

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 10, 2012

M. Nottingham
Rackspace
May 9, 2012

Home Documents for HTTP APIs
draft-nottingham-json-home-00

Abstract

This document proposes a "home document" format for non-browser HTTP clients.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on November 10, 2012.

Copyright Notice

Copyright (c) 2012 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. Requirements](#) [3](#)
- [3. JSON Home Documents](#) [3](#)
- [4. Resource Objects](#) [5](#)
 - [4.1. Resolving Templated Links](#) [5](#)
- [5. Resource Hints](#) [6](#)
 - [5.1. allow](#) [7](#)
 - [5.2. representations](#) [7](#)
 - [5.3. accept-patch](#) [7](#)
 - [5.4. accept-post](#) [7](#)
 - [5.5. accept-ranges](#) [7](#)
- [6. Creating and Serving Home Documents](#) [7](#)
 - [6.1. Managing Change in Home Documents](#) [8](#)
 - [6.2. Evolving and Mixing APIs with Home Documents](#) [8](#)
 - [6.3. Documenting APIs that use Home Documents](#) [8](#)
- [7. Consuming Home Documents](#) [9](#)
- [8. Security Considerations](#) [9](#)
- [9. IANA Considerations](#) [9](#)
- [10. References](#) [9](#)
 - [10.1. Normative References](#) [9](#)
 - [10.2. Informative References](#) [10](#)
- [Appendix A. Frequently Asked Questions](#) [10](#)
 - [A.1. Why not Microformats?](#) [10](#)
 - [A.2. What about authentication?](#) [11](#)
 - [A.3. What about "Faults" \(i.e., errors\)?](#) [11](#)
 - [A.4. How Do I find the XML Schema / JSON Schema / etc. for
a particular media type?](#) [11](#)
- [Appendix B. Open Issues](#) [11](#)
- [Author's Address](#) [12](#)

Nottingham

Expires November 10, 2012

[Page 2]

1. Introduction

There is an emerging preference for non-browser Web applications (colloquially, "HTTP APIs") to use a link-driven approach to their interactions to assure loose coupling, thereby enabling extensibility and API evolution.

This is based upon experience with previous APIs which specified static URI paths (such as "http://api.example.com/v1.0/widgets/abc123/properties") have resulted in brittle, tight coupling between clients and servers.

Sometimes, these APIs are documented by a document format like WADL [[1](#)] that is used as a "design-time" description; i.e., the URIs and other information they describe are "baked into" client implementations.

In contrast, a "follow your nose" API advertises the resources available to clients using link relations [[RFC5988](#)] and the formats they support using internet media types [[RFC4288](#)]. A client can then decide - at "run time" - which resources to interact with based upon its capabilities (as described by link relations), and the server can safely add new resources and formats without disturbing clients that are not yet aware of them.

As such, the client needs to be able to discover this information quickly and efficiently use it to interact with the server. Just as with a human-targeted "home page" for a site, we can create a "home document" for a HTTP API (often, at the same URI, found through server-driven content negotiation) that describes it to non-browser clients.

Of course, an HTTP API might use any format to do so; however, there are advantages to having a standard home document format. This specification suggests one for consideration, using the JSON format [[RFC4627](#)].

2. Requirements

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

3. JSON Home Documents

A JSON Home Document uses the format described in [[RFC4627](#)] and has

Nottingham

Expires November 10, 2012

[Page 3]

the media type "application/json-home".

Its content consists of a root object with a "resources" property, whose names are link relation types (as defined by [RFC5988](#)), and values are Resource Objects, defined below.

For example:

```
GET / HTTP/1.1
```

```
Host: example.org
```

```
Accept: application/json-home
```

```
HTTP/1.1 200 OK
```

```
Content-Type: application/json-home
```

```
Cache-Control: max-age=3600
```

```
Connection: close
```

```
{
  "resources": {
    "http://example.org/rel/widgets": {
      "href": "/widgets/"
    },
    "http://example.org/rel/widget": {
      "href-template": "/widgets/{widget_id}",
      "href-vars": {
        "widget_id": "http://example.org/param/widget"
      },
      "hints": {
        "allow": ["GET", "PUT", "DELETE", "PATCH"],
        "representations": ["application/json"],
        "accept-patch": ["application/json-patch"],
        "accept-post": ["application/xml"],
        "accept-ranges": ["bytes"]
      }
    }
  }
}
```

Here, we have a home document that links two a resource, "/widgets/" with the relation "http://example.org/rel/widgets", and also links to one or more with the relation type "http://example.org/rel/widget" using a URI Template [RFC6570](#).

It also gives several hints about interacting with the latter "widget" resources, including the HTTP methods usable with them, the patch formats they accept, and the fact that they support partial requests [RFC2616](#) using the "bytes" range-specifier.

