SUN BINARY ALIGNMENT PROPOSAL TO OASIS KMIP TC

April 24, 2009

Matt Ball
Sun Microsystems, Inc

Proposal Overview

- Sun proposes to change the KMIP draft so that the binary encoding aligns all values on 4-byte boundaries
- This optimizes parsing in hard-aligned processors, like the ARM

The Problem

- The KMIP binary encoding, as proposed in version 0.98, does not align fields to 4-byte boundaries
- This causes big performance penalties on ARM processors, which has to manually break up accesses into 32-bit boundary accesses, or take expensive traps to handle unaligned accesses

Examples from KMIP Proposal 0.98

- An Integer containing the decimal value 8:
  - > 42 00 00 20 | 01 | 00 00 00 04 | 00 00 00 08
- A Long Integer containing the decimal value 123456789000000000:
  - > 42 00 00 20 | 02 | 00 00 00 08 | 01 B6 9B 4B A5 74 92 00
- A Big Integer containing the decimal value 123456789000000000000000000:
  - > 42 00 00 20 | 03 | 00 00 00 0C | 03 FD 35 EB 6B C2 DF 46 18 08 00 00
- An Enumeration with value 255:
  - > 42 00 00 20 | 04 | 00 00 00 04 | 00 00 00 FF

The Proposed Solution

1) Merge the Item Tag and Item Type fields into a single 4-byte word.
   1) Could replace ‘42’ in Tag field with the Type field
   2) Could reduce existing Tag enumerations to 3 bytes and append Type in the remaining byte.
2) Require all Item Value fields to be aligned to 4 bytes.
   1) If the Item Length indicates a non-multiple of 4 bytes, then add pad bytes to reach a 4-byte boundary.
   2) Padding could be zeros, FF, incrementing bytes or other pattern (like 42)

Examples of Aligned Encoding

All examples assume that ‘42’ is replaced with Type field.

- An Integer containing the decimal value 8:
  - 01 | 00 00 20 | 00 00 00 04 | 00 00 00 08
- A Boolean with the value True:
  - 05 | 00 00 20 | 00 00 00 01 | 01 XX XX XX
- A Text String:
  - 06 | 00 00 20 | 00 00 00 0B |
    48 65 6C 6C 6F 20 57 6F 72 6C 64 XX
- An Octet String:
  - 07 | 00 00 20 | 00 00 00 03 | 01 02 03 XX

Case Study: IPsec ESP (RFC 4303)

- Padding ensures payloads are aligned to 32 bits

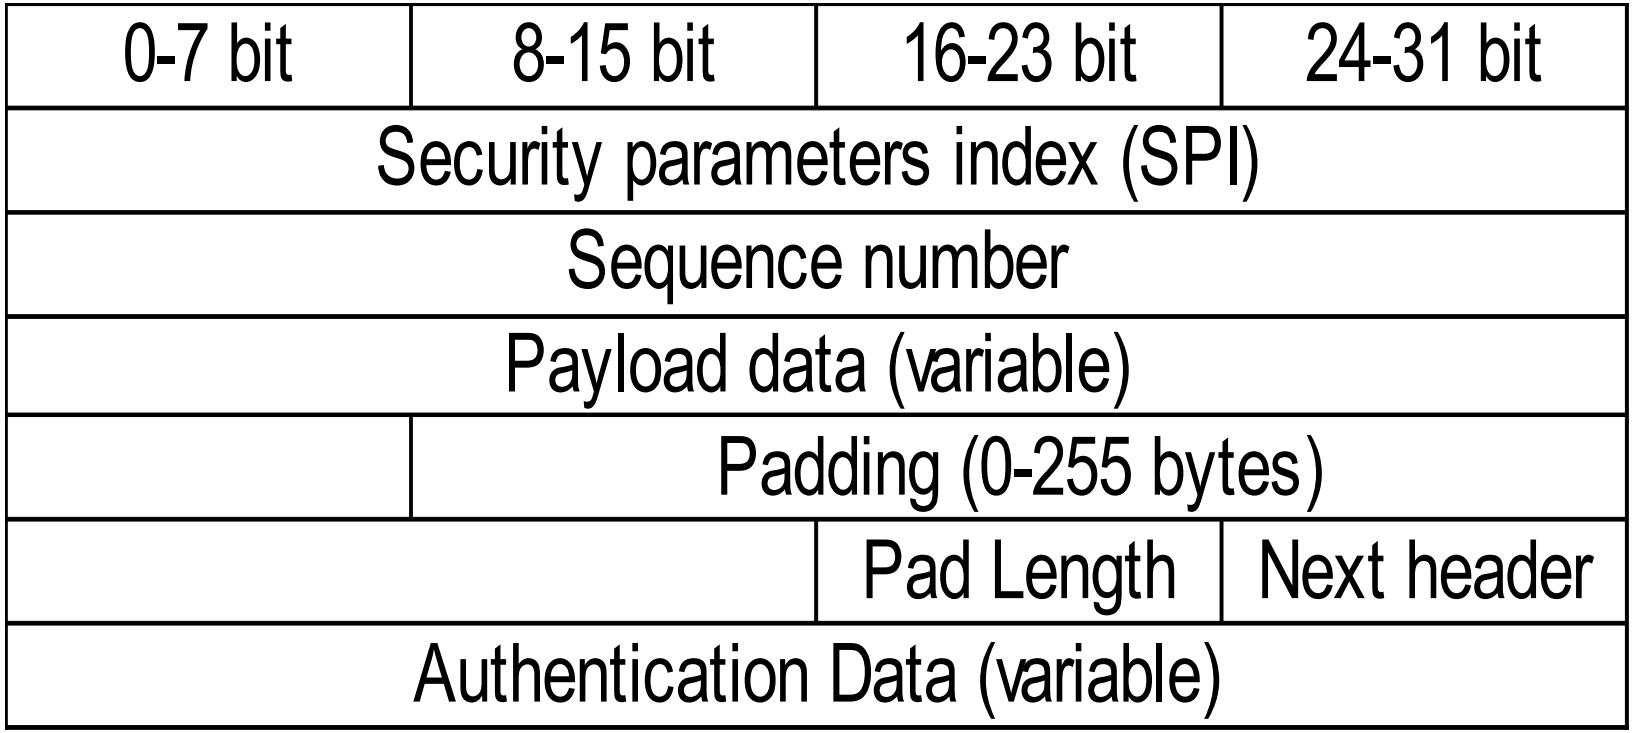

Other Ideas

- Change length of Boolean from 1 to 4 bytes
  - It will get padded to 4 bytes anyway
- Require Big Integers to be multiples of 4 bytes and pad the most significant bytes with zeros
  - Optimizes BIGNUM handling because number is already right-justified

QUESTIONS?

Matt Ball
matthew.ball@sun.com