CoRE Working Group
Internet-Draft
Intended status: Informational
Expires: November 27, 2012

C. Bormann
K. Hartke
Universitaet Bremen TZI

May 26, 2012

Miscellaneous additions to CoAP
draft-bormann-coap-misc-17

Abstract

This short I-D makes a number of partially interrelated proposals how to solve certain problems in the CoRE WG's main protocol, the Constrained Application Protocol (CoAP). The current version has been resubmitted to keep information about these proposals available; the proposals are not all fleshed out at this point in time.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 27, 2012.

Table of Contents

1. Introduction 4
2. Getting rid of artificial limitations 5
2.1. Option Length encoding beyond 270 bytes 5
3. Vendor-defined Option 8
4. Patience, Leisure, and Pledge 9
4.1. Patience 9
4.2. Leisure 9
4.3. Pledge 10
4.4. Option Formats 10
5. CONNECT 12
5.1. Requesting a Tunnel with CONNECT 12
5.2. Using a CONNECT Tunnel 12
5.3. Closing down a CONNECT Tunnel 13
6. Envelope Options 14
6.1. Example envelope option: solving #230 15
6.2. Example envelope option: proxy-elective options 15
7. Protocol Constants and Time Constants 17
7.1. Protocol Constants 17
7.2. Time Constants derived from Protocol Constants 18
8. IANA Considerations 21
9. Security Considerations 22
10. Acknowledgements 23
11. References 24
11.1. Normative References 24
11.2. Informative References 24
Appendix A. The Nursery (Things that still need to ripen a bit) 26
A.1. Payload-Length Option 26
A.2. URI Authorities with Binary Adresses 26
A.3. Length-aware number encoding (o256) 27
A.4. SMS encoding 29
A.4.1. ASCII-optimized SMS encoding 30
Appendix B. The Cemetery (Things we won't do) 33
B.1. Stateful URI compression 33
B.2. Beyond 270 bytes in a single option 34
B.3. Beyond 15 options 35
B.3.1. Implementation considerations 36
B.3.2. What should we do now? 37
B.3.3. Alternatives 37
B.3.4. Alternative: Going to a delimiter model 37
B.4. Implementing the option delimiter for 15 or more options 37
Appendix C. Experimental Options 39
C.1. Options indicating absolute time 39
C.2. Representing Durations 40
C.3. Rationale 41

C.4. Pseudo-Floating Point 42
C.5. A Duration Type for CoAP 43
Authors' Addresses 50

1. Introduction

The CoRE WG is tasked with standardizing an Application Protocol for Constrained Networks/Nodes, CoAP [I-D.ietf-core-coap]. This protocol is intended to provide RESTful [REST] services not unlike HTTP [RFC2616], while reducing the complexity of implementation as well as the size of packets exchanged in order to make these services useful in a highly constrained network of themselves highly constrained nodes.

This objective requires restraint in a number of sometimes conflicting ways:

- reducing implementation complexity in order to minimize code size,
- reducing message sizes in order to minimize the number of fragments needed for each message (in turn to maximize the probability of delivery of the message), the amount of transmission power needed and the loading of the limited-bandwidth channel,
- reducing requirements on the environment such as stable storage, good sources of randomness or user interaction capabilities.

This draft attempts to address a number of problems not yet adequately solved in [I-D.ietf-core-coap]. The solutions proposed to these problems are somewhat interrelated and are therefore presented in one draft.

The appendix contains the "CoAP cemetery" (possibly later to move into its own draft), documenting roads that the WG decided not to take, in order to spare readers from reinventing them in vain.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The term "byte" is used in its now customary sense as a synonym for "octet".

2. Getting rid of artificial limitations

Artificial limitations are limitations of a protocol or system that are not rooted in limitations of actual capabilities, but in arbitrary design decisions. Proper system design tries to avoid artificial limitations, as these tend to cause complexity in systems that need to work with these limitations.

E.g., the original UNIX filesystem had an artificial limitation of the length of a path name component to 14 bytes. This led to all kinds of workarounds in programs that manipulate file names: E.g., systematically replacing a ".el" extension in a filename with a ".elc" for the compiled file might exceed the limit, so all ".el" files were suddenly limited to 13-byte filenames.

Note that, today, there still is a limitation in most file system implementations, typically at 255. This just happens to be high enough to rarely be of real-world concern; we will refer to this case as a "painless" artificial limitation.

CoAP-08 had two highly recognizable artificial limitations in its protocol encoding

- The number of options in a single message is limited to 15 max.

- The length of an option is limited to 270 max.

It has been argued that the latter limitation causes few problems, just as the 255-byte path name component limitation in filenames today causes few problems. Appendix B.2 provided a design to extend this; as a precaution to future extensions of this kind, the current encoding for length 270 (eight ones in the extension byte) could be marked as reserved today. Since, Matthias Kovatsch has proposed a simpler scheme that seems to gain favor in the WG, see Section 2.1.

The former limitation has been solved in CoAP-09. A historical discussion of other approaches for going beyond 15 options is in Appendix B.3. Appendix B.4 discusses implementation.

2.1. Option Length encoding beyond 270 bytes

For option lengths beyond 270 bytes, we reserve the value 255 of an extension byte to mean "add 255, read another extension byte" Figure 1. While this causes the length of the option header to grow linearly with the size of the option value, only 0.4 % of that size is used. With a focus on short options, this encoding is justified.

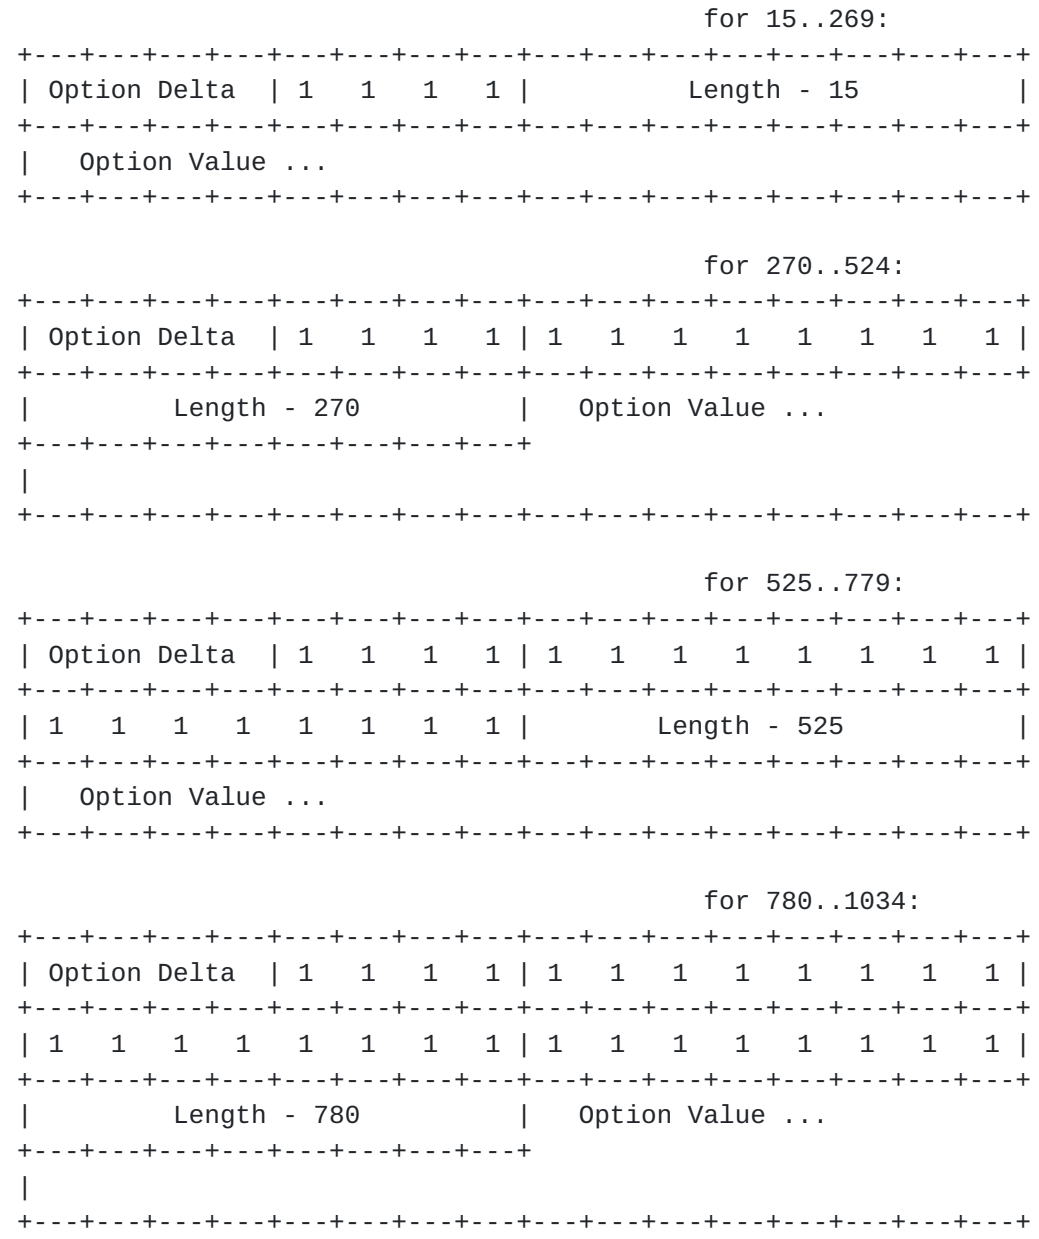

Figure 1: Options beyond 270 bytes

Options that are longer than 1034 bytes MUST NOT be sent; an option that has 255 (all one bits) in the field called "Length - 780" MUST be rejected upon reception as an invalid option.

