

Thing-to-Thing Research Group
Internet-Draft
Intended status: Informational
Expires: April 25, 2019

K. Hartke
Ericsson
October 22, 2018

CoRE Applications
draft-hartke-core-apps-08

Abstract

The application programmable interfaces of RESTful, hypermedia-driven Web applications consist of a number of reusable components such as Internet media types and link relation types. This document proposes "CoRE Applications", a convention for application designers to build the interfaces of their applications in a structured way, so that implementers can easily build interoperable clients and servers, and other designers can reuse the components in their own applications.

Note to Readers

This Internet-Draft should be discussed on the Thing-to-Thing Research Group (T2TRG) mailing list <t2trg@irtf.org> <<https://www.irtf.org/mailman/listinfo/t2trg>>.

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 <https://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 April 25, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	CoRE Applications	3
2.1.	Communication Protocols	4
2.1.1.	URI Schemes	4
2.2.	Representation Formats	5
2.2.1.	Internet Media Types	5
2.3.	Links	7
2.3.1.	Link Relation Types	8
2.3.2.	Template Variable Names	8
2.4.	Forms	8
2.4.1.	Form Relation Types	9
2.4.2.	Form Field Names	9
2.5.	Well-Known Locations	10
3.	CoRE Application Descriptions	10
3.1.	Template	11
4.	URI Design Considerations	12
5.	Security Considerations	14
6.	IANA Considerations	14
7.	References	14
7.1.	Normative References	14
7.2.	Informative References	15
	Acknowledgements	16
	Author's Address	17

1. Introduction

Representational State Transfer (REST) [16] is an architectural style for distributed hypermedia systems. Over the years, REST has gained popularity not only as an approach for large-scale information dissemination, but also as the basic principle for designing and building Internet-based applications in general.

In the coming years, the size and scope of the Internet is expected to increase greatly as physical-world objects become smart enough to communicate over the Internet -- a phenomenon known as the Internet of Things (IoT). As things learn to speak the languages of the net,

the idea of applying REST principles to the design of IoT application architectures suggests itself. To this end, the Constrained Application Protocol (CoAP) [23] was created, an application-layer protocol that enables RESTful applications in constrained-node networks [10], giving rise to a new setting for Internet-based applications: the Constrained RESTful Environment (CoRE).

To realize the full benefits and advantages of the REST architectural style, a set of constraints needs to be maintained when designing applications and their application programming interfaces (APIs). One of the fundamental principles of REST is that "REST APIs must be hypertext-driven" [17]. However, this principle is often ignored by application designers. Instead, APIs are specified out-of-band in terms of fixed URI patterns (e.g., in the API documentation or in a machine-readable format that facilitates code generation). Although this approach may appear easy for clients to use, the fixed resource names and data formats lead to a tight coupling between client and server implementations and make the system less flexible [5]. Violations of REST design principles like this result in APIs that may not be as scalable, extensible, and interoperable as promised by REST.

REST is intended for network-based applications that are long-lived and span multiple organizations [17]. Principled REST APIs require some design effort, since application designers do not only have to take current requirements into consideration, but also have to anticipate changes that may be required in the future -- years or even decades after the application has been deployed for the first time. The reward is long-term stability and evolvability, both of which are very desirable features in the Internet of Things.

To aid application designers in the design process, this document proposes "CoRE Applications", a convention for building the APIs of RESTful, hypermedia-driven Web applications. The goal is to help application designers avoid common mistakes by focusing almost all of the descriptive effort on defining the Internet media type(s) that are used for representing resources and driving application state.

A template for a "CoRE Application Description" provides a consistent format for the description of APIs so that implementers can easily build interoperable clients and servers, and other application designers can reuse the components in their own applications.

2. CoRE Applications

A CoRE Application API is a named set of reusable components. It describes a contract between a server hosting an instance of the

described application and clients that wish to interface with that instance.

The API is generally comprised of:

- o communication protocols, identified by URI schemes,
- o representation formats, identified by Internet media types,
- o link relation types,
- o form relation types,
- o template variables in templated links,
- o form field names in forms, and
- o well-known locations.

Together, these components provide the specific, in-band instructions to a client for interfacing with a given application.

2.1. Communication Protocols

The foundation of a hypermedia-driven REST API are the communication protocol(s) spoken between a client and a server. Although HTTP/1.1 [14] is by far the most common communication protocol for REST APIs, a REST API should typically not be dependent on any specific communication protocol.

2.1.1. URI Schemes

The usage of a particular protocol by a client is guided by URI schemes [7]. URI schemes specify the syntax and semantics of URI references [1] that the server includes in hypermedia controls such as links and forms.

A URI scheme refers to a family of protocols, typically distinguished by a version number. For example, the "http" URI scheme refers to the two members of the HTTP family of protocols: HTTP/1.1 [14] and HTTP/2 [8] (as well as some predecessors). The specific HTTP version used is negotiated between a client and a server in-band using the version indicator in the HTTP request-line or the TLS Application-Layer Protocol Negotiation (ALPN) extension [18].

IANA maintains a list of registered URI schemes at <http://www.iana.org/assignments/uri-schemes>.

2.2. Representation Formats

In RESTful applications, clients and servers exchange representations that capture the current or intended state of a resource and that are labeled with a media type. A representation is a sequence of bytes whose structure and semantics are specified by a representation format: a set of rules for encoding information.

Representation formats should generally allow clients with different goals, so they can do different things with the same data. The specification of a representation format "describes a problem space, not a prescribed relationship between client and server. Client and server must share an understanding of the representations they're passing back and forth, but they don't need to have the same idea of what the problem is that needs to be solved." [21]

Representation formats and their specifications frequently evolve over time. It is part of the responsibility of the designer of a new version to insure both forward and backward compatibility: new representations should work reasonably (with some fallback) with old processors and old representations should work reasonably with new processors [20].

Representation formats enable hypermedia-driven applications when they support the expression of hypermedia controls such as links (Section 2.3) and forms (Section 2.4).

2.2.1. Internet Media Types

One of the most important aspect of hypermedia-driven communications is the concept of Internet media types [2]. Media types are used to label representations so that it is known how the representation should be interpreted and how it is encoded. The centerpiece of a CoRE Application Description should be one or more media types.

Note: The terms media type and representation format are often used interchangeably. In this document, the term "media type" refers specifically to a string of characters such as "application/xml" that is used to label representations; the term "representation format" refers to the definition of the syntax and semantics of representations, such as XML 1.0 [12] or XML 1.1 [13].

A media type identifies a versioned series of representation formats (Section 2.2): a media type does not identify a particular version of a representation format; rather, the media type identifies the family, and includes provisions for version indicator(s) embedded in the representations themselves to determine more precisely the nature

of how the data is to be interpreted [20]. A new media type is only needed to designate a completely incompatible format [20].

Media types consist of a top-level type and a subtype, structured into trees [2]. Optionally, media types can have parameters. For example, the media type "text/plain; charset=utf-8" is a subtype for plain text under the "text" top-level type in the standards tree and has a parameter "charset" with the value "utf-8".

Media types can be further refined by

- o structured type name suffixes (e.g., "+xml" appended to the base subtype name; see Section 4.2.8 of RFC 6838 [2]),
- o a "profile" parameter (see Section 3.1 of RFC 6906 [24]),
- o subtype information embedded in the representations themselves (e.g., "xmlns" declarations in XML documents [11]),

or a similar annotation. An annotation directly in the media type is generally preferable, since subtype information embedded in representations can typically not be negotiated during content negotiation (e.g., using the CoAP Accept option).

In CoAP, media types are paired with a content coding [15] to indicate the "content format" [23] of a representation. Each content format is assigned a numeric identifier that can be used instead of the (more verbose) textual name of the media type in representation formats with size constraints. The flat number space loses the structural information that the textual names have, however.

The media type of a representation must be determined from in-band information (e.g., from the CoAP Content-Format option). Clients must not assume a structure from the application context or other out-of-band information.

IANA maintains a list of registered Internet media types at <http://www.iana.org/assignments/media-types>.

IANA maintains a list of registered structured suffixes at <http://www.iana.org/assignments/media-type-structured-suffix>.

IANA maintains a list of registered CoAP content formats at <http://www.iana.org/assignments/core-parameters>.

2.3. Links

As defined in RFC 8288 [6], a link is a typed connection between two resources. Additionally, a link is the primary means for a client to navigate from one resource to another.

A link is comprised of:

- o a link context,
- o a link relation type that identifies the semantics of the link (see Section 2.3.1),
- o a link target, identified by a URI, and
- o optionally, target attributes that further describe the link or the link target.

A link can be viewed as a statement of the form "{link context} has a {link relation type} resource at {link target}, which has {target attributes}" [6]. For example, the resource <http://example.com/> could have a "terms-of-service" resource at <http://example.com/tos>, which has a representation with the media type "text/html".

There are two special kinds of links:

- o An embedding link is a link with an additional hint: when the link is processed, it should be substituted with the representation of the referenced resource rather than cause the client to navigate away from the current resource. Thus, traversing an embedding link adds to the current state rather than replacing it.

The most well known example for an embedding link is the HTML element. When a Web browser processes this element, it automatically dereferences the "src" and renders the resulting image in place of the element.

- o A templated link is a link where the client constructs the link target URI from provided in-band instructions. The specific rules for such instructions are described by the representation format. URI Templates [3] provide a generic way to construct URIs through variable expansion.

Templated links allow a client to construct resource URIs without being coupled to the resource structure at the server, provided that the client learns the template from a representation sent by the server and does not have the template hard-coded.

2.3.1. Link Relation Types

A link relation type identifies the semantics of a link [6]. For example, a link with the relation type "copyright" indicates that the resource identified by the target URI is a statement of the copyright terms applying to the link context.

Relation types are not to be confused with media types; they do not identify the format of the representation that results when the link is dereferenced [6]. Rather, they only describe how the link context is related to another resource [6].

IANA maintains a list of registered link relation types at <http://www.iana.org/assignments/link-relations>.

Applications that don't wish to register a link relation type can use an extension link relation type [6]: a URI that uniquely identifies the link relation type. For example, an application can use the string "http://example.com/foo" as link relation type without having to register it. Using a URI to identify an extension link relation type, rather than a simple string, reduces the probability of different link relation types using the same identifiers.

2.3.2. Template Variable Names

A templated link enables clients to construct the target URI of a link, for example, when the link refers to a space of resources rather than a single resource. The most prominent mechanisms for this are URI Templates [3] and the HTML <form> element with a submission method of GET.

To enable an automated client to construct an URI reference from a URI Template, the name of the variable in the template can be used to identify the semantics of the variable. For example, when retrieving the representation of a collection of temperature readings, a variable named "threshold" could indicate the variable for setting a threshold of the readings to retrieve.

Template variable names are scoped to link relation types, i.e., two variables with the same name can have different semantics if they appear in links with different link relation types.

2.4. Forms

A form is the primary means for a client to submit information to a server, typically in order to change resource state.

A form is comprised of:

- o a form context,
- o a form relation type that identifies the semantics of the form (see Section 2.4.1),
- o a request method (e.g., PUT, POST, DELETE),
- o a submission URI,
- o a description of a representation that the server expects as part of the form submission, and
- o optionally, target attributes that further describe the form or the form target.

A form can be viewed as an instruction of the form "To perform a {form relation type} operation on {form context}, make a {request method} request to {submission URI}, which has {target attributes}". For example, to "update" the resource <http://example.com/config>, a client would make a PUT request to <http://example.com/config>. (In many cases, the target of a form is the same resource as the context, but this is not required.)

The description of the expected representation can be a set of form fields (see Section 2.4.2) or simply a list of acceptable media types.

Note: A form with a submission method of GET is, strictly speaking, a templated link, since it provides a way to construct a URI and does not submit a representation to the server.

2.4.1. Form Relation Types

A form relation type identifies the semantics of a form. For example, a form with the form relation type "create" indicates that a new item can be created within the form context by making a request to the resource identified by the target URI.

Similarly to extension link relation types, applications can use extension form relation types when they don't wish to register a form relation type.

2.4.2. Form Field Names

Forms can have a detailed description of the representation expected by the server as part of form submission. This description typically consists of a set of form fields where each form field is comprised

of a field name, a field type, and optionally a number of attributes such as a default value, a validation rule or a human-readable label.

To enable an automated client to fill out a form, the field name can be used to identify the semantics of the form field. For example, when controlling a smart light bulb, the field name "brightness" could indicate the field for setting the desired brightness of the light bulb.

Field names are scoped to form relation types, i.e., two form fields with the same name can have different semantics if they appear in forms with different form relation types.

The type of a form field is a data type such as "an integer between 1 and 100" or "an RGB color". The type is orthogonal to the field name, i.e., the type should not be determined from the field name even though the client can identify the semantics of the field from the name. This separation makes it easy to change the set of acceptable values in the future.

2.5. Well-Known Locations

Some applications may require the discovery of information about a host, known as "site-wide metadata" in RFC 5785 [4]. For example, RFC 6415 [19] defines a metadata document format for describing a host; similarly, RFC 6690 [22] defines a link format for the discovery of resources hosted by a server.

Applications that need to define a resource for this kind of metadata can register new "well-known locations". RFC 5785 [4] defines the path prefix `"/.well-known/"` in "http" and "https" URIs for this purpose. RFC 7252 [23] extends this convention to "coap" and "coaps" URIs.

IANA maintains a list of registered well-known URIs at <http://www.iana.org/assignments/well-known-uris>.

3. CoRE Application Descriptions

As applications are implemented and deployed, it becomes important to describe them in some structured way. This section provides a simple template for CoRE Application Descriptions. A uniform structure allows implementers to easily determine the components that make up the interface of an application.

The template below lists all components of applications that both the client and the server implementation of the application need to understand in order to interoperate. Crucially, items not listed in

the template are not part of the contract between clients and servers -- they are implementation details. This includes in particular the URIs of resources (see Section 4).

CoRE Application Descriptions are intended to be published in human-readable format by designers of applications and by operators of deployed application instances. Application designers may publish an application description as a general specification of all application instances, so that implementers can create interoperable clients and servers. Operators of application instances may publish an application description as part of the API documentation of the service, which should also include instructions how the service can be located and which communication protocols and security modes are used.

3.1. Template

The fields of the template are as follows:

Application name:

Name of the application. The name is not used to negotiate capabilities; it is purely informational. A name may include a version number or, for example, refer to a living standard that is updated continuously.

URI schemes:

URI schemes identifying the communication protocols that need to be understood by clients and servers. This information is mostly relevant for deployed instances of the application rather than for the general specification of the application.

Media types:

Internet media types that identify the representation formats that need to be understood by clients and servers. An application description must comprise at least one media type. Additional media types may be required or optional.

Link relation types:

Link relation types that identify the semantics of links. An application description may comprise IANA-registered link relation types and extension link relation types. Both may be required or optional.

Template variable names:

For each link relation type, variable names that identify the semantics of variables in templated links with that link relation type. Whether a template variable is required or optional is indicated in-band inside the templated link.

Form relation types:

Form relation types that identify the semantics of forms and, for each form relation type, the submission method(s) to be used. An application description may comprise IANA-registered form relation types and extension form relation types. Both may be required or optional.

Form field names:

For each form relation type, form field names that identify the semantics of form fields in forms with that form relation type. Whether a form field is required or optional is indicated in-band inside the form.

Well-known locations:

Well-known locations in the resource identifier space of servers that clients can use to discover information given the DNS name or IP address of a server.

Interoperability considerations:

Any issues regarding the interoperable use of the components of the application should be given here.

Security considerations:

Security considerations for the security of the application must be specified here.

Contact:

Person (including contact information) to contact for further information.

Author/Change controller:

Person (including contact information) authorized to change this application description.

Each field should include full citations for all specifications necessary to understand the application components.

4. URI Design Considerations

URIs [1] are a cornerstone of RESTful applications. They enable uniform identification of resources via URI schemes [7] and are used every time a client interacts with a particular resource or when a resource representation references another resource.

URIs often include structured application data in the path and query components, such as paths in a filesystem or keys in a database. It is common for many RESTful applications to use these structures not only as an implementation detail but also make them part of the

public REST API, prescribing a fixed format for this data. However, there are a number of problems with this practice [5], in particular if the application designer and the server owner are not the same entity.

In hypermedia-driven applications, URIs are therefore not included in the application interface. A CoRE Application Description must not mandate any particular form of URI substructure.

RFC 7320 [5] describes the problematic practice of fixed URI structures in detail and provides some acceptable alternatives.

Nevertheless, the design of the URI structure on a server is an essential part of implementing a RESTful application, even though it is not part of the application interface. The server implementer is responsible for binding the resources identified by the application designer to URIs.

A good RESTful URI is:

- o Short. Short URIs are easier to remember and cause less overhead in requests and representations.
- o Meaningful. A URI should describe the resource in a way that is meaningful and useful to humans.
- o Consistent. URIs should follow a consistent pattern to make it easy to reason about the application.
- o Bookmarkable. Cool URIs don't change [9]. However, in practice, application resource structures do change. That should cause URIs to change as well so they better reflect reality. Implementations should not depend on unchanging URIs.
- o Shareable. A URI should not be context sensitive, e.g., to the currently logged-in user. It should be possible to share a URI with third parties so they can access the same resource.
- o Extension-less. Some applications return different data for different extensions, e.g., for "contacts.xml" or "contacts.json". But different URIs imply different resources. RESTful URIs should identify a single resource. Different representations of the resource can be negotiated (e.g., using the CoAP Accept option).

5. Security Considerations

The security considerations of RFC 3986 [1], RFC 5785 [4], RFC 6570 [3], RFC 6838 [2], RFC 7320 [5], RFC 7595 [7], and RFC 8288 [6] are inherited.

All components of an application description are expected to contain clear security considerations. CoRE Application Descriptions should furthermore contain security considerations that need to be taken into account for the security of the overall application.

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

- [1] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [2] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/info/rfc6838>>.
- [3] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", RFC 6570, DOI 10.17487/RFC6570, March 2012, <<https://www.rfc-editor.org/info/rfc6570>>.
- [4] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known Uniform Resource Identifiers (URIs)", RFC 5785, DOI 10.17487/RFC5785, April 2010, <<https://www.rfc-editor.org/info/rfc5785>>.
- [5] Nottingham, M., "URI Design and Ownership", BCP 190, RFC 7320, DOI 10.17487/RFC7320, July 2014, <<https://www.rfc-editor.org/info/rfc7320>>.
- [6] Nottingham, M., "Web Linking", RFC 8288, DOI 10.17487/RFC8288, October 2017, <<https://www.rfc-editor.org/info/rfc8288>>.

- [7] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", BCP 35, RFC 7595, DOI 10.17487/RFC7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.

7.2. Informative References

- [8] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <<https://www.rfc-editor.org/info/rfc7540>>.
- [9] Berners-Lee, T., "Cool URIs don't change", 1998, <<http://www.w3.org/Provider/Style/URI.html>>.
- [10] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [11] Bray, T., Hollander, D., Layman, A., Tobin, R., and H. Thompson, "Namespaces in XML 1.0 (Third Edition)", World Wide Web Consortium Recommendation REC-xml-names-20091208, December 2009, <<http://www.w3.org/TR/2009/REC-xml-names-20091208>>.
- [12] Bray, T., Paoli, J., Sperberg-McQueen, M., Maler, E., and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)", World Wide Web Consortium Recommendation REC-xml-20081126, November 2008, <<http://www.w3.org/TR/2008/REC-xml-20081126>>.
- [13] Bray, T., Paoli, J., Sperberg-McQueen, M., Maler, E., Yergeau, F., and J. Cowan, "Extensible Markup Language (XML) 1.1 (Second Edition)", World Wide Web Consortium Recommendation REC-xml11-20060816, August 2006, <<http://www.w3.org/TR/2006/REC-xml11-20060816>>.
- [14] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [15] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.

- [16] Fielding, R., "Architectural Styles and the Design of Network-based Software Architectures", Ph.D. Dissertation, University of California, Irvine, 2000, <http://www.ics.uci.edu/~fielding/pubs/dissertation/fielding_dissertation.pdf>.
- [17] Fielding, R., "REST APIs must be hypertext-driven", October 2008, <<http://roy.gbiv.com/untangled/2008/rest-apis-must-be-hypertext-driven>>.
- [18] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <<https://www.rfc-editor.org/info/rfc7301>>.
- [19] Hammer-Lahav, E., Ed. and B. Cook, "Web Host Metadata", RFC 6415, DOI 10.17487/RFC6415, October 2011, <<https://www.rfc-editor.org/info/rfc6415>>.
- [20] Masinter, L., "MIME and the Web", draft-masinter-mime-web-info-02 (work in progress), January 2011.
- [21] Richardson, L. and M. Amundsen, "RESTful Web APIs", O'Reilly Media, ISBN 978-1-4493-5806-8, September 2013.
- [22] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [23] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [24] Wilde, E., "The 'profile' Link Relation Type", RFC 6906, DOI 10.17487/RFC6906, March 2013, <<https://www.rfc-editor.org/info/rfc6906>>.

Acknowledgements

Jan Algermissen, Mike Amundsen, Mike Kelly, Julian Reschke, and Erik Wilde provided valuable input on link and form relation types.

Thanks to Olaf Bergmann, Carsten Bormann, Stefanie Gerdes, Ari Keranen, Michael Koster, Matthias Kovatsch, Teemu Savolainen, and Bilhanan Silverajan for helpful comments and discussions that have shaped the document.

Some of the text in this document has been borrowed from [5], [6], [17], and [20]. All errors are my own.

This work was funded in part by Nokia.

Author's Address

Klaus Hartke
Ericsson
Torshamnsgatan 23
Stockholm SE-16483
Sweden

Email: klaus.hartke@ericsson.com

Thing-to-Thing Research Group
Internet-Draft
Intended status: Experimental
Expires: May 4, 2017

K. Hartke
Universitaet Bremen TZI
October 31, 2016

CBOR-encoded Form Data
draft-hartke-t2trg-cbor-forms-00

Abstract

This document describes a media type to encode form data in CBOR format.

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 May 4, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	2
2. Form Body Syntax	2
3. Form Data Syntax	3
4. Security Considerations	3
5. IANA Considerations	4
5.1. Media Type	4
5.2. CoAP Content-Format	4
6. References	4
6.1. Normative References	4
6.2. Informative References	4
Author's Address	5

1. Introduction

This document describes a media type to encode form data in CBOR [RFC7049] format, similar to the well-known "application/x-www-form-urlencoded" [W3C.REC-html5-20141028] and "multipart/form-data" [RFC7578] media types. The use of a compact, binary representation format enables the processing of form submissions on systems with very limited memory, processor power, and instruction sets [RFC7228].

1.1. Terminology

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

2. Form Body Syntax

CoRAL [I-D.hartke-t2trg-coral] provides a representation format for forms that is suitable for constrained systems. However, it does not provide a syntax for specifying the fields of a form in the form body and leaves this task to other specifications. This section defines such a syntax.

A form body in this syntax consists of one or more form fields. Each form field consists of a field name and a type. The field name is an arbitrary text string. The field type is a text string that matches the ABNF rule "type" from [I-D.greevenbosch-appsawg-cbor-cddl].

Form fields can be combined with two combinators, ALL and ANY. The ALL combinator indicates that the form submitter MUST submit all the members of the expression. The ANY combinator indicates that the form submitter MUST submit exactly one member of the expression.

Using CDDL notation [I-D.greevenbosch-appsawg-cbor-cddl], the form body syntax can be expressed as follows:

```
form-body    = expr
expr         = all-expr / any-expr / field
all-expr     = [1, +expr]
any-expr     = [2, +expr]
field        = [3, field-name, field-type]
field-name   = text
field-type   = text
```

Example:

```
[1, [3, "http://xmlns.com/foaf/0.1/firstName", "text"],
     [3, "http://xmlns.com/foaf/0.1/lastName", "text"],
     [3, "http://xmlns.com/foaf/0.1/age", "uint"]]
```

(using the FOAF [FOAF] vocabulary to indicate the semantics of the individual form fields to an automated agent).

3. Form Data Syntax

Form data is serialized as a CBOR table where each entry consists of a form field name and a value. The field name MUST be the name of one of the fields defined in the form body. The value MUST be a CBOR data item that matches the type specification in the field definition. The media type is "application/form-data+cbor".

Using CDDL notation [I-D.greevenbosch-appsawg-cbor-cddl], the form data syntax can be expressed as follows:

```
form-data    = {*field-name => field-value}
field-name   = text
field-value  = any
```

Example:

```
{"http://xmlns.com/foaf/0.1/firstName": "Jane",
 "http://xmlns.com/foaf/0.1/lastName": "Doe",
 "http://xmlns.com/foaf/0.1/age": 32}
```

4. Security Considerations

TODO.

5. IANA Considerations

5.1. Media Type

This document registers the media type "application/form-data+cbor" in the "Media Types" registry.

TODO.

5.2. CoAP Content-Format

This document registers a content format for the "application/form-data+cbor" media type in the "CoAP Content-Formats" registry.

- o Media Type: application/form-data+cbor
 - Encoding: -
 - ID: 61
 - Reference: [RFCXXXX]

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<http://www.rfc-editor.org/info/rfc7049>>.

6.2. Informative References

- [FOAF] Brickley, D. and L. Miller, "FOAF Vocabulary Specification 0.99", January 2014, <<http://xmlns.com/foaf/spec/20140114.html>>.
- [I-D.greevenbosch-appsawg-cbor-cddl] Vigano, C. and H. Birkholz, "CBOR data definition language (CDDL): a notational convention to express CBOR data structures", draft-greevenbosch-appsawg-cbor-cddl-09 (work in progress), September 2016.
- [I-D.hartke-t2trg-coral] Hartke, K., "The Constrained RESTful Application Language (CoRAL)", draft-hartke-t2trg-coral-01 (work in progress), October 2016.

- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.
- [RFC7578] Masinter, L., "Returning Values from Forms: multipart/form-data", RFC 7578, DOI 10.17487/RFC7578, July 2015, <<http://www.rfc-editor.org/info/rfc7578>>.
- [W3C.REC-html5-20141028]
Hickson, I., Berjon, R., Faulkner, S., Leithead, T., Navara, E., O'Connor, T., and S. Pfeiffer, "HTML5", World Wide Web Consortium Recommendation REC-html5-20141028, October 2014, <<http://www.w3.org/TR/2014/REC-html5-20141028>>.

Author's Address

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

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

Thing-to-Thing Research Group
Internet-Draft
Intended status: Experimental
Expires: January 9, 2020

K. Hartke
Ericsson
July 8, 2019

The Constrained RESTful Application Language (CoRAL)
draft-hartke-t2trg-coral-09

Abstract

The Constrained RESTful Application Language (CoRAL) defines a data model and interaction model as well as two specialized serialization formats for the description of typed connections between resources on the Web ("links"), possible operations on such resources ("forms"), as well as simple resource metadata.

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 <https://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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Notational Conventions	4
2.	Data and Interaction Model	4
2.1.	Browsing Context	4
2.2.	Documents	5
2.3.	Links	5
2.4.	Forms	6
2.4.1.	Form Fields	7
2.5.	Embedded Representations	7
2.6.	Navigation	7
2.7.	History Traversal	9
3.	Binary Format	9
3.1.	Data Structure	10
3.1.1.	Documents	10
3.1.2.	Links	10
3.1.3.	Forms	11
3.1.4.	Embedded Representations	12
3.1.5.	Directives	13
3.2.	Dictionaries	13
3.2.1.	Dictionary References	13
3.2.2.	Media Type Parameter	14
4.	Textual Format	14
4.1.	Lexical Structure	15
4.1.1.	Line Terminators	15
4.1.2.	White Space	15
4.1.3.	Comments	15
4.1.4.	Identifiers	15
4.1.5.	IRIs and IRI References	16
4.1.6.	Literals	16
4.1.7.	Punctuators	20
4.2.	Syntactic Structure	20
4.2.1.	Documents	20
4.2.2.	Links	20
4.2.3.	Forms	21
4.2.4.	Embedded Representations	22
4.2.5.	Directives	23
5.	Usage Considerations	24
5.1.	Specifying CoRAL-based Applications	24
5.1.1.	Application Interfaces	24
5.1.2.	Resource Names	25
5.1.3.	Implementation Limits	25
5.2.	Minting Vocabulary	26
5.3.	Expressing Registered Link Relation Types	27
5.4.	Expressing Simple RDF Statements	27
5.5.	Expressing Language-Tagged Strings	27
5.6.	Embedding CoRAL in CBOR Data	28

5.7. Submitting CoRAL Documents	28
5.7.1. PUT Requests	28
5.7.2. POST Requests	29
6. Security Considerations	29
7. IANA Considerations	30
7.1. Media Type "application/coral+cbor"	30
7.2. Media Type "text/coral"	32
7.3. CoAP Content Formats	33
7.4. CBOR Tag	33
8. References	34
8.1. Normative References	34
8.2. Informative References	36
Appendix A. Core Vocabulary	38
A.1. Base	38
A.2. Collections	39
A.3. HTTP	39
A.4. CoAP	40
Appendix B. Default Dictionary	41
Acknowledgements	41
Author's Address	42

1. Introduction

The Constrained RESTful Application Language (CoRAL) is a language for the description of typed connections between resources on the Web ("links"), possible operations on such resources ("forms"), as well as simple resource metadata.

CoRAL is intended for driving automated software agents that navigate a Web application based on a standardized vocabulary of link relation types and operation types. It is designed to be used in conjunction with a Web transfer protocol such as the Hypertext Transfer Protocol (HTTP) [RFC7230] or the Constrained Application Protocol (CoAP) [RFC7252].

This document defines the CoRAL data and interaction model, as well as two specialized CoRAL serialization formats.

The CoRAL data and interaction model is a superset of the Web Linking model of RFC 8288 [RFC8288]. The data model consists of two primary elements: "links" that describe the relationship between two resources and the type of that relationship, and "forms" that describe a possible operation on a resource and the type of that operation. Additionally, the data model can describe simple resource metadata in a way similar to the Resource Description Framework (RDF) [W3C.REC-rdf11-concepts-20140225]. In contrast to RDF, the focus of CoRAL however is on the interaction with resources, not just the relationships between them. The interaction model derives from HTML

5 [W3C.REC-html52-20171214] and specifies how an automated software agent can navigate between resources by following links and perform operations on resources by submitting forms.

The primary CoRAL serialization format is a compact, binary encoding of links and forms in Concise Binary Object Representation (CBOR) [RFC7049]. It is intended for environments with constraints on power, memory, and processing resources [RFC7228] and shares many similarities with the message format of the Constrained Application Protocol (CoAP) [RFC7252]: For example, it uses numeric identifiers instead of verbose strings for link relation types and operation types, and pre-parses Uniform Resource Identifiers (URIs) [RFC3986] into (what CoAP considers to be) their components, which simplifies URI processing for constrained nodes a lot. As a result, link serializations in CoRAL are often much more compact than equivalent serializations in CoRE Link Format [RFC6690].

The secondary CoRAL serialization format is a lightweight, textual encoding of links and forms that is intended to be easy to read and write for humans. The format is loosely inspired by the syntax of Turtle [W3C.REC-turtle-20140225] and is mainly intended for giving examples.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Terms defined in this document appear in *_cursive_* where they are introduced.

2. Data and Interaction Model

The Constrained RESTful Application Language (CoRAL) is designed for building Web-based applications [W3C.REC-webarch-20041215] in which automated software agents navigate between resources by following links and perform operations on resources by submitting forms.

2.1. Browsing Context

Borrowing from HTML 5 [W3C.REC-html52-20171214], each such agent maintains a *_browsing context_* in which the representations of Web resources are processed. (In HTML 5, the browsing context typically corresponds to a tab or window in a Web browser.)

At any time, one representation in each browsing context is designated the `_active_` representation.

2.2. Documents

A resource representation in one of the CoRAL serialization formats is called a CoRAL `_document_`. The Internationalized Resource Identifier (IRI) [RFC3987] that was used to retrieve such a document is called the document's `_retrieval context_`.

A CoRAL document consists of a list of zero or more links, forms, and embedded resource representations, collectively called `_elements_`. CoRAL serialization formats may define additional types of elements for efficiency or convenience, such as a base for relative IRI references [RFC3987].

