IPv6 Configuration in IKEv2

draft-eronen-ipsec-ikev2-ipv6-config-04

pasi.eronen@nokia.com

Background: IPv4

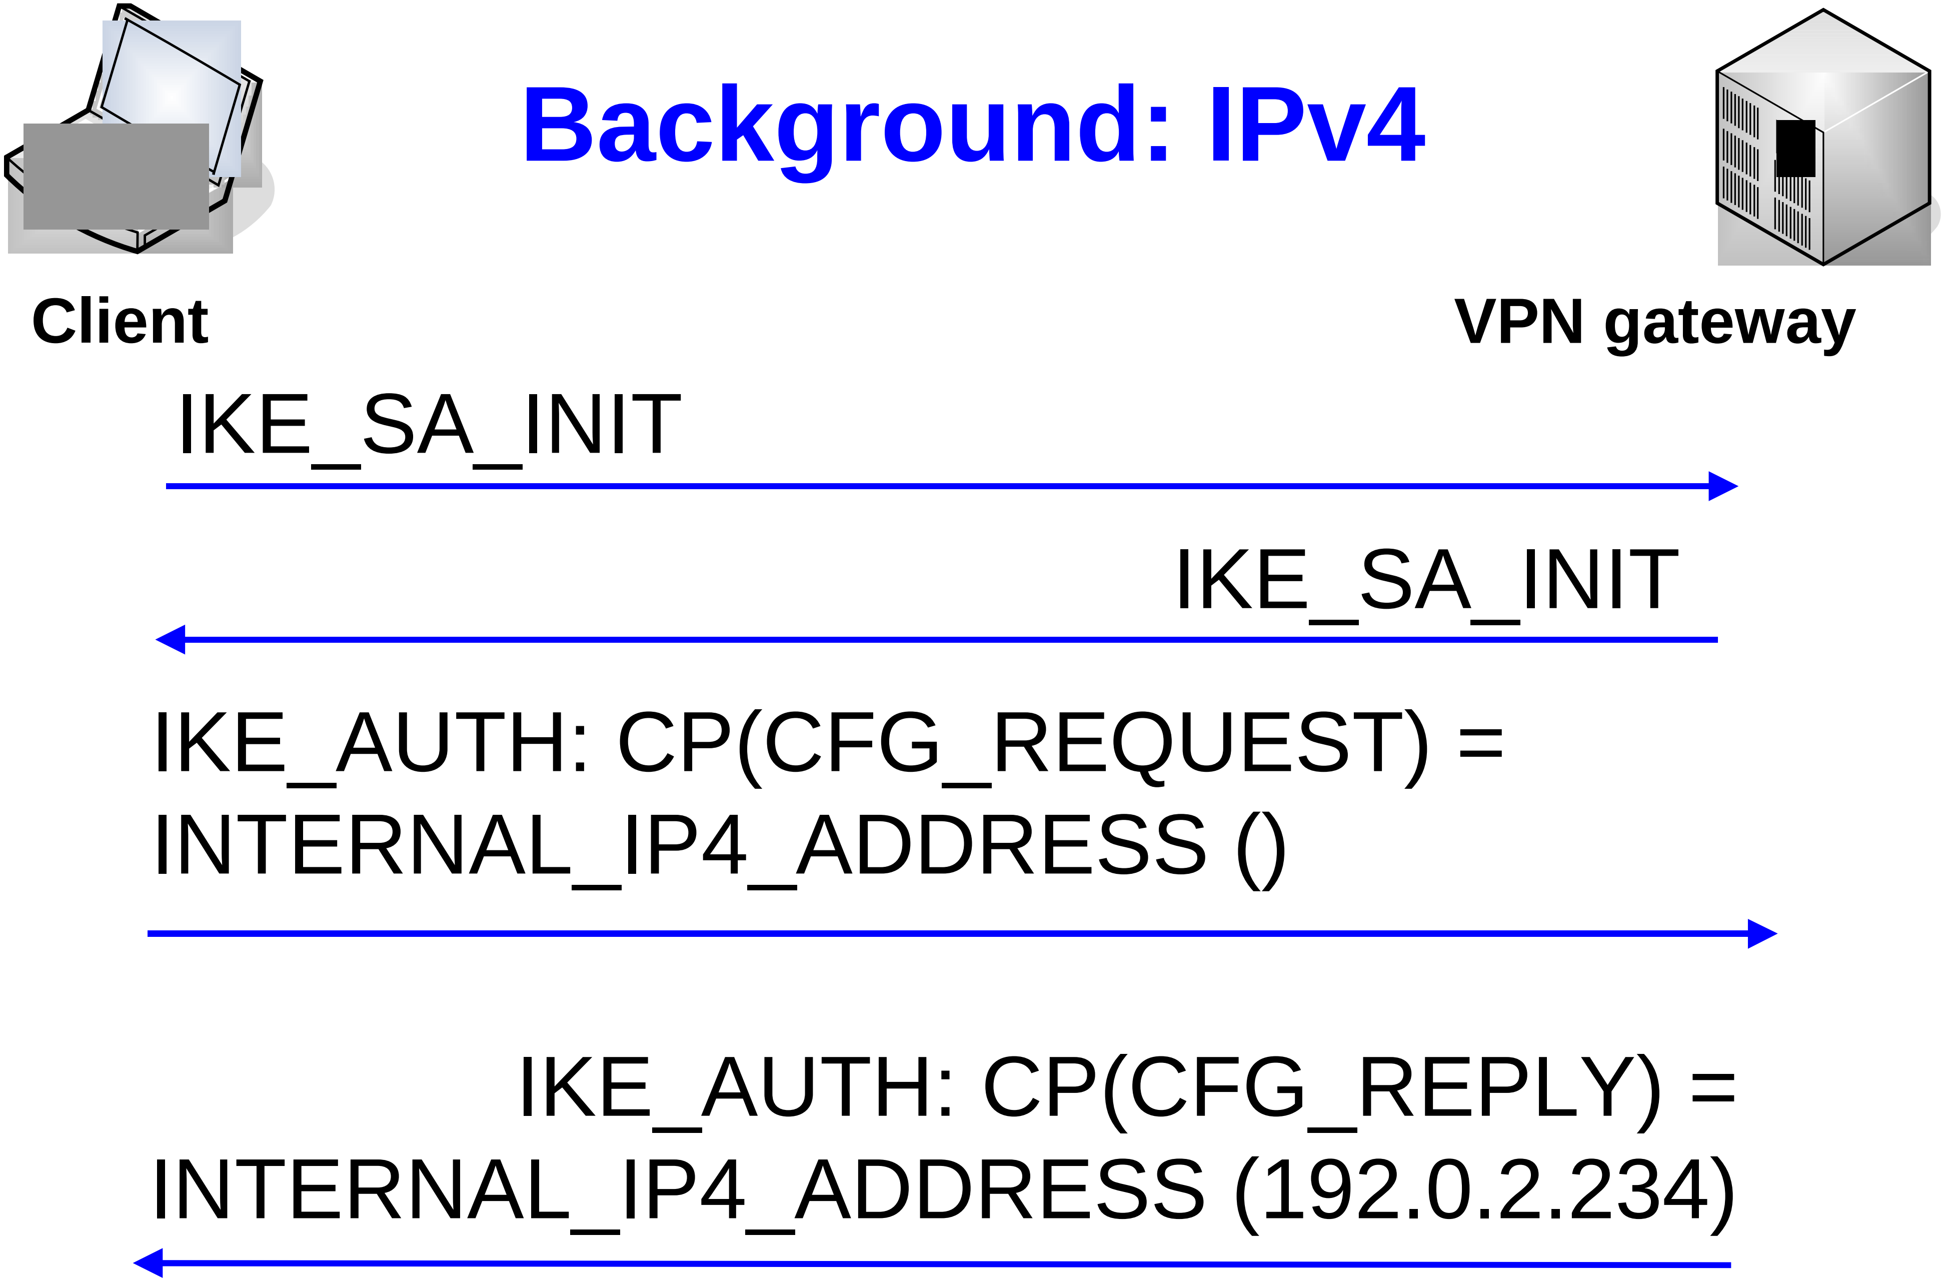

Client

VPN gateway

IKE_SA_INIT

IKE_SA_INIT

IKE_AUTH: CP(CFG_REQUEST) =
INTERNAL_IP4_ADDRESS ()

IKE_AUTH: CP(CFG_REPLY) =
INTERNAL_IP4_ADDRESS (192.0.2.234)

Behind the scenes: gateway

IKE_AUTH: CP(CFG_REQUEST) =
INTERNAL_IP4_ADDRESS ()

- Pick an unused address (from internal pool, DHCP, or AAA)
- Create PAD entries authorizing IDi to create CHILD_SAs for this address
- (If needed, update SPD)
- Narrow TSi/TSr using PAD/SPD

