Secure L3VPN over Public Infrastructure

- Draft-rosen-bess-secure-l3vpn-00
  - Eric Rosen, Ron Bonica
- Goals:
  - augment RFC4364 technology for use over “public” backbone network
    - untrusted
    - no MPLS transport
  - retain RFC4364 “multi tenancy” features

Basic Concept

- Customer Premises Equipment (CPE) provides PE functionality (“C-PE”)
- C-PE control plane: IPsec-protected BGP to secure RR
  - Private routes advertised as VPN-IP routes with private C-PE loopback as next hop
  - Private C-PE loopback advertised as IP route with:
    - Public C-PE address as next hop
    - Tunnel Encapsulation attribute that specifies a C-PE to C-PE IPsec Security Association (SA)
- Data plane uses these IPsec SAs

C-PE Red Routes

- C-PEs have red interfaces (to private sites) and black interfaces (to public net)
- Each C-PE has a red loopback (private) and a black loopback (public)
- C-PEs originate two kinds of red route:
  - VPN-IP routes pointing out red interfaces
    - **Set Next Hop to red loopback**
  - IP route to red loopback (see next slide)
- Red routes only advertised over IPsec-protected (red) BGP sessions (IBGP or EBGP)

Use of Tunnel Encapsulation Attribute

- The red IP route whose NLRI is the red loopback carries a Tunnel Encapsulation attribute (TEA):
  - tunnel type = MPLS-in-IPsec (RFC 4023)
  - remote endpoint = black loopback
    - Note: does not change when route is propagated, even when propagated via EBGP
  - TLVs with whatever other information is needed to set up the IPsec SA

Resolution of Red VPN-IP Routes

- How does C-PE1 forward a packet it receives over a local red interface?
  - Suppose packet's IP DA, interpreted in proper VRF context, matches <NLRI=X, NH=C-PE2-red>
  - Recursive resolution of C-PE2-red finds TEA:
    - Tunnel type = MPLS-in-IPsec
    - Remote endpoint = black loopback of C-PE2
  - IPsec SA gets set up over public backbone between C-PE black loopbacks
    - Remember, black loopbacks are public addresses
  - Therefore the packet gets sent to C-PE2 through the MPLS-in-IPsec tunnel

Cautions

- MUST NOT:
  - accept VPN-IP route from insecure BGP session
  - transfer data between red and black interfaces unless protected by IPsec on the black interfaces
- MUST:
  - Resolve next hop of VPN-IP route via route (with appropriate TEA) received over secure BGP session

Setting up the Secure BGP Sessions

- RRs have red loopback address, black loopback address, black interface addresses
- BGP sessions to C-PEs run through *IPsec transport mode SAs* between the black addresses
- RRs:
  - may be provisioned with pre-shared secrets of C-PEs,
  - or may use certificates to authenticate C-PEs,
  - have no prior knowledge of C-PE black addresses, so C-PEs can move
- C-PEs initiate the sessions

The Data Plane IPsec SAs

- Can be set up when route with TEA is received
- Or can be set up when needed for data
- Granularity: C-PE to C-PE
- BTW, what is MPLS-in-IPsec?
  - Same as MPLS-in-IP in IPsec transport mode
  - On wire, IPsec header followed by label followed by user payload
  - Only black C-PE addresses are in the clear

Next Steps

- Call for adoption