2.3. Links

A `_link_` describes a relationship between two resources on the Web [RFC8288]. As defined in RFC 8288, it consists of a `_link context_`, a `_link relation type_`, and a `_link target_`. In CoRAL, a link can additionally have a nested list of zero or more elements, which take the place of link target attributes.

A link can be viewed as a statement of the form "{link context} has a {link relation type} resource at {link target}" where the link target may be further described by nested elements.

The link relation type identifies the semantics of a link. In HTML 5 and RFC 8288, link relation types are typically denoted by an IANA-registered name, such as "stylesheet" or "type". In CoRAL, they are denoted by an IRI such as <http://www.iana.org/assignments/relation/stylesheet> or <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>. This allows for the creation of new link relation types without the risk of collisions when from different organizations or domains of knowledge. An IRI also can lead to documentation, schema, and other information about the link relation type. These IRIs are only used as identity tokens, though, and are compared using Simple String Comparison (Section 5.1 of RFC 3987).

The link context and the link target are both either by an IRI or (similarly to RDF) a literal. If the IRI scheme indicates a Web transfer protocol such as HTTP or CoAP, then an agent can dereference the IRI and navigate the browsing context to the referenced resource; this is called `_following the link_`. A literal directly identifies a value. This can be a Boolean value, an integer, a floating-point number, a date/time value, a byte string, or a text string.

A link can occur as a top-level element in a document or as a nested element within a link. When a link occurs as a top-level element, the link context implicitly is the document's retrieval context. When a link occurs nested within a link, the link context of the inner link is the link target of the outer link.

There are no restrictions on the cardinality of links; there can be multiple links to and from a particular target, and multiple links of the same or different types between a given link context and target. However, the nested data structure constrains the description of a resource graph to a tree: Links between linked resources can only be described by further nesting links.

2.4. Forms

A `_form_` provides instructions to an agent for performing an operation on a Web resource. It consists of a `_form context_`, an `_operation type_`, a `_request method_`, and a `_submission target_`. Additionally, a form may be accompanied by a list of `_form fields_`.

A form can be viewed as an instruction of the form "To perform an {operation type} operation on {form context}, make a {request method} request to {submission target}" where the request may be further described by form fields.

The operation type identifies the semantics of the operation. Operation types are denoted like link relation types by an IRI.

The form context is the resource on which an operation is ultimately performed. To perform the operation, an agent needs to construct a request with the specified method and the specified submission target as the request IRI. Usually, the submission target is the same resource as the form context, but it may be a different resource. Constructing and sending the request is called `_submitting the form_`.

Form fields, specified in the next section, can be used to provide more detailed instructions to the agent for constructing the request. For example, form fields can instruct the agent to include a payload or certain headers in the request that must match the specifications of the form fields.

A form can occur as a top-level element in a document or as a nested element within a link. When a form occurs as a top-level element, the form context implicitly is the document's retrieval context. When a form occurs nested within a link, the form context is the link target of the enclosing link.

2.4.1. Form Fields

Form fields provide further instructions to agents for constructing a request.

For example, a form field could identify one or more data items that need to be included in the request payload or reference another resource (such as a schema) that describes the structure of the payload. A form field could also provide other kinds of information, such as acceptable media types for the payload or expected request headers. Form fields may be specific to the protocol used for submitting the form.

A form field is the pair of a `_form field type_` and a `_form field value_`.

The form field type identifies the semantics of the form field. Form field types are denoted like link relation types and operation types by an IRI.

The form field value can be either an IRI, a Boolean value, an integer, a floating-point number, a date/time value, a byte string, or a text string.

2.5. Embedded Representations

When a document contains links to many resources and an agent needs a representation of each link target, it may be inefficient to retrieve each of these representations individually. To alleviate this, documents can directly embed representations of resources.

An `_embedded representation_` consists of a sequence of bytes, labeled with `_representation metadata_`.

An embedded representation may be a full, partial, or inconsistent version of the representation served from the IRI of the resource.

An embedded representation can occur as a top-level element in a document or as a nested element within a link. When it occurs as a top-level element, it provides an alternate representation of the document's retrieval context. When it occurs nested within a link, it provides a representation of link target of the enclosing link.

2.6. Navigation

An agent begins interacting with an application by performing a GET request on an `_entry point IRI_`. The entry point IRI is the only IRI an agent is expected to know before interacting with an application.

From there, the agent is expected to make all requests by following links and submitting forms provided by the server in responses. The entry point IRI can be obtained by manual configuration or through some discovery process.

If dereferencing the entry point IRI yields a CoRAL document (or any other representation that implements the CoRAL data and interaction model), then the agent makes this document the active representation in the browsing context and proceeds as follows:

1. The first step for the agent is to decide what to do next, i.e., which type of link to follow or form to submit, based on the link relation types and operation types it understands.
2. The agent then finds the link(s) or form(s) with the respective type in the active representation. This may yield one or more candidates, from which the agent will have to select the most appropriate one. The set of candidates may be empty, for example, when a transition is not supported or not allowed.
3. The agent selects one of the candidates based on the metadata associated with each of these. Metadata includes the content type of the target resource representation, the IRI scheme, the request method, and other information that is provided as nested elements in a link or form fields in a form.

If the selected candidate contains an embedded representation, the agent MAY skip the following steps and immediately proceed with step 8.

4. The agent obtains the `_request IRI_` from the link target or submission target. Fragment identifiers are not part of the request IRI and MUST be separated from the rest of the IRI prior to a dereference.
5. The agent constructs a new request with the request IRI. If the agent is following a link, then the request method MUST be GET. If the agent is submitting a form, then the request method MUST be the one specified in the form. The request IRI may need to be converted to a URI (Section 3.1 of RFC 3987) for protocols that do not support IRIs.

The agent should set HTTP header fields and CoAP request options according to metadata associated with the link or form (e.g., set the HTTP Accept header field or the CoAP Accept option when the media type of the target resource is provided). Depending on the operation type of a form, the agent may also need to include a

request payload that matches the specifications of one or more form fields.

6. The agent sends the request and receives the response.
7. If a fragment identifier was separated from the request IRI, the agent dereferences the fragment identifier within the received representation.
8. The agent `_updates the browsing context_` by making the (embedded or received) representation the active representation.
9. Finally, the agent processes the representation according to the semantics of the content type. If the representation is a CoRAL document (or any other representation that implements the CoRAL data and interaction model), this means the agent has the choice of what to do next again -- and the cycle repeats.

2.7. History Traversal

A browsing context MAY entail a `_session history_` that lists the resource representations that the agent has processed, is processing, or will process.

An entry in the session history consists of a resource representation and the request IRI that was used to retrieve the representation. New entries are added to the session history as the agent navigates from resource to resource.

An agent can navigate a browsing context by `_traversing the session history_` in addition to following links and submitting forms. For example, if an agent received a representation that doesn't contain any further links or forms, it can revert the active representation back to one it has visited earlier.

Traversing the history should take advantage of caches to avoid new requests. An agent MAY reissue a safe request (e.g., a GET request) when it doesn't have a fresh representation in its cache. An agent MUST NOT reissue an unsafe request (e.g., a PUT or POST request) unless it intends to perform that operation again.

3. Binary Format

This section defines the encoding of documents in the CoRAL binary format.

A document in the binary format is a data item in Concise Binary Object Representation (CBOR) [RFC7049]. The structure of this data

item is presented in the Concise Data Definition Language (CDDL) [RFC8610]. The media type is "application/coral+cbor".

The following restrictions are placed on CBOR encoders: Byte strings and text strings MUST be encoded with definite length. Integers and floating-point values MUST be encoded as such (e.g., a floating-point value of 0.0 must not be encoded as the integer 0).

3.1. Data Structure

The data structure of a document in the binary format is made up of four kinds of elements: links, forms, embedded representations, and (as an extension to the CoRAL data model) base directives. Base directives provide a way to encode IRIs with a common base more efficiently.

Elements are processed in the order they appear in the document. Document processors need to maintain an `_environment_` while iterating an array of elements. The environment consists of two variables: the `_current context_` and the `_current base_`. Both the current context and the current base are initially set to the document's retrieval context.

3.1.1. Documents

The body of a document in the binary format is encoded as an array of zero or more links, forms, embedded representations, and directives.

```
document = body
```

```
body = [*(link / form / representation / directive)]
```

3.1.2. Links

A link is encoded as an array that consists of the unsigned integer 2, followed by the link relation type and the link target, optionally followed by a link body that contains nested elements.

```
link = [2, relation-type, link-target, ?body]
```

The link relation type is encoded as a text string that conforms to the syntax of an IRI [RFC3987].

```
relation-type = text
```

The link target is denoted by an IRI reference or represented by a literal value. An IRI reference MUST be resolved against the current base. The encoding of and resolution process for IRI references in

the binary format is described in RFC XXXX [I-D.hartke-t2trg-ciri]. The link target may be null, which indicates that the link target is an unidentified resource.

```
link-target = ciri / literal
```

```
ciri = <Defined in Section X of RFC XXXX>
```

```
literal = bool / int / float / time / bytes / text / null
```

The array of elements in the link body, if any, MUST be processed in a fresh environment. Both the current context and the current base in the new environment are initially set to the link target of the enclosing link.

3.1.3. Forms

A form is encoded as an array that consists of the unsigned integer 3, followed by the operation type and the submission target, optionally followed by a list of form fields.

```
form = [3, operation-type, submission-target, ?form-fields]
```

The operation type is defined in the same way as a link relation type (Section 3.1.2).

```
operation-type = text
```

The request method is either implied by the operation type or encoded as a form field. If there are both, the form field takes precedence over the operation type. Either way, the method MUST be defined for the Web transfer protocol identified by the scheme of the submission target.

The submission target is denoted by an IRI reference. This IRI reference MUST be resolved against the current base.

```
submission-target = ciri
```

3.1.3.1. Form Fields

A list of form fields is encoded as an array of zero or more type-value pairs.

```
form-fields = [*(form-field-type, form-field-value)]
```

The list, if any, MUST be processed in a fresh environment. Both the current context and the current base in the new environment are initially set to the submission target of the enclosing form.

A form field type is defined in the same way as a link relation type (Section 3.1.2).

```
form-field-type = text
```

A form field value can be an IRI reference, a Boolean value, an integer, a floating-point number, a date/time value, a byte string, a text string, or null. An IRI reference MUST be resolved against the current base.

```
form-field-value = ciri / literal
```

3.1.4. Embedded Representations

An embedded representation is encoded as an array that consists of the unsigned integer 0, followed by a byte string containing the representation data, optionally followed by representation metadata.

```
representation = [0, bytes, ?representation-metadata]
```

Representation metadata is encoded as an array of zero or more name-value pairs.

```
representation-metadata = [(metadata-name, metadata-value)]
```

The metadata, if any, MUST be processed in a fresh environment. All variables in the new environment are initially set to a copy of the variables in the current environment.

The metadata name is defined in the same way as a link relation type (Section 3.1.2).

```
metadata-name = text
```

A metadata value can be an IRI reference, a Boolean value, an integer, a floating-point number, a date/time value, a byte string, a text string, or null. An IRI reference MUST be resolved against the current base.

```
metadata-value = ciri / literal
```

3.1.5. Directives

Directives provide the ability to manipulate the environment when processing a list of elements. There is one type of directives available: the Base directive.

```
directive = base-directive
```

3.1.5.1. Base Directives

A Base directive is encoded as an array that consists of the unsigned integer 1, followed by a base.

```
base-directive = [1, base]
```

The base is denoted by an IRI reference. This IRI reference MUST be resolved against the current context (not the current base).

```
base = ciri
```

The directive is processed by resolving the IRI reference against the current context and assigning the result to the current base.

3.2. Dictionaries

The binary format can reference values from a dictionary to reduce representation size and processing cost. Dictionary references can be used in place of link relation types, link targets, operation types, submission targets, form field types, form field values, representation metadata names, and representation metadata values.

3.2.1. Dictionary References

A dictionary reference is encoded as an unsigned integer. Where a dictionary reference cannot be expressed unambiguously, the unsigned integer is tagged with CBOR tag TBD6.

```
relation-type /= uint
```

```
link-target /= #6.TBD6(uint)
```

```
operation-type /= uint
```

```
submission-target /= #6.TBD6(uint)
```

```
form-field-type /= uint
```

```
form-field-value /= #6.TBD6(uint)
```

```
metadata-name /= uint
```

```
metadata-value /= #6.TBD6(uint)
```

3.2.2. Media Type Parameter

The "application/coral+cbor" media type is defined to have a "dictionary" parameter that specifies the dictionary in use. The dictionary is identified by a URI [RFC3986]. For example, a CoRAL document that uses the dictionary identified by the URI <http://example.com/dictionary> can use the following content type:

```
application/coral+cbor;dictionary="http://example.com/dictionary"
```

The URI serves only as an identifier; it does not necessarily have to be dereferencable (or even use a dereferencable URI scheme). It is permissible, though, to use a dereferencable URI and to serve a representation that provides information about the dictionary in a human- or machine-readable way. (The format of such a representation is outside the scope of this document.)

For simplicity, a CoRAL document can reference values only from one dictionary; the value of the "dictionary" parameter MUST be a single URI. If the "dictionary" parameter is absent, the default dictionary specified in Appendix B of this document is assumed.

Once a dictionary has made an assignment, the assignment MUST NOT be changed or removed. A dictionary, however, may contain additional information about an assignment, which may change over time.

In CoAP [RFC7252], media types (including specific values for media type parameters) are encoded as an unsigned integer called "content format". For use with CoAP, each new CoRAL dictionary MUST register a new content format in the IANA CoAP Content-Formats Registry.

4. Textual Format

This section defines the syntax of documents in the CoRAL textual format using two grammars: The lexical grammar defines how Unicode characters are combined to form line terminators, white space, comments, and tokens. The syntactic grammar defines how tokens are combined to form documents. Both grammars are presented in Augmented Backus-Naur Form (ABNF) [RFC5234].

A document in the textual format is a Unicode string in a Unicode encoding form [UNICODE]. The media type for such documents is "text/coral". The "charset" parameter is not used; charset information is transported inside the document in the form of an OPTIONAL Byte Order

Mark (BOM). The use of the UTF-8 encoding scheme [RFC3629], without a BOM, is RECOMMENDED.

4.1. Lexical Structure

The lexical structure of a document in the textual format is made up of four basic elements: line terminators, white space, comments, and tokens. Of these, only tokens are significant in the syntactic grammar. There are five kinds of tokens: identifiers, IRIs, IRI references, literals, and punctuators.

token = identifier / iri / iriref / literal / punctuator

When several lexical grammar rules match a sequence of characters in a document, the longest match takes priority.

4.1.1. Line Terminators

Line terminators divide text into lines. A line terminator is any Unicode character with Line_Break class BK, CR, LF, or NL. However, any CR character that immediately precedes a LF character is ignored. (This affects only the numbering of lines in error messages.)

4.1.2. White Space

White space is a sequence of one or more white space characters. A white space character is any Unicode character with the White_Space property.

4.1.3. Comments

Comments are sequences of characters that are ignored when parsing text into tokens. Single-line comments begin with the characters `"//"` and extend to the end of the line. Delimited comments begin with the characters `"/*"` and end with the characters `"*/"`. Delimited comments can occupy a portion of a line, a single line, or multiple lines.

Comments do not nest. The character sequences `"/*"` and `"*/"` have no special meaning within a single-line comment; the character sequences `"//"` and `"/*"` have no special meaning within a delimited comment.

4.1.4. Identifiers

An identifier token is a user-defined symbolic name. The rules for identifiers correspond to those recommended by the Unicode Standard Annex #31 [UNICODE-UAX31] using the following profile:

```
identifier = START *CONTINUE *(MEDIAL 1*CONTINUE)

START = <Any character with the XID_Start property>

CONTINUE = <Any character with the XID_Continue property>

MEDIAL = "-" / "." / "~" / %x58A / %xF0B

MEDIAL =/ %x2010 / %x2027 / %x30A0 / %x30FB
```

All identifiers MUST be converted into Unicode Normalization Form C (NFC), as defined by the Unicode Standard Annex #15 [UNICODE-UAX15]. Comparison of identifiers is based on NFC and is case-sensitive (unless otherwise noted).

4.1.5. IRIs and IRI References

IRIs and IRI references are Unicode strings that conform to the syntax defined in RFC 3987 [RFC3987]. An IRI reference can be either an IRI or a relative reference. Both IRIs and IRI references are enclosed in angle brackets ("`<`" and "`>`").

```
iri = "<" IRI ">"

iriref = "<" IRI-reference ">"

IRI = <Defined in Section 2.2 of RFC 3987>

IRI-reference = <Defined in Section 2.2 of RFC 3987>
```

4.1.6. Literals

A literal is a textual representation of a value. There are seven types of literals: Boolean, integer, floating-point, date/time, byte string, text string, and null.

```
literal = boolean / integer / float / datetime / bytes / text

literal =/ null
```

4.1.6.1. Boolean Literals

The case-insensitive tokens "`true`" and "`false`" denote the Boolean values true and false, respectively.

```
boolean = "true" / "false"
```

4.1.6.2. Integer Literals

Integer literals denote an integer value of unspecified precision. By default, integer literals are expressed in decimal, but they can also be specified in an alternate base using a prefix: Binary literals begin with "0b", octal literals begin with "0o", and hexadecimal literals begin with "0x".

Decimal literals contain the digits "0" through "9". Binary literals contain "0" and "1", octal literals contain "0" through "7", and hexadecimal literals contain "0" through "9" as well as "A" through "F" in upper- or lowercase.

Negative integers are expressed by prepending a minus sign ("-").

```
integer = ["+" / "-"] (decimal / binary / octal / hexadecimal)
```

```
decimal = 1*DIGIT
```

```
binary = %x30 (%x42 / %x62) 1*BINDIG
```

```
octal = %x30 (%x4F / %x6F) 1*OCTDIG
```

```
hexadecimal = %x30 (%x58 / %x78) 1*HEXDIG
```

```
DIGIT = %x30-39
```

```
BINDIG = %x30-31
```

```
OCTDIG = %x30-37
```

```
HEXDIG = %x30-39 / %x41-46 / %x61-66
```

4.1.6.3. Floating-point Literals

Floating-point literals denote a floating-point number of unspecified precision.

Floating-point literals consist of a sequence of decimal digits followed by a fraction, an exponent, or both. The fraction consists of a decimal point (".") followed by a sequence of decimal digits. The exponent consists of the letter "e" in upper- or lowercase, followed by an optional sign and a sequence of decimal digits that indicate a power of 10 by which the value preceding the "e" is multiplied.

Negative floating-point values are expressed by prepending a minus sign ("-").

```
float = ["+" / "-"] 1*DIGIT [fraction] [exponent]
```

```
fraction = "." 1*DIGIT
```

```
exponent = (%x45 / %x65) ["+" / "-"] 1*DIGIT
```

A floating-point literal can additionally denote either the special "Not-a-Number" (NaN) value, positive infinity, or negative infinity. The NaN value is produced by the case-insensitive token "NaN". The two infinite values are produced by the case-insensitive tokens "+Infinity" (or simply "Infinity") and "-Infinity".

```
float =/ "NaN"
```

```
float =/ ["+" / "-"] "Infinity"
```

4.1.6.4. Date/Time Literals

Date/time literals denote an instant in time.

A date/time literal consists of the prefix "dt" and a sequence of Unicode characters in Internet Date/Time Format [RFC3339], enclosed in single quotes.

```
datetime = %x64.74 SQUOTE date-time SQUOTE
```

```
date-time = <Defined in Section 5.6 of RFC 3339>
```

```
SQUOTE = %x27
```

4.1.6.5. Byte String Literals

Byte string literals denote an ordered sequence of bytes.

A byte string literal consists of a prefix and zero or more bytes encoded in Base16, Base32, or Base64 [RFC4648], enclosed in single quotes. Byte string literals encoded in Base16 begin with "h" or "b16", byte string literals encoded in Base32 begin with "b32", and byte string literals encoded in Base64 begin with "b64".

```
bytes = base16 / base32 / base64
```

```
base16 = (%x68 / %x62.31.36) SQUOTE <Base16 encoded data> SQUOTE
```

```
base32 = %x62.33.32 SQUOTE <Base32 encoded data> SQUOTE
```

```
base64 = %x62.36.34 SQUOTE <Base64 encoded data> SQUOTE
```


4.1.6.6. Text String Literals

Text string literals denote a Unicode string.

A text string literal consists of zero or more Unicode characters enclosed in double quotes. It can include simple escape sequences (such as `\t` for the tab character) as well as hexadecimal and Unicode escape sequences.

```
text = DQUOTE *(char / %x5C escape) DQUOTE
```

```
char = <Any character except %x22, %x5C, and line terminators>
```

```
escape = simple-escape / hexadecimal-escape / unicode-escape
```

```
simple-escape = %x30 / %x62 / %x74 / %x6E / %x76
```

```
simple-escape =/ %x66 / %x72 / %x22 / %x27 / %x5C
```

```
hexadecimal-escape = (%x78 / %x58) 2HEXDIG
```

```
unicode-escape = %x75 4HEXDIG / %x55 8HEXDIG
```

```
DQUOTE = %x22
```

An escape sequence denotes a single Unicode code point. For hexadecimal and Unicode escape sequences, the code point is expressed by the hexadecimal number following the `"\x"`, `"\X"`, `"\u"`, or `"\U"` prefix. Simple escape sequences indicate the code points listed in Table 1.

Escape Sequence	Code Point	Character Name
<code>\0</code>	U+0000	Null
<code>\b</code>	U+0008	Backspace
<code>\t</code>	U+0009	Character Tabulation
<code>\n</code>	U+000A	Line Feed
<code>\v</code>	U+000B	Line Tabulation
<code>\f</code>	U+000C	Form Feed
<code>\r</code>	U+000D	Carriage Return
<code>\"</code>	U+0022	Quotation Mark
<code>\'</code>	U+0027	Apostrophe
<code>\\</code>	U+005C	Reverse Solidus

Table 1: Simple Escape Sequences

4.1.6.7. Null Literal

The case-insensitive tokens "null" and "_" denote the intentional absence of any value.

```
null = "null" / "_"
```

4.1.7. Punctuators

Punctuator tokens are used for grouping and separating.

```
punctuator = "#" / ":" / "*" / "[" / "]" / "{" / "}" / "=" / "->"
```

4.2. Syntactic Structure

The syntactic structure of a document in the textual format is made up of four kinds of elements: links, forms, embedded representations, and (as an extension to the CoRAL data model) directives. Directives provide a way to make documents easier to read and write by setting a base for relative IRI references and introducing shorthands for IRIs.

Elements are processed in the order they appear in the document. Document processors need to maintain an `_environment_` while iterating a list of elements. The environment consists of three variables: the `_current context_`, the `_current base_`, and the `_current mapping from identifiers to IRIs_`. Both the current context and the current base are initially set to the document's retrieval context. The current mapping from identifiers to IRIs is initially empty.

4.2.1. Documents

The body of a document in the textual format consists of zero or more links, forms, embedded representations, and directives.

```
document = body
```

```
body = *(link / form / representation / directive)
```

4.2.2. Links

A link consists of the link relation type, followed by the link target, optionally followed by a link body enclosed in curly brackets ("{" and "}").

```
link = relation-type link-target [{" body "}]
```

The link relation type is denoted by either an IRI, a simple name, or a qualified name.

relation-type = iri / simple-name / qualified-name

A simple name consists of an identifier. It is resolved to an IRI by looking up the empty string in the current mapping from identifiers to IRIs and appending the specified identifier to the result. It is an error if the empty string is not present in the current mapping.

simple-name = identifier

A qualified name consists of two identifiers separated by a colon (":"). It is resolved to an IRI by looking up the identifier on the left hand side in the current mapping from identifiers to IRIs and appending the identifier on the right hand side to the result. It is an error if the identifier on the left hand side is not present in the current mapping.

qualified-name = identifier ":" identifier

The link target is denoted by an IRI reference or represented by a value literal. An IRI reference **MUST** be resolved against the current base. If the link target is null, the link target is an unidentified resource.

link-target = iriref / literal

The list of elements in the link body, if any, **MUST** be processed in a fresh environment. Both the current context and current base in this environment are initially set to the link target of the enclosing link. The mapping from identifiers to IRIs is initially set to a copy of the mapping from identifiers to IRIs in the current environment.

4.2.3. Forms

A form consists of the operation type, followed by a "->" token and the submission target, optionally followed by a list of form fields enclosed in square brackets "[" and "]").

form = operation-type "->" submission-target ["[" form-fields "]"]

The operation type is defined in the same way as a link relation type (Section 4.2.2).

operation-type = iri / simple-name / qualified-name

The request method is either implied by the operation type or encoded as a form field. If there are both, the form field takes precedence over the operation type. Either way, the method **MUST** be defined for

the Web transfer protocol identified by the scheme of the submission target.

The submission target is denoted by an IRI reference. This IRI reference MUST be resolved against the current base.

```
submission-target = iriref
```

4.2.3.1. Form Fields

A list of form fields consists of zero or more type-value pairs.

```
form-fields = *(form-field-type form-field-value)
```

The list, if any, MUST be processed in a fresh environment. Both the current context and the current base in this environment are initially set to the submission target of the enclosing form. The mapping from identifiers to IRIs is initially set to a copy of the mapping from identifiers to IRIs in the current environment.

The form field type is defined in the same way as a link relation type (Section 4.2.2).

```
form-field-type = iri / simple-name / qualified-name
```

The form field value can be an IRI reference, Boolean literal, integer literal, floating-point literal, byte string literal, text string literal, or null. An IRI reference MUST be resolved against the current base.

```
form-field-value = iriref / literal
```

4.2.4. Embedded Representations

An embedded representation consists of a "*" token, followed by the representation data, optionally followed by representation metadata enclosed in square brackets "[" and "]").

```
representation = "*" bytes ["[" representation-metadata "]" ]
```

Representation metadata consists of zero or more name-value pairs.

```
representation-metadata = *(metadata-name metadata-value)
```

The metadata, if any, MUST be processed in a fresh environment. All variables in the new environment are initially set to a copy of the variables in the current environment.

The metadata name is defined in the same way as a link relation type (Section 4.2.2).

```
metadata-name = iri / simple-name / qualified-name
```

The metadata value can be an IRI reference, Boolean literal, integer literal, floating-point literal, byte string literal, text string literal, or null. An IRI reference MUST be resolved against the current base.

```
metadata-value = iriref / literal
```

4.2.5. Directives

Directives provide the ability to manipulate the environment when processing a list of elements. All directives start with a number sign ("#") followed by a directive identifier. Directive identifiers are case-insensitive and constrained to Unicode characters in the Basic Latin block.

The following two types of directives are available: the Base directive and the Using directive.

```
directive = base-directive / using-directive
```

4.2.5.1. Base Directives

A Base directive consists of a number sign ("#"), followed by the case-insensitive identifier "base", followed by a base.

```
base-directive = "#" "base" base
```

The base is denoted by an IRI reference. The IRI reference MUST be resolved against the current context (not the current base).

```
base = iriref
```

The directive is processed by resolving the IRI reference against the current context and assigning the result to the current base.

4.2.5.2. Using Directives

A Using directive consists of a number sign ("#"), followed by the case-insensitive identifier "using", optionally followed by an identifier and an equals sign ("="), finally followed by an IRI. If the identifier is not specified, it is assumed to be the empty string.

```
using-directive = "#" "using" [identifier "="] iri
```

The directive is processed by adding the specified identifier and IRI to the current mapping from identifiers to IRIs. It is an error if the identifier is already present in the mapping.

5. Usage Considerations

This section discusses some considerations in creating CoRAL-based applications and vocabularies.

5.1. Specifying CoRAL-based Applications

CoRAL-based applications naturally implement the Web architecture [W3C.REC-webarch-20041215] and thus are centered around orthogonal specifications for identification, interaction, and representation:

- o Resources are identified by IRIs or represented by value literals.
- o Interactions are based on the hypermedia interaction model of the Web and the methods provided by the Web transfer protocol. The semantics of possible interactions are identified by link relation types and operation types.
- o Representations are CoRAL documents encoded in the binary format defined in Section 3 or the textual format defined in Section 4. Depending on the application, additional representation formats may be used.

5.1.1. Application Interfaces

Specifications for CoRAL-based applications need to list the specific components used in the application interface and their identifiers. This should include the following items:

- o IRI schemes that identify the Web transfer protocol(s) used in the application.
- o Internet media types that identify the representation format(s) used in the application, including the media type(s) of the CoRAL serialization format(s).
- o Link relation types that identify the semantics of links.
- o Operation types that identify the semantics of forms. Additionally, for each operation type, the permissible request method(s).

- o Form field types that identify the semantics of form fields. Additionally, for each form field type, the permissible form field values.
- o Metadata names that identify the semantics of representation metadata. Additionally, for each metadata name, the permissible metadata values.

5.1.2. Resource Names

Resource names -- i.e., URIs [RFC3986] and IRIs [RFC3987] -- are a cornerstone of Web-based applications. They enable the uniform identification of resources and are used every time a client interacts with a server or a resource representation needs to refer to another resource.

URIs and IRIs often include structured application data in the path and query components, such as paths in a filesystem or keys in a database. It is a common practice in many HTTP-based application programming interfaces (APIs) to make this part of the application specification, i.e., to prescribe fixed URI templates that are hard-coded in implementations. There are a number of problems with this practice [RFC7320], though.

In CoRAL-based applications, resource names are therefore not part of the application specification -- they are an implementation detail. The specification of a CoRAL-based application MUST NOT mandate any particular form of resource name structure. BCP 190 [RFC7320] describes the problematic practice of fixed URI structures in more detail and provides some acceptable alternatives.

5.1.3. Implementation Limits

This document places no restrictions on the number of elements in a CoRAL document or the depth of nested elements. Applications using CoRAL (in particular those running in constrained environments) may wish to limit these numbers and specify implementation limits that an application implementation must at least support to be interoperable.

Applications may also mandate the following and other restrictions:

- o use of only either the binary format or the text format;
- o use of only either HTTP or CoAP as supported Web transfer protocol;
- o use of only dictionary references in the binary format for certain vocabulary;

- o use of only either content type strings or content format IDs;
- o use of IRI references only up to a specific string length;
- o use of CBOR in a canonical format (see Section 3.9 of RFC 7049).

5.2. Minting Vocabulary

New link relation types, operation types, form field types, and metadata names can be minted by defining an IRI [RFC3987] that uniquely identifies the item. Although the IRI can point to a resource that contains a definition of the semantics, clients SHOULD NOT automatically access that resource to avoid overburdening its server. The IRI SHOULD be under the control of the person or party defining it, or be delegated to them.

To avoid interoperability problems, it is RECOMMENDED that only IRIs are minted that are normalized according to Section 5.3 of RFC 3987. Non-normalized forms that are best avoided include:

- o Uppercase characters in scheme names and domain names
- o Percent-encoding of characters where it is not required by the IRI syntax
- o Explicitly stated HTTP default port (e.g., <http://example.com/> is preferable over <http://example.com:80/>)
- o Completely empty path in HTTP IRIs (e.g., <http://example.com/> is preferable over <http://example.com>)
- o Dot segments ("./" or "../") in the path component of an IRI
- o Lowercase hexadecimal letters within percent-encoding triplets (e.g., "%3F" is preferable over "%3f")
- o Punycode-encoding of Internationalized Domain Names in IRIs
- o IRIs that are not in Unicode Normalization Form C [UNICODE-UAX15]

IRIs that identify vocabulary do not need to be registered. The inclusion of domain names in IRIs allows for the decentralized creation of new IRIs without the risk of collisions.

However, IRIs can be relatively verbose and impose a high overhead on a representation. This can be a problem in constrained environments [RFC7228]. Therefore, CoRAL alternatively allows the use of unsigned integers to reference CBOR data items from a dictionary, as specified

in Section 3.2. These impose a much smaller overhead but instead need to be assigned by an authority to avoid collisions.

5.3. Expressing Registered Link Relation Types

Link relation types registered in the IANA Link Relations Registry, such as "collection" [RFC6573] or "icon" [W3C.REC-html52-20171214], can be used in CoRAL by appending the registered name to the IRI `<http://www.iana.org/assignments/relation/>`:

```
#using iana = <http://www.iana.org/assignments/relation/>

iana:collection </items>
iana:icon       </favicon.png>
```

Note that registered link relation types are required to be lowercased, as per Section 3.3 of RFC 8288 [RFC8288].

