Why IPsec and BGP don’t play well together in real networks

Brian Weis

Overview

- Of course IKE/IPsec can be configured on routers to provide transport security for BGP and other similar “control plane” protocols.
  - And certainly there are networks where the use of IPsec meets expectations
  - This presentation isn’t intended to marginalize the use of IKE/IPsec
- However there are operational considerations that make IPsec protection inadvisable on many routers speaking BGP
  - Re-keying Issues
  - DoS issues
- But first, consider how a typical BGP router works.

Typical BGP Router

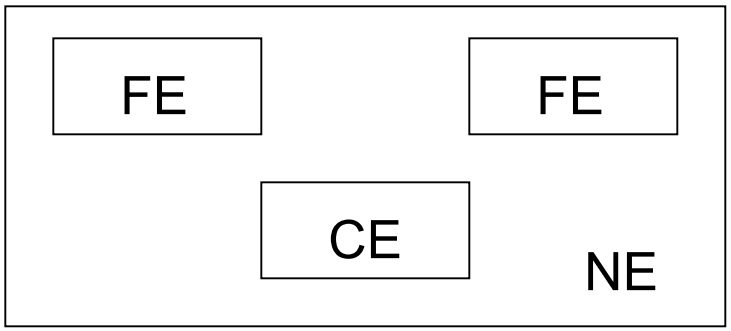

- Using ForCES terminology (RFC 3746)
  - Network Element (NE)
    - In this example, the entire router
  - Forwarding Element (FE)
    - The blades containing network ports
  - Control Element (CE)
    - The blade processing the BGP protocol

Typical BGP Router Architecture

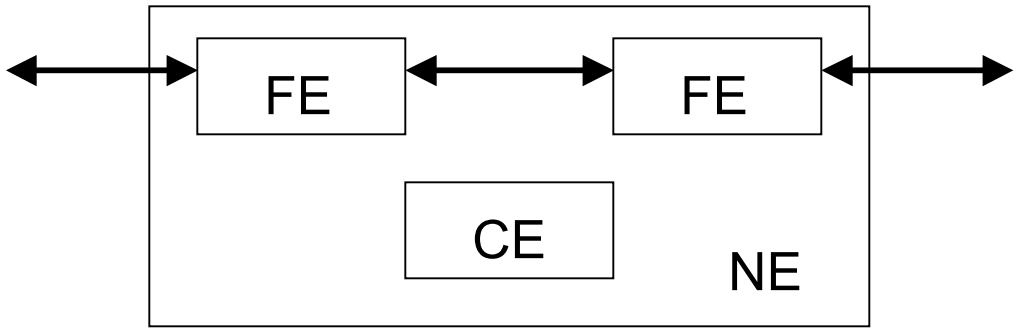

- BGP Router is a device optimized for routing data packets between forwarding elements
  - Data packets are switched either in H/W or in S/W at a high interrupt level
  - This level of packet switching is often called the “fast path”
  - The router has many OC-192 (10 Gbits/Sec) ports.

Typical BGP Router Architecture

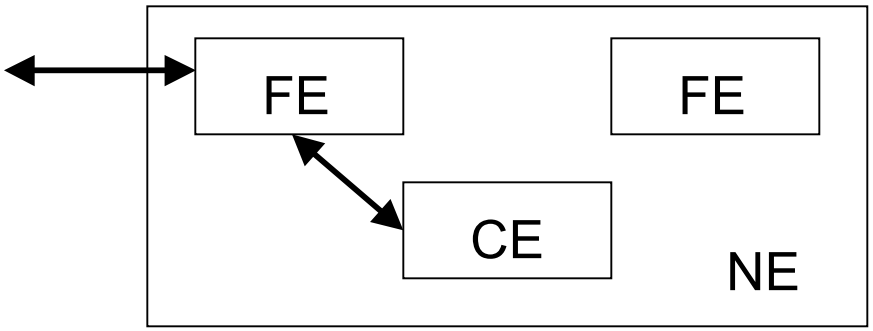

- Control traffic (e.g., BGP) is forwarded to the control element for processing
  - Note that the packets may cross a bus, but in any case must be enqueued on the CE
  - The CE may include a H/W assist (e.g., ASIC) processing the packet before queuing to enforce anti-DoS measures of control traffic.

Typical BGP Configuration

- A BGP router has many configured peers
- Each peer is probably a unique administrative entity, each employing an independent set of operators, and often times owned by a competitor
  - This makes for a difficult operational and security model

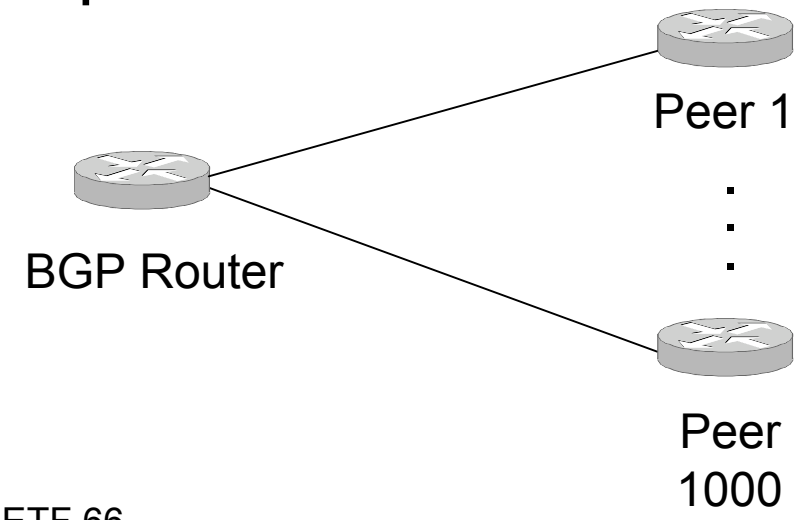

Re-keying

- Why re-keying is important
- IKE/IPsec re-keying issues
- Economic cost of taking an outage due to a missed re-key