In the process, the maximum length of all options that are currently set at 270 should now be set to a carefully chosen value. With the purely encoding-based limit gone, Uri-Proxy should now be restored to be a non-repeatable option.

A first proposal for a new set of per-option length restrictions follows:

number	name	min	max	type	repeat
1	content_type	0	2	uint	no
2	max_age	0	4	uint	no
3	proxy_uri	1	1023	string	no
4	etag	1	8	opaque	yes
5	uri_host	1	255	string	no
6	location_path	0	255	string	yes
7	uri_port	0	2	uint	no
8	location_query	0	255	string	yes
9	uri_path	0	255	string	yes
10	observe	0	2	uint	no
11	token	1	8	opaque	no
12	accept	0	2	uint	yes
13	if_match	0	8	opaque	yes
14	vendor	0	1023	opaque	yes
15	uri_query	1	255	string	yes
17	block2	0	3	uint	no
18	size	0	4	uint	no
19	block1	0	3	uint	no
21	if_none_match	0	0	-	no

(Option 14 with a length of 0 is a fencepost only.)

3. Vendor-defined Option

To enable experimentation and community-specific options, option number 14 (the first NOP option) can also be used as a vendor-defined option. For this application, the option value has one or more bytes, the semantics of which are defined by prior agreement between the communicating partners.

It is RECOMMENDED to start the option value with a unique identifier, e.g., the SDNV [RFC5050] of the enterprise number of the organisation defining the option, possibly followed by additional discriminating bits or bytes as defined by the organisation.

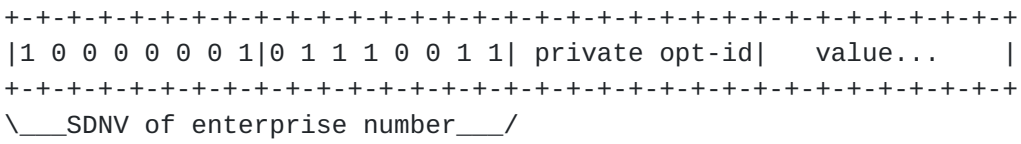

Figure 2: Example option value for vendor-defined option

Note that a Vendor-defined Option cannot be empty, not only because there would be no space for the SDNV, but in particular as the empty option 14 is reserved for fenceposting ([I-D.ietf-core-coap], section 3.2). (Obviously, once a Vendor-defined Option is in use, there is never a need for a fence-post for option number 14.)

Vendor-defined Options are elective.

The Vendor-defined Option is repeatable.

No.	C/E	Name	Format	Length	Default
14	Elective	Vendor	(see below)	1-270 B	(empty)

4. Patience, Leisure, and Pledge

A number of options might be useful for controlling the timing of interactions.

(This section also addresses core-coap ticket #177.)

4.1. Patience

A client may have a limited time period in which it can actually make use of the response for a request. Using the Patience option, it can provide an (elective) indication how much time it is willing to wait for the response from the server, giving the server license to ignore or reject the request if it cannot fulfill it in this period.

If the server knows early that it cannot fulfill the request in the time requested, it MAY indicate this with a 5.04 "Timeout" response. For non-safe methods (such as PUT, POST, DELETE), the server SHOULD indicate whether it has fulfilled the request by either responding with 5.04 "Timeout" (and not further processing the request) or by processing the request normally.

Note that the value of the Patience option should be chosen such that the client will be able to make use of the result even in the presence of the expected network delays for the request and the response. Similarly, when a proxy receives a request with a Patience option and cannot fulfill that request from its cache, it may want to adjust the value of the option before forwarding it to an upstream server.

(TBD: The various cases that arise when combining Patience with Observe.)

The Patience option is elective. Hence, a client MUST be prepared to receive a normal response even after the chosen Patience period (plus an allowance for network delays) has elapsed.

4.2. Leisure

Servers generally will compute an internal value that we will call Leisure, which indicates the period of time that will be used for responding to a request. A Patience option, if present, can be used as an upper bound for the Leisure. Leisure may be non-zero for congestion control reasons, in particular for responses to multicast requests. For these, the server should have a group size estimate G, a target rate R (which both should be chosen conservatively) and an estimated response size S; a rough lower bound for Leisure can then be computed as follows:

```
lb_Leisure = S * G / R
```

Figure 3: Computing a lower bound for the Leisure

E.g., for a multicast request with link-local scope on an 2.4 GHz IEEE 802.15.4 (6LoWPAN) network, G could be (relatively conservatively) set to 100, S to 100 bytes, and the target rate to 8 kbit/s = 1 kB/s. The resulting lower bound for the Leisure is 10 seconds.

To avoid response implosion, responses to multicast requests SHOULD be dithered within a Leisure period chosen by the server to fall between these two bounds.

Currently, we don't foresee a need to signal a value for Leisure from client to server (beyond the signalling provided by Patience) or from server to client, but an appropriate Option might be added later.

4.3. Pledge

In a basic observation relationship [I-D.ietf-core-observe], the server makes a pledge to keep the client in the observation relationship for a resource at least until the max-age for the resource is reached.

To save the client some effort in re-establishing observation relationships each time max-age is reached, the server MAY want to extend its pledge beyond the end of max-age by signalling in a response/notification an additional time period using the Pledge Option, in parallel to the Observe Option.

The Pledge Option MUST NOT be used unless the server can make a reasonable promise not to lose the observation relationship in this time frame.

Currently, we don't foresee a need to signal a value for Pledge from client to server, but an appropriate behavior might be added later for this option when sent in a request.

4.4. Option Formats

No.	C/E	Name	Format	Length	Default
22	Elective	Patience	Duration in mis	1 B	(none)
24	Elective	Pledge	Duration in s	1 B	0

All timing options use the Duration data type (see Appendix C.2), however Patience (and Leisure, if that ever becomes an option) uses a timebase of mibiseconds (mis = 1/1024 s) instead of seconds. (This reduces the range of the Duration from ~ 91 days to 128 minutes.)

Implementation note: As there are no strong accuracy requirements on the clocks employed, making use of any existing time base of milliseconds is a valid implementation approach (2.4 % off).

None of the options may be repeated.

5. CONNECT

[RFC2817] defines the HTTP CONNECT method to establish a TCP tunnel through a proxy so that end-to-end TLS connections can be made through the proxy. Recently, a requirement for similar functionality has been discussed for CoAP. This section defines a strawman CONNECT method and related methods and response codes for CoAP.

(IANA considerations for this section TBD.)

5.1. Requesting a Tunnel with CONNECT

CONNECT is allocated as a new method code in the "CoAP Method Codes" registry. When a client makes a CONNECT request to an intermediary, the intermediary evaluates the Uri-Host, Uri-Port, and/or the authority part of the Proxy-Uri Options in a way that is defined by the security policy of the intermediary. If the security policy allows the allocation of a tunnel based on these parameters, the method returns an empty payload and a response code of 2.30 Tunnel Established. Other possible response codes include 4.03 Forbidden.

It may be the case that the intermediary itself can only reach the requested origin server through another intermediary. In this case, the first intermediary SHOULD make a CONNECT request of that next intermediary, requesting a tunnel to the authority. A proxy MUST NOT respond with any 2.xx status code unless it has either a direct or tunnel connection established to the authority.

An origin server which receives a CONNECT request for itself MAY respond with a 2.xx status code to indicate that a tunnel is established to itself.

Code 2.30 "Tunnel Established" is allocated as a new response code in the "CoAP Response Codes" registry.

5.2. Using a CONNECT Tunnel

Any successful (2.xx) response to a CONNECT request indicates that the intermediary has established a tunnel to the requested host and port. The tunnel is bound to the requesting end-point and the Token supplied in the request (as always, the default Token is admissable). The tunnel can be used by the client by making a DATAGRAM request.

DATAGRAM is allocated as a new method code in the "CoAP Method Codes" registry. When a client makes a DATAGRAM request to an intermediary, the intermediary looks up the tunnel bound to the client end-point and Token supplied in the DATAGRAM request (no other Options are permitted). If a tunnel is found and the intermediary's security

policy permits, the intermediary forwards the payload of the DATAGRAM request as the UDP payload towards the host and port established for the tunnel. No response is defined for this request (note that the request can be given as a CON or NON request; for CON, there will be an ACK on the message layer if the tunnel exists).

The security policy on the intermediary may restrict the allowable payloads based on its security policy, possibly considering host and port. An inadmissable paylaod SHOULD cause a 4.03 Forbidden response with a diagnostic message as payload.

The UDP payload of any datagram received from the tunnel and admitted by the security policy is forwarded to the client as the payload of a 2.31 "Datagram Received" response. The response does not carry any Option except for Token, which identifies the tunnel towards the client.

Code 2.31 "Datagram Received" is allocated as a new response code in the "CoAP Response Codes" registry.

An origin server that has established a tunnel to itself processes the CoAP payloads of related DATAGRAM requests as it would process an incoming UDP payload, and forwards what would be outgoing UDP payloads in 2.31 "Datagram Received" responses.

5.3. Closing down a CONNECT Tunnel

A 2.31 "Datagram Received" response may be replied to with a RST, which closes down the tunnel. Similarly, the Token used in the tunnel may be reused by the client for a different purpose, which also closes down the tunnel.

