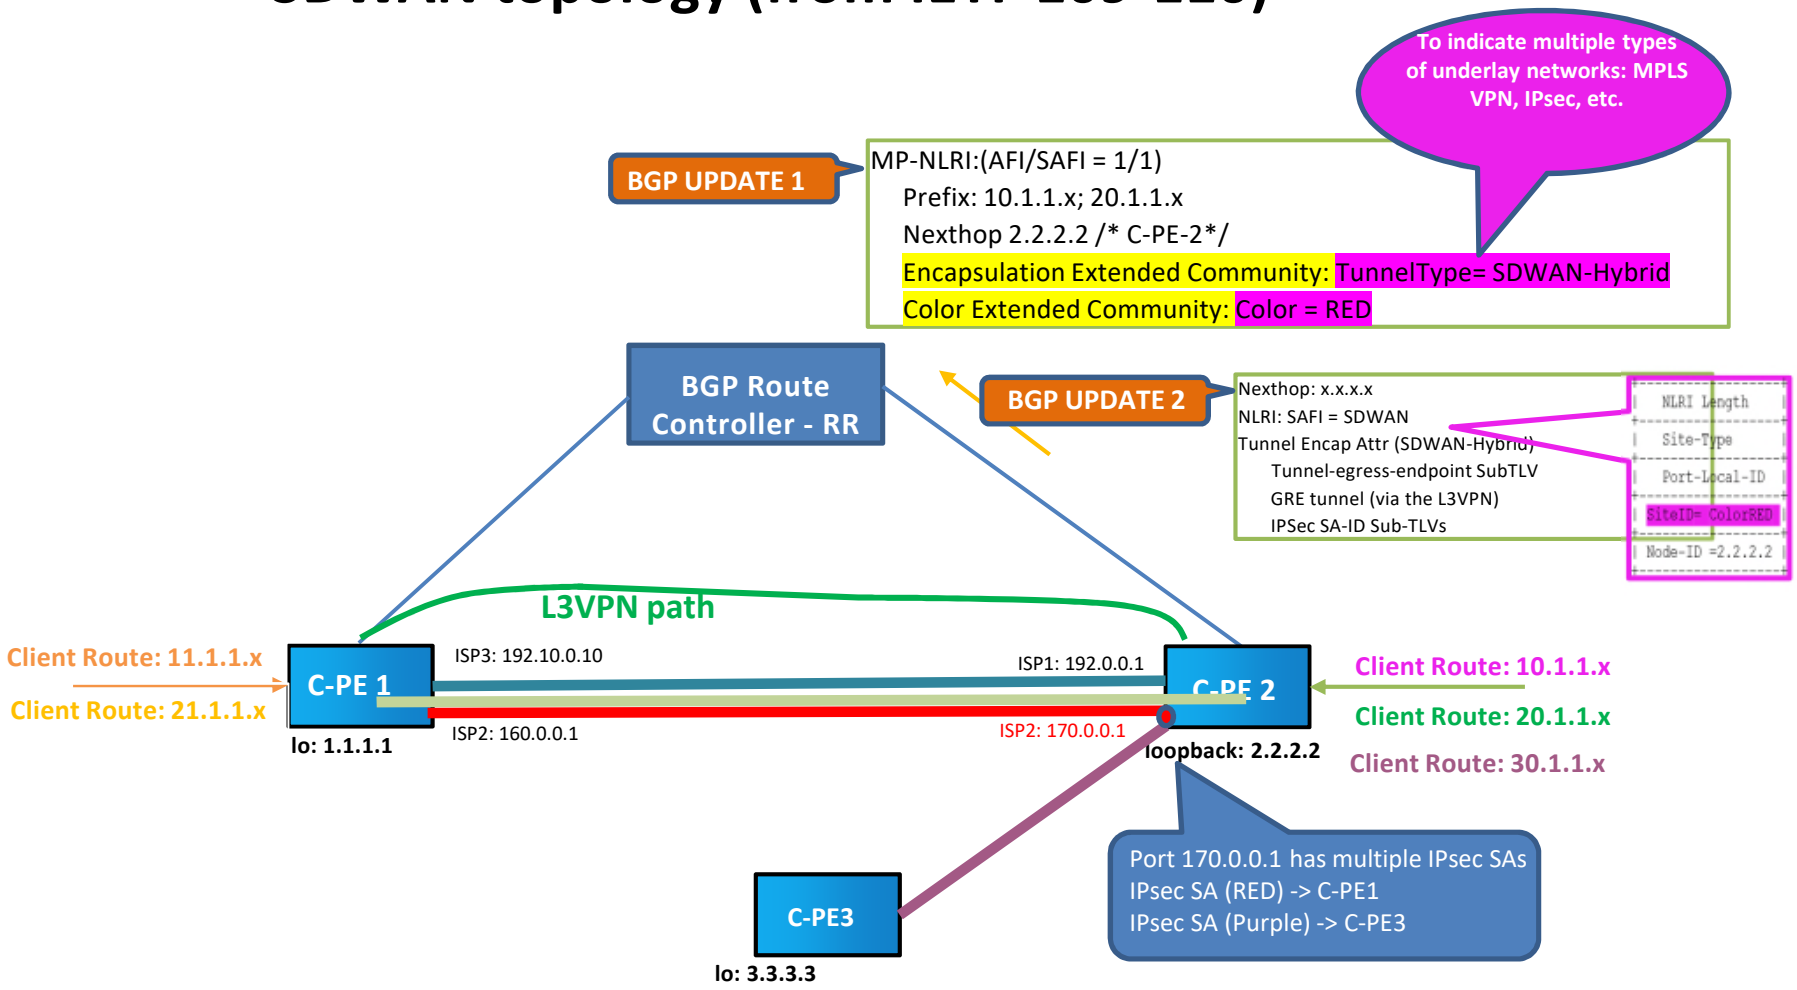

Hybrid Tunnels: with Pre-configured IPsec SA IDs

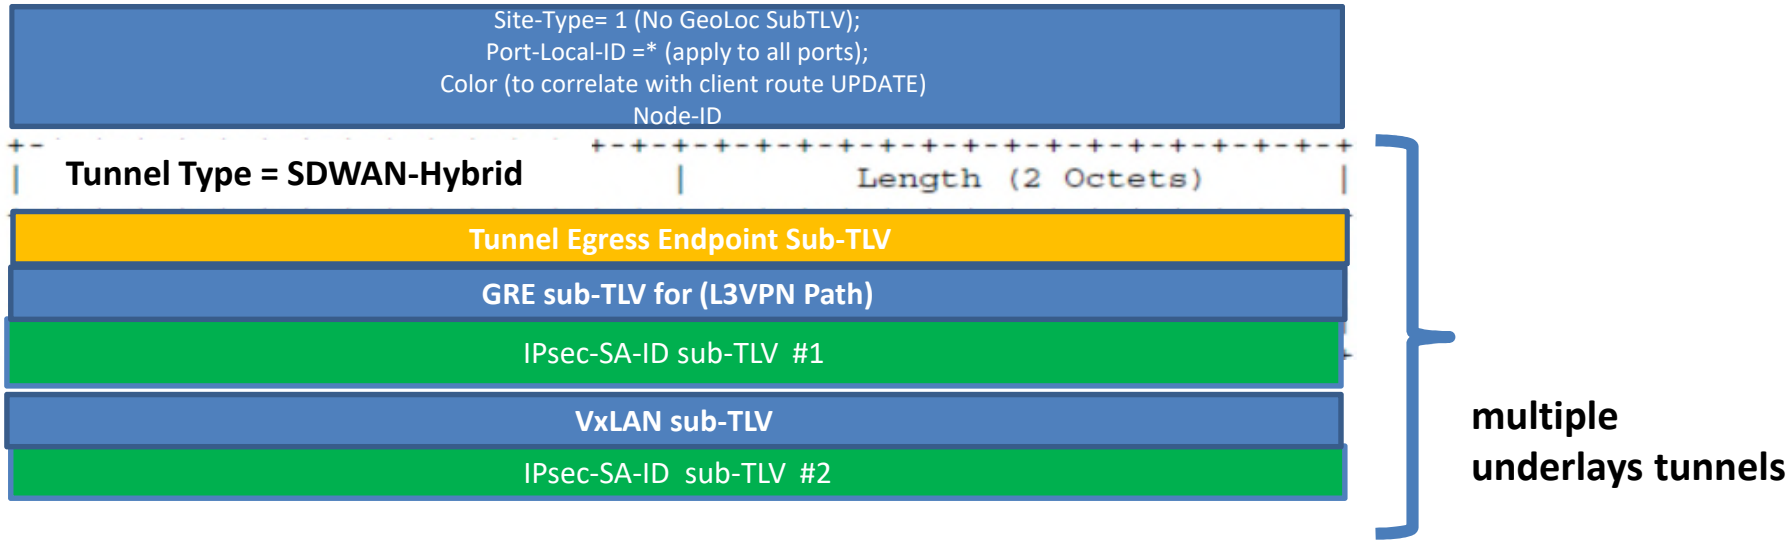


```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Type= IPsec-SA-ID subTLV      |  Length (2 Octets)            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      IPsec SA Identifier #1                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      IPsec SA Identifier #2                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Hybrid Tunnels: with detailed IPsec SA sub-TLVs

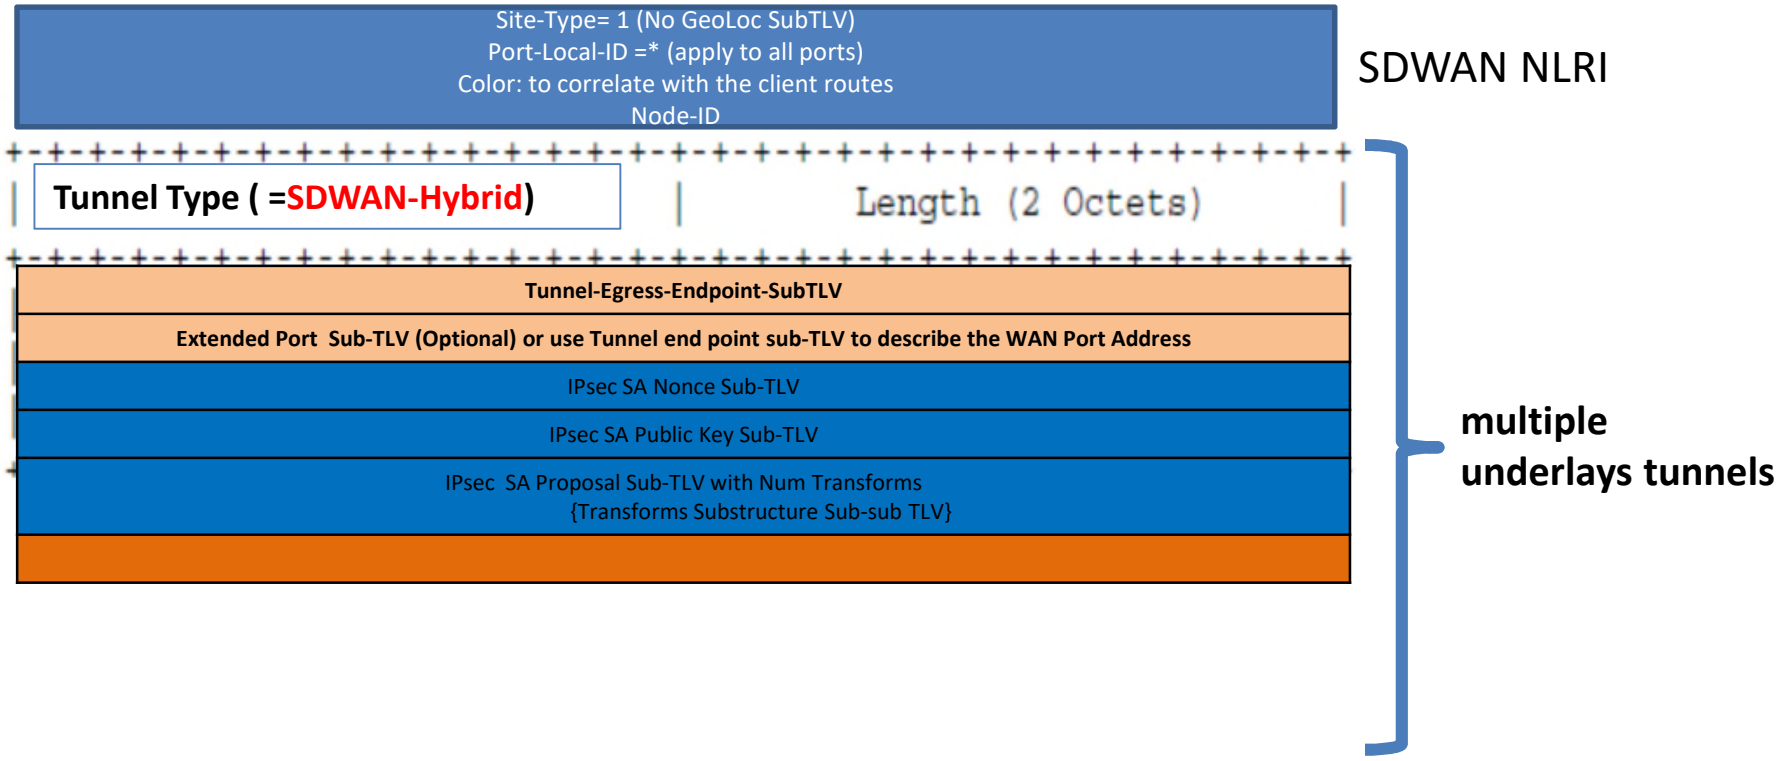

NLRI: SDWAN-Hybrid SAFI = 74

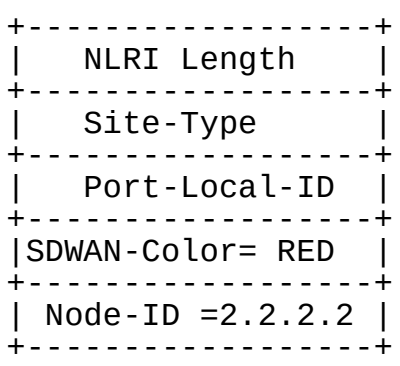

- Site Type:

  Site-Type = 1: For simple deployment, with node ID to identify the precise geolocation.