Why re-keying is important

- It has been posited in some quarters that there's no need to ever rekey BGP sessions
  - That's true only if you have pair-wise trusted links to all peers, and appropriate ingress filtering.
- Reasons to re-key:
  - It's *always* bad practice to assume that keys won't be overused, be leaked, etc. When one of those events happen, no orderly recovery is possible, and the result is using a bad key.
  - Service Provider's have operational staff turnover, and BCP is to change keys to reduce the risk from a disgruntled ex-employee
- Re-keying is important whether the keys are session keys or authentication keys.

IKE/IPsec re-keying issues

IKE Public key pairs:

- Using certificates for public key distribution
  - Re-keying is achieved by issuing a new certificate
  - But how many trust anchors does a BGP router need to communicate with 1000 peers? It isn't reasonable today to expect all ISPs to belong to a single common PKI.
- Distributing raw public keys
  - Results in twice as many keys being exchanged in the secret key case
    - Routers may have a limited amount of NVRAM or flash available, and public keys are relatively large
- Public key operations are computationally expensive
  - Supporting public key operations from many peers requires customers to buy blades/boxes that they don't otherwise need.

Therefore, public keys pairs are generally considered infeasible for BGP routers today.

IKE/IPsec re-keying issues

Pre-shared secret keys:

- IKE Pre-shared secret keys
  - IKE pre-shared keys are associated with an IKE identity. Strictly speaking, there can only be one IKE pre-shared key per peer
  - Re-keying happens by exchanging new keys.
    - There is some synchronization of the key change needed
    - The IKE key lifetime provides a window in which IKE pre-shared keys can be different on different peers (unless one side needs to unexpectedly re-negotiate IPsec SAs before the key is changed on both peers)
- IPsec manually configured SAs
  - It is possible to have multiple IPsec manually configured SAs for a particular peer.

One of these secret key solutions will have to do….

BGP Re-key operational procedures

- Consider two independent operations staffs, each on their own schedule.
  - The logistics are such that the change of a single key *cannot* be tightly synchronized.
    - There may be a period of hours or days between the time Provider 1 and Provider 2 install the new key
  - An IKE pre-shared key lifetime is not guaranteed to be long enough to avoid a synchronization error.
- In the meantime, the BGP session must not go down!

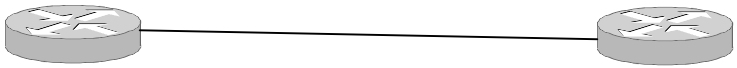

Provider 1 Provider 2

Economic cost of taking an outage due to a missed re-key

- The Service Provider business model cause them to prefer reliable up-times over security
  - In many cases they have made Service Level Agreement (SLA) promises to customers, and a single BGP session reset could break one or even many SLAs
  - Many SPs have to report customer perceivable outages in their network to their Telecommunications Regulator
  - Idling an OC-192 connection is just too costly due to the secondary effects of re-routing that traffic across other links.
- Therefore, re-keying must be reliable!

Secret Key Rollover

- In either the case of IKE pre-shared keys or IPsec manual keys we need some strategy for synchronizing a rekey. E.g.,
  - draft-bonica-tcp-auth-04 style key lists
  - draft-ramaiah-key-rollover-00 methods
  - draft-bellovin-keyroll2385-00 method
- Some of these methods assume the use of multiple keys per peer.
  - That's reasonable for IPsec manually configured SAs because each SA has a unique SPI
  - What about IKE pre-shared keys for a single peer?
    - This seems like a fairly heretical idea, because we've always assumed a PKI should be used instead
    - The receiver would need to try all candidate keys in turn, which can be used as a DoS method

DoS issues

- By nature of their position in the network BGP routers are available to attack
- Fewer control plane protocols available to be attacked is better.
  - In this sense, IKE has the disadvantage of being an additional attack vector (although IKEv2 may be less susceptible than IKEv1)
- BGP routers have layered DoS protections that IPsec-encapsulated packets may weaken.

Layered DoS Protection

- BGP routers protect against
  - Receive queue saturation
  - Bus or backplane saturation
  - CPU saturation
- Mitigation requires looking at the packet
  - Queuing based on protocol type or other parameters (e.g., IP precedence)
  - ACL checks
- These mitigation measures may be foiled by AH/ESP encapsulation
  - Filtering on a SPI isn't as fine-grained as filtering on protocols and ports.
  - The router must de-capsulate the packet before it can do full filtering.

Summary

- IKE/IPsec can be configured to protect BGP, and indeed is sufficient in *some* situations.
- But there are operational issues that discourage its use by most BGP routers used in the Internet today.