6. Envelope Options

As of [I-D.ietf-core-coap], options can take one of three types, two of which are mostly identical:

- uint: A non-negative integer which is represented in network byte order using a variable number of bytes (see [I-D.ietf-core-coap] Appendix A);

- string: a sequence of bytes that is nominally a Net-Unicode string [RFC5198];

- opaque: a sequence of bytes.

It turns out some options would benefit from some internal structure. Also, it may be a good idea to be able to bundle multiple options into one, in order to ensure consistency for a set of elective options that need to be processed all or nothing (i.e., the option becomes critical as soon as another option out of the set is processed, too).

In this section, we introduce a fourth CoAP option type: Envelope options.

An envelope option is a sequence of bytes that looks and is interpreted exactly like a CoAP sequence of options. Instead of an option count or an end-of-option marker, the sequence of options is terminated by the end of the envelope option.

The nested options (options inside the envelope option) may come from the same number space as the top-level CoAP options, or the envelope option may define its own number space - this choice needs to be defined for each envelope option.

If the top-level number space is used, the envelope option typically will restrict the set of options that actually can be used in the envelope. In particular, it is unlikely that an envelope option will allow itself inside the envelope (this would be a recursive option).

Envelope options are a general, but simple mechanism. Some of its potential uses are illustrated by two examples below. (Each of these examples also has its own merits and demerits, which should be discussed separately from the concept of Envelope options employed in the examples.)

6.1. Example envelope option: solving #230

Ticket #230 [CoRE230] points out a design flaw of [I-D.ietf-core-coap]: When we split the elective Location option of draft -01 into multiple elective options, we made it possible that an implementation might process some of these and ignore others, leading to an incorrect interpretation of the Location expressed by the server.

There are several more or less savory solutions to #230.

Each of the elective options that together make up the Location could be defined in such a way that it makes a requirement on the processing of the related option (essentially revoking their elective status once the option under consideration is actually processed). This falls flat as soon as another option is defined that would also become part of the Location: existing implementations would not know that the new option is also part of the cluster that is re-interpreted as critical. The potential future addition of Location-Host and Location-Port makes this a valid consideration.

A better solution would be to define an elective Envelope Option called Location. Within a Location Option, the following top-level options might be allowed (now or in the future):

- Uri-Host
- Uri-Port
- Uri-Path
- Uri-Query

This would unify the code for interpreting the top-level request options that indicate the request URI with the code that interprets the Location URI.

The four options listed are all critical, while the envelope is elective. This gives exactly the desired semantics: If the envelope is processed at all (which is elective), the nested options are critical and all need to be processed.

6.2. Example envelope option: proxy-elective options

Another potential application of envelope options is motivated by the observation that new critical options might not be implemented by all proxies on the CoAP path to an origin server. So that this does not become an obstacle to introducing new critical options that are of

interest only to client and origin server, the client might want to mark some critical options proxy-elective, i.e. elective for a proxy but still critical for the origin server.

One way to do this would be an Envelope option, the Proxy-Elective Option. A client might bundle a number of critical options into a critical Proxy-Elective Option. A proxy that processes the message is obliged to process the envelope (or reject the message), where processing means passing on the nested options towards the origin server (preferably again within a Proxy-Elective option). It can pass on the nested options, even ones unknown to the proxy, knowing that the client is happy with proxies not processing all of them.

(The assumption here is that the Proxy-Elective option becomes part of the base standard, so all but the most basic proxies would know how to handle it.)

7. Protocol Constants and Time Constants

CoAP defines a number of time-related protocol constants. There are also a number of assumptions about the timing behavior of the network. This section attempts to provide input to a possible update, addressing #201 [CoRE201].

7.1. Protocol Constants

Section 9 of [I-D.ietf-core-coap] defines three protocol constants:

name	default value
RESPONSE_TIMEOUT	2 seconds
RESPONSE_RANDOM_FACTOR	1.5
MAX_RETRANSMIT	4

Section 9 gives license to a configuration to modify these protocol constants, without specifying a configuration method. It also does not give any further guidance. Both Cullen Jennings (msg03072by) and Michael Scharf (msg03280e) have argued we do need some additional guidance here.

In particular, Michael Scharf notes that

- a decrease of RESPONSE_TIMEOUT below 1 second would violate the guidelines of RFC 5405, and it would not be safe for use in the Internet;

- significantly decreasing the RESPONSE_TIMEOUT would require an adaptive RTT measurement to be robust, see RFC 5405 or draft-allman-tcpm-rto-consider. [...] the draft should strongly discourage any decrease of RESPONSE_TIMOUT until such more advanced mechanisms exist, and it should explain why.

We should therefore add the following text at the end of Section 9:

> The protocol constants have been chosen to achieve a behavior in the presence of congestion that is safe in the Internet. If a configuration desires to use different values, the onus is on the configureation to ensure these congestion control properties are not violated. In particular, a decrease of RESPONSE_TIMEOUT below 1 second would violate the guidelines of [RFC5405]. ([I-D.allman-tcpm-rto-consider] provides some additional

background.) CoAP was designed to enable implementations that do not maintain round-trip-time (RTT) measurements. However, where it is desired to decrease the RESPONSE_TIMEOUT significantly, this can only be done safely when maintaining such measurements. Configurations MUST NOT decrease RESPONSE_TIMEOUT without using mechanisms that ensure congestion control safety, either defined in the configuration or in future standards documents.

RESPONSE_RANDOM_FACTOR MUST NOT be decreased below 1.0, and it SHOULD have a value that is sufficiently different from 1.0 to provide some protection from synchronization effects.

MAX_RETRANSMIT can be freely adjusted, but a too small value will reduce the probability that a confirmed message is actually received, while a larger value will require further adjustments in the time constants (see discussion below).

If the choice of protocol constants leads to an increase of derived time constants (see below), the configuration mechanism MUST ensure the adjusted value is available to the corresponding end-points, too.

7.2. Time Constants derived from Protocol Constants

(This section might become a new section 9.2 in [I-D.ietf-core-coap]. The various places where these calculations are included in the text, e.g., section 4.1, can then just use these numbers.)

The combination of RESPONSE_TIMEOUT, RESPONSE_RANDOM_FACTOR and MAX_RETRANSMIT influences the timing of retransmissions, which in turn influences how long certain information items need to be kept by an implementation. To be able to unambiguously reference these derived time constants, we give them names as follows:

- MAX_TRANSMIT_SPAN is the maximum time from the first transmission of a confirmed message to its last retransmission. For the default protocol constants, the value is (2+4+8+16)*1.5 = 45 seconds, or more generally:

  ```
  RESPONSE_TIMEOUT * (2 ** MAX_RETRANSMIT - 1) *
  RESPONSE_RANDOM_FACTOR
  ```

- MAX_TRANSMIT_WAIT is the maximum time from the first transmission of a confirmed message to the time when the sender gives up on receiving a response. For the default protocol constants, the value is (2+4+8+16+32)*1.5 = 93 seconds, or more generally:

```
RESPONSE_TIMEOUT * (2 ** (MAX_RETRANSMIT + 1) - 1) *
RESPONSE_RANDOM_FACTOR
```

In addition, some assumptions need to be made on the characteristics of the network and the nodes.

- MAX_LATENCY is the maximum time a datagram is expected to take from the start of its transmission to the completion of its reception. This constant is related to the MSL (Maximum Segment Lifetime) of [RFC0793], which is "arbitrarily defined to be 2 minutes" ([RFC0793] glossary, page 81). Note that this is not necessarily smaller than MAX_TRANSMIT_WAIT, as MAX_LATENCY is not intended to describe a situation when the protocol works well, but the worst case situation against which the protocol has to guard. We, also arbitrarily, define MAX_LATENCY to be 100 seconds. Apart from being reasonably realistic for the bulk of configurations as well as close to the historic choice for TCP, this value also allows message ID lifetime timers to be represented in 8 bits (when measured in seconds). In these calculations, there is no assumption that the direction of the transmission is irrelevant (i.e. that the network is symmetric), just that the same value can reasonably be used as a maximum value for both directions. If that is not the case, the following calculations become only slightly more complex.

- PROCESSING_DELAY is the time a node takes to turn around a confirmed message into an acknowledgement. We assume the node will attempt to send an ACK before having the sender time out, so as a conservative assumption we set it equal to RESPONSE_TIMEOUT.

- MAX_RTT is the maximum round-trip time, or:

```
2 * MAX_LATENCY + PROCESSING_DELAY
```

From these constants, we can derive the following values relevant to the protocol operation:

- EXCHANGE_LIFETIME is the time from starting to send a confirmed message to the time when a response is no longer expected, i.e. message layer information about the message exchange can be purged. EXCHANGE_LIFETIME includes a MAX_TRANSMIT_SPAN, a MAX_LATENCY forward, PROCESSING_DELAY, and a MAX_LATENCY for the way back. Note that there is no need to consider MAX_TRANSMIT_WAIT if the configuration is chosen such that the last waiting period (RESPONSE_TIMEOUT * (2 ** MAX_RETRANSMIT) or the difference between MAX_TRANSMIT_SPAN and MAX_TRANSMIT_WAIT) is less than MAX_LATENCY -- which is a likely choice, as MAX_LATENCY is a worst case value unlikely to be met in the real world. In

this case, EXCHANGE_LIFETIME simplifies to:

```
(RESPONSE_TIMEOUT * (2 ** MAX_RETRANSMIT) *
RESPONSE_RANDOM_FACTOR) + (2 * MAX_LATENCY)
```

or 248 seconds with the default protocol constants.

