

Beyond Custom TLVs

Joe Hildebrand

Brian Trammell

Start: two types of types

- *Storage* types can be parsed generically
 - uint64_t
 - UTF8-encoded string
 - Array of bytes
 - Name/value map
- *Semantic* types drive behavior, **always** protocol-specific
 - Source/destination address: 4/16 byte array? String?
 - Hop count: unsigned integer

Today: Below layer 7

- Many protocols have custom Type, Length, Value formats
- “Type” often means both storage type *and* semantic type
- Custom parsing required
- Custom type system required
- Example: IPFIX sourceIPv4Address
 - Storage type is a 32-bit integer IP address in network byte order
 - Semantic type is “Source IPv4 Address”

Offsets	Octet	0							1							2							3										
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Set ID: 2														Set Length: 12																	
4	32	ID: 256														Count: 1																	
8	64	Type: sourceIPv4Address														Length: 4																	
12	96	Set ID: 256														Set Length: 4																	
16	128	192.168.1.1																															

Today: Applications use JSON

- No schema for parsing
- Storage types: bool, number, string, object, array, etc.
- Semantic types: key names in object, position in array, etc.
- Parse internal field structure in same pass
- Extensions as new keys in a key/value struct
- Ignore what you don't understand

```
{  
  "sourceIPv4Address": [192, 168, 1, 1]  
}
```

Why JSON might not be a good fit

- See RFC [7159](#), search for “interop”
 - More edge cases than you think
 - My favorite: 53-bit integers
- Parser more complicated than you expect
 - Larger code size, more CPU
 - Example: String un-escaping
- Binary data requires encoding (such as Base64)
- Larger wire size

Why CBOR might be a better fit

- RFC [7049](#): binary encoding of JSON++
- Small wire size: often smaller than TLV
- Small code size (e.g. 880 bytes of ARM code)
- Lower CPU, latency to parse
- Fixes the known issues of JSON (integers, floats, strings, etc.)
- Binary data first-class type
- Defined diagnostic rendering

```
0xa171736f75726365495076344164647265737344c0a8
0101

  a1          -- Map with 1 pairs
    71        -- Map[0].key: UTF-8 string length
17:
    736f...   "sourceIPv4Address"
    44        -- Map[0].value: Byte string
length 4
c0a80101 -- Bytes content: 192, 168, 1, 1
```

How to make CBOR even more suitable

- Profile out the pieces that you don't need
- Allow parse failures if those features arrive
- Potential removals:
 - Indefinite-length types
 - Tagging (bignums, etc)
 - Floats

Benefits to choosing a single approach

- One set of code
 - Smaller
 - Better code coverage
 - Optimization more likely: e.g. hardware
- Potential security benefits – new syntax sometimes a source of bugs
- Time to market
- Better diagnostic tooling

Suggested topics for discussion

- Could [routing, ops, etc.] protocols use a single approach like this?
- What are the potential downsides?
- Is CBOR a potential encoding?
- Is there a protocol that could be used as an experiment?