

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 14, 2017

E. Wilde
CA Technologies
March 13, 2017

Link Relation Types for Web Services
draft-wilde-service-link-rel-03

Abstract

Many resources provided on the Web are part of sets of resources that are provided in a context that is managed by one particular service provider. Often, these sets of resources are referred to as "Web Services" or "Web APIs". This specification defines link relations for representing relationships from those resources to ones that provide documentation or descriptions of the Web services. The difference between these concepts is that documentation is primarily intended for human consumers, whereas descriptions are primarily intended for automated consumers. It also defines a link relation to identify a status resource that is used to represent operational information about a service's status.

Note to Readers

Please discuss this draft on the ART mailing list (<https://www.ietf.org/mailman/listinfo/art>).

Online access to all versions and files is available on GitHub (<https://github.com/dret/I-D/tree/master/service-link-rel>).

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.	Introduction	2
2.	Terminology	3
3.	Web Services	4
3.1.	Documenting Web Services	4
3.2.	Describing Web Services	5
3.3.	Unified Documentation/Description	5
4.	Link Relations for Web Services	5
4.1.	The service-doc Link Relation Type	5
4.2.	The service-desc Link Relation Type	6
5.	Web Service Status Resources	6
6.	IANA Considerations	6
6.1.	Link Relation Type: service-doc	6
6.2.	Link Relation Type: service-desc	7
6.3.	Link Relation Type: status	7
7.	Security Considerations	7
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	7
Appendix A.	Acknowledgements	7
	Author's Address	8

[1.](#) Introduction

One of the defining aspects of the Web is that it is possible to interact with Web resources without any prior knowledge of the specifics of the resource. Following Web Architecture

[[W3C.REC-webarch-20041215](#)] by using URIs, HTTP, and media types, the Web's uniform interface allows interactions with resources without the more complex binding procedures of other approaches.

Many resources on the Web are provided as part of a set of resources that are referred to as a "Web Service" or a "Web API". In many cases, these services or APIs are defined and managed as a whole, and it may be desirable for clients to be able to discover this service information.

Service information can be broadly separated into two categories: One category is primarily targeted for human users and often uses generic representations for human readable documents, such as HTML or PDF. The other category is structured information that follows some more formalized description model, and is primarily intended for consumption by machines, for example for tools and code libraries.

In the context of this memo, the human-oriented variant is referred to as "documentation", and the machine-oriented variant is referred to as "description".

These two categories are not necessarily mutually exclusive, as there are representations that have been proposed that are intended for both human consumption, and for interpretation by machine clients. In addition, a typical pattern for service documentation/description is that there is human-oriented high-level documentation that is intended to put a service in context and explain the general model, which is complemented by a machine-level description that is intended as a detailed technical description of the service. These two resources could be interlinked, but since they are intended for different audiences, it can make sense to provide entry points for both of them.

This memo places no constraints on the specific representations used for either of those two categories. It simply allows providers of a Web service to make the documentation and/or the description of their services discoverable, and defines two link relations that serve that purpose.

In addition, this memo defines a link relation that allows providers

of a Web service to link to a resource that represents status information about the service. This information often represents operational information that allows service consumers to retrieve information about "service health" and related issues.

[2.](#) Terminology

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 [RFC 2119](#) [[RFC2119](#)].

[3.](#) Web Services

"Web Services" or "Web APIs" (sometimes also referred to as "HTTP API" or "REST API") are a way to expose information and services on the Web. Following the principles of Web architecture [[W3C.REC-webarch-20041215](#)], they expose URI-identified resources, which are then accessed and transferred using a specific representation. Many services use representations that contain links, and often these links are typed links.

Using typed links, resources can identify relationship types to other resources. [RFC 5988](#) [[RFC5988](#)] establishes a framework of registered link relation types, which are identified by simple strings and registered in an IANA registry. Any resource that supports typed links according to [RFC 5988](#) can then use these identifiers to represent resource relationships on the Web without having to re-invent registered relation types.

In recent years, Web services as well as their documentation and description languages have gained popularity, due to the general popularity of the Web as a platform for providing information and services. However, the design of documentation and description languages varies with a number of factors, such as the general application domain, the preferred application data model, and the preferred approach for exposing services.

