

Workgroup: HTTP
Updates: [5323](#) (if approved)
Published: 8 June 2021
Intended Status: Standards Track
Expires: 10 December 2021
Authors: J. Reschke A. Malhotra J.M. Snell
 greenbytes

HTTP SEARCH Method

Abstract

This specification updates the definition and semantics of the HTTP SEARCH request method originally defined by RFC 5323.

Editorial Note

This note is to be removed before publishing as an RFC.

Discussion of this draft takes place on the HTTP working group mailing list (ietf-http-wg@w3.org), which is archived at <https://lists.w3.org/Archives/Public/ietf-http-wg/>.

Working Group information can be found at <https://httpwg.org/>; source code and issues list for this draft can be found at <https://github.com/httpwg/http-extensions/labels/safe-method-w-body>.

The changes in this draft are summarized in [Appendix A.1](#).

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 10 December 2021.

Copyright Notice

Copyright (c) 2021 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. SEARCH](#)
- [3. The "Accept-Search" Header Field](#)
- [4. Examples](#)
 - [4.1. Simple SEARCH with a Direct Response](#)
 - [4.2. Simple SEARCH with indirect response \(303 See Other\)](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Normative References](#)
- [Appendix A. Change Log](#)
 - [A.1. Since draft-ietf-httpbis-safe-method-w-body-00](#)
- [Authors' Addresses](#)

1. Introduction

This specification updates the HTTP SEARCH method originally defined in [[RFC5323](#)].

Many existing HTTP-based applications use the HTTP GET and POST methods in various ways to implement the functionality provided by SEARCH.

Using a GET request with some combination of query parameters included within the request URI (as illustrated in the example below) is arguably the most common mechanism for implementing search in web applications. With this approach, implementations are required to parse the request URI into distinct path (everything before the '?') and query elements (everything after the '?'). The path identifies the resource processing the query (in this case 'http://example.org/feed') while the query identifies the specific parameters of the search operation.

A typical use of HTTP GET for requesting a search

```
GET /feed?q=foo&limit=10&sort=-published HTTP/1.1
Host: example.org
```

While there are definite advantages to using GET requests in this manner, the disadvantages should not be overlooked. Specifically:

- *Without specific knowledge of the resource and server to which the GET request is being sent, there is no way for the client to know that a search operation is being requested. Identical requests sent to two different servers can implement entirely different semantics.
- *Encoding query parameters directly into the request URI effectively casts every possible combination of query inputs as distinct resources. For instance, because mechanisms such as HTTP caching handle request URIs as opaque character sequences, queries such as 'http://example.org/?q=foo' and 'http://example.org/?q=Foo' will be treated as entirely separate resources even if they yield identical results.
- *While most modern browser and server implementations allow for long request URIs, there is no standardized minimum or maximum length for URIs in general. Many resource constrained devices enforce strict limits on the maximum number of characters that can be included in a URI. Such limits can prove impractical for large or complex query parameters.
- *Query expressions included within a request URI must either be restricted to relatively simple key value pairs or encoded such that the query can be safely represented in the limited character-set allowed by URL standards. Such encoding can add significant complexity, introduce bugs, or otherwise reduce the overall visibility of the query being requested.

As an alternative to using GET, many implementations make use of the HTTP POST method to perform queries, as illustrated in the example below. In this case, the input parameters to the search operation are passed along within the request payload as opposed to using the request URI.

A typical use of HTTP POST for requesting a search

```
POST /feed HTTP/1.1
Host: example.org
Content-Type: application/x-www-form-urlencoded

q=foo&limit=10&sort=-published
```

This variation, however, suffers from the same basic limitation as GET in that it is not readily apparent -- absent specific knowledge of the resource and server to which the request is being sent -- that a search operation is what is being requested. Web applications use the POST method for a wide variety of uses including the

creation or modification of existing resources. Sending the request above to a different server, or even repeatedly sending the request to the same server could have dramatically different effects.

The SEARCH method provides a solution that spans the gap between the use of GET and POST. As with POST, the input to the query operation is passed along within the payload of the request rather than as part of the request URI. Unlike POST, however the semantics of the SEARCH method are specifically defined.

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.

2. SEARCH

The SEARCH method is used to initiate a server-side search. Unlike the HTTP GET method, which requests that a server return a representation of the resource identified by the target URI (as defined by [Section 7.1](#) of [[RFCHTTP](#)]), the SEARCH method is used to ask the server to perform a query operation (described by the request payload) over some set of data scoped to the effective request URI. The payload returned in response to a SEARCH cannot be assumed to be a representation of the resource identified by the effective request URI.

The body payload of the request defines the query. Implementations **MAY** use a request body of any content type with the SEARCH method; however, for backwards compatibility with existing WebDAV implementations, SEARCH requests that use the text/xml or application/xml media types with a root element ([Section 2.1](#) of [[XML](#)]) in the "DAV:" XML namespace ([[XMLNS](#)]) **MUST** be processed per the requirements established by [[RFC5323](#)].

SEARCH requests are both safe and idempotent with regards to the resource identified by the request URI. That is, SEARCH requests do not alter the state of the targeted resource. However, while processing a search request, a server can be expected to allocate computing and memory resources or even create additional HTTP resources through which the response can be retrieved.

A successful response to a SEARCH request is expected to provide some indication as to the final disposition of the search operation. For instance, a successful search that yields no results can be represented by a 204 No Content response. If the response includes a content, it is expected to describe the results of the search operation. In some cases, the server may choose to respond

indirectly to the SEARCH request by returning a 3xx Redirection with a Location header field specifying an alternate Request URI from which the search results can be retrieved using an HTTP GET request. Various non-normative examples of successful SEARCH responses are illustrated in [Section 4](#).