Nottingham

Expires November 10, 2012

[Page 4]

It gives no such hints about the "widgets" resource. This does not mean that it (for example) doesn't support any HTTP methods; it means that the client will need to discover this by interacting with the resource, and/or examining the documentation for its link relation type.

4. Resource Objects

A Resource Object links to resources of the defined type using one of two mechanisms; either a direct link (in which case there is exactly one resource of that relation type associated with the API), or a templated link, in which case there are zero to many such resources.

Direct links are indicated with an "href" property, whose value is a URI [[RFC3986](#)].

Templated links are indicated with an "href-template" property, whose value is a URI Template [[RFC6570](#)]. When "href-template" is present, the Resource Object MUST have a "href-vars" property; see "Resolving Templated Links".

Resource Objects MUST have exactly one of the "href" and "href-vars" properties.

In both forms, the links that "href" and "href-template" refer to are URI-references [[RFC3986](#)] whose base URI is that of the JSON Home Document itself.

Resource Objects MAY also have a "hints" property, whose value is an object that uses named Resource Hints as its properties.

4.1. Resolving Templated Links

A URI can be derived from a Templated Link by treating the "href-template" value as a Level 3 URI Template [[RFC6570](#)], using the "href-vars" property to fill the template.

The "href-vars" property, in turn, is an object that acts as a mapping between variable names available to the template and absolute URIs that are used as global identifiers for the semantics and syntax of those variables.

For example, given the following Resource Object:

Nottingham

Expires November 10, 2012

[Page 5]

```
"http://example.org/rel/widget": {
  "href-template": "/widgets/{widget_id}",
  "href-vars": {
    "widget_id": "http://example.org/param/widget"
  },
  "hints": {
    "allow": ["GET", "PUT", "DELETE", "PATCH"],
    "representations": ["application/json"],
    "accept-patch": ["application/json-patch"],
    "accept-post": ["application/xml"],
    "accept-ranges": ["bytes"]
  }
}
```

If you understand that "http://example.org/param/widget" is an numeric identifier for a widget (perhaps by dereferencing that URL and reading the documentation), you can then find the resource corresponding to widget number 12345 at "http://example.org/widgets/12345" (assuming that the Home Document is located at "http://example.org/").

5. Resource Hints

Resource hints allow clients to find relevant information about interacting with a resource beforehand, as a means of optimising communications, as well as advertising available behaviours (e.g., to aid in laying out a user interface for consuming the API).

Hints are just that - they are not a "contract", and are to only be taken as advisory. The runtime behaviour of the resource always overrides hinted information.

For example, a resource might hint that the PUT method is allowed on all "widget" resources. this means that generally, the user has the ability to PUT to a particular resource, but a specific resource may reject a PUT based upon access control or other considerations. More fine-grained information might be gathered by interacting with the resource (e.g., via a GET), or by another resource "containing" it (such as a "widgets" collection).

This specification defines a set of common hints, based upon information that's discoverable by directly interacting with resources. A future draft will explain how to define new hints.

Nottingham

Expires November 10, 2012

[Page 6]

[5.1.](#) allow

Hints the HTTP methods that the current client will be able to use to interact with the resource; equivalent to the Allow HTTP response header.

Content MUST be an array of strings, containing HTTP methods.

[5.2.](#) representations

Hints the representation types that the resource produces and consumes, using the GET and PUT methods respectively, subject to the 'allow' hint.

Content MUST be an array of strings, containing media types.

[5.3.](#) accept-patch

Hints the PATCH request formats accepted by the resource for this client; equivalent to the Accept-Patch HTTP response header.

Content MUST be an array of strings, containing media types.

[5.4.](#) accept-post

Hints the POST request formats accepted by the resource for this client.

Content MUST be an array of strings, containing media types.

[5.5.](#) accept-ranges

Hints the range-specifiers available to the client for this resource; equivalent to the Accept-Ranges HTTP response header.

Content MUST be an array of strings, containing HTTP range-specifiers.

6. Creating and Serving Home Documents

When making a home document available, there are a few things to keep in mind:

- o A home document is best located at a memorable URI, because its URI will effectively become the URI for the API itself to clients.

- o Home documents can be personalised, just as "normal" home pages can. For example, you might advertise different URIs, and/or different kinds of link relations, depending on the client's identity.
- o Home documents SHOULD be assigned a freshness lifetime (e.g., "Cache-Control: max-age=3600") so that clients can cache them, to avoid having to fetch it every time the client interacts with the service.
- o Custom link relation types, as well as the URIs for variables, should lead to documentation for those constructs.

6.1. Managing Change in Home Documents

The URIs used in home documents MAY change over time. However, changing them can cause issues for clients that are relying on cached home documents containing old links.

