IPsec and IKE anti-replay sequence number subspaces for multi-path tunnels and multi-core processing

draft-pponchon-ipsecme-anti-replay-subspaces

Paul Ponchon (presenter), Mohsin Shaikh, Pierre Pfister, Guillaume Solignac

IETF 115 @ London

What's the problem ?

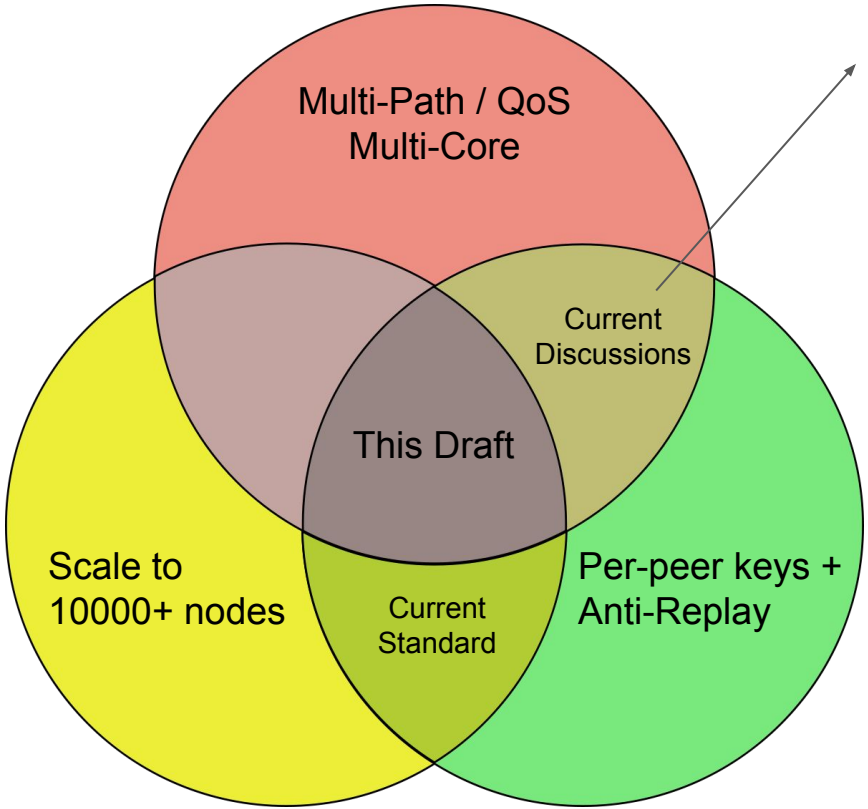

Current drafts suggest using multiple SAs

- (2019) draft-mglt-ipsecme-multiple-child-sa-00
- (2022) draft-pwouters-ipsecme-multi-sa-performance-04
- IETF 108 Presentation

But Multi-Path / QoS requirement multiplies the number of Child SAs:

- e.g., assuming 6 paths, 8 cores, 8 QoS classes: **384 Child SAs per peer**.

10.000 peers become 3.840.000 Child SAs:

- **Unnecessary load on IKE**
- **Fills CPU cache and hardware memory**

The source cause is Anti-Replay: Let’s fix it

Challenges with Anti-Replay:

- Anti-replay material cannot be efficiently shared across cores
- Multiple paths cause out-of-order packets: Packets are dropped !

Proposed solution:

- Single Child SA maintains **multiple anti-replay sequence number and vector**.
- Sender sets the subspace ID in the ESP header
- Receiver uses the subspace ID to use the appropriate anti-replay material

Discussion #1: IKE Negotiation

The draft will describe an IKE-based negotiation.

Input from the working group is welcome !

Discussion #2: subspace ID encoding

Option 1 (in current draft): 8 higher-order bits of the sequence number space

- 24 bits explicit sequence number cycles in 20 seconds at 10Gbps
- Extended Sequence Number (ESN) will be needed
- At very high speeds (e.g. 100Gbps), outages might cause resync
- More than 8 bits if ESN is made explicit

Option 2: Add a new 32 bits field between sequence number and IV

- 16 bits to be used as subspace-ID
- Remaining bits reserved for future use

Discussion #3: AES-GCM FIPS compliance

- AES-GCM for IPsec max usage count is 2^64 and enforced by RFC4106

> Moreover, as with most block cipher modes of operation, the security assurance of GCM degrades as more data is processed with a single key. Therefore, the total number of blocks of plaintext and AAD that are protected by invocations of the authenticated encryption function during the lifetime of the key should be limited. A reasonable limit for most applications would be 2^{64}, consistent with the requirement on the number of invocations in Sec. 8.3.

NIST 800-38d Appendix B

- How to make different cores share the 64 bits IV space is implementation specific
  - Only requirement is that an IV value shall only be used once.
  - A sub-field can be used to encode the sender core ID, or the subspace ID, the rest is used as a counter.

Conclusion

- We want to help the working group to solve anti-replay issues (with multi-core and multi-paths).
- Current proposed solutions rely on creating more SAs, but we believe this will show scalability limitations.

- Let’s find a common ground and work on a single solution:
  - Agree that multiplying the number of Child SAs and keymat will not scale ?
  - Allow multiple anti-replay material in the same Child SA ?
  - Encode the anti-replay subspace ID in sequence number space ? Or in a different field ?
  - Negotiate the option with IKE ?