(The convention of appending the link relation types to the prefix "http://www.iana.org/assignments/relation/" to form IRIs is adopted from Atom [RFC4287]; see also Appendix A.2 of RFC 8288 [RFC8288].)

5.4. Expressing Simple RDF Statements

An RDF statement [W3C.REC-rdf11-concepts-20140225] says that some relationship, indicated by a predicate, holds between two resources. RDF predicates can therefore be good source for vocabulary to provide resource metadata. For example, a CoRAL document could use the FOAF vocabulary [FOAF] to describe the person or software that made it:

```
#using rdf = <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
#using foaf = <http://xmlns.com/foaf/0.1/>

foaf:maker null {
  rdf:type          <http://xmlns.com/foaf/0.1/Person>
  foaf:familyName  "Hartke"
  foaf:givenName   "Klaus"
  foaf:mbox        <mailto:klaus.hartke@ericsson.com>
}
```

5.5. Expressing Language-Tagged Strings

Text strings that are the target of a link can be associated with a language tag [RFC5646] by nesting a link of type `<http://coreapps.org/base#lang>` under the link. The target of this nested link MUST be a text string that conforms to the syntax specified in Section 2.1 of RFC 5646:

```
#using <http://coreapps.org/base#>
#using example = <http://example.org/>
#using iana = <http://www.iana.org/assignments/relation/>

iana:terms-of-service </tos> {
  example:title "Nutzungsbedingungen" { lang "de" }
  example:title "Terms of use"      { lang "en" }
}
```

5.6. Embedding CoRAL in CBOR Data

Data items in the CoRAL binary format (Section 3) may be embedded in other CBOR data [RFC7049] data. Specifications using CDDL [RFC8610] SHOULD reference the following CDDL definitions for this purpose:

```
CoRAL-Document = document
```

```
CoRAL-Link = link
```

```
CoRAL-Form = form
```

For each embedded document, link, and form, the retrieval context, link context, and form context needs to be specified, respectively.

5.7. Submitting CoRAL Documents

By default, a CoRAL document is a representation that captures the current state of a resource. The meaning of a CoRAL document changes when it is submitted in a request. Depending on the request method, the CoRAL document can capture the intended state of a resource (PUT) or be subject to application-specific processing (POST).

5.7.1. PUT Requests

A PUT request with a CoRAL document enclosed in the request payload requests that the state of the target resource be created or replaced with the state described by the CoRAL document. A successful PUT of a CoRAL document generally means that a subsequent GET on that same target resource would result in an equivalent document being sent in a success response.

An origin server SHOULD verify that a submitted CoRAL document is consistent with any constraints the server has for the target resource. When a document is inconsistent with the target resource, the origin server SHOULD either make it consistent (e.g., by removing inconsistent elements) or respond with an appropriate error message containing sufficient information to explain why the document is unsuitable.

The retrieval context of a CoRAL document in a PUT is the request IRI of the request.

5.7.2. POST Requests

A POST request with a CoRAL document enclosed in the request payload requests that the target resource process the CoRAL document according to the resource's own specific semantics.

The retrieval context of a CoRAL document in a POST is the request IRI of the request.

6. Security Considerations

Parsers of CoRAL documents must operate on input that is assumed to be untrusted. This means that parsers **MUST** fail gracefully in the face of malicious inputs (e.g., inputs not adhering to the data structure). Additionally, parsers **MUST** be prepared to deal with resource exhaustion (e.g., resulting from the allocation of big data items) or exhaustion of the call stack (stack overflow).

CoRAL documents intentionally do not feature the equivalent of XML entity references as to preclude the whole class of exponential XML entity expansion ("billion laughs") [CAPEC-197] and improper XML external entity [CAPEC-201] attacks.

Implementers of the CoRAL binary format need to consider the security aspects of processing CBOR with the restrictions described in Section 3. Notably, different number representations for the same numeric value are not equivalent in the CoRAL binary format. See Section 8 of RFC 7049 [RFC7049] for security considerations relating to CBOR.

Implementers of the CoRAL textual format need to consider the security aspects of handling Unicode input. See the Unicode Standard Annex #36 [UNICODE-UAX36] for security considerations relating to visual spoofing and misuse of character encodings. See Section 10 of RFC 3629 [RFC3629] for security considerations relating to UTF-8.

CoRAL makes extensive use of IRIs and URIs. See Section 8 of RFC 3987 [RFC3987] for security considerations relating to IRIs. See Section 7 of RFC 3986 [RFC3986] for security considerations relating to URIs.

The security of applications using CoRAL can depend on the proper preparation and comparison of internationalized strings. For example, such strings can be used to make authentication and authorization decisions, and the security of an application could be

compromised if an entity providing a given string is connected to the wrong account or online resource based on different interpretations of the string. See RFC 6943 [RFC6943] for security considerations relating to identifiers in IRIs and other places.

CoRAL is intended to be used in conjunction with a Web transfer protocol like HTTP or CoAP. See Section 9 of RFC 7230 [RFC7230], Section 9 of RFC 7231 [RFC7231], etc., for security considerations relating to HTTP. See Section 11 of RFC 7252 [RFC7252] for security considerations relating to CoAP.

CoRAL does not define any specific mechanisms for protecting the confidentiality and integrity of CoRAL documents. It relies on application layer or transport layer mechanisms for this, such as Transport Layer Security (TLS) [RFC8446].

CoRAL documents and the structure of a web of resources revealed from automatically following links can disclose personal information and other sensitive information. Implementations need to prevent the unintentional disclosure of such information. See Section of 9 of RFC 7231 [RFC7231] for additional considerations.

Applications using CoRAL ought to consider the attack vectors opened by automatically following, trusting, or otherwise using links and forms in CoRAL documents. Notably, a server that is authoritative for the CoRAL representation of a resource may not necessarily be authoritative for nested elements in the document. See Section 5 of RFC 8288 [RFC8288] for related considerations.

Unless an application mitigates this risk by specifying more specific rules, any link or form in a document where the link or form context and the document's retrieval context don't share the same Web origin [RFC6454] MUST be discarded ("same-origin policy").

7. IANA Considerations

7.1. Media Type "application/coral+cbor"

This document registers the media type "application/coral+cbor" according to the procedures of BCP 13 [RFC6838].

Type name:
application

Subtype name:
coral+cbor

Required parameters:

N/A

Optional parameters:

dictionary - See Section 3.2 of [I-D.hartke-t2trg-coral].

Encoding considerations:

binary - See Section 3 of [I-D.hartke-t2trg-coral].

Security considerations:

See Section 6 of [I-D.hartke-t2trg-coral].

Interoperability considerations:

N/A

Published specification:

[I-D.hartke-t2trg-coral]

Applications that use this media type:

See Section 1 of [I-D.hartke-t2trg-coral].

Fragment identifier considerations:

As specified for "application/cbor".

Additional information:

Deprecated alias names for this type: N/A

Magic number(s): N/A

File extension(s): .coral.cbor

Macintosh file type code(s): N/A

Person & email address to contact for further information:

See the Author's Address section of [I-D.hartke-t2trg-coral].

Intended usage:

COMMON

Restrictions on usage:

N/A

Author:

See the Author's Address section of [I-D.hartke-t2trg-coral].

Change controller:

IESG

Provisional registration?

No

7.2. Media Type "text/coral"

This document registers the media type "text/coral" according to the procedures of BCP 13 [RFC6838] and guidelines in RFC 6657 [RFC6657].

Type name:
text

Subtype name:
coral

Required parameters:
N/A

Optional parameters:
N/A

Encoding considerations:
binary - See Section 4 of [I-D.hartke-t2trg-coral].

Security considerations:
See Section 6 of [I-D.hartke-t2trg-coral].

Interoperability considerations:
N/A

Published specification:
[I-D.hartke-t2trg-coral]

Applications that use this media type:
See Section 1 of [I-D.hartke-t2trg-coral].

Fragment identifier considerations:
N/A

Additional information:
Deprecated alias names for this type: N/A
Magic number(s): N/A
File extension(s): .coral
Macintosh file type code(s): N/A

Person & email address to contact for further information:
See the Author's Address section of [I-D.hartke-t2trg-coral].

Intended usage:
COMMON

Restrictions on usage:

N/A

Author:

See the Author's Address section of [I-D.hartke-t2trg-coral].

Change controller:

IESG

Provisional registration?

No

7.3. CoAP Content Formats

This document registers CoAP content formats for the content types "application/coral+cbor" and "text/coral" according to the procedures of RFC 7252 [RFC7252].

- o Content Type: application/coral+cbor
Content Coding: identity
ID: TBD3
Reference: [I-D.hartke-t2trg-coral]
- o Content Type: text/coral
Content Coding: identity
ID: TBD4
Reference: [I-D.hartke-t2trg-coral]

[[NOTE TO RFC EDITOR: Please replace all occurrences of "TBD3" and "TBD4" in this document with the code points assigned by IANA.]]

[[NOTE TO IMPLEMENTERS: Experimental implementations can use content format ID 65087 for "application/coral+cbor" and content format ID 65343 for "text/coral" until IANA has assigned code points.]]

7.4. CBOR Tag

This document registers a CBOR tag for dictionary references according to the procedures of RFC 7049 [RFC7049].

- o Tag: TBD6
Data Item: unsigned integer
Semantics: Dictionary reference
Reference: [I-D.hartke-t2trg-coral]

[[NOTE TO RFC EDITOR: Please replace all occurrences of "TBD6" in this document with the code point assigned by IANA.]]

8. References

8.1. Normative References

- [I-D.hartke-t2trg-ciri]
Hartke, K., "Constrained Internationalized Resource Identifiers", draft-hartke-t2trg-ciri-03 (work in progress), July 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<https://www.rfc-editor.org/info/rfc3339>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<https://www.rfc-editor.org/info/rfc3987>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<https://www.rfc-editor.org/info/rfc5646>>.
- [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/info/rfc6454>>.

- [RFC6657] Melnikov, A. and J. Reschke, "Update to MIME regarding "charset" Parameter Handling in Textual Media Types", RFC 6657, DOI 10.17487/RFC6657, July 2012, <<https://www.rfc-editor.org/info/rfc6657>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/info/rfc6838>>.
- [RFC6943] Thaler, D., Ed., "Issues in Identifier Comparison for Security Purposes", RFC 6943, DOI 10.17487/RFC6943, May 2013, <<https://www.rfc-editor.org/info/rfc6943>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.
- [UNICODE] The Unicode Consortium, "The Unicode Standard", <<http://www.unicode.org/versions/latest/>>.

Note that this reference is to the latest version of Unicode, rather than to a specific release. It is not expected that future changes in the Unicode specification will have any impact on this document.

- [UNICODE-UAX15]
The Unicode Consortium, "Unicode Standard Annex #15: Unicode Normalization Forms", <<http://unicode.org/reports/tr15/>>.
- [UNICODE-UAX31]
The Unicode Consortium, "Unicode Standard Annex #31: Unicode Identifier and Pattern Syntax", <<http://unicode.org/reports/tr31/>>.

[UNICODE-UAX36]

The Unicode Consortium, "Unicode Standard Annex #36: Unicode Security Considerations",
<<http://unicode.org/reports/tr36/>>.

8.2. Informative References

[CAPEC-197]

MITRE, "CAPEC-197: XML Entity Expansion", July 2018,
<<https://capec.mitre.org/data/definitions/197.html>>.

[CAPEC-201]

MITRE, "CAPEC-201: XML Entity Linking", July 2018,
<<https://capec.mitre.org/data/definitions/201.html>>.

[FOAF]

Brickley, D. and L. Miller, "FOAF Vocabulary Specification 0.99", January 2014,
<<http://xmlns.com/foaf/spec/20140114.html>>.

[RFC4287]

Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, DOI 10.17487/RFC4287, December 2005, <<https://www.rfc-editor.org/info/rfc4287>>.

[RFC5789]

Dusseault, L. and J. Snell, "PATCH Method for HTTP", RFC 5789, DOI 10.17487/RFC5789, March 2010,
<<https://www.rfc-editor.org/info/rfc5789>>.

[RFC6573]

Amundsen, M., "The Item and Collection Link Relations", RFC 6573, DOI 10.17487/RFC6573, April 2012,
<<https://www.rfc-editor.org/info/rfc6573>>.

[RFC6690]

Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012,
<<https://www.rfc-editor.org/info/rfc6690>>.

[RFC7228]

Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014,
<<https://www.rfc-editor.org/info/rfc7228>>.

[RFC7230]

Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014,
<<https://www.rfc-editor.org/info/rfc7230>>.

- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7320] Nottingham, M., "URI Design and Ownership", BCP 190, RFC 7320, DOI 10.17487/RFC7320, July 2014, <<https://www.rfc-editor.org/info/rfc7320>>.
- [RFC8132] van der Stok, P., Bormann, C., and A. Sehgal, "PATCH and FETCH Methods for the Constrained Application Protocol (CoAP)", RFC 8132, DOI 10.17487/RFC8132, April 2017, <<https://www.rfc-editor.org/info/rfc8132>>.
- [RFC8288] Nottingham, M., "Web Linking", RFC 8288, DOI 10.17487/RFC8288, October 2017, <<https://www.rfc-editor.org/info/rfc8288>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [W3C.REC-html52-20171214]
Faulkner, S., Eicholz, A., Leithead, T., Danilo, A., and S. Moon, "HTML 5.2", World Wide Web Consortium Recommendation REC-html52-20171214, December 2017, <<https://www.w3.org/TR/2017/REC-html52-20171214>>.
- [W3C.REC-rdf-schema-20140225]
Brickley, D. and R. Guha, "RDF Schema 1.1", World Wide Web Consortium Recommendation REC-rdf-schema-20140225, February 2014, <<http://www.w3.org/TR/2014/REC-rdf-schema-20140225>>.
- [W3C.REC-rdf11-concepts-20140225]
Cyganiak, R., Wood, D., and M. Lanthaler, "RDF 1.1 Concepts and Abstract Syntax", World Wide Web Consortium Recommendation REC-rdf11-concepts-20140225, February 2014, <<http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225>>.

[W3C.REC-turtle-20140225]

Prud'hommeaux, E. and G. Carothers, "RDF 1.1 Turtle",
World Wide Web Consortium Recommendation REC-turtle-
20140225, February 2014,
<<http://www.w3.org/TR/2014/REC-turtle-20140225>>.

[W3C.REC-webarch-20041215]

Jacobs, I. and N. Walsh, "Architecture of the World Wide
Web, Volume One", World Wide Web Consortium
Recommendation REC-webarch-20041215, December 2004,
<<http://www.w3.org/TR/2004/REC-webarch-20041215>>.

Appendix A. Core Vocabulary

This section defines the core vocabulary for CoRAL: a set of link relation types, operation types, form field types, and metadata names.

A.1. Base

Link Relation Types:

<<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>>

Indicates that the link's context is an instance of the class specified as the link's target, as defined by RDF Schema [W3C.REC-rdf-schema-20140225].

<<http://coreapps.org/base#lang>>

Indicates that the link target is a language tag [RFC5646] that specifies the language of the link context.

The link target MUST be a text string in the format specified in Section 2.1 of RFC 5646 [RFC5646].

Operation Types:

<<http://coreapps.org/base#update>>

Indicates that the state of the form's context can be replaced with the state described by a representation submitted to the server.

This operation type defaults to the PUT method [RFC7231] [RFC7252] for both HTTP and CoAP. Typical overrides by a form field include the PATCH method [RFC5789] [RFC8132] for HTTP and CoAP and the iPATCH method [RFC8132] for CoAP.

<<http://coreapps.org/base#search>>

Indicates that the form's context can be searched by submitting a search query.

This operation type defaults to the POST method [RFC7231] for HTTP and the FETCH method [RFC8132] for CoAP. Typical overrides by a form field include the POST method [RFC7252] for CoAP.

A.2. Collections

Link Relation Types:

<http://www.iana.org/assignments/relation/item>

Indicates that the link's context is a collection and that the link's target is a member of that collection, as defined in Section 2.1 of RFC 6573 [RFC6573].

<http://www.iana.org/assignments/relation/collection>

Indicates that the link's target is a collection and that the link's context is a member of that collection, as defined in Section 2.2 of RFC 6573 [RFC6573].

Operation Types:

<http://coreapps.org/collections#create>

Indicates that the form's context is a collection and that a new item can be created in that collection with the state defined by a representation submitted to the server.

This operation type defaults to the POST method [RFC7231] [RFC7252] for both HTTP and CoAP.

<http://coreapps.org/collections#delete>

Indicates that the form's context is a member of a collection and that the form's context can be removed from that collection.

This operation type defaults to the DELETE method [RFC7231] [RFC7252] for both HTTP and CoAP.

A.3. HTTP

Form Field Types:

<http://coreapps.org/http#method>

Specifies the HTTP method for the request.

The form field value MUST be a text string in the format defined in Section 4.1 of RFC 7231 [RFC7231]. The set of possible values is maintained in the IANA HTTP Method Registry.

A form field of this type MUST NOT occur more than once in a form. If absent, it defaults to the request method implied by the form's operation type.

<http://coreapps.org/http#accept>

Specifies an acceptable HTTP content type for the request payload. There may be multiple form fields of this type. If a form does not include a form field of this type, the server accepts any or no request payload, depending on the operation type.

The form field value MUST be a text string in the format defined in Section 3.1.1.1 of RFC 7231 [RFC7231]. The possible set of media types and their parameters are maintained in the IANA Media Types Registry.

Representation Metadata:

<http://coreapps.org/http#type>

Specifies the HTTP content type of the representation.

The metadata value MUST be specified as a text string in the format defined in Section 3.1.1.1 of RFC 7231 [RFC7231]. The possible set of media types and their parameters are maintained in the IANA Media Types Registry.

Metadata of this type MUST NOT occur more than once for a representation. If absent, its value defaults to content type "application/octet-stream".

A.4. CoAP

Form Field Types:

<http://coreapps.org/coap#method>

Specifies the CoAP method for the request.

The form field value MUST be an integer identifying one of the CoAP request methods maintained in the IANA CoAP Method Codes Registry (e.g., the integer 2 for the POST method).

A form field of this type MUST NOT occur more than once in a form. If absent, it defaults to the request method implied by the form's operation type.

<http://coreapps.org/coap#accept>

Specifies an acceptable CoAP content format for the request payload. There may be multiple form fields of this type. If a form does not include a form field of this type, the server

accepts any or no request payload, depending on the operation type.

The form field value MUST be an integer identifying one of content formats maintained in the IANA CoAP Content-Formats Registry.

Representation Metadata:

<http://coreapps.org/coap#type>

Specifies the CoAP content format of the representation.

The metadata value MUST be an integer identifying one of content formats maintained in the IANA CoAP Content-Formats Registry.

Metadata of this type MUST NOT occur more than once for a representation. If absent, it defaults to content format 42 (i.e., content type "application/octet-stream" without a content coding).

Appendix B. Default Dictionary

This section defines a default dictionary that is assumed when the "application/coral+cbor" media type is used without a "dictionary" parameter.

Key	Value
0	<http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
1	<http://www.iana.org/assignments/relation/item>
2	<http://www.iana.org/assignments/relation/collection>
3	<http://coreapps.org/collections#create>
4	<http://coreapps.org/base#update>
5	<http://coreapps.org/collections#delete>
6	<http://coreapps.org/base#search>
7	<http://coreapps.org/coap#accept>
8	<http://coreapps.org/coap#type>
9	<http://coreapps.org/base#lang>
10	<http://coreapps.org/coap#method>

Table 2: Default Dictionary

Acknowledgements

Thanks to Christian Amsuess, Carsten Bormann, Jaime Jimenez, Sebastian Kaebisch, Ari Keranen, Michael Koster, Matthias Kovatsch,

Jim Schaad, and Niklas Widell for helpful comments and discussions that have shaped the document.

Author's Address

Klaus Hartke
Ericsson
Torshamnsgatan 23
Stockholm SE-16483
Sweden

Email: klaus.hartke@ericsson.com

Thing-to-Thing Research Group
Internet-Draft
Intended status: Experimental
Expires: November 10, 2020

K. Hartke
Ericsson
May 9, 2020

Thing-to-Thing Data Hub
draft-hartke-t2trg-data-hub-06

Abstract

A "Thing-to-Thing Data Hub" is a RESTful, hypermedia-driven Web application that can be used in Thing-to-Thing communications to share data items such as thing descriptions, configurations, resource descriptions, or firmware updates at a central location.

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 <https://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 10, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Notational Conventions	4
2.	Data Hubs	4
2.1.	Data Items	4
2.1.1.	Data Item Representation	4
2.2.	Data Collections	4
2.2.1.	Data Collection Representation	5
2.2.2.	Filter Query Representation	5
2.3.	Data Hub Discovery	5
2.4.	Interactions	6
2.4.1.	Getting All Data Items	6
2.4.2.	Getting Data Items by Metadata	6
2.4.3.	Creating a Data Item	6
2.4.4.	Reading a Data Item	7
2.4.5.	Observing a Data Item	7
2.4.6.	Updating a Data Item	7
2.4.7.	Deleting a Data Item	7
3.	Security Considerations	7
4.	IANA Considerations	8
5.	References	8
5.1.	Normative References	8
5.2.	Informative References	9
	Acknowledgements	10
	Author's Address	10

1. Introduction

In Thing-to-Thing communication, there is often a need to share data items of common interest through a central location. For example, the Resource Directory [I-D.ietf-core-resource-directory] aggregates descriptions of Web resources held on constrained nodes, which enables other nodes to easily discover these resources; a Thing Directory [W3C.CR-wot-architecture-20190516] stores metadata of IoT devices, allowing clients to discover interaction affordances and supported protocol bindings of Things; a Firmware Server [I-D.ietf-suit-architecture] stores firmware images and manifests, making this data available to deployed devices, commissioning tools, and other services.

As more and more Thing-to-Thing applications are implemented, it becomes increasingly important being able to not only share resource descriptions and firmware updates but also many other kinds of data, such as default configurations for new devices, service locations, or certificate revocation lists. Resource directories and firmware servers are not a good fit for these kinds of data, as they're specialized to their use cases and generally not accepting any other

kinds of data. The creation of new, specialized applications for every type of data is not practical in the long term.

This document defines a simple "data hub" application, a RESTful Web application with a machine-understandable hypermedia API. A "data hub" generalizes the concept of a central repository for different applications and is suitable for constrained environments [RFC7228]. Specifically, it enables clients to share data items in any format and provides means for creating, reading, observing, updating, deleting, and finding data items at a data hub server.

Data hubs are primarily intended to be accessible over the Constrained Application Protocol (CoAP) [RFC7252].

Features:

- o General

The data hub generalizes the concept of a directory or repository to data items of any Internet media type. This means that applications using the data hub aren't stuck forever with the same media types or limited to just resource descriptions or firmware updates.

- o Searchable

Clients can retrieve a subset of data items from a data hub based on item metadata.

- o Observable

Data items published to a data hub are exposed as resources. As such, they can be observed for changes [RFC7641] over CoAP. This allows clients to stay informed of information that other clients update over time. As a result, the data hub functions similar to a Publish-Subscribe Broker [I-D.ietf-core-coap-pubsub].

- o Evolvable

The key differentiator of the data hub compared to Resource Directory [I-D.ietf-core-resource-directory] and CoAP Publish-Subscribe Broker [I-D.ietf-core-coap-pubsub] lies in the evolvability of the application -- the ability to respond effectively to the need for changes without negatively impacting existing and new clients.

Data hubs enable fine-grained evolvability by driving all interactions by machine-understandable hypermedia elements.

Features can be added, changed or removed in a safe, backwards-compatible way simply by updating the data hub representation to expose appropriate links and forms.

1.1. Notational Conventions

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

2. Data Hubs

The "Thing-to-Thing Data Hub" application consists of two types of resources: a "data collection" and a number of "data items" that have been shared (Figure 1).




Figure 1: Resources of a Data Hub

2.1. Data Items

A data item is a resource that is a member of the data collection resource.

2.1.1. Data Item Representation

The representation of a data item can be of any media type. However, a data hub can restrict the media types it accepts for publication (see Section 2.2.1).

2.2. Data Collections

A data collection is a collection resource that contains data item resources.

Design Note: In this version of this document, a data hub has only a depth of one level; i.e., all data item resources are organized directly under the top-level data collection resource. This could be extended to multiple levels in a future version.

2.2.1. Data Collection Representation

The representation of a data collection is a CoRAL document [I-D.ietf-core-coral] containing the following elements:

A form of type `<http://coreapps.org/collections#create>` containing the following fields:

For each content format accepted for publication, a form field of type `<http://coreapps.org/coap#accept>` indicating that content format. The absence of any form fields of this type indicates that any content format is accepted.

A form of type `<http://coreapps.org/base#search>`.

For each data item in the data collection, a link of type `<http://www.iana.org/assignments/relation/item>` [RFC6573] targeting the data item and containing the following elements:

A form of type `<http://coreapps.org/base#update>`.

A form of type `<http://coreapps.org/collections#delete>`.

Optionally, a (complete or partial) embedded representation of the data item.

The document MAY additionally contain other links and forms not described in this document. For example, a document could contain a link with the `<http://www.iana.org/assignments/relation/alternate>` link relation type [W3C.REC-html52-20171214] that references an alternate representation of a data item.

Any of the links and forms MUST be omitted if following or submitting it can never lead to a successful outcome, for example, because the client is not authorized or the server does not support the feature.

2.2.2. Filter Query Representation

TODO.

2.3. Data Hub Discovery

In this version of this document, clients are assumed to be pre-configured with the URI of a data collection at a data hub.

2.4. Interactions

2.4.1. Getting All Data Items

A client can list all data items in a data collection by making a GET request to the data collection URI (e.g., after discovering the data hub as described in Section 2.3).

On success, the server returns a 2.05 (Content) response with a representation of the collection. As specified in Section 2.2.1 above, this representation includes links to (and, optionally, representations of) the data items in the data collection as well as forms for creating, updating, deleting, and finding data items.

2.4.2. Getting Data Items by Metadata

If the representation of a data collection contains a form of type `<http://coreapps.org/base#search>`, the client can filter the data collection by submitting this form with a search query (see Section 2.2.2).

Implementations of this version of this document MUST use the method implied by the `<http://coreapps.org/base#search>` operation type, i.e., the FETCH method [RFC8132]. Any form indicating a different method MUST be ignored.

On success, the server returns a 2.05 (Content) response with a representation of a list of data items in the collection (see Section 2.2.1) that match the query.

2.4.3. Creating a Data Item

If the representation of a data collection contains a form of type `<http://coreapps.org/collections#create>`, the client can create a new data item in the data collection by submitting this form with a representation in one of the acceptable media types. The acceptable media types are indicated by form fields of type `<http://coreapps.org/coap#accept>`.

Implementations of this version of this document MUST use the method implied by the `<http://coreapps.org/collections#create>` operation type, i.e., the POST method [RFC7252]. Any form indicating a different method MUST be ignored.

On success, the server returns a 2.01 (Created) response. The location of the created data item is conveyed using the Location-Path and Location-Query options [RFC7252].

2.4.4. Reading a Data Item

A client can read a data item by following a link with the `<http://www.iana.org/assignments/relation/item>` link relation type in the representation of the data collection.

2.4.5. Observing a Data Item

A client can observe a data item by following a link with the `<http://www.iana.org/assignments/relation/item>` link relation type in the representation of the data collection and observing the target resource as specified in RFC 7641 [RFC7641].

2.4.6. Updating a Data Item

If the representation of a data collection includes a form of type `<http://coreapps.org/base#update>` nested within the link to a data item, a client can update the data item by submitting this form with a representation of the updated data item.

Implementations of this version of this document MUST use the method implied by the `<http://coreapps.org/base#update>` operation type, i.e., the PUT method [RFC7252]. Any form indicating a different method MUST be ignored.

On success, the server returns a 2.04 (Changed) response.

2.4.7. Deleting a Data Item

If the representation of a data collection includes a form of type `<http://coreapps.org/collections#delete>` nested within the link to a data item, the client can delete the data item by submitting this form.

Implementations of this version of this document MUST use the method implied by the `<http://coreapps.org/collections#delete>` operation type, i.e., the DELETE method [RFC7252]. Any form indicating a different method MUST be ignored.

On success, the server returns a 2.02 (Deleted) response.

3. Security Considerations

The data hub application relies on a Web transfer protocol like CoAP to exchange representations in a CoRAL serialization format. See Section 11 of RFC 7252 [RFC7252] and Section 7 of RFC 7641 [RFC7641] for security considerations relating to CoAP. See Section XX of RFC

XXXX [I-D.ietf-core-coral] for security considerations relating to CoRAL.

The data hub application does not define any specific mechanisms for protecting the confidentiality and integrity of messages exchanged between a data hub and a client. It is recommended that implementations employ application layer or transport layer mechanisms for interactions with a data hub.

The data hub application does not define any specific mechanisms for protecting the confidentiality and integrity of representations of data items shared through a data hub. For scenarios where end-to-end security matters, such as for firmware updates [I-D.ietf-suit-information-model], implementations should employ an object security mechanism.

4. IANA Considerations

This document has no IANA actions.

5. References

5.1. Normative References

- [I-D.ietf-core-coral]
Hartke, K., "The Constrained RESTful Application Language (CoRAL)", draft-ietf-core-coral-03 (work in progress), March 2020.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6573] Amundsen, M., "The Item and Collection Link Relations", RFC 6573, DOI 10.17487/RFC6573, April 2012, <<https://www.rfc-editor.org/info/rfc6573>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", RFC 7641, DOI 10.17487/RFC7641, September 2015, <<https://www.rfc-editor.org/info/rfc7641>>.

[RFC8132] van der Stok, P., Bormann, C., and A. Sehgal, "PATCH and FETCH Methods for the Constrained Application Protocol (CoAP)", RFC 8132, DOI 10.17487/RFC8132, April 2017, <<https://www.rfc-editor.org/info/rfc8132>>.

5.2. Informative References

- [I-D.ietf-core-coap-pubsub]
Koster, M., Keranen, A., and J. Jimenez, "Publish-Subscribe Broker for the Constrained Application Protocol (CoAP)", draft-ietf-core-coap-pubsub-09 (work in progress), September 2019.
- [I-D.ietf-core-resource-directory]
Shelby, Z., Koster, M., Bormann, C., Stok, P., and C. Amsuess, "CoRE Resource Directory", draft-ietf-core-resource-directory-24 (work in progress), March 2020.
- [I-D.ietf-suit-architecture]
Moran, B., Tschofenig, H., Brown, D., and M. Meriac, "A Firmware Update Architecture for Internet of Things", draft-ietf-suit-architecture-08 (work in progress), November 2019.
- [I-D.ietf-suit-information-model]
Moran, B., Tschofenig, H., and H. Birkholz, "An Information Model for Firmware Updates in IoT Devices", draft-ietf-suit-information-model-05 (work in progress), January 2020.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [W3C.CR-wot-architecture-20190516]
Kovatsch, M., Matsukura, R., Lagally, M., Kawaguchi, T., Toumura, K., and K. Kajimoto, "Web of Things (WoT) Architecture", World Wide Web Consortium Candidate Recommendation CR-wot-architecture-20190516, May 2019, <<https://www.w3.org/TR/2019/CR-wot-architecture-20190516>>.
- [W3C.REC-html52-20171214]
Faulkner, S., Eicholz, A., Leithead, T., Danilo, A., and S. Moon, "HTML 5.2", World Wide Web Consortium Recommendation REC-html52-20171214, December 2017, <<https://www.w3.org/TR/2017/REC-html52-20171214>>.

Acknowledgements

Thanks to Christian Amsuess and Jaime Jimenez for helpful comments and discussions that have shaped the document.