- (others? We need to go through [I-D.ietf-core-coap] and find places where we can substitute in these constants.)

8. IANA Considerations

This draft adds option numbers to Table 2 of [I-D.ietf-core-coap]:

Number	Name	Reference
14	Vendor	[RFCXXXX]
22	Patience	[RFCXXXX]
24	Pledge	[RFCXXXX]

Table 1: New CoAP Option Numbers

9. Security Considerations

TBD.

10. Acknowledgements

This work was partially funded by the Klaus Tschira Foundation.

Of course, much of the content of this draft is the result of discussions with the [I-D.ietf-core-coap] authors.

Patience and Leisure were influenced by a mailing list discussion with Esko Dijk, Kepeng Li, and Salvatore Loreto - thanks!

11. References

11.1. Normative References

[I-D.ietf-core-coap]
Shelby, Z., Hartke, K., Bormann, C., and B. Frank, "Constrained Application Protocol (CoAP)", draft-ietf-core-coap-09 (work in progress), March 2012.

[I-D.ietf-core-observe]
Hartke, K., "Observing Resources in CoAP", draft-ietf-core-observe-05 (work in progress), March 2012.

[I-D.ietf-httpbis-p1-messaging]
Fielding, R., Lafon, Y., and J. Reschke, "HTTP/1.1, part 1: URIs, Connections, and Message Parsing", draft-ietf-httpbis-p1-messaging-19 (work in progress), March 2012.

[I-D.ietf-httpbis-p4-conditional]
Fielding, R., Lafon, Y., and J. Reschke, "HTTP/1.1, part 4: Conditional Requests", draft-ietf-httpbis-p4-conditional-19 (work in progress), March 2012.

[I-D.ietf-httpbis-p6-cache]
Fielding, R., Lafon, Y., Nottingham, M., and J. Reschke, "HTTP/1.1, part 6: Caching", draft-ietf-httpbis-p6-cache-19 (work in progress), March 2012.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.

11.2. Informative References

[CoRE201] "Multiple Location options need to be processed as a unit", CoRE ticket #201, 2012, <http://trac.tools.ietf.org/wg/core/trac/ticket/201>.

[CoRE230] "Multiple Location options need to be processed as a

unit", CoRE ticket #230, 2012,
<http://trac.tools.ietf.org/wg/core/trac/ticket/230>.

[I-D.allman-tcpm-rto-consider]
Allman, M., "Retransmission Timeout Considerations", draft-allman-tcpm-rto-consider-01 (work in progress), May 2012.

[REST] Fielding, R., "Architectural Styles and the Design of Network-based Software Architectures", 2000.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC1924] Elz, R., "A Compact Representation of IPv6 Addresses", RFC 1924, April 1996.

[RFC2817] Khare, R. and S. Lawrence, "Upgrading to TLS Within HTTP/1.1", RFC 2817, May 2000.

[RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.

[RFC5198] Klensin, J. and M. Padlipsky, "Unicode Format for Network Interchange", RFC 5198, March 2008.

[RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, November 2008.

[RFC6256] Eddy, W. and E. Davies, "Using Self-Delimiting Numeric Values in Protocols", RFC 6256, May 2011.

Appendix A. The Nursery (Things that still need to ripen a bit)

A.1. Payload-Length Option

Not all transport mappings may provide an unambiguous length of the CoAP message. For UDP, it may also be desirable to pack more than one CoAP message into one UDP payload (aggregation); in that case, for all but the last message there needs to be a way to delimit the payload of that message.

This can be solved using a new option, the Payload-Length option. If this option is present, the value of this option is an unsigned integer giving the length of the payload of the message (note that this integer can be zero for a zero-length payload, which can in turn be represented by a zero-length option value). (In the UDP aggregation case, what would have been in the payload of this message after "payload-length" bytes is then actually one or more additional messages.)

A.2. URI Authorities with Binary Adresses

One problem with the way URI authorities are represented in the URI syntax is that the authority part can be very bulky if it encodes an IPv6 address in ASCII.

Proposal: Provide an option "Uri-Authority-Binary" that can be an even number of bytes between 2 and 18 except 12 or 14.

- If the number of bytes is 2, the destination IP address of the packet transporting the CoAP message is implied.
- If the number of bytes is 4 or 6, the first four bytes of the option value are an IPv4 address in binary.
- If the number of bytes is 8 or 10, the first eight bytes are the lower 64 bits of an IPv6 address; the upper eight bytes are implied from the destination address of the packet transporting the CoAP message.
- If the number of bytes is 16 or 18, the first 16 bytes are an IPv6 address.
- If two more bytes remain, this is a port number (as always in network byte order).

The resulting authority is (conceptually translated into ASCII and) used in place of an Uri-Authority option, or inserted into a Proxy-Uri. Examples:

Proxy-Uri	Uri-Authority-Bi nary	Uri-Pat h	URI
(none)	(none)	(none)	"/"
(none)	(none)	'temp'	"/temp"
(none)	2 bytes: 61616	'temp'	"coap://[DA]:61616/tem p"
(none)	16 bytes: 2000::1	temp	"coap://[2000::1]/temp "
'http://'	10 bytes: ::123:45 + 616	(none)	"http://[DA::123:45]:6 16"
'http:///te mp'	18 bytes: 2000::1 + 616	(none)	"http://[2000::1]:616/ temp"

A.3. Length-aware number encoding (o256)

The number encoding defined in Appendix A of [I-D.ietf-core-coap] has one significant flaw: Every number has an infinite number of representations, which can be derived by adding leading zero bytes. This runs against the principle of minimizing unnecessary choice. The resulting uncertainty in encoding ultimately leads to unnecessary interoperability failures. (It also wastes a small fraction of the encoding space, i.e., it wastes bytes.)

We could solve the first, but not the second, by outlawing leading zeroes, but then we have to cope with error cases caused by illegal values, another source of interoperability problems.

The number encoding "o256" defined in this section avoids this flaw. The suggestion is not to replace CoAP's "uint" encoding wholesale (CoAP is already too widely implemented for such a change), but to consider this format for new options.

The basic requirements for such an encoding are:

- numbers are encoded as a sequence of zero or more bytes
- each number has exactly one encoding
- for a < b, encoding-size(a) <= encoding-size(b) -- i.e., with larger numbers, the encoding only gets larger, never smaller

again.

- within each encoding size (0 bytes, 1 byte, etc.), lexicographical ordering of the bytes is the same as numeric ordering

Obviously, there is only one encoding that satisfies all these requirements. As illustrated by Figure 4, this is unambiguously derived by

1. enumerating all possible byte sequences, ordered by length and within the same length in lexicographic ordering, and,

2. assigning sequential cardinals.

```
0x'' -> 0
0x'00' -> 1
0x'01' -> 2
0x'02' -> 3
...
0x'fe' -> 255
0x'ff' -> 256
0x'0000' -> 257
0x'0001' -> 258
...
0x'fefd' -> 65534
0x'fefe' -> 65535
0x'feff' -> 65536
...
0x'ffff' -> 65792
0x'000000' -> 65793
0x'000001' -> 65794
```

Figure 4: Enumerating byte sequences by length and then lexicographic order

This results in an exceedingly simple algorithm: each byte is interpreted in the base-256 place-value system, but stands for a number between 1 and 256 instead of 0 to 255. We therefore call this encoding "o256" (one-to-256). 0 is always encoded in zero bytes; 1 to 256 is one byte, 257 (0x101) to 65792 (0x10100) is two bytes, 65793 (0x10101) to 16843008 (0x1010100) is three bytes, etc.

To further illustrate the algorithmic simplicity, pseudocode for encoding and decoding is given in Figure 5 and Figure 6, respectively (in the encoder, "prepend" stands for adding a byte at the _leading_ edge, the requirement for which is a result of the network byte order). Note that this differs only in a single subtraction/addition

(resp.) of one from the canonical algorithm for Appendix A uints.

```
while num > 0
  num -= 1
  prepend(num & 0xFF)
  num >>= 8
end
```

Figure 5: o256 encoder (pseudocode)

```
num = 0
each_byte do |b|
  num <<= 8
  num += b + 1
end
```

Figure 6: o256 decoder (pseudocode)

On a more philosophical note, it can be observed that o256 solves the inverse problem of Self-Delimiting Numeric Values (SDNV) [RFC6256]: SDNV encodes variable-length numbers together with their length (allowing decoding without knowing their length in advance, deriving delimiting information from the number encoding). o256 encodes variable-length numbers when there is a way to separately convey the length (as in CoAP options), encoding (and later deriving) a small part of the numeric value into/from that size information.

A.4. SMS encoding

For use in SMS applications, CoAP messages can be transferred using SMS binary mode. However, there is operational experience showing that some environments cannot successfully send a binary mode SMS.

For transferring SMS in character mode (7-bit characters), base64-encoding [RFC4648] is an obvious choice. 3 bytes of message (24 bits) turn into 4 characters, which cosume 28 bits. The overall overhead is approximately 17 %; the maximum message size is 120 bytes (160 SMS characters).

If a more compact encoding is desired, base85 encoding can be employed (however, probably not the version defined in [RFC1924] -- instead, the version used in tools such as btoa and PDF should be chosen). However, this requires division operations. Also, the base85 character set includes several characters that cannot be transferred in a single 7-bit unit in SMS and/or are known to cause operational problems. A modified base85 character set can be defined to solve the latter problem. 4 bytes of message (32 bits) turn into 5

characters, which consume 35 bits. The overall overhead is approximately 9.3 %; the resulting maximum message size is 128 bytes (160 SMS characters).

