draft-hujun-idr-bgp-ipsec-00

Hu Jun, Nokia

IETF 105

Mechanism Overview

- The draft defines a method of using BGP to signal IPsec tunnel configuration along with NLRI, it uses and extends tunnel encapsulation attribute as specified in [I-D.ietf-idr-tunnel-encaps] for IPsec tunnel.
- BGP is only used to signal certain IPsec configuration, the IPsec tunnel is still created via **IKEv2** between routers after the configuration is learned via BGP UPDATES.

Proposed BGP Tunnel Encapsulation Attribute Changes

Extends tunnel encapsulation attribute specified in **draft-ietf-idr-tunnel-encaps** by:

- Uses existing tunnel type for IPsec tunnel: ESP tunnel mode – type 4
- AH tunnel mode is not in this document.

Adds subTLVs

- **Public routing instance** SubTLV – indicates where IPSEC packet is forwarding in.
  - Could be different from routing instance payload packet belongs.
- **remote address prefix SubTLV** – is a remote traffic select (from receiver Point of View)
  - Another way to do this is to use recursive lookup, but need more updates
- **New Sub-TLV for local address prefix** – is local traffic selector (from receiver's point of view)

It also reuses following existing sub-TLV:

- **Tunnel** Endpoint: IPsec tunnel endpoint address
- **Color:** IPsec configuration attributes like ESP transform; the meaning of this sub-TLV is local to the administrative domain
  - This could also be addressed by new subTLV for color to avoid reuse of existing color subTLV
- **Embedded Label Handling:**
  - For VPN-IPv4 Labeled Unicast or VPN-IPv6 Labeled Unicast routes, these NLRIs contain embedded label, however a labeled payload packet can't be encapsulated in ESP directly
  - However with Ipsec tunnel encapsulation, embedded label could be ignored since CHILD_SA itself could already identify the private routing instance
  - So an UPDATE that include IPsec tunnel encapsulation attribute, which contains value 2 of Embedded Label Handling Sub-TLV, could be used to signal ignoring embedded label

Mapping between BGP signaled config and IPsec needed Config

IPsec Config	BGP signaled Config
IKE config (transform, auth, lifetime…etc)	aggregated by Color Sub-TLV (or a new sub-TLV)
CHILD_SA mode (ESP tunnel)	Tunnel Type field in Tunnel Encap Attr
CHILD_SA transform	aggregated by Color Sub-TLV (or a new sub-TLV)
CHILD_SA local traffic Selector Address Range	a new sub-TLV
CHILD_SA local traffic Selector Address Range	a new sub-TLV
CHILD_SA traffic selector protocol/port-range	not signaled, ANY
CHILD_SA Anti-Replay Config	aggregaed by Color Sub-TLV (or a new sub-TLV)
CHILD_SA ESN config	aggregaed by Color Sub-TLV (or a new sub-TLV)

IP Security Association Set-up		
		1 Domain BGP
Generating initial IPSec SAs		based on the routing lookup, using the Ipsec config signaled in chosen BGP update to trigger IKEv2 to setup IKE_SA and CHILD_SA
Rekey of IPSEC SAs	Rules	IKE_SA and CHILD_SA rekey is done via IKEv2, based on config (e.g. lifetime) signaled by Color sub-TLV
	Single device Rekey	Just normal IKEv2 operation, rekey with its tunnel peer, no involvement of 3rd party
	Simultaneous multiple device rekey	normal IKEv2 operation, rekey with its tunnel peer, no involvement of 3rd party
IPSec DB generation	SPD (Security policy DB)	The SPD is essentially signaled via NLRI, local prefix sub-TLV and remote prefix sub-TLV
	SAD (security association DB) • Key generation • Nonces • SPI • IPSEC	SAD entry is dynamically created when IKEv2 create the tunnel or when CHILD_SA rekey happens; New key/nonce/SPI are generated dynamically/locally (not from BGP) whenever above event happens
	Peer Authorization DB	This is signaled by tunnel endpoint sub-TLV
Policy Distribution	Policy distribution Policy negotiation	Advertising router signals the Ipsec config it requires other to use for reaching the NLRI it advertise; this is could be done with RR or without RR; receiver use normal BGP routing lookup to decide which update to use, and use the Ipsec config in the selected BGP update to create tunnel

Security Issues: BGP Tunnel Attribute

BGP Attribute validation	1 Domain BGP
BGP Origin (RFC6811)	Could be used for further secure BGP update
Filters to stream out BGP security attacks	Could be used for further secure BGP exchange
BGPSEC	Could be used for further secure BGP update
Nested Tunnels	Like other type of tunnels, attention need to be payed to avoid looping as described in section 6.2 of draft-ietf-idr-tunnel-encaps-13

Security issues: Controller to Device issues

Question	1 Doman BGP
How does this draft handle tunneling across untrusted AS or administrative domain	As long as same set of mapping between color and Ipsec config are agreed, this draft could work across admin domain; if there is no such agreement, worst consequence is tunnel won't be able to create, and traffic might get dropped
Who sets the traffic selection policy?	Distribution: by advertising router of the NLRI Turning on: local router policy
Who sets up security DBs?	SPD: by advertising router of the NLRI SAD: dynamically created via IKEv2
Controller Conflict	2BGP preference: if two router advertise same NLRI but with different Ipse config, then it got resolved by normal BGP mechanism; BGP/non-BGP: this is local matter, for example, local configured Ipsec config should be prefer over BGP signaled one for same NLRI
Zero Touch set-up	Supported: This draft doesn't address zero touch; since it still requires some pre-provision (e.g. color to Ipsec config mapping, IPsec auth credential) Impact: however this draft could facilitate zero touch provision since only relative static (same for all NLRI) config need to be pre-provisioned

Questions