Author's Address

Klaus Hartke
Ericsson
Torshamnsgatan 23
Stockholm SE-16483
Sweden

Email: klaus.hartke@ericsson.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 16, 2019

O. Garcia-Morchon
Philips IP&S
S. Kumar
Philips Research
M. Sethi
Ericsson
December 13, 2018

State-of-the-Art and Challenges for the Internet of Things Security
draft-irtf-t2trg-iot-seccons-16

Abstract

The Internet of Things (IoT) concept refers to the usage of standard Internet protocols to allow for human-to-thing and thing-to-thing communication. The security needs for IoT systems are well-recognized and many standardization steps to provide security have been taken, for example, the specification of Constrained Application Protocol (CoAP) secured with Datagram Transport Layer Security (DTLS). However, security challenges still exist, not only because there are some use cases that lack a suitable solution, but also because many IoT devices and systems have been designed and deployed with very limited security capabilities. In this document, we first discuss the various stages in the lifecycle of a thing. Next, we document the security threats to a thing and the challenges that one might face to protect against these threats. Lastly, we discuss the next steps needed to facilitate the deployment of secure IoT systems. This document can be used by implementors and authors of IoT specifications as a reference for details about security considerations while documenting their specific security challenges, threat models, and mitigations.

This document is a product of the IRTF Thing-to-Thing Research Group (T2TRG).

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 <https://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 June 16, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. The Thing Lifecycle	4
3. Security Threats and Managing Risk	7
4. State-of-the-Art	11
4.1. IP-based IoT Protocols and Standards	11
4.2. Existing IP-based Security Protocols and Solutions	14
4.3. IoT Security Guidelines	16
5. Challenges for a Secure IoT	19
5.1. Constraints and Heterogeneous Communication	19
5.1.1. Resource Constraints	19
5.1.2. Denial-of-Service Resistance	20
5.1.3. End-to-end security, protocol translation, and the role of middleboxes	21
5.1.4. New network architectures and paradigm	23
5.2. Bootstrapping of a Security Domain	23
5.3. Operational Challenges	24
5.3.1. Group Membership and Security	24
5.3.2. Mobility and IP Network Dynamics	25
5.4. Secure software update and cryptographic agility	26
5.5. End-of-Life	28
5.6. Verifying device behavior	28
5.7. Testing: bug hunting and vulnerabilities	29
5.8. Quantum-resistance	30
5.9. Privacy protection	31
5.10. Reverse engineering considerations	32
5.11. Trustworthy IoT Operation	33

6. Conclusions and Next Steps	34
7. Security Considerations	34
8. IANA Considerations	34
9. Acknowledgments	35
10. Informative References	35
Authors' Addresses	47

1. Introduction

The Internet of Things (IoT) denotes the interconnection of highly heterogeneous networked entities and networks that follow a number of different communication patterns such as: human-to-human (H2H), human-to-thing (H2T), thing-to-thing (T2T), or thing-to-things (T2Ts). The term IoT was first coined by the Auto-ID center [AUTO-ID] in 1999 which had envisioned a world where every physical object is tagged with a radio-frequency identification (RFID) tag having a globally unique identifier. This would not only allow tracking of objects in real-time but also allow querying of data about them over the Internet. However, since then, the meaning of the Internet of Things has expanded and now encompasses a wide variety of technologies, objects and protocols. It is not surprising that the IoT has received significant attention from the research community to (re)design, apply, and use standard Internet technology and protocols for the IoT.

The things that are part of the Internet of Things are computing devices that understand and react to the environment they reside in. These things are also often referred to as smart objects or smart devices. The introduction of IPv6 [RFC6568] and CoAP [RFC7252] as fundamental building blocks for IoT applications allows connecting IoT hosts to the Internet. This brings several advantages including: (i) a homogeneous protocol ecosystem that allows simple integration with other Internet hosts; (ii) simplified development for devices that significantly vary in their capabilities; (iii) a unified interface for applications, removing the need for application-level proxies. These building blocks greatly simplify the deployment of the envisioned scenarios which range from building automation to production environments and personal area networks.

This document presents an overview of important security aspects for the Internet of Things. We begin by discussing the lifecycle of a thing in Section 2. In Section 3, we discuss security threats for the IoT and methodologies for managing these threats when designing a secure system. Section 4 reviews existing IP-based (security) protocols for the IoT and briefly summarizes existing guidelines and regulations. Section 5 identifies remaining challenges for a secure IoT and discusses potential solutions. Section 6 includes final remarks and conclusions. This document can be used by IoT standards

specifications as a reference for details about security considerations applying to the specified system or protocol.

The first draft version of this document was submitted in March 2011. Initial draft versions of this document were presented and discussed during the CORE meetings at IETF 80 and later. Discussions on security lifecycle at IETF 92 (March 2015) evolved into more general security considerations. Thus, the draft was selected to address the T2TRG work item on the security considerations and challenges for the Internet of Things. Further updates of the draft were presented and discussed during the T2TRG meetings at IETF 96 (July 2016) and IETF 97 (November 2016) and at the joint interim in Amsterdam (March 2017). This document has been reviewed by, commented on, and discussed extensively for a period of nearly six years by a vast majority of T2TRG and related group members; the number of which certainly exceeds 100 individuals. It is the consensus of T2TRG that the security considerations described in this document should be published in the IRTF Stream of the RFC series. This document does not constitute a standard.

2. The Thing Lifecycle

The lifecycle of a thing refers to the operational phases of a thing in the context of a given application or use case. Figure 1 shows the generic phases of the lifecycle of a thing. This generic lifecycle is applicable to very different IoT applications and scenarios. For instance, [RFC7744] provides an overview of relevant IoT use cases.

In this document, we consider a Building Automation and Control (BAC) system to illustrate the lifecycle and the meaning of these different phases. A BAC system consists of a network of interconnected nodes that performs various functions in the domains of HVAC (Heating, Ventilating, and Air Conditioning), lighting, safety, etc. The nodes vary in functionality and a large majority of them represent resource-constrained devices such as sensors and luminaries. Some devices may be battery operated or may rely on energy harvesting. This requires us to also consider devices that sleep during their operation to save energy. In our BAC scenario, the life of a thing starts when it is manufactured. Due to the different application areas (i.e., HVAC, lighting, or safety) nodes/things are tailored to a specific task. It is therefore unlikely that one single manufacturer will create all nodes in a building. Hence, interoperability as well as trust bootstrapping between nodes of different vendors is important.

The thing is later installed and commissioned within a network by an installer during the bootstrapping phase. Specifically, the device

identity and the secret keys used during normal operation may be provided to the device during this phase. Different subcontractors may install different IoT devices for different purposes. Furthermore, the installation and bootstrapping procedures may not be a discrete event and may stretch over an extended period. After being bootstrapped, the device and the system of things are in operational mode and execute the functions of the BAC system. During this operational phase, the device is under the control of the system owner and used by multiple system users. For devices with lifetimes spanning several years, occasional maintenance cycles may be required. During each maintenance phase, the software on the device can be upgraded or applications running on the device can be reconfigured. The maintenance tasks can be performed either locally or from a backend system. Depending on the operational changes to the device, it may be required to re-bootstrap at the end of a maintenance cycle. The device continues to loop through the operational phase and the eventual maintenance phases until the device is decommissioned at the end of its lifecycle. However, the end-of-life of a device does not necessarily mean that it is defective and rather denotes a need to replace and upgrade the network to the next-generation devices for additional functionality. Therefore, the device can be removed and re-commissioned to be used in a different system under a different owner thereby starting the lifecycle all over again.

We note that the presented lifecycle represents to some extent a simplified model. For instance, it is possible to argue that the lifecycle does not start when a tangible device is manufactured but rather when the oldest bit of code that ends up in the device - maybe from an open source project or from the used operating system - was written. Similarly, the lifecycle could also include an on-the-shelf phase where the device is in the supply-chain before an owner/user purchases and installs it. Another phase could involve the device being re-badged by some vendor who is not the original manufacturer. Such phases can significantly complicate other phases such as maintenance and bootstrapping. Finally, other potential end-states can be, e.g., a vendor that no longer supports a device type because it is at end-of-life or a situation in which a device is simply forgotten but remains functional.




Figure 1: The lifecycle of a thing in the Internet of Things

Security is a key requirement in any communication system. However, security is an even more critical requirement in real-world IoT deployments for several reasons. First, compromised IoT systems can not only endanger the privacy and security of a user, but can also cause physical harm. This is because IoT systems often comprise sensors, actuators and other connected devices in the physical environment of the user which could adversely affect the user if they are compromised. Second, a vulnerable IoT system means that an attacker can alter the functionality of a device from a given manufacturer. This not only affects the manufacturer's brand image, but can also leak information that is very valuable for the manufacturer (such as proprietary algorithms). Third, the impact of attacking an IoT system goes beyond a specific device or an isolated system since compromised IoT systems can be misused at scale. For example, they may be used to perform a Distributed Denial of Service (DDoS) attack that limits the availability of other networks and services. The fact that many IoT systems rely on standard IP protocols allows for easier system integration, but this also makes attacks on standard IP protocols widely applicable in other environments. This results in new requirements regarding the implementation of security.

The term security subsumes a wide range of primitives, protocols, and procedures. For instance, the term security includes services such as confidentiality, authentication, integrity, authorization, source authentication, and availability. The term security often also includes augmented services such as duplicate detection and detection of stale packets (timeliness). These security services can be implemented through a combination of cryptographic mechanisms such as block ciphers, hash functions, and signature algorithms; as well as

non-cryptographic mechanisms that implement authorization and other security policy enforcement aspects. For ensuring security in IoT networks, one should not only focus on the required security services, but also pay special attention to how the services are realized in the overall system.

3. Security Threats and Managing Risk

Security threats in related IP protocols have been analyzed in multiple documents including Hypertext Transfer Protocol (HTTP) over Transport Layer Security (TLS) (HTTPS) [RFC2818], Constrained Application Protocol (COAP) [RFC7252], IPv6 over Low-Power Wireless Personal Area Networks (6LoWPAN) [RFC4919], Access Node Control Protocol (ANCP) [RFC5713], Domain Name System (DNS) [RFC3833], IPv6 Neighbor Discovery (ND) [RFC3756], and Protocol for Carrying Authentication and Network Access (PANA) [RFC4016]. In this section, we specifically discuss the threats that could compromise an individual thing or the network as a whole. Some of these threats might go beyond the scope of Internet protocols but we gather them here for the sake of completeness. The threats in the following list are not in any particular order and some threats might be more critical than others depending on the deployment scenario under consideration:

1. **Vulnerable Software/Code:** Things in the Internet of Things rely on software that might contain severe bugs and/or bad design choices. This makes the things vulnerable to many different types of attacks, depending on the criticality of the bugs, e.g., buffer overflows or lack of authentication. This can be considered as one of the most important security threat. The large-scale distributed denial-of-service (DDoS) attack, popularly known as the Mirai botnet [mirai], was caused by things that had well-known or easy-to-guess passwords for configuration.
2. **Privacy threat:** The tracking of a thing's location and usage may pose a privacy risk to people around it. For instance, an attacker can infer privacy sensitive information from the data gathered and communicated by individual things. Such information may subsequently be sold to interested parties for marketing purposes and targeted advertising. In extreme cases, such information might be used to track dissidents in oppressive regimes. Unlawful surveillance and interception of traffic to/from a thing by intelligence agencies is also a privacy threat.
3. **Cloning of things:** During the manufacturing process of a thing, an untrusted factory can easily clone the physical characteristics, firmware/software, or security configuration of

the thing. Deployed things might also be compromised and their software reverse engineered allowing for cloning or software modifications. Such a cloned thing may be sold at a cheaper price in the market, and yet can function normally as a genuine thing. For example, two cloned devices can still be associated and work with each other. In the worst-case scenario, a cloned device can be used to control a genuine device or perform an attack. One should note here, that an untrusted factory may also change functionality of the cloned thing, resulting in degraded functionality with respect to the genuine thing (thereby, inflicting potential damage to the reputation of the original thing manufacturer). Moreover, additional functionality can be introduced in the cloned thing. An example of such functionality is a backdoor.

4. Malicious substitution of things: During the installation of a thing, a genuine thing may be substituted with a similar variant (of lower quality) without being detected. The main motivation may be cost savings, where the installation of lower-quality things (for example, non-certified products) may significantly reduce the installation and operational costs. The installers can subsequently resell the genuine things to gain further financial benefits. Another motivation may be to inflict damage to the reputation of a competitor's offerings.
5. Eavesdropping attack: During the commissioning of a thing into a network, it may be susceptible to eavesdropping, especially if operational keying materials, security parameters, or configuration settings, are exchanged in clear using a wireless medium or if used cryptographic algorithms are not suitable for the envisioned lifetime of the device and the system. After obtaining the keying material, the attacker might be able to recover the secret keys established between the communicating entities, thereby compromising the authenticity and confidentiality of the communication channel, as well as the authenticity of commands and other traffic exchanged over this communication channel. When the network is in operation, T2T communication can be eavesdropped if the communication channel is not sufficiently protected or if a session key is compromised due to protocol weaknesses. An adversary may also be able to eavesdrop if keys are not renewed or updated appropriately. Lastly, messages can also be recorded and decrypted offline at a later point of time. The Venona project [venona-project] is one such example where messages were recorded for offline decryption.
6. Man-in-the-middle attack: Both the commissioning phase and operational phases may also be vulnerable to man-in-the-middle

attacks. For example, when keying material between communicating entities is exchanged in the clear and the security of the key establishment protocol depends on the tacit assumption that no third party can eavesdrop during the execution of this protocol. Additionally, device authentication or device authorization may be non-trivial, or may need support of a human decision process, since things usually do not have a-priori knowledge about each other and cannot always differentiate friends and foes via completely automated mechanisms.

7. Firmware attacks: When a thing is in operation or maintenance phase, its firmware or software may be updated to allow for new functionality or new features. An attacker may be able to exploit such a firmware upgrade by maliciously replacing the thing's firmware, thereby influencing its operational behavior. For example, an attacker could add a piece of malicious code to the firmware that will cause it to periodically report the energy usage of the thing to a data repository for analysis. The attacker can then use this information to determine when a home or enterprise (where the thing is installed) is unoccupied and break in. Similarly, devices whose software has not been properly maintained and updated might contain vulnerabilities that might be exploited by attackers to replace the firmware on the device.
8. Extraction of private information: IoT devices (such as sensors, actuators, etc.) are often physically unprotected in their ambient environment and they could easily be captured by an attacker. An attacker with physical access may then attempt to extract private information such as keys (for example, device's key, private-key, group key), sensed data (for example, healthcare status of a user), configuration parameters (for example, the Wi-Fi key), or proprietary algorithms (for example, algorithm performing some data analytics task). Even when the data originating from a thing is encrypted, attackers can perform traffic analysis to deduce meaningful information which might compromise the privacy of the thing's owner and/or user.
9. Routing attack: As highlighted in [ID-Daniel], routing information in IoT networks can be spoofed, altered, or replayed, in order to create routing loops, attract/repel network traffic, extend/shorten source routes, etc. A non-exhaustive list of routing attacks includes 1) Sinkhole attack (or blackhole attack), where an attacker declares himself to have a high-quality route/path to the base station, thus allowing him to do manipulate all packets passing through it. 2) Selective forwarding, where an attacker may selectively forward

packets or simply drop a packet. 3) Wormhole attack, where an attacker may record packets at one location in the network and tunnel them to another location, thereby influencing perceived network behavior and potentially distorting statistics, thus greatly impacting the functionality of routing. 4) Sybil attack, whereby an attacker presents multiple identities to other things in the network. We refer to [ID-Daniel] for further router attacks and a more detailed description.

10. Elevation of privilege: An attacker with low privileges can misuse additional flaws in the implemented authentication and authorization mechanisms of a thing to gain more privileged access to the thing and its data.
11. Denial-of-Service (DoS) attack: Often things have very limited memory and computation capabilities. Therefore, they are vulnerable to resource exhaustion attack. Attackers can continuously send requests to specific things so as to deplete their resources. This is especially dangerous in the Internet of Things since an attacker might be located in the backend and target resource-constrained devices that are part of a constrained node network [RFC7228]. DoS attack can also be launched by physically jamming the communication channel. Network availability can also be disrupted by flooding the network with a large number of packets. On the other hand, things compromised by attackers can be used to disrupt the operation of other networks or systems by means of a Distributed DoS (DDoS) attack.

To deal with above threats it is required to find and apply suitable security mitigations. However, new threats and exploits appear on a daily basis and products are deployed in different environments prone to different types of threats. Thus, ensuring a proper level of security in an IoT system at any point of time is challenging. To address this challenge, some of the following methodologies can be used:

1. A Business Impact Analysis (BIA) assesses the consequences of the loss of basic security attributes: confidentiality, integrity and availability in an IoT system. These consequences might include the impact from lost data, reduced sales, increased expenses, regulatory fines, customer dissatisfaction, etc. Performing a business impact analysis allows a business to determine the relevance of having a proper security design.
2. A Risk Assessment (RA) analyzes security threats to an IoT system while considering their likelihood and impact. It also includes categorizing each of them with a risk level. Risks classified as

moderate or high must be mitigated, i.e., the security architecture should be able to deal with those threat.

3. A privacy impact assessment (PIA) aims at assessing the Personally Identifiable Information (PII) that is collected, processed, or used in an IoT system. By doing so, the goal is to fulfill applicable legal requirements, determine risks and effects of manipulation and loss of PII.
4. Procedures for incident reporting and mitigation refer to the methodologies that allow becoming aware of any security issues that affect an IoT system. Furthermore, this includes steps towards the actual deployment of patches that mitigate the identified vulnerabilities.

BIA, RA, and PIA should generally be realized during the creation of a new IoT system or when deploying significant system/feature upgrades. In general, it is recommended to re-assess them on a regular basis taking into account new use cases and/or threats. The way a BIA, RA, PIA are performed depends on the environment and the industry. More information can be found in NIST documents such as [NISTSP800-34r1], [NISTSP800-30r1], and [NISTSP800-122].

4. State-of-the-Art

This section is organized as follows. Section 4.1 summarizes state-of-the-art on IP-based IoT systems, within IETF and in other standardization bodies. Section 4.2 summarizes state-of-the-art on IP-based security protocols and their usage. Section 4.3 discusses guidelines and regulations for securing IoT as proposed by other bodies. Note that the references included in this section are a representative of the state-of-the-art at the point of writing and they are by no means exhaustive. The references are also at varying levels of maturity, and thus, it is advisable to review their specific status.

4.1. IP-based IoT Protocols and Standards

Nowadays, there exists a multitude of control protocols for IoT. For BAC systems, the ZigBee standard [ZB], BACNet [BACNET], and DALI [DALI] play key roles. Recent trends, however, focus on an all-IP approach for system control.

In this setting, a number of IETF working groups are designing new protocols for resource-constrained networks of smart things. The 6LoWPAN working group [WG-6LoWPAN] for example has defined methods and protocols for the efficient transmission and adaptation of IPv6 packets over IEEE 802.15.4 networks [RFC4944].

The CoRE working group [WG-CoRE] has specified the Constrained Application Protocol (CoAP) [RFC7252]. CoAP is a RESTful protocol for constrained devices that is modeled after HTTP and typically runs over UDP to enable efficient application-level communication for things.

In many smart object networks, the smart objects are dispersed and have intermittent reachability either because of network outages or because they sleep during their operational phase to save energy. In such scenarios, direct discovery of resources hosted on the constrained server might not be possible. To overcome this barrier, the CoRE working group is specifying the concept of a Resource Directory (RD) [ID-rd]. The Resource Directory hosts descriptions of resources which are located on other nodes. These resource descriptions are specified as CoRE link format [RFC6690].

While CoAP defines a standard communication protocol, a format for representing sensor measurements and parameters over CoAP is required. The Sensor Measurement Lists (SenML) [RFC8428] is a specification that defines media types for simple sensor measurements and parameters. It has a minimalistic design so that constrained devices with limited computational capabilities can easily encode their measurements and, at the same time, servers can efficiently collect large number of measurements.

In many IoT deployments, the resource-constrained smart objects are connected to the Internet via a gateway that is directly reachable. For example, an IEEE 802.11 Access Point (AP) typically connects the client devices to the Internet over just one wireless hop. However, some deployments of smart object networks require routing between the smart objects themselves. The IETF has therefore defined the IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) [RFC6550]. RPL provides support for multipoint-to-point traffic from resource-constrained smart objects towards a more resourceful central control point, as well as point-to-multipoint traffic in the reverse direction. It also supports point-to-point traffic between the resource-constrained devices. A set of routing metrics and constraints for path calculation in RPL are also specified [RFC6551].

The IPv6 over Networks of Resource-constrained Nodes (6lo) [WG-6lo] working group of the IETF has specified how IPv6 packets can be transmitted over various link layer protocols that are commonly employed for resource-constrained smart object networks. There is also ongoing work to specify IPv6 connectivity for a Non-Broadcast Multi-Access (NBMA) mesh network that is formed by IEEE 802.15.4 TimeSlotted Channel Hopping (TSCH) links [ID-6tisch]. Other link layer protocols for which IETF has specified or is currently specifying IPv6 support include Bluetooth [RFC7668], Digital Enhanced

Cordless Telecommunications (DECT) Ultra Low Energy (ULE) air interface [RFC8105], and Near Field Communication (NFC) [ID-6lonfc].

Baker and Meyer [RFC6272] identify which IP protocols can be used in smart grid environments. They give advice to smart grid network designers on how they can decide on a profile of the Internet protocol suite for smart grid networks.

The Low Power Wide-Area Network (LPWAN) working [WG-LPWAN] group is analyzing features, requirements, and solutions to adapt IP-based protocols to networks such as LORA [lora], SigFox [sigfox], NB-IoT [nbiot], etc. These networking technologies enable a smart thing to run for years on a single coin-cell by relying on a star network topology and using optimized radio modulation with frame sizes in the order of tens of bytes. Such networks bring new security challenges since most existing security mechanism do not work well with such resource constraints.

JavaScript Object Notation (JSON) is a lightweight text representation format for structured data [RFC8259]. It is often used for transmitting serialized structured data over the network. IETF has defined specifications for encoding cryptographic keys, encrypted content, signed content, and claims to be transferred between two parties as JSON objects. They are referred to as JSON Web Keys (JWK) [RFC7517], JSON Web Encryption (JWE) [RFC7516], JSON Web Signatures (JWS) [RFC7515] and JSON Web Token (JWT) [RFC7519].

An alternative to JSON, Concise Binary Object Representation (CBOR) [RFC7049] is a concise binary data format that is used for serialization of structured data. It is designed for resource-constrained nodes and therefore it aims to provide a fairly small message size with minimal implementation code, and extensibility without the need for version negotiation. CBOR Object Signing and Encryption (COSE) [RFC8152] specifies how to encode cryptographic keys, message authentication codes, encrypted content, and signatures with CBOR.

The Light-Weight Implementation Guidance (LWIG) working group [WG-LWIG] is collecting experiences from implementers of IP stacks in constrained devices. The working group has already produced documents such as RFC7815 [RFC7815] which defines how a minimal Internet Key Exchange Version 2 (IKEv2) initiator can be implemented.

The Thing-2-Thing Research Group (T2TRG) [RG-T2TRG] is investigating the remaining research issues that need to be addressed to quickly turn the vision of IoT into a reality where resource-constrained nodes can communicate with each other and with other more capable nodes on the Internet.

Additionally, industry alliances and other standardization bodies are creating constrained IP protocol stacks based on the IETF work. Some important examples of this include:

1. Thread [Thread]: Specifies the Thread protocol that is intended for a variety of IoT devices. It is an IPv6-based network protocol that runs over IEEE 802.15.4.
2. Industrial Internet Consortium [IIoT]: The consortium defines reference architectures and security frameworks for development, adoption and widespread use of Industrial Internet technologies based on existing IETF standards.
3. Internet Protocol for Smart Objects IPSO [IPSO]: The alliance specifies a common object model that enables application software on any device to interoperate with other conforming devices.
4. OneM2M [OneM2M]: The standards body defines technical and API specifications for IoT devices. It aims to create a service layer that can run on any IoT device hardware and software.
5. Open Connectivity Foundation (OCF) [OCF]: The foundation develops standards and certifications primarily for IoT devices that use Constrained Application Protocol (CoAP) as the application layer protocol.
6. Fairhair Alliance [Fairhair]: Specifies an IoT middleware to enable a common IP network infrastructure between different application standards used in building automation and lighting systems such as BACnet, KNX and ZigBee.
7. OMA LWM2M [LWM2M]: OMA Lightweight M2M is a standard from the Open Mobile Alliance for M2M and IoT device management. LWM2M relies on CoAP as the application layer protocol and uses a RESTful architecture for remote management of IoT devices.

4.2. Existing IP-based Security Protocols and Solutions

There are three main security objectives for IoT networks: 1. protecting the IoT network from attackers. 2. protecting IoT applications and thus, the things and users. 3. protecting the rest of the Internet and other things from attacks that use compromised things as an attack platform.

In the context of the IP-based IoT deployments, consideration of existing Internet security protocols is important. There are a wide range of specialized as well as general-purpose security solutions for the Internet domain such as IKEv2/IPsec [RFC7296], Transport

Layer Security (TLS) [RFC8446], Datagram Transport Layer Security (DTLS) [RFC6347], Host Identity Protocol (HIP) [RFC7401], PANA [RFC5191], Kerberos ([RFC4120]), Simple Authentication and Security Layer (SASL) [RFC4422], and Extensible Authentication Protocol (EAP) [RFC3748].

TLS provides security for TCP and requires a reliable transport. DTLS secures and uses datagram-oriented protocols such as UDP. Both protocols are intentionally kept similar and share the same ideology and cipher suites. The CoAP base specification [RFC7252] provides a description of how DTLS can be used for securing CoAP. It proposes three different modes for using DTLS: the PreSharedKey mode, where nodes have pre-provisioned keys for initiating a DTLS session with another node, RawPublicKey mode, where nodes have asymmetric-key pairs but no certificates to verify the ownership, and Certificate mode, where public keys are certified by a certification authority. An IoT implementation profile [RFC7925] is defined for TLS version 1.2 and DTLS version 1.2 that offers communication security for resource-constrained nodes.

There is ongoing work to define an authorization and access-control framework for resource-constrained nodes. The Authentication and Authorization for Constrained Environments (ACE) [WG-ACE] working group is defining a solution to allow only authorized access to resources that are hosted on a smart object server and are identified by a URI. The current proposal [ID-aceoauth] is based on the OAuth 2.0 framework [RFC6749] and it comes with profiles intended for different communication scenarios, e.g. DTLS Profile for Authentication and Authorization for Constrained Environments [ID-acedtls].

OSCORE [ID-OSCORE] is a proposal that protects CoAP messages by wrapping them in the CBOR Object Signing and Encryption (COSE) [RFC8152] format. Thus, OSCORE falls in the category of object security and it can be applied wherever CoAP can be used. The advantage of OSCORE over DTLS is that it provides some more flexibility when dealing with end-to-end security. Section 5.1.3 discusses this further.

The Automated Certificate Management Environment (ACME) [WG-ACME] working group is specifying conventions for automated X.509 certificate management. This includes automatic validation of certificate issuance, certificate renewal, and certificate revocation. While the initial focus of working group is on domain name certificates (as used by web servers), other uses in some IoT deployments is possible.

The Internet Key Exchange (IKEv2)/IPsec - as well as the less used Host Identity protocol (HIP) - reside at or above the network layer in the OSI model. Both protocols are able to perform an authenticated key exchange and set up the IPsec for secure payload delivery. Currently, there are also ongoing efforts to create a HIP variant coined Diet HIP [ID-HIP-DEX] that takes constrained networks and nodes into account at the authentication and key exchange level.

Migault et al. [ID-dietesp] are working on a compressed version of IPsec so that it can easily be used by resource-constrained IoT devices. They rely on the Internet Key Exchange Protocol version 2 (IKEv2) for negotiating the compression format.

The Extensible Authentication Protocol (EAP) [RFC3748] is an authentication framework supporting multiple authentication methods. EAP runs directly over the data link layer and, thus, does not require the deployment of IP. It supports duplicate detection and retransmission, but does not allow for packet fragmentation. The Protocol for Carrying Authentication for Network Access (PANA) is a network-layer transport for EAP that enables network access authentication between clients and the network infrastructure. In EAP terms, PANA is a UDP-based EAP lower layer that runs between the EAP peer and the EAP authenticator.

4.3. IoT Security Guidelines

Attacks on and from IoT devices have become common in the last years, for instance, large scale Denial of Service (DoS) attacks on the Internet Infrastructure from compromised IoT devices. This fact has prompted many different standards bodies and consortia to provide guidelines for developers and the Internet community at large to build secure IoT devices and services. A subset of the different guidelines and ongoing projects are as follows:

1. Global System for Mobile Communications (GSM) Association (GSMA) IoT security guidelines [GSMAsecurity]: GSMA has published a set of security guidelines for the benefit of new IoT product and service providers. The guidelines are aimed at device manufacturers, service providers, developers and network operators. An enterprise can complete an IoT Security Self-Assessment to demonstrate that its products and services are aligned with the security guidelines of the GSMA.
2. Broadband Internet Technical Advisory Group (BITAG) IoT Security and Privacy Recommendations [BITAG]: BITAG has published recommendations for ensuring security and privacy of IoT device users. BITAG observes that many IoT devices are shipped from the factory with software that is already outdated and

vulnerable. The report also states that many devices with vulnerabilities will not be fixed either because the manufacturer does not provide updates or because the user does not apply them. The recommendations include that IoT devices should function without cloud and Internet connectivity, and that all IoT devices should have methods for automatic secure software updates.

3. United Kingdom Department for Digital, Culture, Media and Sport (DCMS) [DCMS]: UK DCMS has released a report that includes a list of 13 steps for improving IoT security. These steps, for example, highlight the need for implementing a vulnerability disclosure policy and keeping software updated. The report is aimed at device manufacturers, IoT service providers, mobile application developers and retailers.
4. Cloud Security Alliance (CSA) New Security Guidance for Early Adopters of the IoT [CSA]: CSA recommendations for early adopters of IoT encourages enterprises to implement security at different layers of the protocol stack. It also recommends implementation of an authentication/authorization framework for IoT deployments. A complete list of recommendations is available in the report [CSA].
5. United States Department of Homeland Security [DHS]: DHS has put forth six strategic principles that would enable IoT developers, manufacturers, service providers and consumers to maintain security as they develop, manufacture, implement or use network-connected IoT devices.
6. National Institute of Standards and Technology (NIST) [NIST-Guide]: The NIST special publication urges enterprise and US federal agencies to address security throughout the systems engineering process. The publication builds upon the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 15288 standard and augments each process in the system lifecycle with security enhancements.
7. National Institute of Standards and Technology (NIST) [nist-lightweight-project]: NIST is running a project on lightweight cryptography with the purpose of: (i) identifying application areas for which standard cryptographic algorithms are too heavy, classifying them according to some application profiles to be determined; (ii) determining limitations in those existing cryptographic standards; and (iii) standardizing lightweight algorithms that can be used in specific application profiles.

8. Open Web Application Security Project (OWASP) [OWASP]: OWASP provides security guidance for IoT manufactures, developers and consumers. OWASP also includes guidelines for those who intend to test and analyze IoT devices and applications.
9. IoT Security foundation [IoTSecFoundation]: IoT security foundation has published a document that enlists various considerations that need to be taken into account when developing IoT applications. For example, the document states that IoT devices could use hardware-root of trust to ensure that only authorized software runs on the devices.
10. National Highway Traffic Safety Administration (NHTSA) [NHTSA]: The US NHTSA provides guidance to the automotive industry for improving the cyber security of vehicles. While some of the guidelines are general, the document provides specific recommendations for the automotive industry such as how various automotive manufacturer can share cyber security vulnerabilities discovered.
11. Best Current Practices (BCP) for IoT devices [ID-Moore]: This document provides a list of minimum requirements that vendors of Internet of Things (IoT) devices should to take into account while developing applications, services and firmware updates in order to reduce the frequency and severity of security incidents that arise from compromised IoT devices.
12. European Union Agency for Network and Information Security (ENISA) [ENISA-ICS]: ENISA published a document on communication network dependencies for Industrial Control Systems (ICS)/Supervisory Control And Data Acquisition (SCADA) systems in which security vulnerabilities, guidelines and general recommendations are summarized.
13. Internet Society Online Trust Alliance [ISOC-OTA]: The Internet Society's IoT Trust Framework identifies the core requirements manufacturers, service providers, distributors, purchasers and policymakers need to understand, assess and embrace for effective security and privacy as part of the Internet of Things.