Base64 and base85 do not make use of the fact that much CoAP data will be ASCII-based. Therefore, we define the following experimental SMS encoding.

A.4.1. ASCII-optimized SMS encoding

Not all 128 theoretically possible SMS characters are operationally free of problems. We therefore define:

Shunned code characters: @ sign, as it maps to 0x00

LF and CR signs (0x0A, 0x0D)

uppercase C cedilla (0x09), as it is often mistranslated in gateways

ESC (0x1B), as it is used in certain character combinations only

Some ASCII characters cannot be transferred in the base SMS character set, as their code positions are taken by non-ASCII characters. These are simply encoded with their ASCII code positions, e.g., an underscore becomes a section mark (even though underscore has a different code position in the SMS character set).

Equivalently translated input bytes: $, @, [, \,], ^, _, `, {, |, }, ~, DEL

In other words, bytes 0x20 to 0x7F are encoded into the same code positions in the 7-bit character set.

Out of the remaining code characters, the following SMS characters are available for encoding:

Non-equivalently translated (NET) code characters: 0x01 to 0x08, (8 characters)

0x0B, 0x0C, (2 characters)

0x0E to 0x1A, (13 characters)

0x1C to 0x1F, (4 characters)

Of the 27 NET code characters, 18 are taken as prefix characters (see below), and 8 are defined as directly translated characters:

Directly translated bytes: Equivalently translated input bytes are represented as themselves

0x00 to 0x07 are represented as 0x01 to 0x08

This leaves 0x08 to 0x1F and 0x80 to 0xFF. Of these, the bytes 0x80 to 0x87 and 0xA0 to 0xFF are represented as the bytes 0x00 to 0x07 (represented by characters 0x01 to 0x08) and 0x20 to 0x7F, with a prefix of 1 (see below). The characters 0x08 to 0x1F are represented as the characters 0x28 to 0x3F with a prefix of 2 (see below). The characters 0x88 to 0x9F are represented as the characters 0x48 to 0x5F with a prefix of 2 (see below). (Characters 0x01 to 0x08, 0x20 to 0x27, 0x40 to 0x47, and 0x60 to 0x7f with a prefix of 2 are reserved for future extensions, which could be used for some backreferencing or run-length compression.)

Bytes that do not need a prefix (directly translated bytes) are sent as is. Any byte that does need a prefix (i.e., 1 or 2) is preceded by a prefix character, which provides a prefix for this and the following two bytes as follows:

0x0B	100		0x15	200
0x0C	101		0x16	201
0x0E	102		0x17	202
0x0F	110		0x18	210
0x10	111		0x19	211
0x11	112		0x1A	212
0x12	120		0x1C	220
0x13	121		0x1D	221
0x14	122		0x1E	222

(This leaves one non-shunned character, 0x1F, for future extension.)

The coding overhead of this encoding for random bytes is similar to Base85, without the need for a division/multiplication. For bytes that are mostly ASCII characters, the overhead can easily become negative. (Conversely, for bytes that are more likely to be non-ASCII than in a random sequence of bytes, the overhead becomes

greater.)

So, for instance, for the CoAP message in Figure 7:

```
ver     tt      code    mid
1       ack     2.05    17033
content_type    40
token           sometok
3c 2f 3e 3b 74 69 74 6c 65 3d 22 47 65 6e 65 72  |</>;title="Gener|
61 6c 20 49 6e 66 6f 22 3b 63 74 3d 30 2c 3c 2f  |al Info";ct=0,</|
74 69 6d 65 3e 3b 69 66 3d 22 63 6c 6f 63 6b 22  |time>;if="clock"|
3b 72 74 3d 22 54 69 63 6b 73 22 3b 74 69 74 6c  |;rt="Ticks";titl|
65 3d 22 49 6e 74 65 72 6e 61 6c 20 43 6c 6f 63  |e="Internal Cloc|
6b 22 3b 63 74 3d 30 2c 3c 2f 61 73 79 6e 63 3e  |k";ct=0,</async>|
3b 63 74 3d 30                                   |;ct=0           |
```

Figure 7: CoAP response message as captured and decoded

The 116 byte unencoded message is shown as ASCII characters in Figure 8 (\xDD stands for the byte with the hex digits DD):

```
bEB\x89\x11(\xA7sometok</>;title="General Info";ct=0,</time>
;if="clock";rt="Ticks";title="Internal Clock";ct=0,</async>;ct=0
```

Figure 8: CoAP response message shown as unencoded characters

The equivalent SMS encoding is shown as equivalent-coded SMS characters in Figure 9 (7 bits per character, \x12 is a 220 prefix and \x0B is a 100 prefix, the rest is shown in equivalent encoding), adding two characters of prefix overhead, for a total length of 118 7-bit characters or 104 (103.25 plus padding) bytes:

```
bEB\x12I1(\x0B'sometok</>;title="General Info";ct=0,</time>
;if="clock";rt="Ticks";title="Internal Clock";ct=0,</async>;ct=0
```

Figure 9: CoAP response message shown as SMS-encoded characters

Appendix B. The Cemetery (Things we won't do)

This annex documents roads that the WG decided not to take, in order to spare readers from reinventing them in vain.

B.1. Stateful URI compression

Is the approximately 25 % average saving achievable with Huffman-based URI compression schemes worth the complexity? Probably not, because much higher average savings can be achieved by introducing state.

Henning Schulzrinne has proposed for a server to be able to supply a shortened URI once a resource has been requested using the full-length URI. Let's call such a shortened referent a _Temporary Resource Identifier_, _TeRI_ for short. This could be expressed by a response option as shown in Figure 10.

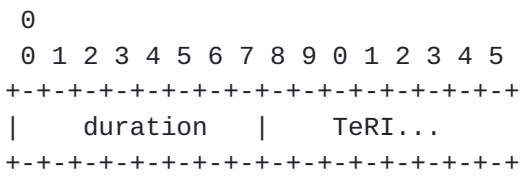

Figure 10: Option for offering a TeRI in a response

The TeRI offer option indicates that the server promises to offer this resources under the TeRI given for at least the time given as the duration. Another TeRI offer can be made later to extend the duration.

Once a TeRI for a URI is known (and still within its lifetime), the client can supply a TeRI instead of a URI in its requests. The same option format as an offer could be used to allow the client to indicate how long it believes the TeRI will still be valid (so that the server can decide when to update the lifetime duration). TeRIs in requests could be distinguished from URIs e.g. by using a different option number.

Proposal: Add a TeRI option that can be used in CoAP requests and responses.

Add a way to indicate a TeRI and its duration in a link-value.

Do not add any form of stateless URI encoding.

Benefits: Much higher reduction of message size than any stateless URI encoding could achieve.

As the use of TeRIs is entirely optional, minimal complexity nodes can get by without implementing them.

Drawbacks: Adds considerable state and complexity to the protocol.

It turns out that real CoAP URIs are short enough that TeRIs are not needed.

(Discuss the security implications of TeRIs.)

B.2. Beyond 270 bytes in a single option

The authors would argue that 270 as the maximum length of an option is already beyond the "painless" threshold.

If that is not the consensus of the WG, the scheme can easily be extended as in Figure 11:

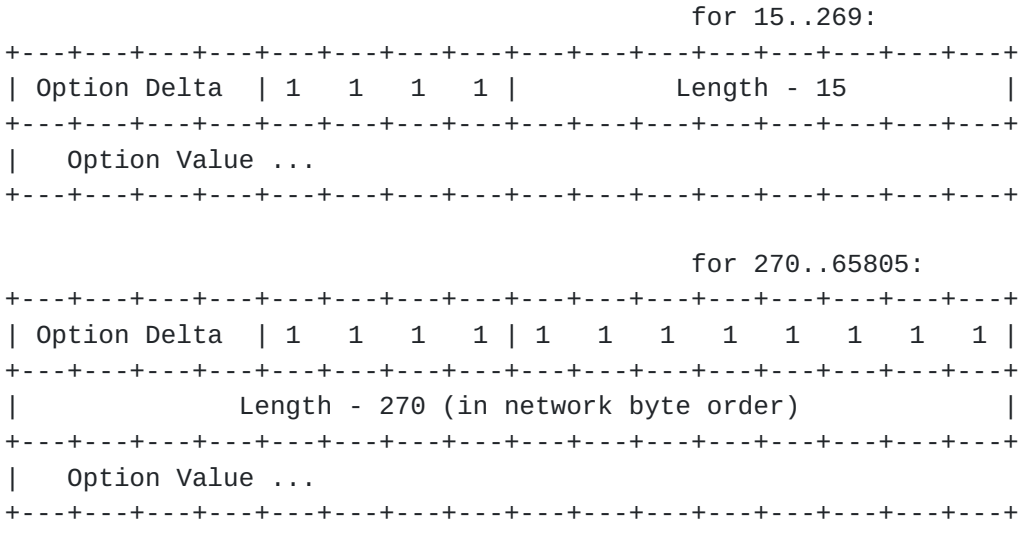

Figure 11: Ridiculously Long Option Header

The infinite number of obvious variations on this scheme are left as an exercise to the reader.

Again, as a precaution to future extensions, the current encoding for length 270 (eight ones in the extension byte) could be marked as reserved today.

