BGP IPSec: Links and VPNs

Susan Hares
Hickory Hill Consulting

Drafts

- Drafts considered:
  - draft-sajassi-bess-secure-evpn-02.txt,
  - draft-hujun-idr-bgp-ipsec-00.txt,
  - draft-dunbar-idr-sdwan-port-safi-03.txt
- Supporting drafts:
  - draft-carrel-ipsecme-controller-ike-00.txt
  - draft-ietf-i2nsf-sdn-ipsec-flow-protection-04.txt
  - draft-ietf-idr-tunnel-encaps-12.txt
  - Draft-dunbar-bess-bgp-sdwan-usage-01.txt

Why this meeting?

- Multiple overlapping proposals on IPSEC links and VPNS in Bess, IDR, and I2NSF with lots in common
- BESS and IDR Chairs agreed
  - Common TLVs for draft-idr-tunnel-encaps agreed upon by IDR
  - SA mechanisms need to be harmonized across the 3 drafts, but RTG chairs need input from Security
  - Determine if NLRI request in draft-dunbar-idr-sdwan-port-safi – should be looked at separately
- IDR/BESS know routing but need Security Area aid on IP Security methodology
  - Security area people agreed to meet us today – Thanks

Basic topology- Multiplied by 10,000

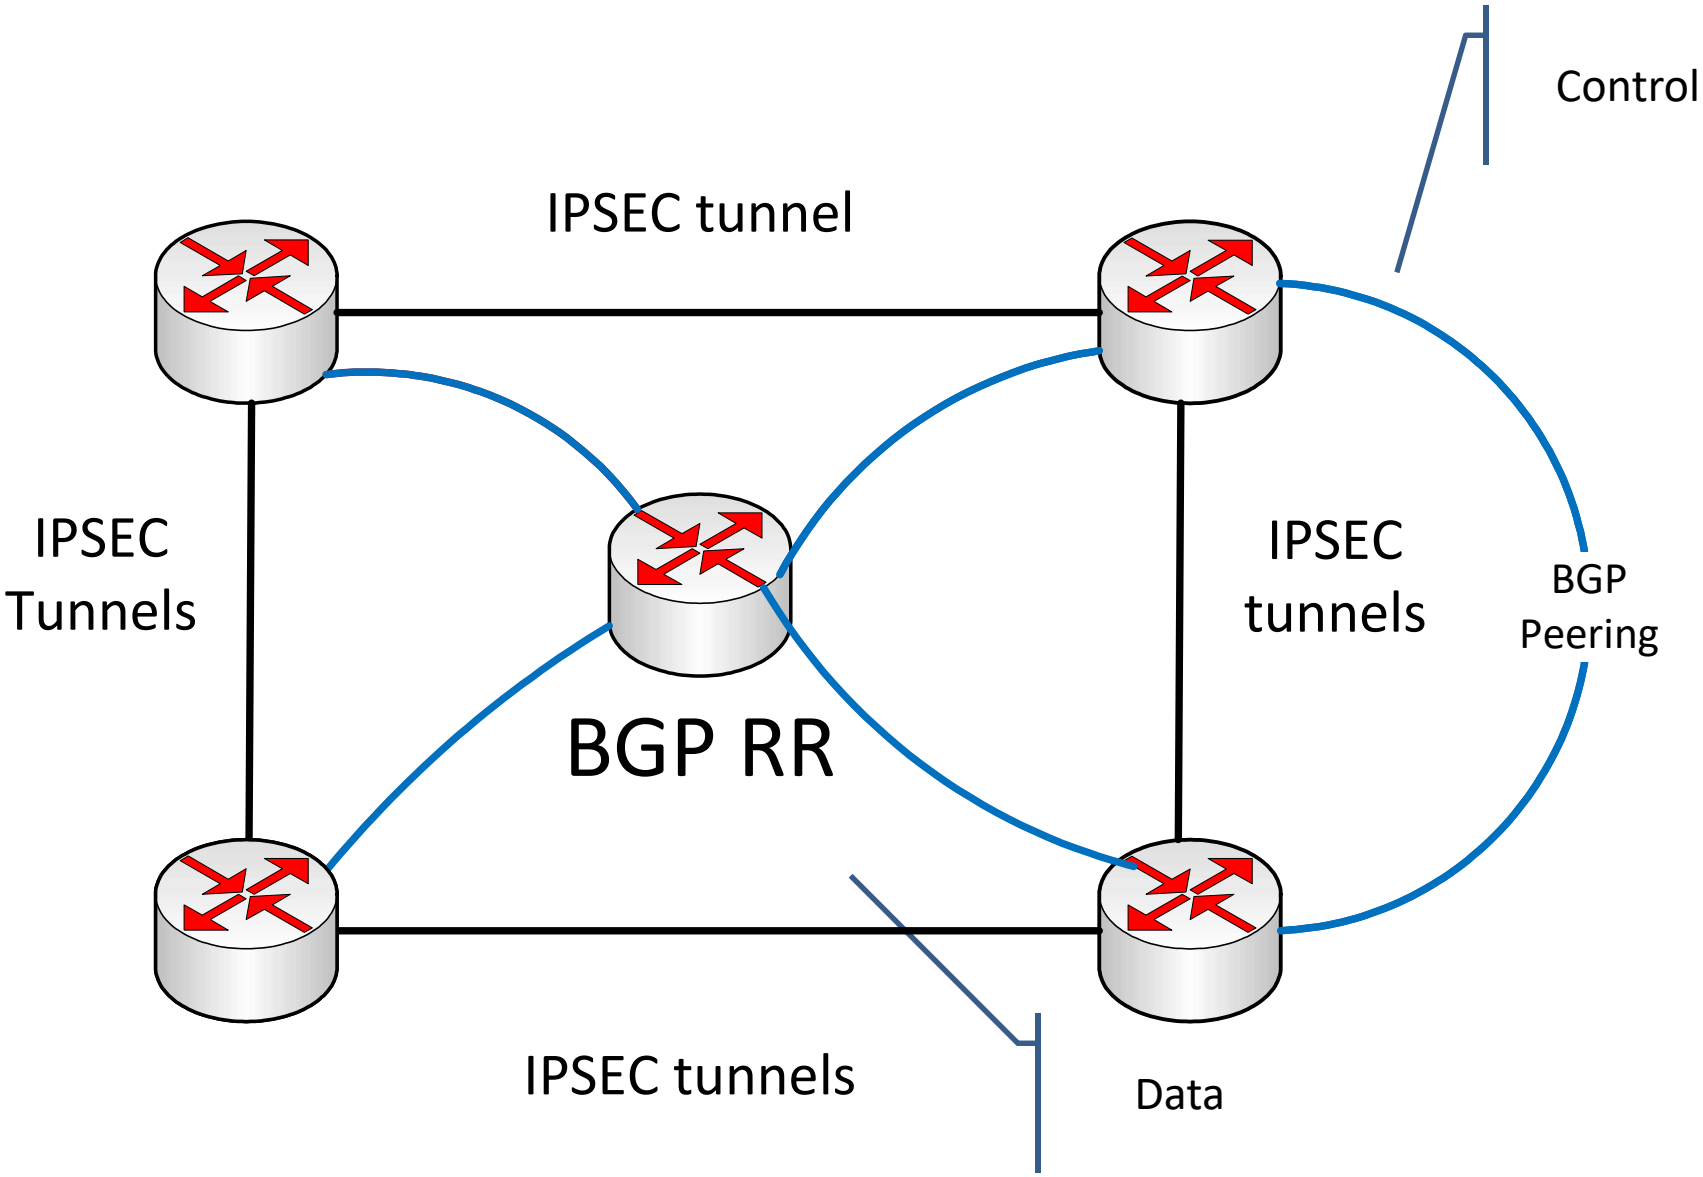

Overview of Drafts

draft-sajaassi-secure-evpn-03.txt

- Secure EVPN as part of the EVPN services from BESS

draft-dunbar-idr-sdwan-port-safi-03.txt

- SDWAN: from EVPN services from BESS which provides secure VPN for WANs mixing private secure VPNs and public VPNS

draft-hujun-idr-bgp-ipsec-00.txt

- To make provision & management of large number of IPsec mesh tunnels simpler and more efficient;
- Specially in a network without central controller for BGP

Personal Caveat

- I am a co-author on one of the drafts proposals.
- For this session, I will acting as WG chair
  - My only comments on the SDWAN draft will be to point out errors.