Other guideline and recommendation documents may exist or may later be published. This list should be considered non-exhaustive. Despite the acknowledgment that security in the Internet is needed and the existence of multiple guidelines, the fact is that many IoT devices and systems have very limited security. There are multiple reasons for this. For instance, some manufactures focus on delivering a product without paying enough attention to security.

This may be because of lack of expertise or limited budget. However, the deployment of such insecure devices poses a severe threat on the privacy and safety of users. The vast amount of devices and their inherent mobile nature also implies that an initially secure system can become insecure if a compromised device gains access to the system at some point in time. Even if all other devices in a given environment are secure, this does not prevent external attacks caused by insecure devices. Recently the Federal Communications Commission (FCC) [FCC] has stated the need for additional regulation of IoT systems. It is possible that we may see other such regional regulations in the future.

5. Challenges for a Secure IoT

In this section, we take a closer look at the various security challenges in the operational and technical features of IoT and then discuss how existing Internet security protocols cope with these technical and conceptual challenges through the lifecycle of a thing. This discussion should neither be understood as a comprehensive evaluation of all protocols, nor can it cover all possible aspects of IoT security. Yet, it aims at showing concrete limitations and challenges in some IoT design areas rather than giving an abstract discussion. In this regard, the discussion handles issues that are most important from the authors' perspectives.

5.1. Constraints and Heterogeneous Communication

Coupling resource-constrained networks and the powerful Internet is a challenge because the resulting heterogeneity of both networks complicates protocol design and system operation. In the following we briefly discuss the resource constraints of IoT devices and the consequences for the use of Internet Protocols in the IoT domain.

5.1.1. Resource Constraints

IoT deployments are often characterized by lossy and low-bandwidth communication channels. IoT devices are also often constrained in terms of CPU, memory, and energy budget available [RFC7228]. These characteristics directly impact the design of protocols for the IoT domain. For instance, small packet size limits at the physical layer (127 Bytes in IEEE 802.15.4) can lead to (i) hop-by-hop fragmentation and reassembly or (ii) small IP-layer maximum transmission unit (MTU). In the first case, excessive fragmentation of large packets that are often required by security protocols may open new attack vectors for state exhaustion attacks. The second case might lead to more fragmentation at the IP layer which commonly downgrades the overall system performance due to packet loss and the need for retransmission.

The size and number of messages should be minimized to reduce memory requirements and optimize bandwidth usage. In this context, layered approaches involving a number of protocols might lead to worse performance in resource-constrained devices since they combine the headers of the different protocols. In some settings, protocol negotiation can increase the number of exchanged messages. To improve performance during basic procedures such as, for example, bootstrapping, it might be a good strategy to perform those procedures at a lower layer.

Small CPUs and scarce memory limit the usage of resource-expensive cryptographic primitives such as public-key cryptography as used in most Internet security standards. This is especially true if the basic cryptographic blocks need to be frequently used or the underlying application demands low delay.

There are ongoing efforts to reduce the resource consumption of security protocols by using more efficient underlying cryptographic primitives such as Elliptic Curve Cryptography [RFC8446]. The specification of elliptic curve X25519 [ecc25519], stream ciphers such as ChaCha [ChaCha], Diet HIP [ID-HIP-DEX], and ECC groups for IKEv2 [RFC5903] are all examples of efforts to make security protocols more resource efficient. Additionally, most modern security protocols have been revised in the last few years to enable cryptographic agility, making cryptographic primitives interchangeable. However, these improvements are only a first step in reducing the computation and communication overhead of Internet protocols. The question remains if other approaches can be applied to leverage key agreement in these heavily resource-constrained environments.

A further fundamental need refers to the limited energy budget available to IoT nodes. Careful protocol (re)design and usage is required to reduce not only the energy consumption during normal operation, but also under DoS attacks. Since the energy consumption of IoT devices differs from other device classes, judgments on the energy consumption of a particular protocol cannot be made without tailor-made IoT implementations.

5.1.2. Denial-of-Service Resistance

The tight memory and processing constraints of things naturally alleviate resource exhaustion attacks. Especially in unattended T2T communication, such attacks are difficult to notice before the service becomes unavailable (for example, because of battery or memory exhaustion). As a DoS countermeasure, DTLS, IKEv2, HIP, and Diet HIP implement return routability checks based on a cookie mechanism to delay the establishment of state at the responding host

until the address of the initiating host is verified. The effectiveness of these defenses strongly depend on the routing topology of the network. Return routability checks are particularly effective if hosts cannot receive packets addressed to other hosts and if IP addresses present meaningful information as is the case in today's Internet. However, they are less effective in broadcast media or when attackers can influence the routing and addressing of hosts (for example, if hosts contribute to the routing infrastructure in ad-hoc networks and meshes).

In addition, HIP implements a puzzle mechanism that can force the initiator of a connection (and potential attacker) to solve cryptographic puzzles with variable difficulties. Puzzle-based defense mechanisms are less dependent on the network topology but perform poorly if CPU resources in the network are heterogeneous (for example, if a powerful Internet host attacks a thing). Increasing the puzzle difficulty under attack conditions can easily lead to situations where a powerful attacker can still solve the puzzle while weak IoT clients cannot and are excluded from communicating with the victim. Still, puzzle-based approaches are a viable option for sheltering IoT devices against unintended overload caused by misconfiguration or malfunctioning things.

5.1.3. End-to-end security, protocol translation, and the role of middleboxes

The term end-to-end security often has multiple interpretations. Here, we consider end-to-end security in the context end-to-end IP connectivity, from a sender to a receiver. Services such as confidentiality and integrity protection on packet data, message authentication codes or encryption are typically used to provide end-to-end security. These protection methods render the protected parts of the packets immutable as rewriting is either not possible because a) the relevant information is encrypted and inaccessible to the gateway or b) rewriting integrity-protected parts of the packet would invalidate the end-to-end integrity protection.

Protocols for constrained IoT networks are not exactly identical to their larger Internet counterparts for efficiency and performance reasons. Hence, more or less subtle differences between protocols for constrained IoT networks and Internet protocols will remain. While these differences can be bridged with protocol translators at middleboxes, they may become major obstacles if end-to-end security measures between IoT devices and Internet hosts are needed.

If access to data or messages by the middleboxes is required or acceptable, then a diverse set of approaches for handling such a scenario are available. Note that some of these approaches affect

the meaning of end-to-end security in terms of integrity and confidentiality since the middleboxes will be able to either decrypt or modify partially the exchanged messages:

1. Sharing credentials with middleboxes enables them to transform (for example, decompress, convert, etc.) packets and re-apply the security measures after transformation. This method abandons end-to-end security and is only applicable to simple scenarios with a rudimentary security model.
2. Reusing the Internet wire format for IoT makes conversion between IoT and Internet protocols unnecessary. However, it can lead to poor performance in some use cases because IoT specific optimizations (for example, stateful or stateless compression) are not possible.
3. Selectively protecting vital and immutable packet parts with a message authentication code or with encryption requires a careful balance between performance and security. Otherwise this approach might either result in poor performance or poor security depending on which parts are selected for protection, where they are located in the original packet, and how they are processed. [ID-OSCORE] proposes a solution in this direction by encrypting and integrity protecting most of the message fields except those parts that a middlebox needs to read or change.
4. Homomorphic encryption techniques can be used in the middlebox to perform certain operations. However, this is limited to data processing involving arithmetic operations. Furthermore, performance of existing libraries, for example, SEAL [SEAL] is still too limited and homomorphic encryption techniques are not widely applicable yet.
5. Message authentication codes that sustain transformation can be realized by considering the order of transformation and protection (for example, by creating a signature before compression so that the gateway can decompress the packet without recalculating the signature). Such an approach enables IoT specific optimizations but is more complex and may require application-specific transformations before security is applied. Moreover, the usage of encrypted or integrity-protected data prevents middleboxes from transforming packets.
6. Mechanisms based on object security can bridge the protocol worlds, but still require that the two worlds use the same object security formats. Currently the object security format based on CBOR Object Signing and Encryption (COSE) [RFC8152] is different from JSON Object Signing and Encryption (JOSE) [RFC7520] or

Cryptographic Message Syntax (CMS) [RFC5652]. Legacy devices relying on traditional Internet protocols will need to update to the newer protocols for constrained environments to enable real end-to-end security. Furthermore, middleboxes do not have any access to the data and this approach does not prevent an attacker who is capable of modifying relevant message header fields that are not protected.

To the best of our knowledge, none of the mentioned security approaches that focus on the confidentiality and integrity of the communication exchange between two IP end-points provide the perfect solution in this problem space.

5.1.4. New network architectures and paradigm

There is a multitude of new link layer protocols that aim to address the resource-constrained nature of IoT devices. For example, the IEEE 802.11 ah [IEEE802ah] has been specified for extended range and lower energy consumption to support Internet of Things (IoT) devices. Similarly, Low-Power Wide-Area Network (LPWAN) protocols such as LoRa [lora], Sigfox [sigfox], NarrowBand IoT (NB-IoT) [nbiot] are all designed for resource-constrained devices that require long range and low bit rates. [RFC8376] provides an informational overview of the set of LPWAN technologies being considered by the IETF. It also identifies the potential gaps that exist between the needs of those technologies and the goal of running IP in such networks. While these protocols allow IoT devices to conserve energy and operate efficiently, they also add additional security challenges. For example, the relatively small MTU can make security handshakes with large X509 certificates a significant overhead. At the same time, new communication paradigms also allow IoT devices to communicate directly amongst themselves with or without support from the network. This communication paradigm is also referred to as Device-to-Device (D2D) or Machine-to-Machine (M2M) or Thing-to-Thing (T2T) communication and it is motivated by a number of features such as improved network performance, lower latency and lower energy requirements.

5.2. Bootstrapping of a Security Domain

Creating a security domain from a set of previously unassociated IoT devices is a key operation in the lifecycle of a thing in an IoT network. This aspect is further elaborated and discussed in the T2TRG draft on bootstrapping [ID-bootstrap].

5.3. Operational Challenges

After the bootstrapping phase, the system enters the operational phase. During the operational phase, things can use the state information created during the bootstrapping phase in order to exchange information securely. In this section, we discuss the security challenges during the operational phase. Note that many of the challenges discussed in Section 5.1 apply during the operational phase.

5.3.1. Group Membership and Security

Group key negotiation is an important security service for IoT communication patterns in which a thing sends some data to multiple things or data flows from multiple things towards a thing. All discussed protocols only cover unicast communication and therefore, do not focus on group-key establishment. This applies in particular to (D)TLS and IKEv2. Thus, a solution is required in this area. A potential solution might be to use the Diffie-Hellman keys - that are used in IKEv2 and HIP to setup a secure unicast link - for group Diffie-Hellman key-negotiations. However, Diffie-Hellman is a relatively heavy solution, especially if the group is large.

Symmetric and asymmetric keys can be used in group communication. Asymmetric keys have the advantage that they can provide source authentication. However, doing broadcast encryption with a single public/private key pair is also not feasible. Although a single symmetric key can be used to encrypt the communication or compute a message authentication code, it has inherent risks since the capture of a single node can compromise the key shared throughout the network. The usage of symmetric-keys also does not provide source authentication. Another factor to consider is that asymmetric cryptography is more resource-intensive than symmetric key solutions. Thus, the security risks and performance trade-offs of applying either symmetric or asymmetric keys to a given IoT use case need to be well-analyzed according to risk and usability assessments. [ID-multicast] is looking at a combination of symmetric (for encryption) and asymmetric (for authentication) in the same packet.

Conceptually, solutions that provide secure group communication at the network layer (IPsec/IKEv2, HIP/Diet HIP) may have an advantage in terms of the cryptographic overhead when compared to application-focused security solutions (TLS/DTLS). This is due to the fact that application-focused solutions require cryptographic operations per group application, whereas network layer approaches may allow sharing secure group associations between multiple applications (for example, for neighbor discovery and routing or service discovery). Hence, implementing shared features lower in the communication stack can

avoid redundant security measures. However, it is important to note that sharing security contexts among different applications involves potential security threats, e.g., if one of the applications is malicious and monitors exchanged messages or injects fake messages. In the case of OSCORE, it provides security for CoAP group communication as defined in RFC7390, i.e., based on multicast IP. If the same security association is reused for each application, then this solution does not seem to have more cryptographic overhead compared to IPsec.

Several group key solutions have been developed by the MSEC working group [WG-MSEC] of the IETF. The MIKEY architecture [RFC4738] is one example. While these solutions are specifically tailored for multicast and group broadcast applications in the Internet, they should also be considered as candidate solutions for group key agreement in IoT. The MIKEY architecture for example describes a coordinator entity that disseminates symmetric keys over pair-wise end-to-end secured channels. However, such a centralized approach may not be applicable in a distributed IoT environment, where the choice of one or several coordinators and the management of the group key is not trivial.

5.3.2. Mobility and IP Network Dynamics

It is expected that many things (for example, wearable sensors, and user devices) will be mobile in the sense that they are attached to different networks during the lifetime of a security association. Built-in mobility signaling can greatly reduce the overhead of the cryptographic protocols because unnecessary and costly re-establishments of the session (possibly including handshake and key agreement) can be avoided. IKEv2 supports host mobility with the MOBIKE [RFC4555] and [RFC4621] extension. MOBIKE refrains from applying heavyweight cryptographic extensions for mobility. However, MOBIKE mandates the use of IPsec tunnel mode which requires the transmission of an additional IP header in each packet.

HIP offers a simple yet effective mobility management by allowing hosts to signal changes to their associations [RFC8046]. However, slight adjustments might be necessary to reduce the cryptographic costs, for example, by making the public-key signatures in the mobility messages optional. Diet HIP does not define mobility yet but it is sufficiently similar to HIP and can use the same mechanisms. DTLS provides some mobility support by relying on a connection ID (CID). The use of connection IDs can provide all the mobility functionality described in [ID-Williams], except, sending the updated location. The specific need for IP-layer mobility mainly depends on the scenario in which the nodes operate. In many cases, mobility supported by means of a mobile gateway may suffice to enable

mobile IoT networks, such as body sensor networks. Using message based application-layer security solutions such as OSCORE [ID-OSCORE] can also alleviate the problem of re-establishing lower-layer sessions for mobile nodes.

5.4. Secure software update and cryptographic agility

IoT devices are often expected to stay functional for several years and decades even though they might operate unattended with direct Internet connectivity. Software updates for IoT devices are therefore not only required for new functionality, but also to eliminate security vulnerabilities due to software bugs, design flaws, or deprecated algorithms. Software bugs might remain even after careful code review. Implementations of security protocols might contain (design) flaws. Cryptographic algorithms can also become insecure due to advances in cryptanalysis. Therefore, it is necessary that devices which are incapable of verifying a cryptographic signature are not exposed to the Internet (even indirectly).

Schneier [SchneierSecurity] in his essay highlights several challenges that hinder mechanisms for secure software update of IoT devices. First, there is a lack of incentives for manufactures, vendors and others on the supply chain to issue updates for their devices. Second, parts of the software running on IoT devices is simply a binary blob without any source code available. Since the complete source code is not available, no patches can be written for that piece of code. Lastly Schneier points out that even when updates are available, users generally have to manually download and install them. However, users are never alerted about security updates and at many times do not have the necessary expertise to manually administer the required updates.

The FTC staff report on Internet of Things - Privacy & Security in a Connected World [FTCreport] and the Article 29 Working Party Opinion 8/2014 on the Recent Developments on the Internet of Things [Article29] also document the challenges for secure remote software update of IoT devices. They note that even providing such a software update capability may add new vulnerabilities for constrained devices. For example, a buffer overflow vulnerability in the implementation of a software update protocol (TR69) [TR69] and an expired certificate in a hub device [wink] demonstrate how the software update process itself can introduce vulnerabilities.

Powerful IoT devices that run general purpose operating systems can make use of sophisticated software update mechanisms known from the desktop world. However, resource-constrained devices typically do not have any operating system and are often not equipped with a

memory management unit or similar tools. Therefore, they might require more specialized solutions.

An important requirement for secure software and firmware updates is source authentication. Source authentication requires the resource-constrained things to implement public-key signature verification algorithms. As stated in Section 5.1.1, resource-constrained things have limited amount of computational capabilities and energy supply available which can hinder the amount and frequency of cryptographic processing that they can perform. In addition to source authentication, software updates might require confidential delivery over a secure (encrypted) channel. The complexity of broadcast encryption can force the usage of point-to-point secure links - however, this increases the duration of a software update in a large system. Alternatively, it may force the usage of solutions in which the software update is delivered to a gateway, and then distributed to the rest of the system with a network key. Sending large amounts of data that later needs to be assembled and verified over a secure channel can consume a lot of energy and computational resources. Correct scheduling of the software updates is also a crucial design challenge. For example, a user of connected light bulbs would not want them to update and restart at night. More importantly, the user would not want all the lights to update at the same time.

Software updates in IoT systems are also needed to update old and insecure cryptographic primitives. However, many IoT systems, some of which are already deployed, are not designed with provisions for cryptographic agility. For example, many devices come with a wireless radio that has an AES128 hardware co-processor. These devices solely rely on the co-processor for encrypting and authenticating messages. A software update adding support for new cryptographic algorithms implemented solely in software might not fit on these devices due to limited memory, or might drastically hinder its operational performance. This can lead to the use of old and insecure software. Therefore, it is important to account for the fact that cryptographic algorithms would need to be updated and consider the following when planning for cryptographic agility:

1. Would it be secure to use the existing cryptographic algorithms available on the device for updating with new cryptographic algorithms that are more secure?
2. Will the new software-based implementation fit on the device given the limited resources?
3. Would the normal operation of existing IoT applications on the device be severely hindered by the update?

Finally, we would like to highlight the previous and ongoing work in the area of secure software and firmware updates at the IETF. [RFC4108] describes how Cryptographic Message Syntax (CMS) [RFC5652] can be used to protect firmware packages. The IAB has also organized a workshop to understand the challenges for secure software update of IoT devices. A summary of the recommendations to the standards community derived from the discussions during that workshop have been documented [RFC8240]. A working group called Software Updates for Internet of Things (suit) [WG-SUIT] is currently working on a new version [RFC4108] to reflect the best current practices for firmware update based on experience from IoT deployments. It is specifically working on describing an IoT firmware update architecture and specifying a manifest format that contains meta-data about the firmware update package. Finally, the Trusted Execution Environment Provisioning working group [WG-TEEP] aims at developing a protocol for lifecycle management of trusted applications running on the secure area of a processor (Trusted Execution Environment (TEE)).

5.5. End-of-Life

Like all commercial devices, IoT devices have a given useful lifetime. The term end-of-life (EOL) is used by vendors or network operators to indicate the point of time in which they limit or end support for the IoT device. This may be planned or unplanned (for example when the manufacturer goes bankrupt, when the vendor just decides to abandon a product, or when a network operator moves to a different type of networking technology). A user should still be able to use and perhaps even update the device. This requires for some form of authorization handover.

Although this may seem far-fetched given the commercial interests and market dynamics, we have examples from the mobile world where the devices have been functional and up-to-date long after the original vendor stopped supporting the device. CyanogenMod for Android devices, and OpenWrt for home routers are two such instances where users have been able to use and update their devices even after the official EOL. Admittedly it is not easy for an average user to install and configure their devices on their own. With the deployment of millions of IoT devices, simpler mechanisms are needed to allow users to add new root-of-trusts and install software and firmware from other sources once the device is EOL.

5.6. Verifying device behavior

Users using new IoT appliances such as Internet-connected smart televisions, speakers and cameras are often unaware that these devices can undermine their privacy. Recent revelations have shown that many IoT device vendors have been collecting sensitive private

data through these connected appliances with or without appropriate user warnings [cctv].

An IoT device user/owner would like to monitor and verify its operational behavior. For instance, the user might want to know if the device is connecting to the server of the manufacturer for any reason. This feature - connecting to the manufacturer's server - may be necessary in some scenarios, such as during the initial configuration of the device. However, the user should be kept aware of the data that the device is sending back to the vendor. For example, the user might want to know if his/her TV is sending data when he/she inserts a new USB stick.

Providing such information to the users in an understandable fashion is challenging. This is because IoT devices are not only resource-constrained in terms of their computational capability, but also in terms of the user interface available. Also, the network infrastructure where these devices are deployed will vary significantly from one user environment to another. Therefore, where and how this monitoring feature is implemented still remains an open question.

Manufacturer Usage Description (MUD) files [ID-MUD] are perhaps a first step towards implementation of such a monitoring service. The idea behind MUD files is relatively simple: IoT devices would disclose the location of their MUD file to the network during installation. The network can then retrieve those files, and learn about the intended behavior of the devices stated by the device manufacturer. A network monitoring service could then warn the user/owner of devices if they don't behave as expected.

Many devices and software services that automatically learn and monitor the behavior of different IoT devices in a given network are commercially available. Such monitoring devices/services can be configured by the user to limit network traffic and trigger alarms when unexpected operation of IoT devices is detected.

5.7. Testing: bug hunting and vulnerabilities

Given that IoT devices often have inadvertent vulnerabilities, both users and developers would want to perform extensive testing on their IoT devices, networks, and systems. Nonetheless, since the devices are resource-constrained and manufactured by multiple vendors, some of them very small, devices might be shipped with very limited testing, so that bugs can remain and can be exploited at a later stage. This leads to two main types of challenges:

1. It remains to be seen how the software testing and quality assurance mechanisms used from the desktop and mobile world will be applied to IoT devices to give end users the confidence that the purchased devices are robust. Bodies such as the European Cyber Security Organization (ECSO) [ECSO] are working on processes for security certification of IoT devices.
2. It is also an open question how the combination of devices from multiple vendors might actually lead to dangerous network configurations. For example, if combination of specific devices can trigger unexpected behavior. It is needless to say that the security of the whole system is limited by its weakest point.

5.8. Quantum-resistance

Many IoT systems that are being deployed today will remain operational for many years. With the advancements made in the field of quantum computers, it is possible that large-scale quantum computers are available in the future for performing cryptanalysis on existing cryptographic algorithms and ciphersuites. If this happens, it will have two consequences. First, functionalities enabled by means of primitives such as RSA or ECC - namely key exchange, public-key encryption and signature - would not be secure anymore due to Shor's algorithm. Second, the security level of symmetric algorithms will decrease, for example, the security of a block cipher with a key size of b bits will only offer $b/2$ bits of security due to Grover's algorithm.

The above scenario becomes more urgent when we consider the so called "harvest and decrypt" attack in which an attacker can start to harvest (store) encrypted data today, before a quantum-computer is available, and decrypt it years later, once a quantum computer is available. Such "harvest and decrypt" attacks are not new and were used in the Venona project [venona-project]. Many IoT devices that are being deployed today will remain operational for a decade or even longer. During this time, digital signatures used to sign software updates might become obsolete making the secure update of IoT devices challenging.

This situation would require us to move to quantum-resistant alternatives, in particular, for those functionalities involving key exchange, public-key encryption and signatures. [ID-c2pq] describes when quantum computers may become widely available and what steps are necessary for transition to cryptographic algorithms that provide security even in presence of quantum computers. While future planning is hard, it may be a necessity in certain critical IoT deployments which are expected to last decades or more. Although increasing the key-size of the different algorithms is definitely an

option, it would also incur additional computational overhead and network traffic. This would be undesirable in most scenarios. There have been recent advancements in quantum-resistant cryptography. We refer to [ETSI-GR-QSC-001] for an extensive overview of existing quantum-resistant cryptography and [RFC7696] provides guidelines for cryptographic algorithm agility.

5.9. Privacy protection

People will eventually be surrounded by hundreds of connected IoT devices. Even if the communication links are encrypted and protected, information about people might still be collected or processed for different purposes. The fact that IoT devices in the vicinity of people might enable more pervasive monitoring can negatively impact their privacy. For instance, imagine the scenario where a static presence sensor emits a packet due to the presence or absence of people in its vicinity. In such a scenario, anyone who can observe the packet, can gather critical privacy-sensitive information.

Such information about people is referred to as personal data in the European Union (EU) or Personally identifiable information (PII) in the United States (US). In particular, the General Data Protection Regulation (GDPR) [GDPR] defines personal data as: 'any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person'.

Ziegeldorf [Ziegeldorf] defines privacy in IoT as a threefold guarantee:

1. Awareness of the privacy risks imposed by IoT devices and services. This awareness is achieved by means of transparent practices by the data controller, i.e., the entity that is providing IoT devices and/or services.
2. Individual control over the collection and processing of personal information by IoT devices and services.
3. Awareness and control of the subsequent use and dissemination of personal information by data controllers to any entity outside the subject's personal control sphere. This point implies that the data controller must be accountable for its actions on the personal information.

Based on this definition, several threats to the privacy of users have been documented [Ziegelendorf] and [RFC6973], in particular considering the IoT environment and its lifecycle:

1. Identification - refers to the identification of the users, their IoT devices, and generated data.
2. Localization - relates to the capability of locating a user and even tracking them, e.g., by tracking MAC addresses in Wi-Fi or Bluetooth.
3. Profiling - is about creating a profile of the user and their preferences.
4. Interaction - occurs when a user has been profiled and a given interaction is preferred, presenting (for example, visually) some information that discloses private information.
5. Lifecycle transitions - take place when devices are, for example, sold without properly removing private data.
6. Inventory attacks - happen if specific information about IoT devices in possession of a user is disclosed.
7. Linkage - is about when information of two or more IoT systems (or other data sets) is combined so that a broader view of the personal data captured can be created.

When IoT systems are deployed, the above issues should be considered to ensure that private data remains private. These issues are particularly challenging in environments in which multiple users with different privacy preferences interact with the same IoT devices. For example, an IoT device controlled by user A (low privacy settings) might leak private information about another user B (high privacy settings). How to deal with these threats in practice is an area of ongoing research.

5.10. Reverse engineering considerations

Many IoT devices are resource-constrained and often deployed in unattended environments. Some of these devices can also be purchased off-the-shelf or online without any credential-provisioning process. Therefore, an attacker can have direct access to the device and apply advanced techniques to retrieve information that a traditional black box model does not consider. Example of those techniques are side-channel attacks or code disassembly. By doing this, the attacker can try to retrieve data such as:

1. long term keys. These long term keys can be extracted by means of a side-channel attack or reverse engineering. If these keys are exposed, then they might be used to perform attacks on devices deployed in other locations.
2. source code. Extraction of source code might allow the attacker to determine bugs or find exploits to perform other types of attacks. The attacker might also just sell the source code.
3. proprietary algorithms. The attacker can analyze these algorithms gaining valuable know-how. The attacker can also create copies of the product (based on those proprietary algorithms) or modify the algorithms to perform more advanced attacks.
4. configuration or personal data. The attacker might be able to read personal data, e.g., healthcare data, that has been stored on a device.

One existing solution to prevent such data leaks is the use of a secure element, a tamper-resistant device that is capable of securely hosting applications and their confidential data. Another potential solution is the usage of Physical Unclonable Function (PUFs) that serves as unique digital fingerprint of a hardware device. PUFs can also enable other functionalities such as secure key storage. Protection against such data leakage patterns is non-trivial since devices are inherently resource-constrained. An open question is whether there are any viable techniques to protect IoT devices and the data in the devices in such an adversarial model.

5.11. Trustworthy IoT Operation

Flaws in the design and implementation of IoT devices and networks can lead to security vulnerabilities. A common flaw is the use of well-known or easy-to-guess passwords for configuration of IoT devices. Many such compromised IoT devices can be found on the Internet by means of tools such as Shodan [shodan]. Once discovered, these compromised devices can be exploited at scale, for example, to launch DDoS attacks. Dyn, a major DNS , was attacked by means of a DDoS attack originating from a large IoT botnet composed of thousands of compromised IP-cameras [dyn-attack]. There are several open research questions in this area:

1. How to avoid vulnerabilities in IoT devices that can lead to large-scale attacks?
2. How to detect sophisticated attacks against IoT devices?

3. How to prevent attackers from exploiting known vulnerabilities at a large scale?

Some ideas are being explored to address this issue. One of the approaches relies on the use of Manufacturer Usage Description (MUD) files [ID-MUD]. As explained earlier, this proposal requires IoT devices to disclose the location of their MUD file to the network during installation. The network can then (i) retrieve those files, (ii) learn from the manufacturers the intended usage of the devices, for example, which services they need to access, and then (iii) create suitable filters and firewall rules.

6. Conclusions and Next Steps

This Internet Draft provides IoT security researchers, system designers and implementers with an overview of security requirements in the IP-based Internet of Things. We discuss the security threats, state-of-the-art, and challenges.

Although plenty of steps have been realized during the last few years (summarized in Section 4.1) and many organizations are publishing general recommendations (Section 4.3) describing how IoT should be secured, there are many challenges ahead that require further attention. Challenges of particular importance are bootstrapping of security, group security, secure software updates, long-term security and quantum-resistance, privacy protection, data leakage prevention - where data could be cryptographic keys, personal data, or even algorithms - and ensuring trustworthy IoT operation.

Authors of new IoT specifications and implementors need to consider how all the security challenges discussed in this draft (and those that emerge later) affect their work. The authors of IoT specifications not only need to put in a real effort towards addressing the security challenges, but also clearly documenting how the security challenges are addressed. This would reduce the chances of security vulnerabilities in the code written by implementors of those specifications.

7. Security Considerations

This entire memo deals with security issues.

8. IANA Considerations

This document contains no request to IANA.

9. Acknowledgments

We gratefully acknowledge feedback and fruitful discussion with Tobias Heer, Robert Moskowitz, Thorsten Dahm, Hannes Tschofenig, Carsten Bormann, Barry Raveendran, Ari Keranen, Goran Selander, Fred Baker, Vicent Roca, Thomas Fossati and Eliot Lear. We acknowledge the additional authors of the previous version of this document Sye Loong Keoh, Rene Hummen and Rene Struik.

10. Informative References

- [Article29] "Opinion 8/2014 on the on Recent Developments on the Internet of Things", Web http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp223_en.pdf, n.d..
- [AUTO-ID] "AUTO-ID LABS", Web <http://www.autoidlabs.org/>, September 2010.
- [BACNET] "BACnet", Web <http://www.bacnet.org/>, February 2011.
- [BITAG] "Internet of Things (IoT) Security and Privacy Recommendations", Web <http://www.bitag.org/report-internet-of-things-security-privacy-recommendations.php>, n.d..
- [cctv] "Backdoor In MVPower DVR Firmware Sends CCTV Stills To an Email Address In China", Web <https://hardware.slashdot.org/story/16/02/17/0422259/backdoor-in-mvpower-dvr-firmware-sends-cctv-stills-to-an-email-address-in-china>, n.d..
- [ChaCha] Bernstein, D., "ChaCha, a variant of Salsa20", Web <http://cr.yp.to/chacha/chacha-20080128.pdf>, n.d..
- [CSA] "Security Guidance for Early Adopters of the Internet of Things (IoT)", Web https://downloads.cloudsecurityalliance.org/whitepapers/Security_Guidance_for_Early_Adopters_of_the_Internet_of_Things.pdf, n.d..
- [DALI] "DALI", Web <http://www.dalibydesign.us/dali.html>, February 2011.

- [DCMS] "Secure by Design: Improving the cyber security of consumer Internet of Things Report", Web <https://www.gov.uk/government/publications/secure-by-design>, n.d..
- [DHS] "Strategic Principles For Securing the Internet of Things (IoT)", Web https://www.dhs.gov/sites/default/files/publications/Strategic_Principles_for_Securing_the_Internet_of_Things-2016-1115-FINAL....pdf, n.d..
- [dyn-attack] "Dyn Analysis Summary Of Friday October 21 Attack", Web <https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/>, n.d..
- [ecc25519] Bernstein, D., "Curve25519: new Diffie-Hellman speed records", Web <https://cr.yp.to/ecdh/curve25519-20060209.pdf>, n.d..
- [ECSO] "European Cyber Security Organization", Web <https://www.ecs-org.eu/>, n.d..
- [ENISA-ICS] "Communication network dependencies for ICS/SCADA Systems", European Union Agency For Network And Information Security , February 2017.
- [ETSI-GR-QSC-001] "Quantum-Safe Cryptography (QSC);Quantum-safe algorithmic framework", European Telecommunications Standards Institute (ETSI) , June 2016.
- [Fairhair] "Fairhair Alliance", Web <https://www.fairhair-alliance.org/>, n.d..
- [FCC] "Federal Communications Comssion Response 12-05-2016", FCC , February 2016.
- [FTCreport] "FTC Report on Internet of Things Urges Companies to Adopt Best Practices to Address Consumer Privacy and Security Risks", Web <https://www.ftc.gov/news-events/press-releases/2015/01/ftc-report-internet-things-urges-companies-adopt-best-practices>, n.d..