  Site-Type = 2: For large SDWAN heterogeneous deployment where a Geo-Loc Sub-TLV [LISP-GEOLoc]is used to identify the precise geolocation

- Port local ID: SDWAN edge node Port identifier, which is locally significant. If the SDWAN NLRI applies to multiple ports, this field is NULL.
- SDWAN-Color: to correlate with the Color-Extended-community included in the client routes UPDATE.
- Node-ID: The node’s IPv4 or IPv6 address

Backup slides
Tunnel Path Attributes and Sub-TLVs inside the SDWAN NLRI

Extended Port (NAT) Sub-TLV

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Port Ext Type  |  EncapExt subTLV Length       |I|O|R|R|R|R|R|R|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| NAT Type      |  Encap-Type   |Trans networkID|     RD ID     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                  Local  IP Address                            |
             32-bits for IPv4, 128-bits for Ipv6
                    ~~~~~~~~~~~~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Local  Port                                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Public IP                                      |
             32-bits for IPv4, 128-bits for Ipv6
                    ~~~~~~~~~~~~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Public Port                                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                ISP-Sub-TLV                                    |
~                                                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

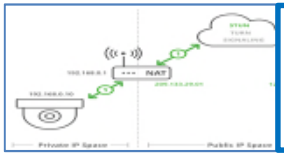

Edge node get NAT properties via STUN requests/responds.
Peers may not be able to access the STUN server.

- 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

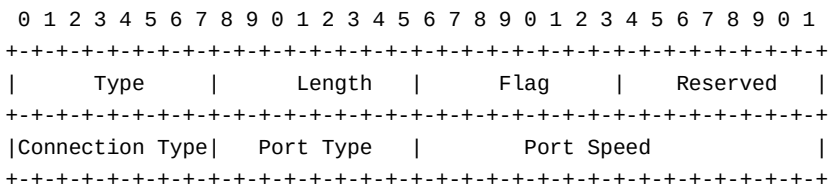

- Type: To be assigned by IANA
- Length: 6 bytes.
- Flag: a 1 octet value.
- Reserved: 1 octet of reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.
- Connection Type: There are two different types of WAN Connectivity. They are listed below as:
  - Wired – 1
  - Wireless – 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.