Behind the scenes: client

IKE_AUTH: CP(CFG_REPLY) = INTERNAL_IP4_ADDRESS (192.0.2.234)

- Create “virtual interface” with this address
- Update source address selection information (e.g., routing table) so that this address gets used by apps (for new TCP connections etc.)
- Create PAD entries authorizing IDr to create CHILD_SAs for this address
- (If needed, update SPD so that all traffic from this address/interface is sent to the gateway)

IPv6 version

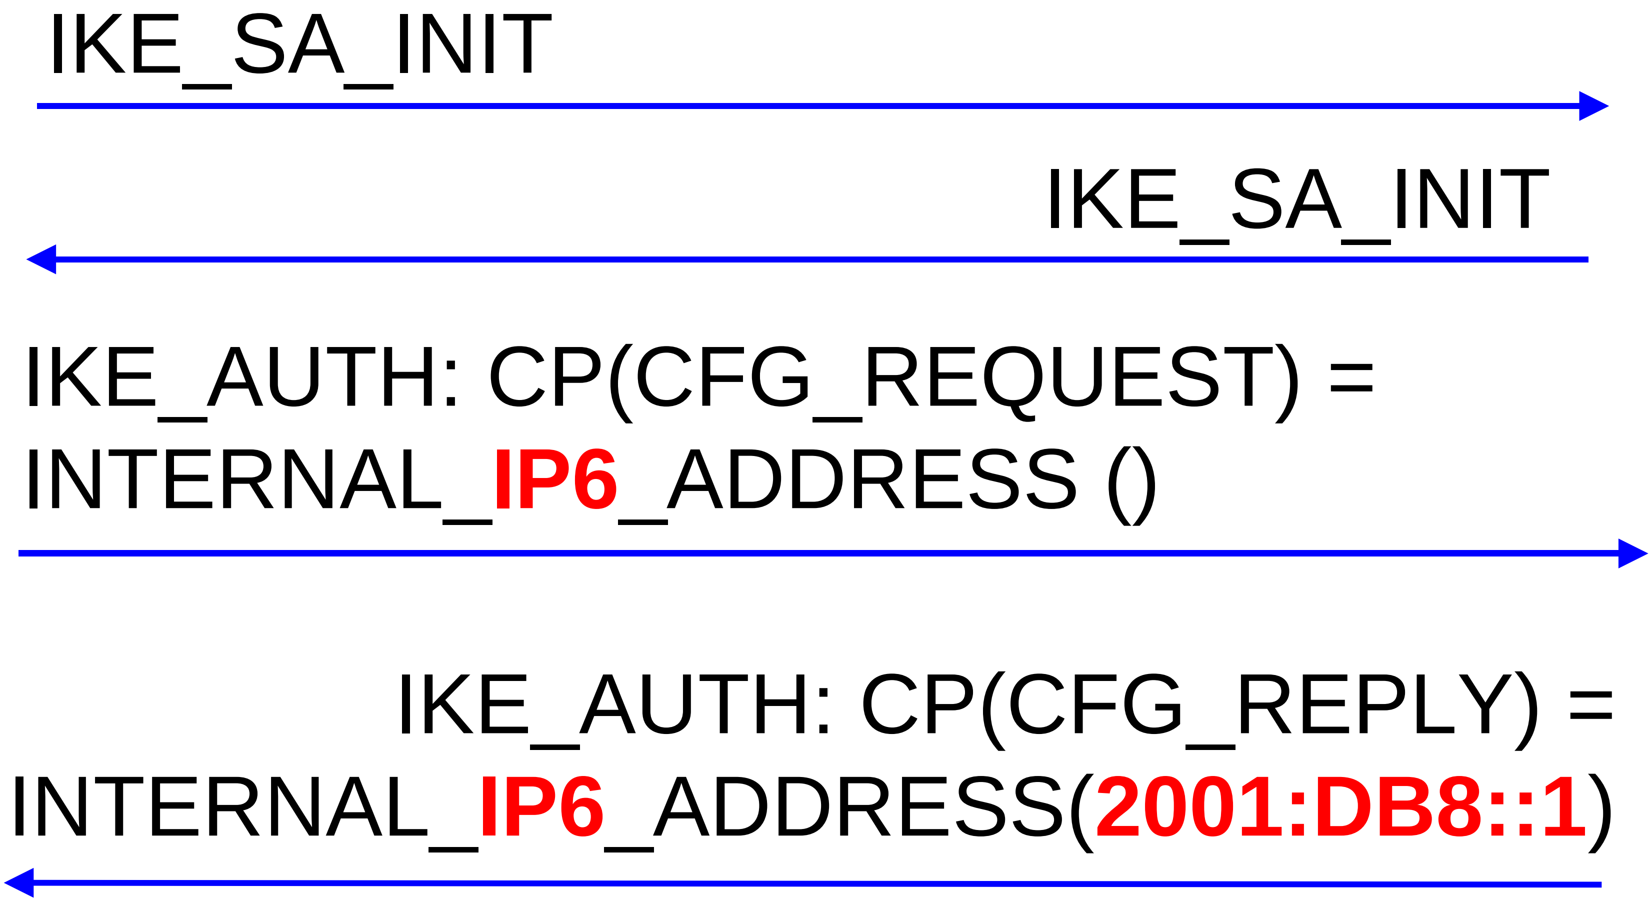

IKE_SA_INIT

IKE_SA_INIT

IKE_AUTH: CP(CFG_REQUEST) =
INTERNAL_**IP6**_ADDRESS ()

IKE_AUTH: CP(CFG_REPLY) =
INTERNAL_**IP6**_ADDRESS(**2001:DB8::1**)

Problems

- No multiple prefixes (renumbering, host-based site multihoming, …)
- No link-local addresses (violates MUST in RFC 4291)
- Interface ID selection (CGAs, HBAs)
- Additional references
  - Why this was bad idea for 3GPP: RFC 3314
  - Why multilink subnets are complex: RFC 4903

Solution space (1 of 3): Link/subnet model

- Point-to-point
  - Every client gets its own prefix
- Multi-access
  - Multiple VPN clients on same “virtual link” (“like Ethernet”)
- “Router aggregation” (NBMA)
  - Shared prefix, but not shared link (multi-link subnet)

Solution space (2 of 3): Layer 3 Access Control

(How gateway drops packets
with wrong source address)

- IPsec traffic selectors in SAD/SPD
- Ingress filtering outside IPsec

Solution space (3 of 3): Where address/prefix is sent

- IKEv2 messages (configuration payloads)
- ND inside tunnel
- DHCPv6 inside tunnel

Solution space (extras)

- **Reauthentication:** When same IDi opens second IKE_SA, same address(es) or different ones?
- **Compatibility with other IPsec uses:** When creating CHILD_SA, is it for the virtual interface or the interface IKE packets are sent over?
- (See draft for details and discussion)

Solution discussion

- Current draft proposes one combination (next slides)