- [GDPR] "The EU General Data Protection Regulation",
Web <https://www.eugdpr.org/>, n.d..
- [GSMAsecurity]
"GSMA IoT Security Guidelines", Web
<http://www.gsma.com/connectedliving/future-iot-networks/iot-security-guidelines/>, n.d..
- [ID-6lonfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi,
"Transmission of IPv6 Packets over Near Field
Communication", draft-ietf-6lo-nfc-12 (work in progress),
November 2018.
- [ID-6tisch]
Thubert, P., "An Architecture for IPv6 over the TSCH mode
of IEEE 802.15.4", draft-ietf-6tisch-architecture-18 (work
in progress), December 2018.
- [ID-acedtls]
Gerdes, S., Bergmann, O., Bormann, C., Selander, G., and
L. Seitz, "Datagram Transport Layer Security (DTLS)
Profile for Authentication and Authorization for
Constrained Environments (ACE)", draft-ietf-ace-dtls-
authorize-05 (work in progress), October 2018.
- [ID-aceoauth]
Seitz, L., Selander, G., Wahlstroem, E., Erdtman, S., and
H. Tschofenig, "Authentication and Authorization for
Constrained Environments (ACE) using the OAuth 2.0
Framework (ACE-OAuth)", draft-ietf-ace-oauth-authz-17
(work in progress), November 2018.
- [ID-bootstrap]
Sarikaya, B., Sethi, M., and D. Garcia-Carillo, "Secure
IoT Bootstrapping: A Survey", draft-sarikaya-t2trg-
sbootstrapping-05 (work in progress), September 2018.
- [ID-c2pq]
Hoffman, P., "The Transition from Classical to Post-
Quantum Cryptography", draft-hoffman-c2pq-04 (work in
progress), August 2018.
- [ID-Daniel]
Park, S., Kim, K., Haddad, W., Chakrabarti, S., and J.
Laganier, "IPv6 over Low Power WPAN Security Analysis",
draft-daniel-6lowpan-security-analysis-05 (work in
progress), March 2011.

- [ID-dietesp] Migault, D., Guggemos, T., and C. Bormann, "Diet-ESP: a flexible and compressed format for IPsec/ESP", draft-mglt-6lo-diet-esp-02 (work in progress), July 2016.
- [ID-HIP-DEX] Moskowitz, R., "HIP Diet EXchange (DEX)", draft-moskowitz-hip-rg-dex-06 (work in progress), May 2012.
- [ID-Moore] Moore, K., Barnes, R., and H. Tschofenig, "Best Current Practices for Securing Internet of Things (IoT) Devices", draft-moore-iot-security-bcp-01 (work in progress), July 2017.
- [ID-MUD] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", draft-ietf-opsawg-mud-25 (work in progress), June 2018.
- [ID-multicast] Tiloca, M., Selander, G., Palombini, F., and J. Park, "Group OSCORE - Secure Group Communication for CoAP", draft-ietf-core-oscore-groupcomm-03 (work in progress), October 2018.
- [ID-OSCORE] Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security for Constrained RESTful Environments (OSCORE)", draft-ietf-core-object-security-15 (work in progress), August 2018.
- [ID-rd] Shelby, Z., Koster, M., Bormann, C., Stok, P., and C. Amsuess, "CoRE Resource Directory", draft-ietf-core-resource-directory-17 (work in progress), October 2018.
- [ID-Williams] Williams, M. and J. Barrett, "Mobile DTLS", draft-barrett-mobile-dtls-00 (work in progress), March 2009.
- [IEEE802ah] "Status of Project IEEE 802.11ah, IEEE P802.11- Task Group AH-Meeting Update.",
Web http://www.ieee802.org/11/Reports/tgah_update.htm, n.d..
- [IIoT] "Industrial Internet Consortium",
Web <http://www.iiconsortium.org/>, n.d..

- [IoTSecFoundation]
"Establishing Principles for Internet of Things Security",
Web <https://iotsecurityfoundation.org/establishing-principles-for-internet-of-things-security/>, n.d..
- [IPSO] "IPSO Alliance", Web <http://www.ipso-alliance.org/>, n.d..
- [ISOC-OTA]
"Internet Society's Online Trust Alliance (OTA)",
Web <https://www.internetsociety.org/ota/>, n.d..
- [loral] "LoRa - Wide Area Networks for IoT", Web <https://www.lora-alliance.org/>, n.d..
- [LWM2M] "OMA LWM2M", Web
<http://openmobilealliance.org/iot/lightweight-m2m-lwm2m/>,
n.d..
- [mirai] Koliass, C., Kambourakis, G., Stavrou, A., and J. Voas,,
"DDoS in the IoT: Mirai and Other Botnets", IEEE
Computer , 2017.
- [nbiot] "NarrowBand IoT", Web
http://www.3gpp.org/ftp/tsg_ran/TSG_RAN/TSGR_69/Docs/RP-151621.zip, n.d..
- [NHTSA] "Cybersecurity Best Practices for Modern Vehicles", Web
https://www.nhtsa.gov/staticfiles/nvs/pdf/812333_CybersecurityForModernVehicles.pdf, n.d..
- [NIST-Guide]
Ross, R., McEvelley, M., and J. Oren, "Systems Security
Engineering", Web
<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160.pdf>, n.d..
- [nist-lightweight-project]
"NIST lightweight Project", Web www.nist.gov/programs-projects/lightweight-cryptography,
www.nist.gov/sites/default/files/documents/2016/10/17/sonmez-turan-presentation-lwc2016.pdf, n.d..
- [NISTSP800-122]
Erika McCallister, ., Tim Grance, ., and . Karen Scarfone,
"NIST SP800-122 - Guide to Protecting the Confidentiality
of Personally Identifiable Information", Web
<https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-122.pdf>, n.d..

- [NISTSP800-30r1] "NIST SP 800-30r1 - Guide for Conducting Risk Assessments", Web <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf>, n.d..
- [NISTSP800-34r1] Marianne Swanson, ., Pauline Bowen, ., Amy Wohl Phillips, ., Dean Gallup, ., and . David Lynes, "NIST SP800-34r1 - Contingency Planning Guide for Federal Information Systems", Web <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-34r1.pdf>, n.d..
- [OCF] "Open Connectivity Foundation", Web <https://openconnectivity.org/>, n.d..
- [OneM2M] "OneM2M", Web <http://www.onem2m.org/>, n.d..
- [OWASP] "IoT Security Guidance", Web https://www.owasp.org/index.php/IoT_Security_Guidance, n.d..
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<https://www.rfc-editor.org/info/rfc2818>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, DOI 10.17487/RFC3748, June 2004, <<https://www.rfc-editor.org/info/rfc3748>>.
- [RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <<https://www.rfc-editor.org/info/rfc3756>>.
- [RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", RFC 3833, DOI 10.17487/RFC3833, August 2004, <<https://www.rfc-editor.org/info/rfc3833>>.
- [RFC4016] Parthasarathy, M., "Protocol for Carrying Authentication and Network Access (PANA) Threat Analysis and Security Requirements", RFC 4016, DOI 10.17487/RFC4016, March 2005, <<https://www.rfc-editor.org/info/rfc4016>>.

- [RFC4108] Housley, R., "Using Cryptographic Message Syntax (CMS) to Protect Firmware Packages", RFC 4108, DOI 10.17487/RFC4108, August 2005, <<https://www.rfc-editor.org/info/rfc4108>>.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, DOI 10.17487/RFC4120, July 2005, <<https://www.rfc-editor.org/info/rfc4120>>.
- [RFC4422] Melnikov, A., Ed. and K. Zeilenga, Ed., "Simple Authentication and Security Layer (SASL)", RFC 4422, DOI 10.17487/RFC4422, June 2006, <<https://www.rfc-editor.org/info/rfc4422>>.
- [RFC4555] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, DOI 10.17487/RFC4555, June 2006, <<https://www.rfc-editor.org/info/rfc4555>>.
- [RFC4621] Kivinen, T. and H. Tschofenig, "Design of the IKEv2 Mobility and Multihoming (MOBIKE) Protocol", RFC 4621, DOI 10.17487/RFC4621, August 2006, <<https://www.rfc-editor.org/info/rfc4621>>.
- [RFC4738] Ignjatic, D., Dondeti, L., Audet, F., and P. Lin, "MIKEY-RSA-R: An Additional Mode of Key Distribution in Multimedia Internet KEYing (MIKEY)", RFC 4738, DOI 10.17487/RFC4738, November 2006, <<https://www.rfc-editor.org/info/rfc4738>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC5191] Forsberg, D., Ohba, Y., Ed., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, DOI 10.17487/RFC5191, May 2008, <<https://www.rfc-editor.org/info/rfc5191>>.

- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <<https://www.rfc-editor.org/info/rfc5652>>.
- [RFC5713] Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", RFC 5713, DOI 10.17487/RFC5713, January 2010, <<https://www.rfc-editor.org/info/rfc5713>>.
- [RFC5903] Fu, D. and J. Solinas, "Elliptic Curve Groups modulo a Prime (ECP Groups) for IKE and IKEv2", RFC 5903, DOI 10.17487/RFC5903, June 2010, <<https://www.rfc-editor.org/info/rfc5903>>.
- [RFC6272] Baker, F. and D. Meyer, "Internet Protocols for the Smart Grid", RFC 6272, DOI 10.17487/RFC6272, June 2011, <<https://www.rfc-editor.org/info/rfc6272>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.
- [RFC6568] Kim, E., Kaspar, D., and JP. Vasseur, "Design and Application Spaces for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6568, DOI 10.17487/RFC6568, April 2012, <<https://www.rfc-editor.org/info/rfc6568>>.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.

- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC7401] Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <<https://www.rfc-editor.org/info/rfc7401>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", RFC 7515, DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<https://www.rfc-editor.org/info/rfc7516>>.
- [RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.

- [RFC7520] Miller, M., "Examples of Protecting Content Using JSON Object Signing and Encryption (JOSE)", RFC 7520, DOI 10.17487/RFC7520, May 2015, <<https://www.rfc-editor.org/info/rfc7520>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.
- [RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<https://www.rfc-editor.org/info/rfc7696>>.
- [RFC7744] Seitz, L., Ed., Gerdes, S., Ed., Selander, G., Mani, M., and S. Kumar, "Use Cases for Authentication and Authorization in Constrained Environments", RFC 7744, DOI 10.17487/RFC7744, January 2016, <<https://www.rfc-editor.org/info/rfc7744>>.
- [RFC7815] Kivinen, T., "Minimal Internet Key Exchange Version 2 (IKEv2) Initiator Implementation", RFC 7815, DOI 10.17487/RFC7815, March 2016, <<https://www.rfc-editor.org/info/rfc7815>>.
- [RFC7925] Tschofenig, H., Ed. and T. Fossati, "Transport Layer Security (TLS) / Datagram Transport Layer Security (DTLS) Profiles for the Internet of Things", RFC 7925, DOI 10.17487/RFC7925, July 2016, <<https://www.rfc-editor.org/info/rfc7925>>.
- [RFC8046] Henderson, T., Ed., Vogt, C., and J. Arkko, "Host Mobility with the Host Identity Protocol", RFC 8046, DOI 10.17487/RFC8046, February 2017, <<https://www.rfc-editor.org/info/rfc8046>>.
- [RFC8105] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt, M., and D. Barthel, "Transmission of IPv6 Packets over Digital Enhanced Cordless Telecommunications (DECT) Ultra Low Energy (ULE)", RFC 8105, DOI 10.17487/RFC8105, May 2017, <<https://www.rfc-editor.org/info/rfc8105>>.
- [RFC8152] Schaad, J., "CBOR Object Signing and Encryption (COSE)", RFC 8152, DOI 10.17487/RFC8152, July 2017, <<https://www.rfc-editor.org/info/rfc8152>>.

- [RFC8240] Tschofenig, H. and S. Farrell, "Report from the Internet of Things Software Update (IoTSU) Workshop 2016", RFC 8240, DOI 10.17487/RFC8240, September 2017, <<https://www.rfc-editor.org/info/rfc8240>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.
- [RFC8376] Farrell, S., Ed., "Low-Power Wide Area Network (LPWAN) Overview", RFC 8376, DOI 10.17487/RFC8376, May 2018, <<https://www.rfc-editor.org/info/rfc8376>>.
- [RFC8428] Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Sensor Measurement Lists (SenML)", RFC 8428, DOI 10.17487/RFC8428, August 2018, <<https://www.rfc-editor.org/info/rfc8428>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RG-T2TRG] "IRTF Thing-to-Thing (T2TRG) Research Group", Web <https://datatracker.ietf.org/rg/t2trg/charter/>, n.d..
- [SchneierSecurity] "The Internet of Things Is Wildly Insecure--And Often Unpatchable", Web https://www.schneier.com/essays/archives/2014/01/the_internet_of_thin.html, n.d..
- [SEAL] "Simple Encrypted Arithmetic Library - SEAL", Web <https://www.microsoft.com/en-us/research/publication/simple-encrypted-arithmetic-library-seal-v2-0/>, n.d..
- [shodan] "Shodan", Web <https://www.shodan.io/>, n.d..
- [sigfox] "Sigfox - The Global Communications Service Provider for the Internet of Things (IoT)", Web <https://www.sigfox.com/>, n.d..
- [Thread] "Thread Group", Web <http://threadgroup.org/>, n.d..

- [TR69] "Too Many Cooks - Exploiting the Internet-of-TR-069-Things", Web https://media.ccc.de/v/31c3_-_6166_-_en_-_saal_6_-_201412282145_-_too_many_cooks_-_exploiting_the_internet-of-tr-069-things_-_lior_oppenheim_-_shahar_tal, n.d..
- [venona-project] "Venona Project", Web <https://www.nsa.gov/news-features/declassified-documents/venona/index.shtml>, n.d..
- [WG-6lo] "IETF IPv6 over Networks of Resource-constrained Nodes (6lo) Working Group", Web <https://datatracker.ietf.org/wg/6lo/charter/>, n.d..
- [WG-6LoWPAN] "IETF IPv6 over Low power WPAN (6lowpan) Working Group", Web <http://tools.ietf.org/wg/6lowpan/>, n.d..
- [WG-ACE] "IETF Authentication and Authorization for Constrained Environments (ACE) Working Group", Web <https://datatracker.ietf.org/wg/ace/charter/>, n.d..
- [WG-ACME] "Automated Certificate Management Environment Working Group", Web <https://datatracker.ietf.org/wg/acme/about/>, n.d..
- [WG-CoRE] "IETF Constrained RESTful Environment (CoRE) Working Group", Web <https://datatracker.ietf.org/wg/core/charter/>, n.d..
- [WG-LPWAN] "IETF Low Power Wide-Area Networks Working Group", Web <https://datatracker.ietf.org/wg/lpwan/>, n.d..
- [WG-LWIG] "IETF Light-Weight Implementation Guidance (LWIG) Working Group", Web <https://datatracker.ietf.org/wg/lwig/charter/>, n.d..
- [WG-MSEC] "IETF MSEC Working Group", Web <https://datatracker.ietf.org/wg/msec/>, n.d..
- [WG-SUIT] "IETF Software Updates for Internet of Things (suit)", Web <https://datatracker.ietf.org/group/suit/about/>, n.d..
- [WG-TEEP] "IETF Trusted Execution Environment Provisioning (teep)", Web <https://datatracker.ietf.org/wg/teep/about/>, n.d..

[wink] "Wink's Outage Shows Us How Frustrating Smart Homes Could Be", Web <http://www.wired.com/2015/04/smart-home-headaches/>, n.d..

[ZB] "ZigBee Alliance", Web <http://www.zigbee.org/>, February 2011.

[Ziegeldorf]
Ziegeldorf, J., Garcia-Morchon, O., and K. Wehrle,,
"Privacy in the Internet of Things: Threats and Challenges", Security and Communication Networks - Special Issue on Security in a Completely Interconnected World , 2013.

Authors' Addresses

Oscar Garcia-Morchon
Philips IP&S
High Tech Campus 5
Eindhoven, 5656 AA
The Netherlands

Email: oscar.garcia-morchon@philips.com

Sandeep S. Kumar
Philips Research
High Tech Campus
Eindhoven, 5656 AA
The Netherlands

Email: sandeep.kumar@philips.com

Mohit Sethi
Ericsson
Hirsalantie 11
Jorvas, 02420
Finland

Email: mohit@piuha.net

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 18, 2018

A. Keranen
Ericsson
M. Kovatsch
ETH Zurich
K. Hartke
Universitaet Bremen TZI
September 14, 2017

RESTful Design for Internet of Things Systems
draft-keranen-t2trg-rest-iot-05

Abstract

This document gives guidance for designing Internet of Things (IoT) systems that follow the principles of the Representational State Transfer (REST) architectural style.

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 March 18, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Basics	6
3.1. Architecture	6
3.2. System design	8
3.3. Uniform Resource Identifiers (URIs)	9
3.4. Representations	10
3.5. HTTP/CoAP Methods	10
3.5.1. GET	11
3.5.2. POST	11
3.5.3. PUT	12
3.5.4. DELETE	12
3.6. HTTP/CoAP Status/Response Codes	12
4. REST Constraints	13
4.1. Client-Server	13
4.2. Stateless	14
4.3. Cache	14
4.4. Uniform Interface	14
4.5. Layered System	15
4.6. Code-on-Demand	15
5. Hypermedia-driven Applications	16
5.1. Motivation	16
5.2. Knowledge	17
5.3. Interaction	18
6. Design Patterns	18
6.1. Collections	18
6.2. Calling a Procedure	19
6.2.1. Instantly Returning Procedures	19
6.2.2. Long-running Procedures	19
6.2.3. Conversion	20
6.2.4. Events as State	20
6.3. Server Push	21
7. Security Considerations	22
8. Acknowledgement	23
9. References	23
9.1. Normative References	23
9.2. Informative References	25
Appendix A. Future Work	26
Authors' Addresses	26

1. Introduction

The Representational State Transfer (REST) architectural style [REST] is a set of guidelines and best practices for building distributed hypermedia systems. At its core is a set of constraints, which when fulfilled enable desirable properties for distributed software systems such as scalability and modifiability. When REST principles are applied to the design of a system, the result is often called RESTful and in particular an API following these principles is called a RESTful API.

Different protocols can be used with RESTful systems, but at the time of writing the most common protocols are HTTP [RFC7230] and CoAP [RFC7252]. Since RESTful APIs are often simple and lightweight, they are a good fit for various IoT applications. The goal of this document is to give basic guidance for designing RESTful systems and APIs for IoT applications and give pointers for more information. Design of a good RESTful IoT system has naturally many commonalities with other Web systems. Compared to other systems, the key characteristics of many IoT systems include:

- o data formats, interaction patterns, and other mechanisms that minimize, or preferably avoid, the need for human interaction
- o preference for compact and simple data formats to facilitate efficient transfer over (often) constrained networks and lightweight processing in constrained nodes

2. Terminology

This section explains some of the common terminology that is used in the context of RESTful design for IoT systems. For terminology of constrained nodes and networks, see [RFC7228].

Cache: A local store of response messages and the subsystem that controls storage, retrieval, and deletion of messages in it.

Client: A node that sends requests to servers and receives responses. In RESTful IoT systems it's common for nodes to have more than one role (e.g., both server and client; see Section 3.1).

Client State: The state kept by a client between requests. This typically includes the currently processed representation, the set of active requests, the history of requests, bookmarks (URIs stored for later retrieval), and application-specific state (e.g., local variables). (Note that this is called "Application State" in [REST], which has some ambiguity in modern (IoT) systems where

the overall state of the distributed application (i.e., application state) is reflected in the union of all Client States and Resource States of all clients and servers involved.)

Content Negotiation: The practice of determining the "best" representation for a client when examining the current state of a resource. The most common forms of content negotiation are Proactive Content Negotiation and Reactive Content Negotiation.

Form: A hypermedia control that enables a client to change the state of a resource or to construct a query locally.

Forward Proxy: An intermediary that is selected by a client, usually via local configuration rules, and that can be tasked to make requests on behalf of the client. This may be useful, for example, when the client lacks the capability to make the request itself or to service the response from a cache in order to reduce response time, network bandwidth, and energy consumption.

Gateway: A reverse proxy that provides an interface to a non-RESTful system such as legacy systems or alternative technologies such as Bluetooth ATT/GATT. See also "Reverse Proxy".

Hypermedia Control: A component, such as a link or a form, embedded in a representation that identifies a resource for future hypermedia interactions. If the client engages in an interaction with the identified resource, the result may be a change to resource state and/or client state.

Idempotent Method: A method where multiple identical requests with that method lead to the same visible resource state as a single such request.

Link: A hypermedia control that enables a client to navigate between resources and thereby change the client state.

Link Relation Type: An identifier that describes how the link target resource relates to the current resource (see [RFC5988]).

Media Type: A string such as "text/html" or "application/json" that is used to label representations so that it is known how the representation should be interpreted and how it is encoded.

Method: An operation associated with a resource. Common methods include GET, PUT, POST, and DELETE (see Section 3.5 for details).

Origin Server: A server that is the definitive source for representations of its resources and the ultimate recipient of any

request that intends to modify its resources. In contrast, intermediaries (such as proxies caching a representation) can assume the role of a server, but are not the source for representations as these are acquired from the origin server.

Proactive Content Negotiation: A content negotiation mechanism where the server selects a representation based on the expressed preference of the client. For example, an IoT application could send a request to a sensor with preferred media type "application/senml+json".

Reactive Content Negotiation: A content negotiation mechanism where the client selects a representation from a list of available representations. The list may, for example, be included by a server in an initial response. If the user agent is not satisfied by the initial response representation, it can request one or more of the alternative representations, selected based on metadata (e.g., available media types) included in the response.

Representation: A serialization that represents the current or intended state of a resource and that can be transferred between clients and servers. REST requires representations to be self-describing, meaning that there must be metadata that allows peers to understand which representation format is used. Depending on the protocol needs and capabilities, there can be additional metadata that is transmitted along with the representation.

Representation Format: A set of rules for serializing resource state. On the Web, the most prevalent representation format is HTML. Other common formats include plain text and formats based on JSON [RFC7159], XML, or RDF. Within IoT systems, often compact formats based on JSON, CBOR [RFC7049], and EXI [W3C.REC-exi-20110310] are used.

Representational State Transfer (REST): An architectural style for Internet-scale distributed hypermedia systems.

Resource: An item of interest identified by a URI. Anything that can be named can be a resource. A resource often encapsulates a piece of state in a system. Typical resources in an IoT system can be, e.g., a sensor, the current value of a sensor, the location of a device, or the current state of an actuator.

Resource State: A model of a resource's possible states that is represented in a supported representation type, typically a media type. Resources can change state because of REST interactions with them, or they can change state for reasons outside of the REST model.

Resource Type: An identifier that annotates the application- semantics of a resource (see Section 3.1 of [RFC6690]).

Reverse Proxy: An intermediary that appears as a server towards the client but satisfies the requests by forwarding them to the actual server (possibly via one or more other intermediaries). A reverse proxy is often used to encapsulate legacy services, to improve server performance through caching, and to enable load balancing across multiple machines.

Safe Method: A method that does not result in any state change on the origin server when applied to a resource.

Server: A node that listens for requests, performs the requested operation and sends responses back to the clients.

Uniform Resource Identifier (URI): A global identifier for resources. See Section 3.3 for more details.

3. Basics

3.1. Architecture

The components of a RESTful system are assigned one or both of two roles: client or server. Note that the terms "client" and "server" refer only to the roles that the nodes assume for a particular message exchange. The same node might act as a client in some communications and a server in others. Classic user agents (e.g., Web browsers) are always in the client role and have the initiative to issue requests. Origin servers always have the server role and govern over the resources they host.




Figure 1: Client-Server Communication

Intermediaries (such as forward proxies, reverse proxies, and gateways) implement both roles, but only forward requests to other intermediaries or origin servers. They can also translate requests to different protocols, for instance, as CoAP-HTTP cross-proxies.




Figure 2: Communication with Forward Proxy

Reverse proxies are usually imposed by the origin server. In addition to the features of a forward proxy, they can also provide an interface for non-RESTful services such as legacy systems or alternative technologies such as Bluetooth ATT/GATT. In this case, reverse proxies are usually called gateways. This property is enabled by the Layered System constraint of REST, which says that a client cannot see beyond the server it is connected to (i.e., it is left unaware of the protocol/paradigm change).




Figure 3: Communication with Reverse Proxy

Nodes in IoT systems often implement both roles. Unlike intermediaries, however, they can take the initiative as a client (e.g., to register with a directory, such as CoRE Resource Directory [I-D.ietf-core-resource-directory], or to interact with another thing) and act as origin server at the same time (e.g., to serve sensor values or provide an actuator interface).




Figure 4: Constrained RESTful environments

3.2. System design

When designing a RESTful system, the primary effort goes into modeling the state of the distributed application and assigning it to the different components (i.e., clients and servers). How clients can navigate through the resources and modify state to achieve their goals is defined through hypermedia controls, that is, links and forms. Hypermedia controls span a kind of a state machine where the nodes are resources and the transitions are links or forms. Clients run this state machine (i.e., the application) by retrieving representations, processing the data, and following the included hypermedia controls. In REST, remote state is changed by submitting forms. This is usually done by retrieving the current state, modifying the state on the client side, and transferring the new state to the server in the form of new representations - rather than calling a service and modifying the state on the server side.

Client state encompasses the current state of the described state machine and the possible next transitions derived from the hypermedia controls within the currently processed representation (see Section 2). Furthermore, clients can have part of the state of the distributed application in local variables.

Resource state includes the more persistent data of an application (i.e., independent of individual clients). This can be static data such as device descriptions, persistent data such as system configurations, but also dynamic data such as the current value of a sensor on a thing.

It is important to distinguish between "client state" and "resource state" and keep them separate. Following the Stateless constraint, the client state must be kept only on clients. That is, there is no establishment of shared information about past and future interactions between client and server (usually called a session). On the one hand, this makes requests a bit more verbose since every request must contain all the information necessary to process it. On the other hand, this makes servers efficient and scalable, since they do not have to keep any state about their clients. Requests can easily be distributed over multiple worker threads or server instances. For IoT systems, this constraint lowers the memory requirements for server implementations, which is particularly important for constrained servers (e.g., sensor nodes) and servers serving large amount of clients (e.g., Resource Directory).

3.3. Uniform Resource Identifiers (URIs)

An important part of RESTful API design is to model the system as a set of resources whose state can be retrieved and/or modified and where resources can be potentially also created and/or deleted.

Uniform Resource Identifiers (URIs) are used to indicate a resource for interaction, to reference a resource from another resource, to advertise or bookmark a resource, or to index a resource by search engines.

```

foo://example.com:8042/over/there?name=ferret#nose
  \_/      \_____/ \_____/ \_____/ \_____/
  |         |         |         |         |
scheme  authority  path   query   fragment

```

A URI is a sequence of characters that matches the syntax defined in [RFC3986]. It consists of a hierarchical sequence of five components: scheme, authority, path, query, and fragment (from most significant to least significant). A scheme creates a namespace for resources and defines how the following components identify a resource within that namespace. The authority identifies an entity that governs part of the namespace, such as the server "www.example.org" in the "http" scheme. A host name (e.g., a fully qualified domain name) or an IP address, potentially followed by a transport layer port number, are usually used in the authority component for the "http" and "coap" schemes. The path and query contain data to identify a resource within the scope of the URI's scheme and naming authority. The fragment allows to refer to some portion of the resource, such as a Record in a SenML Pack. However, fragments are processed only at client side and not sent on the wire. [RFC7320] provides more details on URI design and ownership with best current practices for establishing URI structures, conventions, and formats.

For RESTful IoT applications, typical schemes include "https", "coaps", "http", and "coap". These refer to HTTP and CoAP, with and without Transport Layer Security (TLS) [RFC5246]. (CoAP uses Datagram TLS (DTLS) [RFC6347], the variant of TLS for UDP.) These four schemes also provide means for locating the resource; using the HTTP protocol for "http" and "https", and with the CoAP protocol for "coap" and "coaps". If the scheme is different for two URIs (e.g., "coap" vs. "coaps"), it is important to note that even if the rest of the URI is identical, these are two different resources, in two distinct namespaces.

The query parameters can be used to parametrize the resource. For example, a GET request may use query parameters to request the server

to send only certain kind data of the resource (i.e., filtering the response). Query parameters in PUT and POST requests do not have such established semantics and are not commonly used. Whether the order of the query parameters matters in URIs is unspecified and they can be re-ordered e.g., by proxies. Therefore applications should not rely on their order; see Section 3.3 of [RFC6943] for more details.

3.4. Representations

Clients can retrieve the resource state from an origin server or manipulate resource state on the origin server by transferring resource representations. Resource representations have a media type that tells how the representation should be interpreted by identifying the representation format used.

Typical media types for IoT systems include:

- o "text/plain" for simple UTF-8 text
- o "application/octet-stream" for arbitrary binary data
- o "application/json" for the JSON format [RFC7159]
- o "application/senml+json" [I-D.ietf-core-senml] for Sensor Markup Language (SenML) formatted data
- o "application/cbor" for CBOR [RFC7049]
- o "application/exi" for EXI [W3C.REC-exi-20110310]

A full list of registered Internet Media Types is available at the IANA registry [IANA-media-types] and numerical media types registered for use with CoAP are listed at CoAP Content-Formats IANA registry [IANA-CoAP-media].

3.5. HTTP/CoAP Methods

Section 4.3 of [RFC7231] defines the set of methods in HTTP; Section 5.8 of [RFC7252] defines the set of methods in CoAP. As part of the Uniform Interface constraint, each method can have certain properties that give guarantees to clients.

Safe methods do not cause any state change on the origin server when applied to a resource. For example, the GET method only returns a representation of the resource state but does not change the resource. Thus, it is always safe for a client to retrieve a representation without affecting server-side state.

Idempotent methods can be applied multiple times to the same resource while causing the same visible resource state as a single such request. For example, the PUT method replaces the state of a resource with a new state; replacing the state multiple times with the same new state still results in the same state for the resource. However, the response from the server can be different when the same idempotent method is used multiple times. For example when DELETE is used twice on an existing resource, the first request would remove the association and return success acknowledgement whereas the second request would likely result in error response due to non-existing resource.

The following lists the most relevant methods and gives a short explanation of their semantics.

3.5.1. GET

The GET method requests a current representation for the target resource, while the origin server must ensure that there are no side-effects on the resource state. Only the origin server needs to know how each of its resource identifiers corresponds to an implementation and how each implementation manages to select and send a current representation of the target resource in a response to GET.

A payload within a GET request message has no defined semantics.

The GET method is safe and idempotent.

3.5.2. POST

The POST method requests that the target resource process the representation enclosed in the request according to the resource's own specific semantics.

If one or more resources has been created on the origin server as a result of successfully processing a POST request, the origin server sends a 201 (Created) response containing a Location header field (with HTTP) or Location-Path and/or Location-Query Options (with CoAP) that provide an identifier for the resource created. The server also includes a representation that describes the status of the request while referring to the new resource(s).

The POST method is not safe nor idempotent.

3.5.3. PUT

The PUT method requests that the state of the target resource be created or replaced with the state defined by the representation enclosed in the request message payload. A successful PUT of a given representation would suggest that a subsequent GET on that same target resource will result in an equivalent representation being sent.

The fundamental difference between the POST and PUT methods is highlighted by the different intent for the enclosed representation. The target resource in a POST request is intended to handle the enclosed representation according to the resource's own semantics, whereas the enclosed representation in a PUT request is defined as replacing the state of the target resource. Hence, the intent of PUT is idempotent and visible to intermediaries, even though the exact effect is only known by the origin server.

The PUT method is not safe, but is idempotent.

3.5.4. DELETE

The DELETE method requests that the origin server remove the association between the target resource and its current functionality.

If the target resource has one or more current representations, they might or might not be destroyed by the origin server, and the associated storage might or might not be reclaimed, depending entirely on the nature of the resource and its implementation by the origin server.