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:

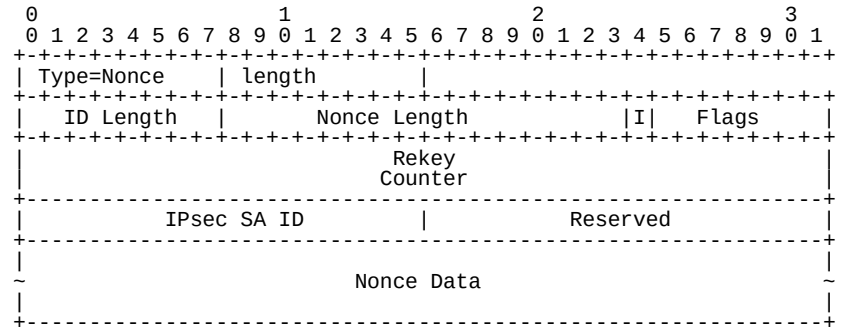

IPsec SA ID - The 2 bytes IPSec SA ID could 0 or non-zero values. It is cross referenced by client route's IPSec Tunnel Encap IPSec-SA-ID or IPSec-SA-Group Sub-TLV in Section 5 of the Draft. When there are multiple IPsec SAs terminated at one address, such as WAN port address or the node address, they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

- Public Sub-TLV:

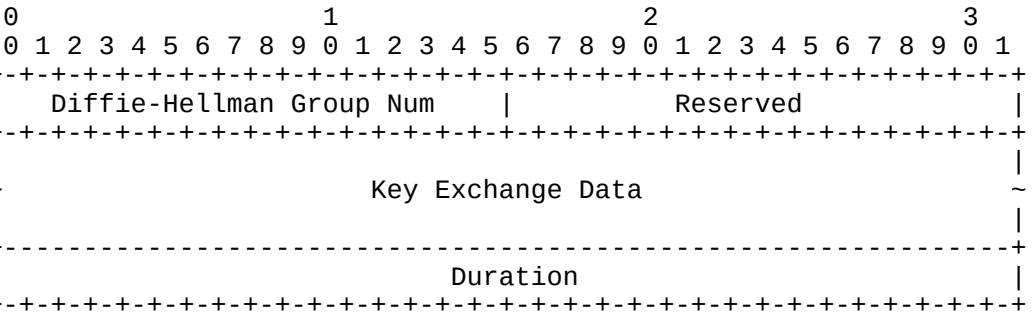

Simplified IPsec SA attributes advertisement

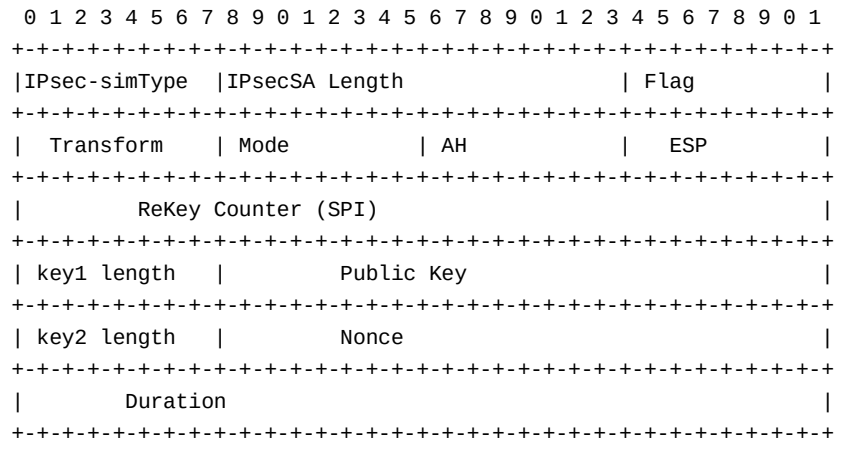

- IPsec-simType: to be assigned by IANA.
- Flags: for future usage.
- Transform (1 Byte): the value can be AH, ESP, or AH+ESP.
- Mode (1 byte): Indicate Tunnel Mode or Transport mode
- AH (1 byte): AH authentication algorithms supported, which can be md5 | sha1 | sha2-256 | sha2-384 | sha2-512 | sm3.
- 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.