Topics

- Use Case and architecture
- Security issues
- Hierarchy Needed
- BGP Mechanisms
  - draft-ietf-idr-tunnel-attribute replaces Encapsulation Extended Community

Architecture – Device

	1 Admin Domain BGP-SEC	Secure E-VPN	SDWAN
Zero Touch Bring up	n/a	Yes -	Yes
Configuration Management	Pre-configured – central or OPS	Yes – device controller	yes – device controller
Orchestration	Uses Color to orchestrate pre-define configuration	Yes	Yes
Signaling	BGP with Tunnel Encapsulation	BGP with Tunnel encapsulation	BGP with tunnel encapsulation

Open Security Issues (TBD)

- Controller to Device
  - Assume RR can security identify the other BGP node
  - Sets up the traffic selection policy (policy distribution)
  - Sets up the Security Databases
    - Security Policy Database (on controller, no
    - Security Association Database (SADB)
- Conflict could occur between 2 mechanisms (I2NSF vs BGP, or 2 BGP) needs Resolution
  - Note: Goal is to either have non-overlapping policy roles for I2NSF and BGP.

- BGP Tunnel attribute (~Extended community) – sent over IPSEC, but BGP Data content is also validated via the following options:
  - Validating BGP Origin (RFC6811) + filtering
  - BGPSEC signature

Hierarchy

Level	1 Domain BGP	E-VPN	SD-WAN	IP VPN
PE group	n/a	n/a	**Site-ID**	(peer group)
PE /CPE level (BGP Peer)	**CTL:** BGP Peers **Tunnel:** peer-peer at If or loopback	**CTL**: PE-RR **Tunnel:** PE-PE or PE-CPE (v4/v6) at loopback	**CTL:** PE-RR **Tunnel:** CPE-CPE Route (v4/v6) or Loopback	**CTL:** PE-RR **Tunnel:** **PE-PE**
Tenant	**CTL:** BGP Peers [prefix]	**CTL:** PE-RR **Tunnel:** EVPN IMET	**CTL**: PE-RR **Tunnel:** EVPN IMET	n/a
Tenant subnet	Specify subnet Prefix (src/dst)	**CTL:** PE-RR EVPN IMET	**CTL:** PE-RR **Tunnel:** EVPN IMET	VPLS AD (~subnet)
Port group	n/a	No equivalent concept	**Port Distinguisher**	
Per IP	VPN IP prefix + color	**CTL:** PE-RR EVPN RT2/RT5	**CTL:** PE-RR Local IP address	**CTL:** PE-RR VPN IP RT
Per MAC	n/a	**CTL:** PE-RR EVPN RT2	**CTL:** PE-RR EVPN RT2	n/a

IPSec Data in BGP-TLVs

Information	1 Domain BGP	EVPN	SD-WAN
Tunnel Identifier	**Tunnel type:** 4 **Sub-TLVs:** Public-routing, Local/remote prefixes,	**Tunnel type: multiple** **DIM sub-TLV** Originator ID + (Tenant ID) + Subnet ID + Tenant Address)	**Tunnel type: multiple** **DIM sub-TLV** Originator ID + (Tenant ID) + Subnet ID + Tenant Address)
Port distinguisher	Private/public	n/a	**SD-WAN NLRI/SAFI:** Port Distinguisher SITE-ID, Node-ID **In Tunnel Attribute:** EncapExt sub-tLV (includes public/private)
Nounce data	Local, auto	DIM sub-TLV: 32 bits	DIM Sub-TLV: 32 bits
Rekey info	Dynamic	Dim sub-TLV: 32 bits	DIM sub-TLV: 32 Bits
Key Exchange	Pre-configured	Key exchange sub-TLV	Key exchange sub-TLV
SA transforms	Pre-configured	ESP SA sub-TLV	IPsec-SA sub-TLV
Not used Sub-TLVs	all EVP	n/a	Remote Endpoint