- Sketches 5 others in Appendix A (and explains why I felt they're less desirable)
- Depends on how you prioritize pros and cons
  - E.g., implementation impact on IKEv2 vs. per-packet IPsec processing (kernel space) vs. rest of IPv6 stack
- Not all combinations make sense

Current proposal

- Point-to-point link model
  - + Each client gets its own /64 prefix, can use (almost) any interface identifiers
  - + Simplest, no complexity of multi-link subnets, or overhead of multi-access
  - - VPN gateway needs larger address pool (not problem for enterprise/ISP, possibly for homes if ISPs don’t follow RFC 3177)

Current proposal

- L3 access control with IPsec SAD/SPD
  - + Aligned with overall IPsec architecture
  - + Same as in IPv4 case
- IKEv2 configuration payloads
  - + Same as in IPv4 case
  - + IKE knows about addresses → easier to do L3 access control with IPsec
  - - Specific to IKE (but can use stateless DHCPv6 for other configuration than address)

Other combinations (quick overview only)

#1: Stateless autoconfiguration (inside tunnel) + point-to-point link

- + Looks elegant (on paper, at least)
- – Implementation impact for kernel-side IPsec and rest of IPv6 stack?
- – L3 access control outside IPsec → not aligned with IPsec architecture
- – Very different from IPv4 case

#2 and #3: Stateless autoconfiguration + NBMA

- + Allows sharing prefixes
- – Non-standard processing of ND messages on gateway?
- – Multi-link subnet
- – L3 access control outside IPsec
- – Very different from IPv4 case

#4: “As close to IPv4 configuration payloads as possible”

- + Similar to IPv4 case
- + L3 access control with IPsec SAD/SPD
- – Potentially more complex Interface ID selection (CGAs, HBAs)
- – Multi-link subnet

#5: “RFC 3456” with DHCPv6 (instead of DHCPv4)

- – RFC 3456 wasn’t really succesful…
- – Multi-link subnet

Next steps

- Editor / second author?
- More discussion