The response to a SEARCH request is not cacheable. It ought to be noted, however, that because SEARCH requests are safe and idempotent, responses to a SEARCH **MUST NOT** invalidate previously cached responses to other requests directed at the same effective request URI. By default, that is. We need to figure out under which conditions we can make the result cacheable.

The semantics of the SEARCH method change to a "conditional SEARCH" if the request message includes an If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match, or If-Range header field ([\[RFCHTTP\]](#), [Section 13](#)). A conditional SEARCH requests that the query be performed only under the circumstances described by the conditional header field(s). It is important to note, however, that such conditions are evaluated against the state of the target resource itself as opposed to the collected results of the search operation.

3. The "Accept-Search" Header Field

The "Accept-Search" response header field **MAY** be used by a server to directly signal support for the SEARCH method while identifying the specific query format media types that may be used.

Accept-Search = 1#media-type

The Accept-Search header field specifies a comma-separated listing of media types (with optional parameters) as defined by [Section 8.3.1](#) of [\[RFCHTTP\]](#).

The order of types listed by the Accept-Search header field is insignificant.

4. Examples

The non-normative examples in this section make use of a simple, hypothetical plain-text based query syntax based on SQL with results returned as comma-separated values. This is done for illustration purposes only. Implementations are free to use any format they wish on both the request and response.

4.1. Simple SEARCH with a Direct Response

A simple query with a direct response:

SEARCH /contacts HTTP/1.1
Host: example.org
Content-Type: example/query
Accept: text/csv

select surname, givenname, email limit 10

Response:

HTTP/1.1 200 OK
Content-Type: text/csv

surname, givenname, email
Smith, John, john.smith@example.org
Jones, Sally, sally.jones@example.com
Dubois, Camille, camille.dubois@example.net

4.2. Simple SEARCH with indirect response (303 See Other)

A simple query with an Indirect Response (303 See Other):

SEARCH /contacts HTTP/1.1
Host: example.org
Content-Type: example/query
Accept: text/csv

select surname, givenname, email limit 10

Response:

HTTP/1.1 303 See Other
Location: http://example.org/contacts/query123

Fetch Query Response:

GET /contacts/query123 HTTP/1.1
Host: example.org

Response:

HTTP/1.1 200 OK
Content-Type: text/csv

surname, givenname, email
Smith, John, john.smith@example.org
Jones, Sally, sally.jones@example.com
Dubois, Camille, camille.dubois@example.net

5. Security Considerations

The SEARCH method is subject to the same general security considerations as all HTTP methods as described in [RFCHTTP].

6. IANA Considerations

IANA is requested to update the registration of the SEARCH method in the permanent registry at <<http://www.iana.org/assignments/http-methods>> (see [Section 16.1.1](#) of [RFCHTTP]).

Method Name	Safe	Idempotent	Specification
SEARCH	Yes	Yes	Section 2

Table 1

7. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5323] Reschke, J., Ed., Reddy, S., Davis, J., and A. Babich, "Web Distributed Authoring and Versioning (WebDAV) SEARCH", RFC 5323, DOI 10.17487/RFC5323, November 2008, <<https://www.rfc-editor.org/info/rfc5323>>.
- [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>>.
- [RFCHTTP] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", Work in Progress, Internet-Draft, draft-ietf-httpbis-semantics-16, 27 May 2021, <<https://tools.ietf.org/html/draft-ietf-httpbis-semantics-16>>.
- [XML] Bray, T., Paoli, J., Sperberg-McQueen, M., Maler, E., and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)", W3C Recommendation REC-xml-20081126, 26 November 2008, <<https://www.w3.org/TR/2008/REC-xml-20081126/>>. Latest version available at <https://www.w3.org/TR/xml/>.
- [XMLNS] Bray, T., Hollander, D., Layman, A., Tobin, R., and H. Thompson, "Namespaces in XML 1.0 (Third Edition)", W3C Recommendation REC-xml-names-20091208, 8 December 2009, <<https://www.w3.org/TR/2009/REC-xml-names-20091208/>>. Latest version available at <https://www.w3.org/TR/xml-names/>.

Appendix A. Change Log

This section is to be removed before publishing as an RFC.

(see <https://trac.tools.ietf.org/tools/xml2rfc/trac/ticket/622>)

A.1. Since draft-ietf-httpbis-safe-method-w-body-00

*Use "example/query" media type instead of undefined "text/query"
(<https://github.com/httpwg/http-extensions/issues/1450>)

*In [Section 3](#), adjust the grammar to just define the field value
(<https://github.com/httpwg/http-extensions/issues/1470>)

*Update to latest HTTP core spec, and adjust terminology accordingly
(<https://github.com/httpwg/http-extensions/issues/1473>)

*Reference RFC 8174 and markup bcp14 terms (<https://github.com/httpwg/http-extensions/issues/1497>)

*Update HTTP reference (<https://github.com/httpwg/http-extensions/issues/1524>)

*Relax restriction of generic XML media type in request body
(<https://github.com/httpwg/http-extensions/issues/1535>)

Authors' Addresses

Julian Reschke
greenbytes GmbH
Hafenweg 16
48155 Münster
Germany

Email: julian.reschke@greenbytes.de
URI: <https://greenbytes.de/tech/webdav/>

Ashok Malhotra

Email: malhotrasahib@gmail.com

James M Snell

Email: jasnell@gmail.com