IKEV2 SUPPORT FOR PER-QUEUE CHILD SA

DRAFT-PWOUTERS-IPSECME-MULTI-SA-PERFORMANCE-00

IPsec, IETF 111

July 2021

Antony Antony, Tobia Brunner, Steffen Klassert, Paul Wouters

Goal of the draft:

- 40-100 Gbps wire speed IPsec using multiple CPU cores
- An unencrypted link of 10 Gbps or more is commonly reduced to 2-5 Gbps when IPsec is used to encrypt the link using AES-GCM.

- By using the implementation specified in this draft, aggregate throughput increased from 5Gbps using 1 CPU to 40-60 Gbps using 25-30 CPUs

Changes since last version

From draft-pwouters-multi-sa-performance
To draft-pwouters-ipsecme-multi-sa-performance

- Due to fixed name, diff not linked, see manual diff
- Separate info notifies for QoS and CPU case
- Clarified terms: Initial Child SA → Fallback Child SA
- Always require an INFO notify for an Additional Child SA
- Negotiate the maximum number (not minimum)
- Attempt to clarify QoS case
- Added some operational considerations
- Clarified case when not having per-queue ACQUIRE

Questions for WG: QoS

- Should we remove/split QoS in separate draft?
  - Authors have little QoS experience, no p-QoS code.
- Negotiate “all” QoS / flows at once? There is no variable number – the number is “all the different ones”
- Need to request ALL combinations? Or just ones you want (eg “bulk” and “voip”)
- How does IPv4 QoS & IPv6 flow label combine in TS?
- Do we need a new “reject this QoS/flow” TS_ error code?
- Can one combine per-CPU and per-QoS ? We don’t really know.

Questions for WG: Other

- When too many Child SA's, return TS_UNACCEPTABLE, ~~NO_ADDITIONAL_SAS~~, or a new error code?

- Considerable effort (two years) made by implementations (Linux, libreswan, strongswan).
  - Need to know if WG wants to move forward or not.