The DELETE method is not safe, but is idempotent.

3.6. HTTP/CoAP Status/Response Codes

Section 6 of [RFC7231] defines a set of Status Codes in HTTP that are used by application to indicate whether a request was understood and satisfied, and how to interpret the answer. Similarly, Section 5.9 of [RFC7252] defines the set of Response Codes in CoAP.

The status codes consist of three digits (e.g., "404" with HTTP or "4.04" with CoAP) where the first digit expresses the class of the code. Implementations do not need to understand all status codes, but the class of the code must be understood. Codes starting with 1 are informational; the request was received and being processed. Codes starting with 2 indicate a successful request. Codes starting with 3 indicate redirection; further action is needed to complete the

request. Codes starting with 4 and 5 indicate errors. The codes starting with 4 mean client error (e.g., bad syntax in the request) whereas codes starting with 5 mean server error; there was no apparent problem with the request, but server was not able to fulfill the request.

Responses may be stored in a cache to satisfy future, equivalent requests. HTTP and CoAP use two different patterns to decide what responses are cacheable. In HTTP, the cacheability of a response depends on the request method (e.g., responses returned in reply to a GET request are cacheable). In CoAP, the cacheability of a response depends on the response code (e.g., responses with code 2.04 are cacheable). This difference also leads to slightly different semantics for the codes starting with 2; for example, CoAP does not have a 2.00 response code whereas 200 ("OK") is commonly used with HTTP.

4. REST Constraints

The REST architectural style defines a set of constraints for the system design. When all constraints are applied correctly, REST enables architectural properties of key interest [REST]:

- o Performance
- o Scalability
- o Reliability
- o Simplicity
- o Modifiability
- o Visibility
- o Portability

The following sub-sections briefly summarize the REST constraints and explain how they enable the listed properties.

4.1. Client-Server

As explained in the Architecture section, RESTful system components have clear roles in every interaction. Clients have the initiative to issue requests, intermediaries can only forward requests, and servers respond requests, while origin servers are the ultimate recipient of requests that intent to modify resource state.

This improves simplicity and visibility, as it is clear which component started an interaction. Furthermore, it improves modifiability through a clear separation of concerns.

4.2. Stateless

The Stateless constraint requires messages to be self-contained. They must contain all the information to process it, independent from previous messages. This allows to strictly separate the client state from the resource state.

This improves scalability and reliability, since servers or worker threads can be replicated. It also improves visibility because message traces contain all the information to understand the logged interactions.

Furthermore, the Stateless constraint enables caching.

4.3. Cache

This constraint requires responses to have implicit or explicit cache-control metadata. This enables clients and intermediary to store responses and re-use them to locally answer future requests. The cache-control metadata is necessary to decide whether the information in the cached response is still fresh or stale and needs to be discarded.

Cache improves performance, as less data needs to be transferred and response times can be reduced significantly. Less transfers also improves scalability, as origin servers can be protected from too many requests. Local caches furthermore improve reliability, since requests can be answered even if the origin server is temporarily not available.

4.4. Uniform Interface

All RESTful APIs use the same, uniform interface independent of the application. This simple interaction model is enabled by exchanging representations and modifying state locally, which simplifies the interface between clients and servers to a small set of methods to retrieve, update, and delete state - which applies to all applications.

In contrast, in a service-oriented RPC approach, all required ways to modify state need to be modeled explicitly in the interface resulting in a large set of methods - which differs from application to application. Moreover, it is also likely that different parties come up with different ways how to modify state, including the naming of

the procedures, while the state within an application is a bit easier to agree on.

A REST interface is fully defined by:

- o URIs to identify resources
- o representation formats to represent (and retrieve and manipulate) resource state
- o self-descriptive messages with a standard set of methods (e.g., GET, POST, PUT, DELETE with their guaranteed properties)
- o hypermedia controls within representations

The concept of hypermedia controls is also known as HATEOAS: Hypermedia As The Engine Of Application State. The origin server embeds controls for the interface into its representations and thereby informs the client about possible next requests. The mostly used control for RESTful systems is Web Linking [RFC5590]. Hypermedia forms are more powerful controls that describe how to construct more complex requests, including representations to modify resource state.

While this is the most complex constraints (in particular the hypermedia controls), it improves many different key properties. It improves simplicity, as uniform interfaces are easier to understand. The self-descriptive messages improve visibility. The limitation to a known set of representation formats fosters portability. Most of all, however, this constraint is the key to modifiability, as hypermedia-driven, uniform interfaces allow clients and servers to evolve independently, and hence enable a system to evolve.

4.5. Layered System

This constraint enforces that a client cannot see beyond the server with which it is interacting.

A layered system is easier to modify, as topology changes become transparent. Furthermore, this helps scalability, as intermediaries such as load balancers can be introduced without changing the client side. The clean separation of concerns helps with simplicity.

4.6. Code-on-Demand

This principle enables origin servers to ship code to clients.

Code-on-Demand improves modifiability, since new features can be deployed during runtime (e.g., support for a new representation format). It also improves performance, as the server can provide code for local pre-processing before transferring the data.

5. Hypermedia-driven Applications

Hypermedia-driven applications take advantage of hypermedia controls, i.e., links and forms, embedded in the resource representations. A hypermedia client is a client that is capable of processing these hypermedia controls. Hypermedia links can be used to give additional information about a resource representation (e.g., the source URI of the representation) or pointing to other resources. The forms can be used to describe the structure of the data that can be sent (e.g., with a POST or PUT method) to a server, or how a data retrieval (e.g., GET) request for a resource should be formed. In a hypermedia-driven application the client interacts with the server using only the hypermedia controls, instead of selecting methods and/or constructing URIs based on out-of-band information, such as API documentation.

5.1. Motivation

The advantage of this approach is increased evolvability and extensibility. This is important in scenarios where servers exhibit a range of feature variations, where it's expensive to keep evolving client knowledge and server knowledge in sync all the time, or where there are many different client and server implementations. Hypermedia controls serve as indicators in capability negotiation. In particular, they describe available resources and possible operations on these resources using links and forms, respectively.

There are multiple reasons why a server might introduce new links or forms:

- o The server implements a newer version of the application. Older clients ignore the new links and forms, while newer clients are able to take advantage of the new features by following the new links and submitting the new forms.
- o The server offers links and forms depending on the current state. The server can tell the client which operations are currently valid and thus help the client navigate the application state machine. The client does not have to have knowledge which operations are allowed in the current state or make a request just to find out that the operation is not valid.

- o The server offers links and forms depending on the client's access control rights. If the client is unauthorized to perform a certain operation, then the server can simply omit the links and forms for that operation.

5.2. Knowledge

A client needs to have knowledge of a couple of things for successful interaction with a server. This includes what resources are available, what representations of resource states are available, what each representation describes, how to retrieve a representation, what state changing operations on a resource are possible, how to perform these operations, and so on.

Some part of this knowledge, such as how to retrieve the representation of a resource state, is typically hard-coded in the client software. For other parts, a choice can often be made between hard-coding the knowledge or acquiring it on-demand. The key to success in either case is the use in-band information for identifying the knowledge that is required. This enables the client to verify that it has all required knowledge and to acquire missing knowledge on-demand.

A hypermedia-driven application typically uses the following identifiers:

- o URI schemes that identify communication protocols,
- o Internet Media Types that identify representation formats,
- o link relation types or resource types that identify link semantics,
- o form relation types that identify form semantics,
- o variable names that identify the semantics of variables in templated links, and
- o form field names that identify the semantics of form fields in forms.

The knowledge about these identifiers as well as matching implementations have to be shared a priori in a RESTful system.

5.3. Interaction

A client begins interacting with an application through a GET request on an entry point URI. The entry point URI is the only URI a client is expected to know before interacting with an application. From there, the client is expected to make all requests by following links and submitting forms that are provided in previous responses. The entry point URI can be obtained, for example, by manual configuration or some discovery process (e.g., DNS-SD [RFC6763] or Resource Directory [I-D.ietf-core-resource-directory]). For Constrained RESTful environments `"/.well-known/core"` relative URI is defined as a default entry point for requesting the links hosted by servers with known or discovered addresses [RFC6690].

6. Design Patterns

Certain kinds of design problems are often recurring in variety of domains, and often re-usable design patterns can be applied to them. Also some interactions with a RESTful IoT system are straightforward to design; a classic example of reading a temperature from a thermometer device is almost always implemented as a GET request to a resource that represents the current value of the thermometer. However, certain interactions, for example data conversions or event handling, do not have as straightforward and well established ways to represent the logic with resources and REST methods.

The following sections describe how common design problems such as different interactions can be modeled with REST and what are the benefits of different approaches.

6.1. Collections

A common pattern in RESTful systems across different domains is the collection. A collection can be used to combine multiple resources together by providing resources that consist of set of (often partial) representations of resources, called items, and links to resources. The collection resource also defines hypermedia controls for managing and searching the items in the collection.

Examples of the collection pattern in RESTful IoT systems are the CoRE Resource Directory [I-D.ietf-core-resource-directory], CoAP pub/sub broker [I-D.ietf-core-coap-pubsub], and resource discovery via `"/.well-known/core"`. Collection+JSON [CollectionJSON] is an example of a generic collection Media Type.

6.2. Calling a Procedure

To modify resource state, clients usually use GET to retrieve a representation from the server, modify that locally, and transfer the resulting state back to the server with a PUT (see Section 4.4). Sometimes, however, the state can only be modified on the server side, for instance, because representations would be too large to transfer or part of the required information shall not be accessible to clients. In this case, resource state is modified by calling a procedure (or "function"). This is usually modeled with a POST request, as this method leaves the behavior semantics completely to the server. Procedure calls can be divided into two different classes based on how long they are expected to execute: "instantly" returning and long-running.

6.2.1. Instantly Returning Procedures

When the procedure can return within the expected response time of the system, the result can be directly returned in the response. The result can either be actual content or just a confirmation that the call was successful. In either case, the response does not contain a representation of the resource, but a so-called action result. Action results can still have hypermedia controls to provide the possible transitions in the application state machine.

6.2.2. Long-running Procedures

When the procedure takes longer than the expected response time of the system, or even longer than the response timeout, it is a good pattern to create a new resource to track the "task" execution. The server would respond instantly with a "Created" status (HTTP code 201 or CoAP 2.01) and indicate the location of the task resource in the corresponding header field (or CoAP option) or as a link in the action result. The created resource can be used to monitor the progress, to potentially modify queued tasks or cancel tasks, and to eventually retrieve the result.

Monitoring information would be modeled as state of the task resource, and hence be retrievable as representation. The result - when available - can be embedded in the representation or given as a link to another sub-resource. Modifying tasks can be modeled with forms that either update sub-resources via PUT or do a partial write using PATCH or POST. Canceling a task would be modeled with a form that uses DELETE to remove the task resource.

6.2.3. Conversion

A conversion service is a good example where REST resources need to behave more like a procedure call. The knowledge of converting from one representation to another is located only at the server to relieve clients from high processing or storing lots of data. There are different approaches that all depend on the particular conversion problem.

As mentioned in the previous sections, POST request are a good way to model functionality that does not necessarily affect resource state. When the input data for the conversion is small and the conversion result is deterministic, however, it can be better to use a GET request with the input data in the URI query part. The query is parameterizing the conversion resource, so that it acts like a look-up table. The benefit is that results can be cached also for HTTP (where responses to POST are not cacheable). In CoAP, cacheability depends on the response code, so that also a response to a POST request can be made cacheable through a 2.05 Content code.

When the input data is large or has a binary encoding, it is better to use POST requests with a proper Media Type for the input representation. A POST request is also more suitable, when the result is time-dependent and the latest result is expected (e.g., exchange rates).

6.2.4. Events as State

In event-centric paradigms such as pub/sub, events are usually represented by an incoming message that might even be identical for each occurrence. Since the messages are queued, the receiver is aware of each occurrence of the event and can react accordingly. For instance, in an event-centric system, ringing a door bell would result in a message being sent that represents the event that it was rung.

In resource-oriented paradigms such as REST, messages usually carry the current state of the remote resource, independent from the changes (i.e., events) that have lead to that state. In a naive yet natural design, a door bell could be modeled as a resource that can have the states unpressed and pressed. There are, however, a few issues with this approach. Polling is not an option, as it is highly unlikely to be able to observe the pressed state with any realistic polling interval. When using CoAP Observe with Confirmable notifications, the server will usually send two notifications for the event that the door bell was pressed: notification for changing from unpressed to pressed and another one for changing back to unpressed. If the time between the state changes is very short, the server might

drop the first notification, as Observe only guarantees only eventual consistency (see Section 1.3 of [RFC7641]).

The solution is to pick a state model that fits better to the application. In the case of the door bell - and many other event-driven resources - the solution could be a counter that counts how often the bell was pressed. The corresponding action is taken each time the client observes a change in the received representation.

In the case of a network outage, this could lead to a ringing sound 10 minutes after the bell was rung. Also including a timestamp of the last counter increment in the state can help to suppress ringing a sound when the event has become obsolete.

6.3. Server Push

Overall, a universal mechanism for server push, that is, change-of-state notifications and stand-alone event notifications, is still an open issue that is being discussed in the Thing-to-Thing Research Group. It is connected to the state-event duality problem and custody transfer, that is, the transfer of the responsibility that a message (e.g., event) is delivered successfully.

A proficient mechanism for change-of-state notifications is currently only available for CoAP: Observing resources [RFC7641]. It offers eventual consistency, which guarantees "that if the resource does not undergo a new change in state, eventually all registered observers will have a current representation of the latest resource state". It intrinsically deals with the challenges of lossy networks, where notifications might be lost, and constrained networks, where there might not be enough bandwidth to propagate all changes.

For stand-alone event notifications, that is, where every single notification contains an identifiable event that must not be lost, observing resources is not a good fit. A better strategy is to model each event as a new resource, whose existence is notified through change-of-state notifications of an index resource (cf. Collection pattern). Large numbers of events will cause the notification to grow large, as it needs to contain a large number of Web links. Blockwise transfers [RFC7959] can help here. When the links are ordered by freshness of the events, the first block can already contain all links to new events. Then, observers do not need to retrieve the remaining blocks from the server, but only the representations of the new event resources.

An alternative pattern is to exploit the dual roles of IoT devices, in particular when using CoAP: they are usually client and server at

the same time. A client observer would subscribe to events by registering a callback URI at the origin server, e.g., using a POST request and receiving the location of a temporary subscription resource as handle. The origin server would then publish events by sending POST requests containing the event to the observer. The cancellation can be modeled through deleting the subscription resource. This pattern makes the origin server responsible for delivering the event notifications. This goes beyond retransmissions of messages; the origin server is usually supposed to queue all undelivered events and to retry until successful delivery or explicit cancellation. In HTTP, this pattern is known as REST Hooks.

In HTTP, there exist a number of workarounds to enable server push, e.g., long polling and streaming [RFC6202] or server-sent events [W3C.REC-html5-20141028]. Long polling as an extension that both server and client need to be aware of. In IoT systems, long polling can introduce a considerable overhead, as the request has to be repeated for each notification. Streaming and server-sent events (in fact an evolved version of streaming) are more efficient, as only one request is sent. However, there is only one response header and subsequent notifications can only have content. There are no means for individual status and metadata, and hence no means for proficient error handling (e.g., when the resource is deleted).

7. Security Considerations

This document does not define new functionality and therefore does not introduce new security concerns. We assume that system designers apply classic Web security on top of the basic RESTful guidance given in this document. Thus, security protocols and considerations from related specifications apply to RESTful IoT design. These include:

- o Transport Layer Security (TLS): [RFC5246] and [RFC6347]
- o Internet X.509 Public Key Infrastructure: [RFC5280]
- o HTTP security: Section 9 of [RFC7230], Section 9 of [RFC7231], etc.
- o CoAP security: Section 11 of [RFC7252]
- o URI security: Section 7 of [RFC3986]

IoT-specific security is mainly work in progress at the time of writing. First specifications include:

- o (D)TLS Profiles for the Internet of Things: [RFC7925]

Further IoT security considerations are available in [I-D.irtf-t2trg-iot-secons].

8. Acknowledgement

The authors would like to thank Mert Ocak, Heidi-Maria Back, Tero Kauppinen, Michael Koster, Robby Simpson, Ravi Subramaniam, Dave Thaler, Erik Wilde, and Niklas Widell for the reviews and feedback.

9. References

9.1. Normative References

- [I-D.ietf-core-object-security] Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security of CoAP (OSCOAP)", draft-ietf-core-object-security-04 (work in progress), July 2017.
- [I-D.ietf-core-resource-directory] Shelby, Z., Koster, M., Bormann, C., Stok, P., and C. Amsuess, "CoRE Resource Directory", draft-ietf-core-resource-directory-11 (work in progress), July 2017.
- [REST] Fielding, R., "Architectural Styles and the Design of Network-based Software Architectures", Ph.D. Dissertation, University of California, Irvine , 2000.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/info/rfc3986>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC5590] Harrington, D. and J. Schoenwaelder, "Transport Subsystem for the Simple Network Management Protocol (SNMP)", STD 78, RFC 5590, DOI 10.17487/RFC5590, June 2009, <<https://www.rfc-editor.org/info/rfc5590>>.

- [RFC5988] Nottingham, M., "Web Linking", RFC 5988, DOI 10.17487/RFC5988, October 2010, <<https://www.rfc-editor.org/info/rfc5988>>.
- [RFC6202] Loreto, S., Saint-Andre, P., Salsano, S., and G. Wilkins, "Known Issues and Best Practices for the Use of Long Polling and Streaming in Bidirectional HTTP", RFC 6202, DOI 10.17487/RFC6202, April 2011, <<https://www.rfc-editor.org/info/rfc6202>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<https://www.rfc-editor.org/info/rfc6690>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<https://www.rfc-editor.org/info/rfc7231>>.
- [RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", RFC 7641, DOI 10.17487/RFC7641, September 2015, <<https://www.rfc-editor.org/info/rfc7641>>.
- [RFC7959] Bormann, C. and Z. Shelby, Ed., "Block-Wise Transfers in the Constrained Application Protocol (CoAP)", RFC 7959, DOI 10.17487/RFC7959, August 2016, <<https://www.rfc-editor.org/info/rfc7959>>.
- [W3C.REC-exi-20110310] Schneider, J. and T. Kamiya, "Efficient XML Interchange (EXI) Format 1.0", World Wide Web Consortium Recommendation REC-exi-20110310, March 2011, <<http://www.w3.org/TR/2011/REC-exi-20110310>>.

[W3C.REC-html5-20141028]

Hickson, I., Berjon, R., Faulkner, S., Leithead, T., Navara, E., O'Connor, T., and S. Pfeiffer, "HTML5", World Wide Web Consortium Recommendation REC-html5-20141028, October 2014, <<http://www.w3.org/TR/2014/REC-html5-20141028>>.

9.2. Informative References

[CollectionJSON]

Amundsen, M., "Collection+JSON - Document Format", February 2013, <<http://amundsen.com/media-types/collection/format/>>.

[I-D.ietf-core-coap-pubsub]

Koster, M., Keranen, A., and J. Jimenez, "Publish-Subscribe Broker for the Constrained Application Protocol (CoAP)", draft-ietf-core-coap-pubsub-02 (work in progress), July 2017.

[I-D.ietf-core-senml]

Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Media Types for Sensor Measurement Lists (SenML)", draft-ietf-core-senml-10 (work in progress), July 2017.

[I-D.irtf-t2trg-iot-seccons]

Garcia-Morchon, O., Kumar, S., and M. Sethi, "State-of-the-Art and Challenges for the Internet of Things Security", draft-irtf-t2trg-iot-seccons-05 (work in progress), September 2017.

[IANA-CoAP-media]

"CoAP Content-Formats", n.d., <<http://www.iana.org/assignments/core-parameters/core-parameters.xhtml#content-formats>>.

[IANA-media-types]

"Media Types", n.d., <<http://www.iana.org/assignments/media-types/media-types.xhtml>>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.

[RFC6943] Thaler, D., Ed., "Issues in Identifier Comparison for Security Purposes", RFC 6943, DOI 10.17487/RFC6943, May 2013, <<https://www.rfc-editor.org/info/rfc6943>>.

- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March 2014, <<https://www.rfc-editor.org/info/rfc7159>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7320] Nottingham, M., "URI Design and Ownership", BCP 190, RFC 7320, DOI 10.17487/RFC7320, July 2014, <<https://www.rfc-editor.org/info/rfc7320>>.
- [RFC7925] Tschofenig, H., Ed. and T. Fossati, "Transport Layer Security (TLS) / Datagram Transport Layer Security (DTLS) Profiles for the Internet of Things", RFC 7925, DOI 10.17487/RFC7925, July 2016, <<https://www.rfc-editor.org/info/rfc7925>>.

Appendix A. Future Work

- o Interface semantics: shared knowledge among system components (URI schemes, media types, relation types, well-known locations; see core-apps)
- o Unreliable (best effort) communication, robust communication in network with high packet loss, 3-way commit
- o Discuss directories, such as CoAP Resource Directory
- o More information on how to design resources; choosing what is modeled as a resource, etc.

Authors' Addresses

Ari Keranen
Ericsson
Jorvas 02420
Finland

Email: ari.keranen@ericsson.com

Matthias Kovatsch
ETH Zurich
Universitaetstrasse 6
Zurich CH-8092
Switzerland

Email: kovatsch@inf.ethz.ch

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

Email: hartke@tzi.org

Thing-to-Thing Research Group
Internet-Draft
Intended status: Experimental
Expires: September 14, 2017

M. Koster
SmartThings
March 13, 2017

Media Types for Hypertext Sensor Markup
draft-koster-t2trg-hsml-01

Abstract

The scale and scope of the worldwide web has been in part driven by the availability of HTML as a common serialization, data model, and interaction model for structured resources on the web. By contrast, the general use of appropriate hypermedia techniques for machine interfaces has been limited by the lack of a common format for serialization and exchange of structured machine resources and sensor/actuator data which includes or embeds standardized hypermedia controls. The IRTF Thing to Thing Research Group [T2TRG] has a charter to investigate the use of REST design style [REST] for machine interactions. The W3C Web of Things Interest Group [W3C-WoT] are investigating abstract hypermedia controls and interaction models for machines. Machine optimized content formats exist for web links [RFC5988] [RFC6690] and for data items [I-D.ietf-core-senml].

Structured data which contains both links and items is known as the collection pattern. This draft describes media types for representation of machine resources structured as collections. A simple, reusable data model is described with a representation format, using a well known set of keywords to expose hypermedia controls, which inform clients how to perform state transfer operations on resources. The underlying assumptions regarding transfer layer processing are specified in this document. The HSML media type described in this document is compatible with SenML and CoRE Link-format by reusing the keyword identifiers and element structures from these content formats. Representations of HSML document content may be obtained in CoRE Link-Format and SenML content formats.

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 September 14, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Scope of this document 4
2. Overview and Use Case Requirements 4
3. Data Model and Interaction Model 4
3.1. Informative Representation Examples 5
3.2. Links 5
3.3. Collections 5
3.4. Link Embedding 5
3.4.1. Batch operations on multiple items in a collection . 6
3.4.2. Collective operation on groups of linked resources . 6
4. Abstract Transfer Model 6
5. Collections 7
5.1. Base element 7
5.2. Link element 7
5.3. Item element 8
5.3.1. Items embedded in the collection 8
5.3.2. Items stored as collections 9
6. Representation Formats 11
6.1. Example Serialization Formats 11
6.1.1. Collection Formats 12
6.1.2. Link Formats 13
6.1.3. Item Formats 13
7. URI and Parameter Processing 15
7.1. URI Path Processing 15
7.2. URI Parameter processing 15

7.2.1. Resource selection	16
8. Transfer Model Mapping to Collections	16
8.1. Target Resource is Collection, Format is Collection . . .	16
8.1.1. RETRIEVE	16
8.1.2. UPDATE	17
8.1.3. CREATE	18
8.1.4. DELETE	20
8.2. Target Resource is Collection, Format is Link	21
8.2.1. RETRIEVE	21
8.2.2. UPDATE	21
8.2.3. CREATE	22
8.2.4. DELETE	23
8.3. Target Resource is Collection, Format is Item	24
8.3.1. Link Embedding Items	24
8.3.2. RETRIEVE	24
8.3.3. UPDATE	25
8.3.4. CREATE	26
8.3.5. DELETE	27
8.4. Target Resource is Item	28
8.4.1. RETRIEVE	28
8.4.2. UPDATE	28
8.4.3. CREATE	29
8.4.4. DELETE	29
8.5. Groups	29
8.5.1. RETRIEVE	30
8.5.2. UPDATE	31
8.5.3. CREATE	32
8.5.4. DELETE	33
9. Link extensions	34
9.1. Actions	34
9.2. Link Bindings and Monitors	35
10. Reserved Identifiers	36
10.1. Default namespace	37
10.2. URI Processing Parameters	37
10.3. Link Keywords	37
10.3.1. Link Relation Types	38
10.3.2. Link Attribute Types	38
10.4. Item Keywords	38
10.5. Link Parameters used in Actions	39
10.6. Link Parameters used in Link Bindings and Monitors . . .	40
10.7. Conditional Observe Parameters used in Monitors	41
10.8. Link Attribute Values	41
11. IANA Considerations	42
11.1. Media Types	42
11.2. CoRE Parameters Content Formats	43
11.3. Link Parameters	43
11.4. Link Relation Types	44
11.5. New CoRE Interface Types	44

Items are data elements, either within the context the collection, or outside the context of the collection. An instance of an item is a resource and is identified by a URI.

An Item may only be in the context of one collection, but may be identified by any number of links in any number of collections.

Items in the collection that use an HSML compatible data model, for example SenML, see [I-D.ietf-core-senml], may be embedded in the collection and transferred either along with the links or separately from links.

3.1. Informative Representation Examples

JSON formatted examples are used in this document to illustrate normative and informative concepts. Representations in other formats may be derived from the JSON representations. For example, compact binary mappings may be defined using available models.

3.2. Links

Links follow the specifications in [RFC5988] and [RFC6690] with extensions to implement actions and monitors as described in this document and any referencing extension documents.

HSML Links may be stored in Resource Directories for discovery using CoRE Resource Directory [I-D.ietf-core-resource-directory].

3.3. Collections

Collections contain one or more links and extended links, and may contain data items referred to by the links. A representation of a collection may contain both links and data items, plus any extended links such as action forms.

3.4. Link Embedding

Link embedding enables the transfer of one or more items in a collection using a single transfer operation. This document describes two types of link embedding for items in the collection. Batch link embedding allows one or more resource instances (item) to each contribute part of an aggregate (collection) representation. Group link embedding allows a particular operation to be repeated for each member (item) of a group (collection).

3.4.1. Batch operations on multiple items in a collection

A collection of items enables operations on more than one item at a time by exposing a structured a representation of multiple resources in the collection.

Applications may select resources by using URI parameters, and transfer representations of multiple named resources using the HSML or SenML multi-item formats.

3.4.2. Collective operation on groups of linked resources

Resource links in the collection may specify group transfer semantics, where transfer operations are routed to each resource in the collection specified by a group link. Group responses are aggregated using a multi-item format which identifies each item by URI.

4. Abstract Transfer Model

The HSML media type assumes a transfer model capable of interacting with representations using a simple CRUD model, allowing for basic life cycle operations on resources and collections.

CREATE

Create an instance of a resource as specified using the payload as a constructor. Optionally return a reference to the created resource. Typically uses POST in CoAP [RFC7252] or HTTP, may use PUBLISH in pubsub protocols.

RETRIEVE

Obtain a representation of the selected resource. Typically uses GET in CoAP or HTTP, could use SUBSCRIBE with message retention in pubsub.

UPDATE

Replace or partially replace the representation of the selected resource. Typically uses PUT or PATCH in CoAP and HTTP, could use PUBLISH in pubsub in the frequent case that CREATE and UPDATE are not needed on the same resource.

DELETE

Remove the representation of the selected resource. Typically uses DELETE in CoAP or HTTP. There is no natural mapping to pubsub if a remove operation is not provided.

OBSERVE

Obtain a sequence of representations of the selected resource, indicating state updates which occur on the resource. Typically uses CoAP Observe, HTTP EventSource, MQTT SUBSCRIBE. OBSERVE is the transfer equivalent of performing a RETRIEVE on the resource immediately following each state change of the resource.

5. Collections

Collection representations in HSML include Base Elements, Link Elements, and Item Elements.

5.1. Base element

A base elements describes the context under which to interpret values embedded in subsequent items within the representation of a collection.

The base identifier element (bi) may contain an absolute URI or an absolute path reference from which to base relative references found in the links. It functions as a base URI embedded in content as per [RFC3986] Section 5.1.1

URI reference follows the definition in [RFC3986] Section 5.

The format of base elements are specified in [I-D.ietf-core-senml]. Figure 1 is an example of a base element.

```
{
  "bi": "/sensors/"
}
```

Figure 1: Example Base URI

Other base items from SenML are permissible, including base time (bt) and base value(bv). If additional senml base values are present, the client MUST interpret the items in the collection in the context of the applicable base elements. For example, if there is a "bv" or base value elment, all of the returned values from items in the collection MUST be added to the base value as per [I-D.ietf-core-senml].

5.2. Link element

A link element is a hyperlink based on the structure and syntax of [RFC6690] and [I-D.ietf-core-links-json]. An example link element is shown in Figure 2.

```
{
  "href": "temp",
  "rt": "some.sensor.temp"
}
```

Figure 2: Example Link Element

5.3. Item element

An item element in a collection is a data element that is referenced by a link in the collection.

Items in the collection are indicated by hyperlink references ("href") that serve as selection variables for matching to URI parameters and resource names ("n") in multi-resource representations. Reference resolution should use the rules defined in [RFC3986].

Items may be embedded in the collection, they may be subresources of the collection, or they may be items in other collections referenced by links in the collection. An example item element is shown in Figure 3

```
{
  "n": "temp",
  "v": 27
}
```

Figure 3: Example Item Element

5.3.1. Items embedded in the collection

Items may be stored as simple sets of key-value pairs in the context of the collection. Links about these items may be obtained from the collection that contains them.

```
[
  {
    "bi": "/sensors/"
  },
  {
    "href": "temp",
    "rel": "item"
  },
  {
    "href": "humid",
    "rel": "item"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 4: Items Embedded in a Collection

5.3.2. Items stored as collections

Alternatively, items themselves may be stored as single-item collections, pointed to by links in another collection. Items stored as collections may contain an item with a zero length href and name, and a self link for the item as shown in the collection representation of the item in Figure 5. Items stored in this way may be augmented by adding additional resources and link content to the collection. Items stored as collections may offer link format and collection format representations.

```
base collection:
[
  {
    "bi": "/sensors/"
  },
  {
    "href": "temp/"
  },
  {
    "href": "humid/"
  }
]

"temp" item:
[
  {
    "bi": "/sensors/temp/"
  },
  {
    "href": "", //may be elided
    "rel": ["self","item"]
  },
  {
    "n": "", //may be elided
    "v": 27
  }
]

"humid" item:
[
  {
    "bi": "/sensors/humid/"
  },
  {
    "rel": ["self","item"]
  },
  {
    "v": 50
  }
]
```

Figure 5: Items as Separate Collections

Items embedded in collections, and items linked and stored as separate collections, will all be returned using the item representation format as shown in Figure 6. A client interacting with the items representation of the example collection at /sensors/

would not need to understand the difference between embedded items and linked items that exposed similar content.

```
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 6: Example Items Representation

6. Representation Formats

The HSML media type includes multiple content types and interface types [I-D.ietf-core-interfaces] to enable the client to select representations that optimize communication for the workflow. Representation formats include links and items together (collection formats), links alone (link formats), or items alone (item formats).

Link formats are useful for discovery workflow, item formats are useful for interaction with resource state machines, and link+item formats are useful for constructing and modifying resource instances.

In addition to HSML native formats, standard CoRE Link-Format [RFC6690] and SenML formats [I-D.ietf-core-senml] may be exposed.