B.3. Beyond 15 options

(This section keeps discussion that is no longer needed as we have agreed to do what is documented in Appendix B.4).

The limit of 15 options is motivated by the fixed four-bit field "OC" that is used for indicating the number of options in the fixed-length CoAP header (Figure 12).

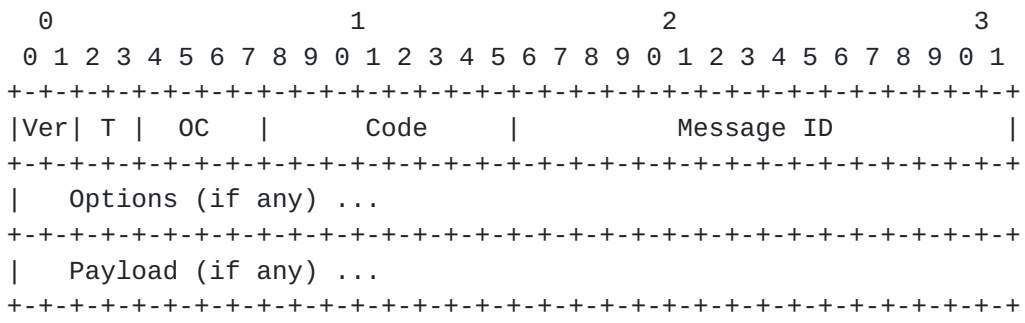

Figure 12: Four-byte fixed header in a CoAP Message

Note that there is another fixed four-bit field in CoAP: the option length (Figure 13 - note that this figure is not to the same scale as the previous figure):

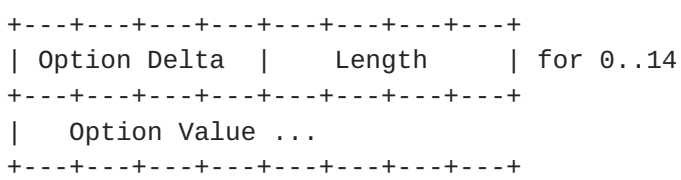

Figure 13: Short Option Header

Since 15 is inacceptable for a maximum option length, the all-ones value (15) was taken out of the set of allowable values for the short header, and a long header was introduced that allows the insertion of an extension byte (Figure 14):

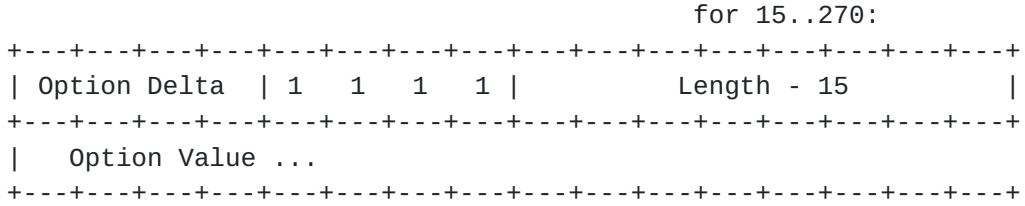

Figure 14: Long Option Header

We might want to use the same technique for the CoAP header as well. There are two obvious places where the extension byte could be placed:

1. right after the byte carrying the OC field, so the structure is the same as for the option header;

2. right after the fixed-size CoAP header.

Both solutions lose the fixed-size-ness of the CoAP header.

Solution 1 has the disadvantage that the CoAP header is also changing in structure: The extension byte is wedged between the first and the second byte of the CoAP header. This is unfortunate, as the number of options only comes into play when the option processing begins, so it is more natural to use solution 2 (Figure 15):

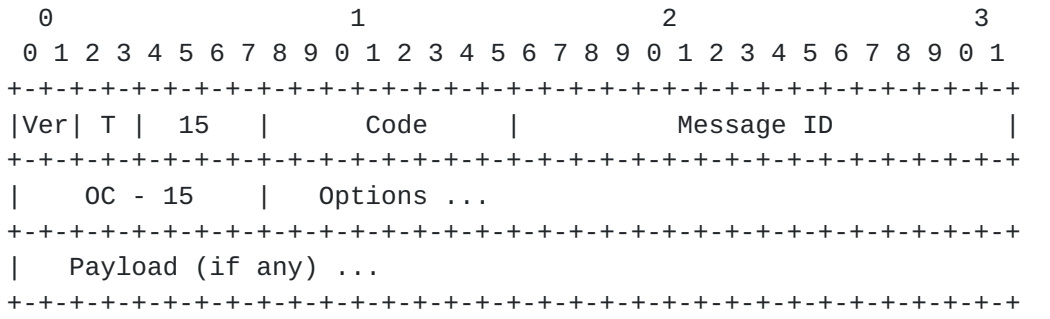

Figure 15: Extended header for CoAP Messages with 15+ options

This would allow for up to 270 options in a CoAP message, which is very likely way beyond the "painless" threshold.

B.3.1. Implementation considerations

For a message decoder, this extension creates relatively little pain, as the number of options only becomes interesting when the encoding turns to the options part of the message, which is then simply lead in by the extension byte if the four-bit field is 15.

For a message encoder, this extension is not so rosy. If the encoder is constructing the message serially, it may not know in advance whether the number of options will exceed 14. None of the following implementation strategies is particularly savory, but all of them do work:

1. Encode the options serially under the assumption that the number of options will be 14 or less. When the 15th option needs to be encoded, abort the option encoding, and restart it from scratch

one byte further to the left.

2. Similar to 1, except that the bytes already encoded are all moved one byte to right, the extension byte is inserted, and the option encoding process is continued.

3. The encoder always leaves space for the extension byte (at least if it can't prove the number will be less thatn 14). If the extension byte is not needed, an Option 0 with length 0 is encoded instead (i.e., one byte is wasted - this option is elective and will be ignored by the receiver).

As a minimum, to enable strategy 3, the option 0 should be reserved at least for the case of length=0.

B.3.2. What should we do now?

As a minimum proposal for the next version of CoAP, the value 15 for OC should be marked as reserved today.

B.3.3. Alternatives

One alternative that has been discussed previously is to have an "Options" Option, which allows the carriage of multiple options in the belly of a single one. This could also be used to carry more than 15 options. However:

- The conditional introduction of an Options option has implementation considerations that are likely to be more severe than the ones listed above;

- since 270 bytes may not be enough for the encoding of _all_ options, the "Options" option would need to be repeatable. This creates many different ways to encode the same message, leading to combinatorial explosion in test cases for ensuring interoperability.

B.3.4. Alternative: Going to a delimiter model

Another alternative is to spend the additional byte not as an extended count, but as an option terminator.

B.4. Implementing the option delimiter for 15 or more options

Implementation note: As can be seen from the proof of concept code in Figure 16, the actual implementation cost for a decoder is around 4 lines of code (or about 8-10 machine code instructions).

```
while numopt > 0
  nextbyte = ... get next byte

  if numopt == 15                  # new
    break if nextbyte == 0xF0      # new
  else                             # new
    numopt -= 1
  end                              # new

  ... decode the delta and length from nextbyte and handle them
end
```

Figure 16: Implementing the Option Terminator

Similarly, creating the option terminator needs about four more lines (not marked "old" in the C code in Figure 17).

```
b0 = 0x40 + (tt << 4);             /* old */
buffer[0] = b0 + 15;               /* guess first byte */

    ....  encode options ....      /* old */

if (option_count >= 15 || first_fragment_already_shipped)
   buffer[pos++] = 0xF0;           /* use delimiter */
else                               /* save a byte: */
   buffer[0] = b0 + option_count;  /* old: backpatch */
```

Figure 17: Creating the Option Terminator

Appendix C. Experimental Options

This annex documents proposals that need significant additional discussion before they can become part of (or go back to) the main CoAP specification. They are not dead, but might die if there turns out to be no good way to solve the problem.

C.1. Options indicating absolute time

HTTP has a number of headers that may indicate absolute time:

- "Date", defined in Section 14.18 in [RFC2616] (Section 9.3 in [I-D.ietf-httpbis-p1-messaging]), giving the absolute time a response was generated;
- "Last-Modified", defined in Section 14.29 in [RFC2616], (Section 6.6 in [I-D.ietf-httpbis-p4-conditional], giving the absolute time of when the origin server believes the resource representation was last modified;
- "If-Modified-Since", defined in Section 14.25 in [RFC2616], "If-Unmodified-Since", defined in Section 14.28 in [RFC2616], and "If-Range", defined in Section 14.27 in [RFC2616] can be used to supply absolute time to gate a conditional request;
- "Expires", defined in Section 14.21 in [RFC2616] (Section 3.3 in [I-D.ietf-httpbis-p6-cache]), giving the absolute time after which a response is considered stale.
- The more obscure headers "Retry-After", defined in Section 14.37 in [RFC2616], and "Warning", defined in section 14.46 in [RFC2616], also may employ absolute time.

[I-D.ietf-core-coap] defines a single "Date" option, which however "indicates the creation time and date of a given resource representation", i.e., is closer to a "Last-Modified" HTTP header. HTTP's caching rules [I-D.ietf-httpbis-p6-cache] make use of both "Date" and "Last-Modified", combined with "Expires". The specific semantics required for CoAP needs further consideration.

In addition to the definition of the semantics, an encoding for absolute times needs to be specified.

In UNIX-related systems, it is customary to indicate absolute time as an integer number of seconds, after midnight UTC, January 1, 1970. Unless negative numbers are employed, this time format cannot represent time values prior to January 1, 1970, which probably is not required for the uses ob absolute time in CoAP.

