Anti-replay sequence number subspaces

draft-ponchon-ipsecme-anti-replay-subspaces

Mohsin Shaikh (presenter), Paul Ponchon, Pierre Pfister, Guillaume Solignac
IETF 116 @ Yokohama

Quick Recap

- Proposal to use multiple sequence number subspaces as an alternative to creating multiple child SAs for multi-core performance
- Additionally support QoS and multi-path

Advantages of multiple sequence subspaces

- Some users require anti-replay and PFS for compliance
- No. of IKEv2 messages exchanged:
  Multiple sequence subspaces : 4
  Multiple child SAs : 2 + (2 x no. of child SAs)
  IKEv2 touts fewer message exchanges as an advantage over IKEv1 [RFC7296 - Appendix A. Summary of Changes from IKEv1]
- Avoids repeating history. IKEv1 protocol allowed a single pair of selectors per CHILD_SA, while IKEv2 improved that by allowing multiple traffic selectors to be negotiated for one child SA. We are repeating the same IKEv1-like behaviour with multiple child SAs per core. [RFC7296 - Appendix A. Summary of Changes from IKEv1]
- Adding too many SAs may slow down per-packet SAD lookup [draft-ietf-ipsecme-multi-sa-performance-00]

Updates in draft-01 since IETF 115

- Increased the sequence number field in ESP header to 64-bits, the subspace ID to 16 bits (from 8 bits in previous draft) and reserve the top 16 bits to store subspace ID.

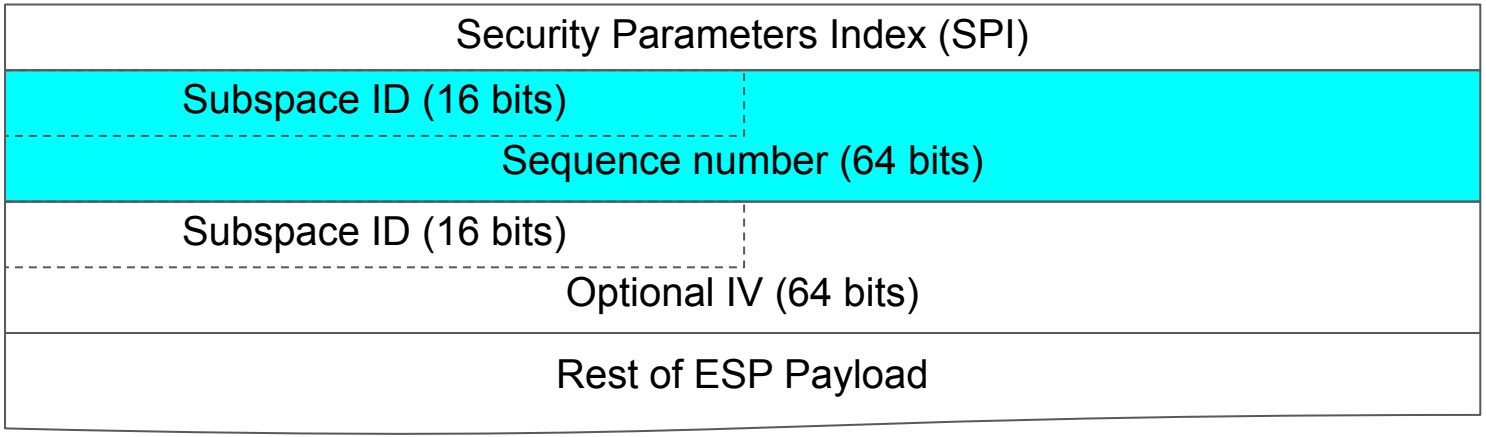

Updates in draft-01 since IETF 115 (contd.)

- IKE negotiation contains a new “Anti-replay subspaces” transform to negotiate the number of subspaces required
- IPR disclosure from SSH Communications Security

Test Data

- Memory usage goes up with multiple child SAs (tested between 2 docker containers using strongSwan with 1 IKE SA).

ESP proposal	aes256gcm8	aes128ctr-sha2_256
0 tunnel	9,040 kB	8,908 kB
1000 tunnels	17,996 kB	19,228 kB
10,000 tunnels	90,360 kB	98,576 kB
100,000 tunnels	824,100 kB	891,568 kB

- Assuming a 64 packet window, the amount of increase in memory per SA is (8+8 bytes) X no. of sequence subspaces.

Should the working group work on this and adopt ?

Q&A