6.1. Example Serialization Formats

Figure 7 shows an example document in hsml+json format. This example contains a base element, three link elements, and two item elements.

```
RETRIEVE /sensors/ accept=application/hsml+json
or
RETRIEVE /sensors/ accept=application/hsml.collection+json
or
RETRIEVE /sensors/
  accept=application/hsml+json;if=hsml.collection
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp"
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 7: Example Collection Format

The HSML media type defines content formats and corresponding CoRE Interface Types that may select partial representations of the resource for interaction.

6.1.1.1. Collection Formats

Collection formats as shown in Figure 7 expose all of the elements of a resource, including items, links, and link extensions.

6.1.2. Link Formats

Link content formats, when used in an "accept" option or "content-type" option in a transfer header, or when selected by the "if=hsml.link" URI parameter, will select the link elements in the collection for interaction as in Figure 8.

```
RETRIEVE /sensors/ accept=application/hsml.link+json
or
RETRIEVE /sensors/ accept=application/hsml+json;if=hsml.link
Response Payload:
[
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp"
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  }
]
```

Figure 8: Example Lnk Format

CoRE link-format content formats, for example application/link-format+json, select RFC6690 compliant links, and may not include representations of extended links (rel=action, rel=monitor).

6.1.3. Item Formats

Item content formats, when used in an "accept" option or "content-type" option in a transfer header, or when selected by the "if=hsml.item" URI parameter, will select the item elements in the collection for interaction as in Figure 9.

```
RETRIEVE /sensors/ accept=application/hsml.item+json
or
RETRIEVE /sensors/ accept=application/hsml+json?if=hsml.item
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 9: Example Item Format

URI Parameters for matching link attributes and relations may be used to select items when item representations are being specified using either content-format (accept) or interface parameters (if=). For example:

```
RETRIEVE /sensors/?if=hsml.item&rt=some.sensor.temp
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 27
  }
]
```

Figure 10: Item Selection Using Link Parameter

SenML content formats select data records and return SenML compliant resource names. "bn" may optionally be returned when compliant resource names "n" may be resolved through simple string concatenation as per [I-D.ietf-core-senml].

7. URI and Parameter Processing

The HSML media type defines URI reference processing and URI Query processing but does not in general define fragment (#) references in URIs.

If fragment references are provided in a particular transfer implementation, they should be used to select single items in collections in accordance with current practice.

7.1. URI Path Processing

The path part of the URI reference used to indicate HSML resources may be used as a reference to a collection or to an item in a collection. Collection references should contain the trailing slash character "/" in accordance with [RFC3986]. Server implementations should return links to collections with the trailing "/", and should attempt to accept references to collections without the trailing "/" if such references can be used to construct unambiguous references.

References to items in a collection should not contain the trailing "/" character. Servers should return items in response to references that do not contain the trailing "/" character, and should attempt to accept references to items in collections with the trailing "/" if such references can be used to construct unambiguous references.

URI references may be routed to collections in the order in which path segments appear in the reference, from left to right reading the path string, separated by "/" characters.

URI references may alternatively be routed as opaque strings to resources. In this case, the resolution of relative references to items in a collection should be possible by concatenating the relative reference to the context URI of the collection. Note that this may enforce certain naming conventions such as the trailing slash in practice.

7.2. URI Parameter processing

URI Parameters, typically mapped as query parameters in HTTP and CoAP, are used for selecting resources, selecting partial representations, and otherwise modifying aspects of the expected or included representation. In this way, they may be considered part of the URI, since they help identify a unique representation to be transferred.

7.2.1. Resource selection

URI Parameters may be used to select resources in a collection for transfer. This is done using the common parameter matching rules specified in [RFC6690].

Resource selection is performed based on matching URI Parameters with Link Parameters of all links in the collection which are exposed by the indicated media type and interface type. URI Parameters listed in Section 10.2 are excluded from the matching process.

The target resource selection depends on the content-format specified in the request or the interface type specified in the URI parameters.

The collection content-formats or interface types select all links and items in the collection, including link extensions. URI parameters included in the request should be matched against link parameters for selecting links and associated items.

The link content formats or interface types select all links in the collection. URI parameters included in the request should be matched against link parameters for link selection.

The item content formats or interface types select all items in the collection. URI parameters included in the request should be matched against link parameters associated with items in the collection for item selection as shown in Figure 10.

8. Transfer Model Mapping to Collections

8.1. Target Resource is Collection, Format is Collection

When the reference of a request targets a collection resource, using a collection format, the representation may contain both links and items. It is implied that operations using this format will interact with both links and items. The collection format is indicated by using a collection content type in the accept or content-type header, or by specifying a collection interface type e.g. if=hsml.collection.

8.1.1. RETRIEVE

Retrieve returns a representation of selected elements, consisting of a list of elements in the collection, including base element, links, and optionally representations of items, as shown in Figure 11. Elements may include link extensions, for example action links and monitor links.

```
RETRIEVE /sensors/ accept=application/hsml.collection+json
Response Payload:
```

```
[
  {
    "bi": "/sensors/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp"
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 11: Retrieve Collection

8.1.2. UPDATE

Update replaces all selected elements in the collection with elements included in the payload. Update operations may include replace (PUT) and partial update (PATCH) operations where supported in the transfer protocol. The server response should indicate that the resource was Changed.

```
UPDATE /sensors/?href=temp
  content-type=application/hsml.collection+json
Payload:
[
  {
    "rt": ["some.sensor.temp", "some.other.type"]
  }
]
Response: Changed

RETRIEVE /sensors/ accept=application/hsml.collection+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": ["some.sensor.temp", "some.other.type"]
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 12: Update Item in Collection

8.1.3. CREATE

Create adds elements to a collection, including links and items, where the elements are specified by representations included in the payload. Hints and directives about the created resource may be included in the payload as link parameters, for example a value for href, specifying the location of the created resource. The server is expected to return the location of created resource instances to the

client in a header or other response parameter. For example, the "Location" option in CoAP or "Location" header in HTTP should be used to identify the created resource. The server response should indicate that a resource was Created.

```
CREATE /sensors/ content-type=application/hsml.collection+json
```

```
Payload:
```

```
[
  {
    "href": "barometer",
    "rt": "some.sensor.mbar"
  },
  {
    "n": "barometer",
    "v": 993
  }
]
```

```
Response: Created
```

```
Location: "barometer"
```

```
RETRIEVE /sensors/ accept=application/hsml.collection+json
```

```
Response Payload:
```

```
[
  {
    "bi": "/sensors/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "barometer",
    "rt": "some.sensor.mbar"
  },
  {
    "href": "temp",
    "rt": ["some.sensor.temp", "some.other.type"]
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  },
  {
    "n": "barometer",
    "v": 993
  },
  {
    "n": "temp",

```

```
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 13: Create Item in Collection

8.1.4. DELETE

Delete removes selected elements from the collection. If no elements are selected, delete removes the entire collection. The server response should indicate that the resource was Deleted.

```
DELETE /sensors/?href=barometer
Response: Deleted
```

```
RETRIEVE /sensors/ accept=application/hsml.collection+json
Response Payload:
```

```
[
  {
    "bi": "/sensors/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": ["some.sensor.temp", "some.other.type"]
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 14: Delete Item in Collection

8.2. Target Resource is Collection, Format is Link

When a collection is referenced and the link format is indicated, using a link content format in the header or specifying a link interface type, e.g. `if=hsml.link`, it is expected that the request will interact with the links in the collection.

8.2.1. RETRIEVE

Retrieve returns a list containing selected links, as shown in Figure 15.

```
RETRIEVE /sensors/ accept=application/hsml.link+json
Response Payload:
[
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": ["some.sensor.temp", "some.other.type"]
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  }
]

RETRIEVE /sensors/?rt=some.sensor.temp
accept=application/hsml.link+json
Response Payload:
[
  {
    "href": "temp",
    "rt": ["some.sensor.temp", "some.other.type"]
  }
]
```

Figure 15: Retrieve Links

8.2.2. UPDATE

Update modifies selected links, replacing link elements with elements included in the payload. Update operations may include replace (PUT) and partial update (PATCH) operations where supported in the transfer

protocol. The server response should indicate that the resource was Changed.

```
UPDATE /sensors/?href=temp
content-type=application/hsml.link+json
Payload:
[
  {
    "rt": "some.sensor.temp"
  }
]

RETRIEVE /sensors/ accept=application/hsml.link+json
Response Payload:
[
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp",
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  }
]
```

Figure 16: Update Links

8.2.3. CREATE

Create adds links to the collection, where the links are included in the payload. The server response should indicate that the resource was Changed.

```
CREATE /sensors/ content-type=application/hsml.link+json
Payload:
[
  {
    "href": "/sensor-group/"
  }
]
Response: Changed

RETRIEVE /sensors/ accept=application/hsml.link+json
Response Payload:
[
  {
    "href": "/sensor-group/"
  },
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp",
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  }
]
```

Figure 17: Create Links

8.2.4. DELETE

Delete removes selected links from the collection. The server response should indicate that the resource was Changed. If links point to items in the context of the collection, either remove the items as well as the links, or leave the collection as is and return a method error (Method Not Allowed).

```
DELETE /sensors/?href=sensor-group
Response: Changed

RETRIEVE /sensors/ accept=application/hsml.link+json
Response Payload:
[
  {
    "anchor": "/sensors/",
    "rel": ["self", "index"]
  },
  {
    "href": "temp",
    "rt": "some.sensor.temp",
  },
  {
    "href": "humid",
    "rt": "some.sensor.humid"
  }
]
```

Figure 18: Delete Links

8.3. Target Resource is Collection, Format is Item

When a collection is referenced and the item format is indicated, either by including an item content type in the request header or using an item interface type, e.g. `if=hsml.item`, it is expected that the request will interact with the items in a collection.

Specifying item interaction with a collection invokes the link embedding operations.

8.3.1. Link Embedding Items

Collective operations on items in collections are invoked by using the URI of the collections, along with URI parameters, to select one or more items in the collection.

Items which are compatible with the HSML item format may be returned with multiple items embedded in a single representation.

8.3.2. RETRIEVE

Retrieve returns a list containing a base element and a composite representation of the selected items as shown in Figure 19.

```
RETRIEVE /sensors/ accept=application/hsml.item+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 27
  },
  {
    "n": "humid",
    "v": 50
  }
]

RETRIEVE /sensors/?href=temp
accept=application/hsml.item+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 27
  }
]
```

Figure 19: Retrieve Items

8.3.3. UPDATE

Update modifies selected items, replacing items in the collection with items included in the payload which match by name "n" value. Update operations may include replace (PUT) and partial update (PATCH) operations where supported in the transfer protocol. The server response should indicate that the resource was Changed.

```
UPDATE /sensors/ content-type=application/hsml.item+json
Payload:
[
  {
    "n": "temp",
    "v": 30
  }
]
Response: Changed

RETRIEVE /sensors/ accept=application/hsml.item+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 30
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 20: Update Items

8.3.4. CREATE

Create adds new items to the collection along with system-constructed links. Link content is determined by the resource type or traits defined by application semantics. Server is expected to return the location of created resource instances to the client in a header or other response parameter. For example, the "Location" option in CoAP or "Location" header in HTTP should be used to identify the created resource. The server response should indicate that a resource was Created.

```
CREATE /sensors/ content-type=application/hsml.item+json
Payload:
[
  {
    "n": "barometer",
    "v": 1002
  }
]
Response: Created

RETRIEVE /sensors/ accept=application/hsml.item+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 30
  },
  {
    "n": "barometer",
    "v": 1002
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 21: Create Items

8.3.5. DELETE

Delete removes selected items and corresponding links from the collection. The server response should indicate that the resource was Deleted. If no items are selected, return a not found error.

```
DELETE /sensors/?href=barometer
Response: Deleted

RETRIEVE /sensors/ accept=application/hsml.item+json
Response Payload:
[
  {
    "bi": "/sensors/"
  },
  {
    "n": "temp",
    "v": 30
  },
  {
    "n": "humid",
    "v": 50
  }
]
```

Figure 22: Delete Items

8.4. Target Resource is Item

When the URI of a reference points to an item in a collection, it is expected that the request will interact with a single item.

8.4.1. RETRIEVE

Retrieve returns a representation of the item in the content type according to the accept option of the RETRIEVE request, or using a system defined content-format if there is no accept option provided.

```
RETRIEVE /sensors/temp accept=text/plain
Response Payload:
30
```

Figure 23: Retrieve One Item

8.4.2. UPDATE

Update replaces the resource state with the state defined in the supplied representation according to the content-type or ct option. Update operations may include replace (PUT) and partial update (PATCH) operations where supported in the transfer protocol. The server response should indicate that the resource was Changed.


```
UPDATE /sensors/temp content-type=text/plain
Payload:
33

RETRIEVE /sensors/temp accept=text/plain
Response Payload:
33
```

Figure 24: Update One Item

8.4.3. CREATE

Not Defined, application dependent.

8.4.4. DELETE

Delete removes any links to the item from the collection, and removes the item. If the item is stored as a collection, delete removes the collection. The server response should indicate that the resource was Deleted.

```
DELETE /sensors/temp

RETRIEVE /sensors/temp accept=text/plain
Response: Not Found
```

Figure 25: Delete One Item

8.5. Groups

Group transfer operations are provided by collections that contain links with the "grp" relation value.

Transfer operations which specify the collection URI as target and use the item content format are routed to the resolved URI of each link in the collection that contains the "grp" relation.

URI Parameters used for resource selection and matching are sent to the target URIs of all links that contain the "grp" relation.

Responses from the selected group resources are aggregated and by default returned as a single response. The group response SHOULD be returned as an outer array where such representation is available, for example a JSON array which contains elements consisting of SenML responses.

Optionally, a chunked response may be specified, if provided by the transfer implementation, in which the response from each group member is returned individually within a sequence of responses.

The return code should be based on successful responses from link targets. An implementation of a group collection may choose to allow some rejected responses from link targets, depending on the composition of the link targets. A group may not be required to be composed of link targets that always accept all requests; this is at the discretion of the resource designer.

No mechanism is provided in this document to enable a client to inspect the separate return codes from each group link target resource. Multiple transfer headers may be supplied in some representations, or mapped to metadata in others.

The following examples assume the prior example from Figure 5 indexed by a group collection as in Figure 26.

```
RETRIEVE /sensor-group/ accept=application/hsml.collection+json
Response Payload:
[
  {
    "bi": "/sensor-group/"
  },
  {
    "anchor": "/sensor-group/",
    "rel": ["self", "index"]
  },
  {
    "href": "/sensors/temp/",
    "rel": "grp"
  },
  {
    "href": "/sensors/humid/",
    "rel": "grp"
  }
]
```

Figure 26: Example Group Collection

8.5.1. RETRIEVE

Retrieve requests are routed to each link in the collection that contains the "grp" relation. The response from each link target is returned as an element in an array representation.

```
RETRIEVE /sensor-group/ accept=application/hsml.item+json
Response Payload:
[
  [
    {
      "bi": "/sensors/temp/"
    },
    {
      "v": 33
    }
  ],
  [
    {
      "bi": "/sensors/humid/"
    },
    {
      "v": 41
    }
  ]
]
```

Figure 27: Group Retrieve

8.5.2. UPDATE

Update requests are routed to each link in the collection that contains the "grp" relation. The target resource of each group link processes the request, including URI parameters and content format. The result code returned should indicate that the resource is Changed if any resource state may have been updated.

```
UPDATE /sensor-group/ content-type=application/hsml.item+json
Payload:
[
  {
    "v": 0
  }
]
Response: Changed

RETRIEVE /sensor-group/ accept=application/hsml.item+json
Response Payload:
[
  [
    {
      "bi": "/sensors/temp/"
    },
    {
      "v": 0
    }
  ],
  [
    {
      "bi": "/sensors/humid/"
    },
    {
      "v": 0
    }
  ]
]
```

Figure 28: Group Update

8.5.3. CREATE

Create requests are routed to each link in the collection that contains the "grp" relation. In the example shown in Figure 29, an additional named resource is being created within each (collection type) item to hold a location value for that item. The result code should indicate that a resource was Created if any resource was created as a result of the create operation.

```
CREATE /sensor-group/ content-type=application/hsml.item+json
Payload:
[
  {
    "n": "location",
    "vs": "living room"
  }
]
Response: Created

RETRIEVE /sensor-group/ accept=application/hsml.item+json
Response Payload:
[
  [
    {
      "bi": "/sensors/temp/"
    },
    {
      "v": 0
    },
    {
      "n": "location",
      "vs": "living room"
    }
  ],
  [
    {
      "bi": "/sensors/humid/"
    },
    {
      "v": 0
    },
    {
      "n": "location",
      "vs": "living room"
    }
  ]
]
```

Figure 29: Group Create

8.5.4. DELETE

Delete requests are routed to each link in the collection that contains the "grp" relation. In the example shown in Figure 30, the URI parameter ?href=location selects the resource at the relative URI reference "location" at each group link target for delete. The

result code should indicate that a resource was Deleted if any resource was deleted as a result of the delete operation.

```
DELETE /sensor-group/?href=location
Response: Deleted
```

```
RETRIEVE /sensor-group/ accept=application/hsml.item+json
```

Response Payload:

```
[
  [
    {
      "bi": "/sensors/temp/"
    },
    {
      "v": 0
    }
  ],
  [
    {
      "bi": "/sensors/humid/"
    },
    {
      "v": 0
    }
  ]
]
```

Figure 30: Group Delete

9. Link extensions

9.1. Actions

Actions are hypermedia controls, indicated by a `rel=action` value in a link, used to construct transfer operations that change the state of resources. The use roughly follows the use of forms in HTML [RFC1866], with semantics more consistent with links. See Section 10.5 for more information.

An example Action element is shown in Figure 31.

```
{
  "rel": "action",
  "type": "st.on",
  "href": "switchcommand",
  "method": "create",
  "accept": "text/plain",
  "schema": {"type": "string", "enum": ["on"]}
}

{
  "rel": "action",
  "type": "st.off",
  "href": "switchcommand",
  "method": "create",
  "accept": "text/plain",
  "schema": {"enum": ["off"]}
}
```

Figure 31: Example Action Element

These Action elements inform the client that to perform a type "st.on" or "st.off" action on the context resource, perform a CREATE method on the "switchcommand" URI relative to the context URI, using the text/plain content type, with a payload as defined by the "schema" parameter. This example uses a free-form fragment of JSON-Schema language to differentiate, by action payloads, the "st.on" and "st.off" actions, which are mapped to the same URI and method.

9.2. Link Bindings and Monitors

Link Bindings and Monitors are hypermedia controls, indicated by a rel=boundto or rel=monitor value in a link, used to construct transfer operations that consume or expose state changes of resources. A monitor invokes a state transfer operation from the link context to a target resource. A Link Binding follows the semantics defined in [I-D.groves-core-dynlink], and invokes a state transfer in the opposite direction, that is from the link target to the link context.

Monitors use the IANA registered link relation "monitor", defined in [RFC5989]. Link Bindings use the link relation type "boundto", defined in [I-D.groves-core-dynlink].

Monitors have a set of accept parameters that indicate how the context resource is being observed, a set of filter parameters that indicate the conditions for generating a state change in the monitor, and a set of target parameters that indicate how state changes are to

be applied to the monitor resource. See Section 10.6 for more information.

An example Monitor element is shown in Figure 32.

```
{
  "rel": "monitor",
  "href": "tank-level-events",
  "content-type": "application/senml+json",
  "transfer-method": "create",
  "pmin": 600,
  "pmax": 3600,
  "nbul": 20,
  "nbl1": 80
}
```

Figure 32: Example Monitor Element

This Monitor element defines a monitor resource at the "tank-level-events" URI relative to the context URI, which OBSERVEs the context URI, and updates the "tank-level-events" resource using the CREATE method to add JSON items to the collection, according to the given conditional parameters no more frequently than once every 600 seconds, at least once every 3600 seconds, when the reading is in the notification band, which has a lower limit of 80 and wraps around zero to an upper limit of 20. This has the effect of defining a low level alert notification and high level alert notification.

10. Reserved Identifiers

This section defines the common reserved identifiers that are expected to be processed by implementations of HSML clients and servers. There are many more relation types and link parameters defined and registered with IANA. Implementations should not restrict processing to the keywords identified in this document; they should accept all IANA registered keywords as valid identifiers.

Many of the keywords listed are defined in other RFCs and IETF documents. This document does not redefine any existing keywords. Where a definition exists, the existing definition will be used. Where multiple conflicting definitions exist, this document will indicate the required definition.

New definitions are summarized in Section 11.

10.1. Default namespace

Identifiers in representations using the HSML media types are assumed to use the default namespace defined in Section 10 of this document. An identifier that does not contain an explicit namespace identifier is assumed to be in the default namespace.

For example, if the identifier "method" is encountered and it doesn't resolve to an IANA registered parameter (reg-parameter in [RFC5988]) resolution should be attempted using the definition of "method" in this document.

10.2. URI Processing Parameters

The following URI Parameters are used to filter representations according to specific processing rules and should not be used to attempt to match link parameters.

"if" Interface type, used to select a partial representation of a collection

"count" Indicates the number of items to be returned from the collection

"start" Indicates the array index of the item in the collection to select as the first item to be returned

"page" Page number, in units of count

10.3. Link Keywords

The following keywords are reserved for use in an HSML serialization to indicate elements of a web link

"anchor"
Overrides the default resource context of the link

"rel"
Link relation type as defined in [RFC5988] and [RFC6690]

"href"
Target of a link reference. This may be a relative path reference in the collection, e.g. "currentValue" or an absolute path reference on the server, "/sensors/temp/currentValue", or an absolute URI, for example "https://example.com/sensors/temp/currentValue"

10.3.1. Link Relation Types

The following keywords are reserved for use in a HSML serialization to indicate types of link relations, and are used for values of "rel".

"self"

Refers to the collection that contains the link

"item"

The link points to an item in the collection, indicating eligibility for collective interaction using link embedding as described in Section 3.4 and Section 8.3.1.

"grp"

The item the link points to is available for collective interaction through the collection URI according to group semantics described in Section 8.5.

10.3.2. Link Attribute Types

The following keywords are reserved for use in a HSML serialization to indicate types of link attributes

"rt"

The resource type(s) of the item

"u"

Units of measure

"ct"

The CoAP content-format number(s) associated with the item

"content-type"

The media type string(s) associated with the item

"obs"

Presence of this attribute indicates that the associated resource is observable

10.4. Item Keywords

The following keywords are reserved for use in a HSML serialization to indicate elements within the serialization. Some of these are defined in [I-D.ietf-core-senml].

"bi"

The base URI of the collection, relative to the service location
e.g. "/sensors/temp/" This is a new definition for HSML

"bt"

The base time that corresponds to the encapsulated state of the collection

"t"

The time stamp that corresponds to the encapsulated state of the item in the collection, relative to the base time "bt"

"n"

The name or URI of the resource, relative to the base name or base URI

"u"

Units of measure

"v"

Number value

"vb"

Boolean value

"vs"

String value

10.5. Link Parameters used in Actions

"anchor"

May override the default context of an action

"rel"

Indicates that this control is an action when rel contains the value "action"

"href"

URI for mapping or invoking the action specified in the action control.

"type"

Additional indicator of the action being exposed, can be used with "rel"

"method"

Transfer method to use on a particular action

"accept"

The Content-Types or CoAP content-formats that are accepted on create and update methods

"content-type"

The media type string(s) that are exposed by retrieve and observe methods

"ct"

The CoAP content-format number(s) exposed

"schema"

Indicates the schema to use for constructing or interpreting transfer payloads, may be a literal value or a URI pointing to an instance of a schema

10.6. Link Parameters used in Link Bindings and Monitors

"anchor"

May override the context URI of a link binding or monitor with any observable resource

"rel"

Indicates that this control is a monitor when rel contains the value "monitor" or a link binding when rel contains the value "boundto"

"href"

The URI of the resource used to monitor context URI, where transfer operations will be sent.

"accept"

The media type string or CoAP content-format to request from the observed resource

"content-type"

The media type string to use in the transfer operation

"ct"

the CoAP content-format number to use in the transfer operation

"accept-method"

(HSML extension) Transfer method to use in request from the observed resource, default is OBSERVE

"transfer-method"

(HSML extension) Transfer method to use for notifications, default is UPDATE

"accept-schema"

(HSML extension) Schema to use in interpreting the observed resource payload, required if transfer-schema is used.

"transfer-schema"

(HSML extension) Schema to use in constructing the notification transfer payload, default is to transfer the accepted payload unmodified to the target resource.

10.7. Conditional Observe Parameters used in Monitors

"pmin"

Minimum time between notifications from a monitor

"pmax"

Maximum time between notifications from a monitor

"gth"

Value to match or exceed to determine notification condition

"lth"

Value to match or be less than to determine notification condition

"st"

Value change +/- from last report to determine notification condition

"eq"

Value to match, or change from, to determine notification condition

"bmn"

Defines a lower limit, at or above which notification is enabled

"bmx"

Defines an upper limit, at or below which notification is enabled

"iv"

Starts the notification state machine with an initial value

10.8. Link Attribute Values

The following keywords are reserved for use in a HSML serialization to indicate values of link attributes

"create"

Transfer layer CREATE operation, value of "method" or "target-method"

"retrieve"
Transfer layer RETRIEVE operation, value of "method" or "accept-method"

"update"
Transfer layer UPDATE operation, value of "method" or "target-method"

"delete"
Transfer layer DELETE operation, value of "method" or "target-method"

"observe"
Transfer layer OBSERVE operation, value of "method" or "accept-method"

11. IANA Considerations

11.1. Media Types

Type

- o application

Subtypes

- o hsml
- o hsml.collection
- o hsml.link
- o hsml.item

Media type strings

- o application/hsml
- o application/hsml.collection
- o application/hsml.link
- o application/hsml.item
- o application/hsml+json
- o application/hsml.collection+json

- o application/hsml.link+json
- o application/hsml.item+json

11.2. CoRE Parameters Content Formats

(subject to Structured Syntax encoding rules TBD)

- o 22000 - application/hsml+json
- o 22001 - application/hsml.link+json
- o 22002 - application/hsml.item+json

11.3. Link Parameters

- o method
- o schema
- o content-type
- o ct
- o accept-method
- o transfer-method
- o accept-schema
- o transfer-schema

The following should be registered in the CoRE dynamic linking draft [I-D.groves-core-dynlink].

- o pmin
- o pmax
- o bmn
- o bmx
- o iv
- o lth
- o gth

- o st

- o eq

11.4. Link Relation Types

- o grp

11.5. New CoRE Interface Types

- o hsml.collection

- o hsml.item

- o hsml.link

11.6. Transfer Layer Methods

These definitions may use the default namespace and do not need to be registered with IANA

- o create

- o retrieve

- o update

- o delete

- o observe

12. Security Considerations

12.1. Object Signing

Collection representations are resource state encapsulations and may be transmitted and stored as signed objects in order to protect the integrity of data and metadata, including time and embedded access control information.

12.2. Signed Embedded Time Stamps

The collection may include time stamps (bt and t) that are signed with the object data and metadata.

12.3. Signed Embedded Access Control

The collection representation may include embedded access control information, also signed with the metadata, that can instruct the server to enforce a particular access policy for transfer requests.

12.4. Secure State Updates

Representations submitted to a server to update the state of a resource (UPDATE, CREATE, DELETE) may also contain embedded signed assertions which may be used by the server to decide whether to apply or reject the update.

12.5. Object Signing and Encryption

Object signing and encryption SHOULD use the mechanisms specified in IETF documents for secure JSON Objects [RFC7516] and CBOR Objects [I-D.ietf-cose-msg] [I-D.selander-ace-object-security].

13. Terminology

Client

Having a client role in a REST operation, transmitting a request and receiving one or more responses.

Server

Having a server role in a REST operation, the origin of data items or proxy for the origin. A server is also an authority for a URI namespace [RFC3986].

Resource

Server endpoint for a REST operation, identified by a URI [RFC3986]

Representation

An encoded form of the state of a resource. The encoding rules may be specified in a media type or content type. Clients and servers exchange representations of resources in order to effect application state changes [REST].

URI

Uniform Resource Identifier, used to identify a resource in a link or as a reference [RFC3986]

Reference

An identifier used to select or identify a particular resource. References are constructed by clients to identify resources when

interacting with servers. Servers match references in client requests against URIs of hosted resources.

Media Type, also Content-Format, Content-Type

A set of rules for encoding, transfer, and processing resource representations

Hypermedia

Design style which uses metadata in the form of hyperlinks to structure resources in relation to each other

Collection

A composite resource that contains links and optionally data items

Link, also Hyperlink

A metadata element as described in [RFC5988] and [RFC6690] that contains a pointer to and description of some data element.

Item

A data item pointed to by one or more links in one or more collections.

Context

The context of a link is the subject of the link or the enclosing scope. In this document the collection is the default context for links in the collection.

Target

The target of a link is the resource being pointed to or described. Links in a collection point to and describe items as link targets.

Transfer Layer

A set of predefined message types used to implement state transfer semantics, for example REST.

Request

A message sent from a client to a server identifying the resource, representation, and method to use for the interaction with the server.

Response

A message sent from a server to a client in response to a request, which communicates the state of the identified resource.

Operation or Method

The state transition type requested by the client for the server to perform on the identified resource. Indicated by the transfer layer method, for example, RETRIEVE, UPDATE, CREATE, DELETE.

Pubsub

A transfer layer semantic interface based on the publish-subscribe paradigm, allowing for asynchronous messages to be routed on demand.

14. Informative References

[I-D.groves-core-dynlink]

Shelby, Z., Vial, M., Koster, M., and C. Groves, "Dynamic Resource Linking for Constrained RESTful Environments", draft-groves-core-dynlink-00 (work in progress), July 2016.

[I-D.ietf-core-interfaces]

Shelby, Z., Vial, M., Koster, M., and C. Groves, "Reusable Interface Definitions for Constrained RESTful Environments", draft-ietf-core-interfaces-08 (work in progress), February 2017.

[I-D.ietf-core-links-json]

Li, K., Rahman, A., and C. Bormann, "Representing CoRE Formats in JSON and CBOR", draft-ietf-core-links-json-06 (work in progress), July 2016.

[I-D.ietf-core-resource-directory]

Shelby, Z., Koster, M., Bormann, C., and P. Stok, "CoRE Resource Directory", draft-ietf-core-resource-directory-09 (work in progress), October 2016.

[I-D.ietf-core-senml]

Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Media Types for Sensor Measurement Lists (SenML)", draft-ietf-core-senml-05 (work in progress), March 2017.

[I-D.ietf-cose-msg]

Schaad, J., "CBOR Object Signing and Encryption (COSE)", draft-ietf-cose-msg-24 (work in progress), November 2016.

[I-D.selander-ace-object-security]

Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security of CoAP (OSCOAP)", draft-selander-ace-object-security-06 (work in progress), October 2016.

- [REST] Fielding, R., "Architectural Styles and the Design of Network-based Software Architectures", Ph.D. Dissertation, University of California, Irvine, 2000, <http://www.ics.uci.edu/~fielding/pubs/dissertation/fielding_dissertation.pdf>.
- [RFC1866] Berners-Lee, T. and D. Connolly, "Hypertext Markup Language - 2.0", RFC 1866, DOI 10.17487/RFC1866, November 1995, <<http://www.rfc-editor.org/info/rfc1866>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5988] Nottingham, M., "Web Linking", RFC 5988, DOI 10.17487/RFC5988, October 2010, <<http://www.rfc-editor.org/info/rfc5988>>.
- [RFC5989] Roach, A., "A SIP Event Package for Subscribing to Changes to an HTTP Resource", RFC 5989, DOI 10.17487/RFC5989, October 2010, <<http://www.rfc-editor.org/info/rfc5989>>.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", RFC 6690, DOI 10.17487/RFC6690, August 2012, <<http://www.rfc-editor.org/info/rfc6690>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", RFC 7516, DOI 10.17487/RFC7516, May 2015, <<http://www.rfc-editor.org/info/rfc7516>>.
- [T2TRG] IRTF, ., "IRTF Thing to Thing Research Group", 2016, <<https://datatracker.ietf.org/rg/t2trg/charter/>>.
- [W3C-WoT] WoT IG, ., "W3C Web of Things Interest Group", 2016, <<https://www.w3.org/WoT/IG/>>.

Author's Address

Michael Koster
SmartThings
1281 Lawrence Station Road
Sunnyvale 94089
USA

Phone: +1-707-502-5136
Email: michael.koster@smarththings.com