If a 32-bit integer is used and allowance is made for a sign-bit in a local implementation, the latest UTC time value that can be represented by the resulting 31 bit integer value is 03:14:07 on January 19, 2038. If the 32-bit integer is used as an unsigned value, the last date is 2106-02-07, 06:28:15.

The reach can be extended by: - moving the epoch forward, e.g. by 40 years (= 1262304000 seconds) to 2010-01-01. This makes it impossible to represent Last-Modified times in that past (such as could be gatewayed in from HTTP). - extending the number of bits, e.g. by one more byte, either always or as one of two formats, keeping the 32-bit variant as well.

Also, the resolution can be extended by expressing time in milliseconds etc., requiring even more bits (e.g., a 48-bit unsigned integer of milliseconds would last well after year 9999.)

For experiments, an experimental "Date" option is defined with the semantics of HTTP's "Last-Modified". It can carry an unsigned integer of 32, 40, or 48 bits; 32- and 40-bit integers indicate the absolute time in seconds since 1970-01-01 00:00 UTC, while 48-bit integers indicate the absolute time in milliseconds since 1970-01-01 00:00 UTC.

However, that option is not really that useful until there is a "If-Modified-Since" option as well.

(Also: Discuss nodes without clocks.)

C.2. Representing Durations

Various message types used in CoAP need the representation of *durations*, i.e. of the length of a timespan. In SI units, these are measured in seconds. CoAP durations represent integer numbers of seconds, but instead of representing these numbers as integers, a more compact single-byte pseudo-floating-point (pseudo-FP) representation is used (Figure 18).

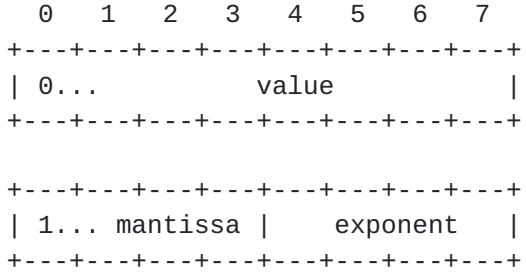

Figure 18: Duration in (8,4) pseudo-FP representation

If the high bit is clear, the entire n-bit value (including the high bit) is the decoded value. If the high bit is set, the mantissa (including the high bit, with the exponent field cleared out but still present) is shifted left by the exponent to yield the decoded value.

The (n,e)-pseudo-FP format can be decoded with a single line of code (plus a couple of constant definitions), as demonstrated in Figure 19.

```
#define N 8
#define E 4
#define HIBIT (1 << (N - 1))
#define EMASK ((1 << E) - 1)
#define MMASK ((1 << N) - 1 - EMASK)

#define DECODE_8_4(r) (r < HIBIT ? r : (r & MMASK) << (r & EMASK))
```

Figure 19: Decoding an (8,4) pseudo-FP value

Note that a pseudo-FP encoder needs to consider rounding; different applications of durations may favor rounding up or rounding down the value encoded in the message.

The highest pseudo-FP value, represented by an all-ones byte (0xFF), is reserved to indicate an indefinite duration. The next lower value (0xEF) is thus the highest representable value and is decoded as 7340032 seconds, a little more than 12 weeks.

C.3. Rationale

Where CPU power and memory is abundant, a duration can almost always be adequately represented by a non-negative floating-point number representing that number of seconds. Historically, many APIs have also used an integer representation, which limits both the resolution (e.g., if the integer represents the duration in seconds) and often the range (integer machine types have range limits that may become relevant). UNIX's "time_t" (which is used for both absolute time and durations) originally was a signed 32-bit value of seconds, but was later complemented by an additional integer to add microsecond ("struct timeval") and then later nanosecond ("struct timespec") resolution.

Three decisions need to be made for each application of the concept of duration:

- the *resolution*. What rounding error is acceptable?
- the *range*. What is the maximum duration that needs to be represented?
- the *number of bits* that can be expended.

Obviously, these decisions are interrelated. Typically, a large range needs a large number of bits, unless resolution is traded. For most applications, the actual requirement for resolution are limited for longer durations, but can be more acute for shorter durations.

C.4. Pseudo-Floating Point

Constrained systems typically avoid the use of floating-point (FP) values, as

- simple CPUs often don't have support for floating-point datatypes
- software floating-point libraries are expensive in code size and slow.

In addition, floating-point datatypes used to be a significant element of market differentiation in CPU design; it has taken the industry a long time to agree on a standard floating point representation.

These issues have led to protocols that try to constrain themselves to integer representation even where the ability of a floating point representation to trade range for resolution would be beneficial.

The idea of introducing _pseudo-FP_ is to obtain the increased range provided by embedding an exponent, without necessarily getting stuck with hardware datatypes or inefficient software floating-point libraries.

For the purposes of this draft, we define an (n,e)-pseudo-FP as a fixed-length value of n bits, e of which may be used for an exponent. Figure 18 illustrates an (8,4)-pseudo-FP value.

If the high bit is clear, the entire n-bit value (including the high bit) is the decoded value. If the high bit is set, the mantissa (including the high bit, but with the exponent field cleared out) is shifted left by the exponent to yield the decoded value.

The (n,e)-pseudo-FP format can be decoded with a single line of code (plus a couple of constant definition), as demonstrated in Figure 19.

Only non-negative numbers can be represented by this format. It is designed to provide full integer resolution for values from 0 to 2^(n-1)-1, i.e., 0 to 127 in the (8,4) case, and a mantissa of n-e bits from 2^(n-1) to (2^n-2^e)*2^(2^e-1), i.e., 128 to 7864320 in the (8,4) case. By choosing e carefully, resolution can be traded against range.

Note that a pseudo-FP encoder needs to consider rounding; different applications of durations may favor rounding up or rounding down the value encoded in the message. This requires a little more than a single line of code (which is left as an exercise to the reader, as the most efficient expression depends on hardware details).

C.5. A Duration Type for CoAP

CoAP needs durations in a number of places. In [I-D.ietf-core-coap], durations occur in the option "Subscription-lifetime" as well as in the option "Max-age". (Note that the option "Date" is not a duration, but a point in time.) Other durations of this kind may be added later.

Most durations relevant to CoAP are best expressed with a minimum resolution of one second. More detailed resolutions are unlikely to provide much benefit.

The range of lifetimes and caching ages are probably best kept below the order of magnitude of months. An (8,4)-pseudo-FP has the maximum value of 7864320, which is about 91 days; this appears to be adequate for a subscription lifetime and probably even for a maximum cache age. Figure 20 shows the values that can be expressed. (If a larger range for the latter is indeed desired, an (8,5)-pseudo-FP could be used; this would last 15 milleniums, at the cost of having only 3 bits of accuracy for values larger than 127 seconds.)

Proposal: A single duration type is used throughout CoAP, based on an (8,4)-pseudo-FP giving a duration in seconds.

Benefits: Implementations can use a single piece of code for managing all CoAP-related durations.

In addition, length information never needs to be managed for durations that are embedded in other data structures: All durations are expressed by a single byte.

It might be worthwhile to reserve one duration value, e.g. 0xFF, for an indefinite duration.

Duration	Seconds	Encoded

00:00:00	0x00000000	0x00
00:00:01	0x00000001	0x01
00:00:02	0x00000002	0x02
00:00:03	0x00000003	0x03
00:00:04	0x00000004	0x04
00:00:05	0x00000005	0x05
00:00:06	0x00000006	0x06
00:00:07	0x00000007	0x07
00:00:08	0x00000008	0x08
00:00:09	0x00000009	0x09
00:00:10	0x0000000a	0x0a
00:00:11	0x0000000b	0x0b
00:00:12	0x0000000c	0x0c
00:00:13	0x0000000d	0x0d
00:00:14	0x0000000e	0x0e
00:00:15	0x0000000f	0x0f
00:00:16	0x00000010	0x10
00:00:17	0x00000011	0x11
00:00:18	0x00000012	0x12
00:00:19	0x00000013	0x13
00:00:20	0x00000014	0x14
00:00:21	0x00000015	0x15
00:00:22	0x00000016	0x16
00:00:23	0x00000017	0x17
00:00:24	0x00000018	0x18
00:00:25	0x00000019	0x19
00:00:26	0x0000001a	0x1a
00:00:27	0x0000001b	0x1b
00:00:28	0x0000001c	0x1c
00:00:29	0x0000001d	0x1d
00:00:30	0x0000001e	0x1e
00:00:31	0x0000001f	0x1f
00:00:32	0x00000020	0x20
00:00:33	0x00000021	0x21
00:00:34	0x00000022	0x22
00:00:35	0x00000023	0x23
00:00:36	0x00000024	0x24
00:00:37	0x00000025	0x25
00:00:38	0x00000026	0x26
00:00:39	0x00000027	0x27
00:00:40	0x00000028	0x28
00:00:41	0x00000029	0x29
00:00:42	0x0000002a	0x2a
00:00:43	0x0000002b	0x2b
00:00:44	0x0000002c	0x2c
00:00:45	0x0000002d	0x2d
00:00:46	0x0000002e	0x2e

```
00:00:47  0x0000002f  0x2f
00:00:48  0x00000030  0x30
00:00:49  0x00000031  0x31
00:00:50  0x00000032  0x32
00:00:51  0x00000033  0x33
00:00:52  0x00000034  0x34
00:00:53  0x00000035  0x35
00:00:54  0x00000036  0x36
00:00:55  0x00000037  0x37
00:00:56  0x00000038  0x38
00:00:57  0x00000039  0x39
00:00:58  0x0000003a  0x3a
00:00:59  0x0000003b  0x3b
00:01:00  0x0000003c  0x3c
00:01:01  0x0000003d  0x3d
00:01:02  0x0000003e  0x3e
00:01:03  0x0000003f  0x3f
00:01:04  0x00000040  0x40
00:01:05  0x00000041  0x41
00:01:06  0x00000042  0x42
00:01:07  0x00000043  0x43
00:01:08  0x00000044  0x44
00:01:09  0x00000045  0x45
00:01:10  0x00000046  0x46
00:01:11  0x00000047  0x47
00:01:12  0x00000048  0x48
00:01:13  0x00000049  0x49
00:01:14  0x0000004a  0x4a
00:01:15  0x0000004b  0x4b
00:01:16  0x0000004c  0x4c
00:01:17  0x0000004d  0x4d
00:01:18  0x0000004e  0x4e
00:01:19  0x0000004f  0x4f
00:01:20  0x00000050  0x50
00:01:21  0x00000051  0x51
00:01:22  0x00000052  0x52
00:01:23  0x00000053  0x53
00:01:24  0x00000054  0x54
00:01:25  0x00000055  0x55
00:01:26  0x00000056  0x56
00:01:27  0x00000057  0x57
00:01:28  0x00000058  0x58
00:01:29  0x00000059  0x59
00:01:30  0x0000005a  0x5a
00:01:31  0x0000005b  0x5b
00:01:32  0x0000005c  0x5c
00:01:33  0x0000005d  0x5d
00:01:34  0x0000005e  0x5e
```