To mitigate these risks, servers changing links SHOULD consider:

- o Reducing the freshness lifetime of home documents before a link change, so that clients are less likely to refer to an "old" document
- o Assure that they handle requests for the "old" URIs appropriately; e.g., with a 404 Not Found, or by redirecting the client to the new URI.
- o Alternatively, considering the "old" and "new" URIs as equally valid references for an "overlap" period.

Generally, servers ought not to change URIs without good cause.

6.2. Evolving and Mixing APIs with Home Documents

Using home documents affords the opportunity to change the "shape" of the API over time, without breaking old clients.

This includes introducing new functions alongside the old ones - by adding new link relation types with corresponding resource objects - as well as adding new template variables, media types, and so on.

It's important to realise that a home document can serve more than one "API" at a time; by listing all relevant relation types, it can effectively "mix" different APIs, allowing clients to work with different resources as they see fit.

6.3. Documenting APIs that use Home Documents

Another use case for "static" API description formats like WSDL and WADL is to generate documentation for the API from them.

An API that uses the home document format correctly won't have a need to do so, provided that the link relation types and media types it uses are well-documented already.

7. Consuming Home Documents

Clients might use home documents in a variety of ways.

In the most common case - actually consuming the API - the client will scan the Resources Object for the link relation(s) that it is interested in, and then to interact with the resource(s) referred to. Resource Hints can be used to optimise communication with the client, as well as to inform as to the permissible actions (e.g., whether PUT is likely to be supported).

Note that the home document is a "living" document; it does not represent a "contract", but rather is expected to be inspected before each interaction. In particular, links from the home document MUST NOT be assumed to be valid beyond the freshness lifetime of the home document, as per HTTP's caching model [[RFC2616](#)].

As a result, clients SHOULD cache the home document (as per [[RFC2616](#)], to avoid fetching it before every interaction (which would otherwise be required).

Likewise, a client encountering a 404 Not Found on a link SHOULD obtain a fresh copy of the home document, to assure that it is up-to-date.

8. Security Considerations

TBD

9. IANA Considerations

TBD

10. References

10.1. Normative References

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

- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", [RFC 4627](#), July 2006.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), October 2010.
- [RFC6570] Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., and D. Orchard, "URI Template", [RFC 6570](#), March 2012.

[10.2. Informative References](#)

- [RFC4288] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 4288](#), December 2005.

URIs

- [1] <http://www.w3.org/Submission/wadl/>
- [2] <http://microformats.org/>

[Appendix A. Frequently Asked Questions](#)

[A.1. Why not Microformats?](#)

Browser-centric Web applications use HTML as their representation format of choice. While it is possible to augment HTML for non-browser clients (using techniques like Microformats [2]), a few issues become evident when doing so:

- o HTML has a very forgiving syntax. While this is appropriate for browsers (especially considering that there are many million HTML authors in the world), it makes for a less-than-precise language for machines, resulting in both overhead (parsing and transmission) as well as lack of precision.
- o HTML is presentation-centric, making it tempting to reformat it from time to time, to improve the "look and feel" of a page. However, doing so can cause comparatively brittle non-browser clients to lose their understanding of the content's semantics, unless very careful controls are in place.

Nottingham

Expires November 10, 2012

[Page 10]

Because of this, it's most practical to define a separate format, and JSON is easily machine-readable, precise, and has a better chance of being managed for stability.

A.2. What about authentication?

In HTTP, authentication is discoverable by interacting with the resource (usually, by getting a 401 Unauthorized response status code, along with one or more challenges). While the home document could hint it, this isn't yet done, to avoid possible security complications.

A.3. What about "Faults" (i.e., errors)?

In HTTP, errors are conveyed by HTTP status codes. While this specification could (and even may) allow enumeration of possible error conditions, there's a concern that this will encourage applications to define many such "faults", leading to tight coupling between the application and its clients.

So, this is an area of possible future development; if any such mechanism appears here, it's likely to be quite restricted.

A.4. How Do I find the XML Schema / JSON Schema / etc. for a particular media type?

That isn't addressed by home documents. Ultimately, it's up to the media type accepted and generated by resources to define and constrain (or not) their syntax.

Appendix B. Open Issues

The following is a list of placeholders for open issues.

- o Refining and extending representation formats - "application/xml", for example, isn't enough. While a media type for every representation is one answer, something like 'profile' might be good too.
- o Object for contact details - do we need an object that describes who's running the API, etc?
- o Defining new hints - guidance is needed on minting new hints. Possibly a registry.
- o Defining new top-level properties - how new ones are minted, registry, etc.
- o Defining new Resource Object properties - how new ones are minted, registry, etc.

Nottingham

Expires November 10, 2012

[Page 11]

Author's Address

Mark Nottingham
Rackspace

Email: mnot@mnot.net

URI: <http://www.mnot.net/>