Current Tunnel Types

- 0 – Reserve [RFC5512]
- L2TPv3 over IP [RFC5512]
- GRE [RFC5512]
- Transmit tunnel endpoint [RFC5566]
- IPSec in Tunnel-mode [RFC5566]
- IP in iP tunnel with IP sec [RFC5512]
- MPLS in IP Tunnel [RFC5566]
- IP in IP [RFC5512]
- VXLAN encapsulation [RFC8365]
- NVGRE encapsulation [RFC8365]
- MPLS Encapsulation [RFC8365]
- VXLAN GPE encapsulation [RFC8365]
- MPLS in UDP Encapsulation [RFC7510] [RFC Errata 4350]
- IPv6 Tunnel [Martin Djernaes]
- SR TE Policy Type [draft-previdi-idr-segment-routing-te-policy]
- Bare [Nicschal Sheth]

draft-ietf-idr-tunnel-encapsulation obsoletes RFC5512.
RFC5566 – depends on RFC5512.
RFC5566must be revised !

draft-ietf-idr-tunnel-encapsulation does not define PMSI (RFC6514) + this idr tunnel attribute

BGP Secure VPN Requirements

- Scalability – 10K nodes, 100K links, 10 million routes, 20 million customers
  - Control traffic needs to be minimized
- Robustness – 99.999% uptime, 99.999% packets get through
- Ready to go key management – SA on the fly within ms
- Rekeying occurs
- Separate path for control vs. Data
- Network Topology with non-bidirectional links

Why BGP as Control Plane (BGP Basics)

- Compelling reasons of using BGP:
  - BGP already widely deployed as sole protocol (see RFC 7938)
  - Reliable transport, Guaranteed in-order delivery over Secure TCP
  - Incremental updates
  - RR Hierarchy reduces full mesh of BGP Peers and Route Table
  - RR already has the capability to apply policies to communications among peers for efficient distribution
  - BGP + RRR supports many logical topologies (hub-spoke, mesh)
- BGP Implementations:
  - Robust, technology widely accepted – minimal learning
  - RR has flexible filtering policies to communications among peers.
  - Deployed in large networks

What IPSec people can help with

- Asked each proposal team to discuss Security portion of their proposal
  - So IPSEC people can comment regarding these proposals
  - Two proposals (Secure EVPN and SD-WAN) use draft-carrel-ipsecme-controller-ike-00
- Perhaps this is beginning of a longer conversation

IP Security Association Set-up

		Domain BGP	Secure EVPN	SD-WAN
Generating initial IPSec SAs				
Rekey of IPSEC SAs	Rules			
	Single device Rekey			
	Simultaneous multiple device rekey			
IPSec DB generation	SPD:security policy DB			
	SAD – security association DB • Key generation • Nonces • SPI • IPSEC			
	Peer Authorization DB			
Policy Distribution	Policy distribution Policy negotiation			

Security Issues:
BGP Tunnel Attribute

BGP Attribute validation	1 Domain BGP	Secure EVPN	SD
BGP Origin (RFC6811)	Support: Y/N	Support: Y/N	
Filters to stream out BGP security attacks			
BGPSEC			
Nested Tunnels			

Extended Communities in BGP can be changed by anyone.
Attributes have a stricter set of rules.
Some proposal for IPSEC use Extended Communities

Security issues:
Controller to Device issues

Question	1 Doman BGP	Secure EVPN	SDWAN
How does this draft handle tunneling across untrusted domain?			
Who sets the traffic selection policy?	Distribution: Turning on:		
Who sets up security DBs?	SPD: SAD:		
Controller Conflict	2BGP preference: BGP/non-BGP:		
Zero Touch set-up	Supported: Y/N Impact:		