00:01:35	0x0000005f	0x5f
00:01:36	0x00000060	0x60
00:01:37	0x00000061	0x61
00:01:38	0x00000062	0x62
00:01:39	0x00000063	0x63
00:01:40	0x00000064	0x64
00:01:41	0x00000065	0x65
00:01:42	0x00000066	0x66
00:01:43	0x00000067	0x67
00:01:44	0x00000068	0x68
00:01:45	0x00000069	0x69
00:01:46	0x0000006a	0x6a
00:01:47	0x0000006b	0x6b
00:01:48	0x0000006c	0x6c
00:01:49	0x0000006d	0x6d
00:01:50	0x0000006e	0x6e
00:01:51	0x0000006f	0x6f
00:01:52	0x00000070	0x70
00:01:53	0x00000071	0x71
00:01:54	0x00000072	0x72
00:01:55	0x00000073	0x73
00:01:56	0x00000074	0x74
00:01:57	0x00000075	0x75
00:01:58	0x00000076	0x76
00:01:59	0x00000077	0x77
00:02:00	0x00000078	0x78
00:02:01	0x00000079	0x79
00:02:02	0x0000007a	0x7a
00:02:03	0x0000007b	0x7b
00:02:04	0x0000007c	0x7c
00:02:05	0x0000007d	0x7d
00:02:06	0x0000007e	0x7e
00:02:07	0x0000007f	0x7f
00:02:08	0x00000080	0x80
00:02:24	0x00000090	0x90
00:02:40	0x000000a0	0xa0
00:02:56	0x000000b0	0xb0
00:03:12	0x000000c0	0xc0
00:03:28	0x000000d0	0xd0
00:03:44	0x000000e0	0xe0
00:04:00	0x000000f0	0xf0
00:04:16	0x00000100	0x81
00:04:48	0x00000120	0x91
00:05:20	0x00000140	0xa1
00:05:52	0x00000160	0xb1
00:06:24	0x00000180	0xc1
00:06:56	0x000001a0	0xd1
00:07:28	0x000001c0	0xe1

```
00:08:00  0x000001e0  0xf1
00:08:32  0x00000200  0x82
00:09:36  0x00000240  0x92
00:10:40  0x00000280  0xa2
00:11:44  0x000002c0  0xb2
00:12:48  0x00000300  0xc2
00:13:52  0x00000340  0xd2
00:14:56  0x00000380  0xe2
00:16:00  0x000003c0  0xf2
00:17:04  0x00000400  0x83
00:19:12  0x00000480  0x93
00:21:20  0x00000500  0xa3
00:23:28  0x00000580  0xb3
00:25:36  0x00000600  0xc3
00:27:44  0x00000680  0xd3
00:29:52  0x00000700  0xe3
00:32:00  0x00000780  0xf3
00:34:08  0x00000800  0x84
00:38:24  0x00000900  0x94
00:42:40  0x00000a00  0xa4
00:46:56  0x00000b00  0xb4
00:51:12  0x00000c00  0xc4
00:55:28  0x00000d00  0xd4
00:59:44  0x00000e00  0xe4
01:04:00  0x00000f00  0xf4
01:08:16  0x00001000  0x85
01:16:48  0x00001200  0x95
01:25:20  0x00001400  0xa5
01:33:52  0x00001600  0xb5
01:42:24  0x00001800  0xc5
01:50:56  0x00001a00  0xd5
01:59:28  0x00001c00  0xe5
02:08:00  0x00001e00  0xf5
02:16:32  0x00002000  0x86
02:33:36  0x00002400  0x96
02:50:40  0x00002800  0xa6
03:07:44  0x00002c00  0xb6
03:24:48  0x00003000  0xc6
03:41:52  0x00003400  0xd6
03:58:56  0x00003800  0xe6
04:16:00  0x00003c00  0xf6
04:33:04  0x00004000  0x87
05:07:12  0x00004800  0x97
05:41:20  0x00005000  0xa7
06:15:28  0x00005800  0xb7
06:49:36  0x00006000  0xc7
07:23:44  0x00006800  0xd7
07:57:52  0x00007000  0xe7
```

```
          08:32:00  0x00007800  0xf7
          09:06:08  0x00008000  0x88
          10:14:24  0x00009000  0x98
          11:22:40  0x0000a000  0xa8
          12:30:56  0x0000b000  0xb8
          13:39:12  0x0000c000  0xc8
          14:47:28  0x0000d000  0xd8
          15:55:44  0x0000e000  0xe8
          17:04:00  0x0000f000  0xf8
          18:12:16  0x00010000  0x89
          20:28:48  0x00012000  0x99
          22:45:20  0x00014000  0xa9
       1d 01:01:52  0x00016000  0xb9
       1d 03:18:24  0x00018000  0xc9
       1d 05:34:56  0x0001a000  0xd9
       1d 07:51:28  0x0001c000  0xe9
       1d 10:08:00  0x0001e000  0xf9
       1d 12:24:32  0x00020000  0x8a
       1d 16:57:36  0x00024000  0x9a
       1d 21:30:40  0x00028000  0xaa
       2d 02:03:44  0x0002c000  0xba
       2d 06:36:48  0x00030000  0xca
       2d 11:09:52  0x00034000  0xda
       2d 15:42:56  0x00038000  0xea
       2d 20:16:00  0x0003c000  0xfa
       3d 00:49:04  0x00040000  0x8b
       3d 09:55:12  0x00048000  0x9b
       3d 19:01:20  0x00050000  0xab
       4d 04:07:28  0x00058000  0xbb
       4d 13:13:36  0x00060000  0xcb
       4d 22:19:44  0x00068000  0xdb
       5d 07:25:52  0x00070000  0xeb
       5d 16:32:00  0x00078000  0xfb
       6d 01:38:08  0x00080000  0x8c
       6d 19:50:24  0x00090000  0x9c
       7d 14:02:40  0x000a0000  0xac
       8d 08:14:56  0x000b0000  0xbc
       9d 02:27:12  0x000c0000  0xcc
       9d 20:39:28  0x000d0000  0xdc
      10d 14:51:44  0x000e0000  0xec
      11d 09:04:00  0x000f0000  0xfc
      12d 03:16:16  0x00100000  0x8d
      13d 15:40:48  0x00120000  0x9d
      15d 04:05:20  0x00140000  0xad
      16d 16:29:52  0x00160000  0xbd
      18d 04:54:24  0x00180000  0xcd
      19d 17:18:56  0x001a0000  0xdd
      21d 05:43:28  0x001c0000  0xed
```

```
      22d 18:08:00  0x001e0000  0xfd
      24d 06:32:32  0x00200000  0x8e
      27d 07:21:36  0x00240000  0x9e
      30d 08:10:40  0x00280000  0xae
      33d 08:59:44  0x002c0000  0xbe
      36d 09:48:48  0x00300000  0xce
      39d 10:37:52  0x00340000  0xde
      42d 11:26:56  0x00380000  0xee
      45d 12:16:00  0x003c0000  0xfe
      48d 13:05:04  0x00400000  0x8f
      54d 14:43:12  0x00480000  0x9f
      60d 16:21:20  0x00500000  0xaf
      66d 17:59:28  0x00580000  0xbf
      72d 19:37:36  0x00600000  0xcf
      78d 21:15:44  0x00680000  0xdf
      84d 22:53:52  0x00700000  0xef
      91d 00:32:00  0x00780000  0xff (reserved)
```

Figure 20

Authors' Addresses

Carsten Bormann
Universitaet Bremen TZI
Postfach 330440
Bremen D-28359
Germany

Phone: +49-421-218-63921
Fax: +49-421-218-7000
Email: cabo@tzi.org

Klaus Hartke
Universitaet Bremen TZI
Postfach 330440
Bremen D-28359
Germany

Phone: +49-421-218-63905
Fax: +49-421-218-7000
Email: hartke@tzi.org