SOCKS Protocol Version 6 (Update)

draft-olteanu-intarea-socks-6-07

Vladimir Olteanu
University Politehnica of Bucharest

New in -07

- Aligned fields
- Address resolution
- Nits galore

Aligned fields

- Refactored messages/options
- Got rid of version minor

- Cleaner/easier to implement (esp. on non-x86)
  - Everything* is aligned after a large recv(), including protocol running on top

  - *Authentication negotiations were not changed and may still cause misalignment. 0-RTT authentication does not cause misalignment.

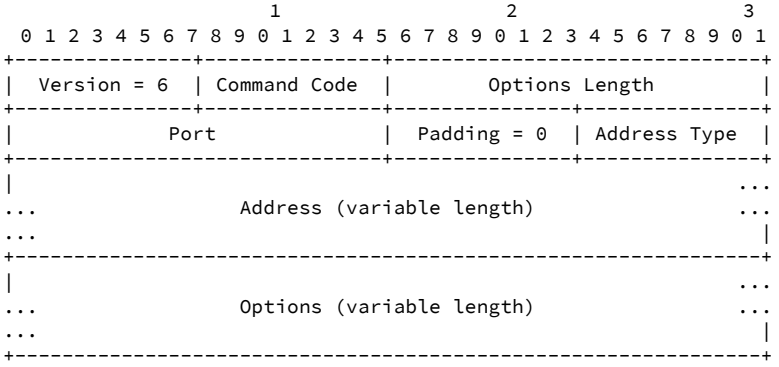

Example: SOCKS v6 Request

Aligned domain names

- Keep existing format, but add padding at the end
- Total length (incl. length byte) MUST be a multiple of 4

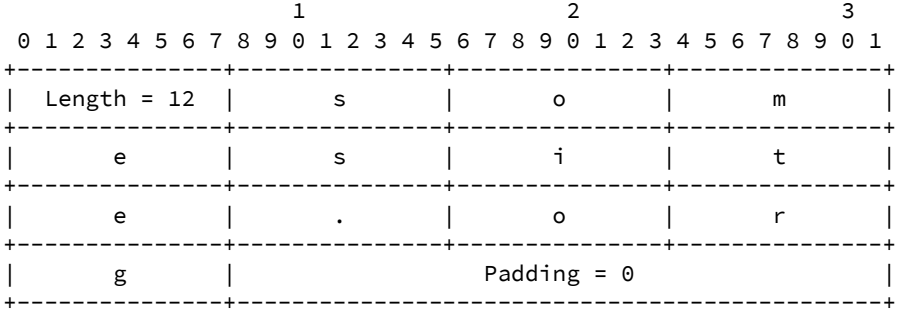

Representation of somesite.org

Option refactoring

- 2-byte Kind field
- Flattened kinds / types / codes
  - Save space
  - Taxonomy useful for socket options, debatable otherwise
  - e.g. Idempotence / Expenditure Reply / Success → Idempotence Accepted

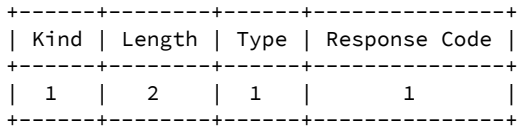

-06: Idempotence / Expenditure Reply / Success

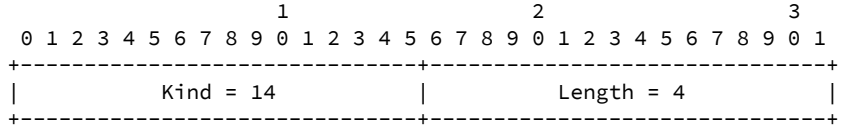

-07: Idempotence Accepted

Address Resolution options

- Have the proxy resolve the address in the Request, relay answer back to client
- Non-standard SOCKS feature used by Tor
- Imitate semantics of gethostbyname() / getaddrinfo()

Address Resolution options

- Client sends Resolution Request
  - Even as part of a NOOP

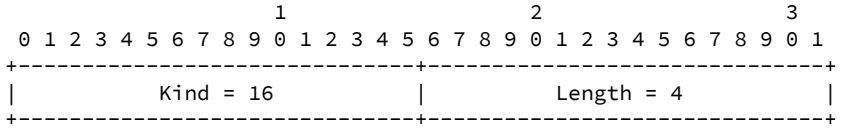

- Proxy replies with {IPv4, IPv6, Domain Name} Resolution options
  - IPv4 + IPv6 if the request contained a domain name
  - Domain name, otherwise

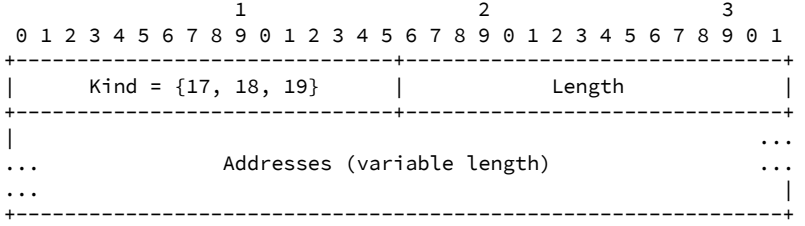

Address Resolution use cases

- UDP
  - Can separately perform address resolution over a separate NOOP (over TCP)
  - Smaller UDP header (no need to embed entire domain name)
    - 4 bytes for an IPv4 address vs. 4+ bytes for the domain name
  - Larger and constant payload size

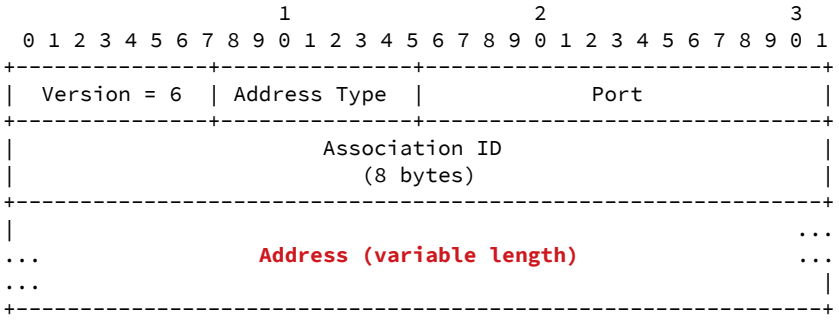

SOCKS v6 UDP Header

Address Resolution use cases

- “Proxified” apps via LD_PRELOAD
  - Intercept socket API calls from a SOCKS-unaware app
  - connect() does not take a domain name
  - Domain name resolved via a separate function call (getaddrinfo() or gethostbyname())
- Motivation
  - Privacy: do not leak DNS requests
  - CDNs: Proxy may have a different vantage point