This specification allows service providers to use a unified way to link to service documentation and/or description. This link should not make any assumptions about the provided type of documentation

and/or description, so that service providers can choose the ones that best fit their services and needs.

[3.1.](#) Documenting Web Services

In the context of this specification, "documentation" refers to information that is primarily intended for human consumption. Typical representations for this kind of documentation are HTML and PDF.

Documentation is often structured, but the exact kind of structure depends on the structure of the service that is documented, as well as on the specific way in which the documentation authors choose to document it.

[3.2.](#) Describing Web Services

In the context of this specification, "description" refers to information that is primarily intended for machine consumption. Typical representations for this are dictated by the technology underlying the service itself, which means that in today's technology landscape, description formats exist that are based on XML, JSON, RDF, and a variety of other structured data models. Also, in each of those technologies, there may be a variety of languages that are defined to achieve the same general purpose of describing a Web service.

Descriptions are always structured, but the structuring principles depend on the nature of the described service. For example, one of the earlier service description approaches, the Web Services Description Language (WSDL), uses "operations" as its core concept, which are essentially identical to function calls, because the underlying model is based on that of the Remote Procedure Call (RPC) model. Other description languages for non-RPC approaches to services will use different structuring approaches.

[3.3.](#) Unified Documentation/Description

If service providers use an approach where there is no distinction of service documentation [Section 3.1](#) and service description [Section 3.2](#), then they may not feel the need to use two separate links. In such a case, an alternative approach is to use the "service" link relation type, which has no indication of whether it links to documentation or description, and thus may be better fit if no such differentiation is required.

[4.](#) Link Relations for Web Services

In order to allow Web services to represent the relation of individual resources to service documentation or description, this specification introduces and registers two new link relation types.

[4.1.](#) The service-doc Link Relation Type

The "service-doc" link relation type is used to represent the fact that a resource is part of a bigger set of resources that are documented at a specific URI. The target resource is expected to provide documentation that is primarily intended for human consumption.

[4.2.](#) The service-desc Link Relation Type

The "service-desc" link relation type is used to represent the fact that a resource is part of a bigger set of resources that are described at a specific URI. The target resource is expected to provide a service description that is primarily intended for machine consumption. In many cases, it is provided in a representation that is consumed by tools, code libraries, or similar components.

[5.](#) Web Service Status Resources

Web services providing access to a set of resources often are hosted and operated in an environment for which status information may be available. This information may be as simple as confirming that a service is operational, or may provide additional information about

different aspects of a service, and/or a history of status information, possibly listing incidents and their resolution.

The "status" link relation type can be used to link to such a status resource, allowing service consumers to retrieve status information about a Web service's status. Such a link may not be available from all resources provided by a Web service, but from key resources such as a Web service's home resource [[I-D.nottingham-json-home](#)].

This memo does not restrict the representation of a status resource in any way. It may be primarily focused on human or machine consumption, or a combination of both. It may be a simple "traffic light" indicator for service health, or a more sophisticated representation conveying more detailed information such as service subsystems and/or a status history.

[6.](#) IANA Considerations

The link relation types below have been registered by IANA per [Section 6.2.1 of RFC 5988](#) [[RFC5988](#)]:

[6.1.](#) Link Relation Type: service-doc

Relation Name: service-doc

Description: Linking to service documentation that is primarily intended for human consumption.

Reference: [[This document]]

[6.2.](#) Link Relation Type: service-desc

Relation Name: service-desc

Description: Linking to service description that is primarily intended for consumption by machines.

Reference: [[This document]]

[6.3.](#) Link Relation Type: status

Relation Name: status

Description: Linking to a resource that represents the status of a Web service or API.

Reference: [[This document]]

[7.](#) Security Considerations

...

[8.](#) References

[8.1.](#) Normative References

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

[RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), October 2010.

[8.2.](#) Informative References

[I-D.nottingham-json-home]
Nottingham, M., "Home Documents for HTTP APIs", [draft-nottingham-json-home-04](#) (work in progress), May 2016.

[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.](#) Acknowledgements

Thanks for comments and suggestions provided by Mike Amundsen, Oliver Gierke, Sebastien Lambla, and Darrell Miller.

Erik Wilde
CA Technologies

Email: erik.wilde@dret.net

URI: <http://dret.net/netdret/>