

Health Check Response Format for HTTP APIs
draft-inadarei-api-health-check-01

Abstract

This document proposes a service health check response format for HTTP APIs.

Note to Readers

RFC EDITOR: please remove this section before publication

The issues list for this draft can be found at <https://github.com/inadarei/rfc-healthcheck/issues> [1].

The most recent draft is at <https://inadarei.github.io/rfc-healthcheck/> [2].

Recent changes are listed at <https://github.com/inadarei/rfc-healthcheck/commits/master> [3].

See also the draft's current status in the IETF datatracker, at <https://datatracker.ietf.org/doc/draft-inadarei-api-health-check/> [4].

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 September 25, 2018.

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.](#) Notational Conventions [3](#)
- [3.](#) API Health Response [3](#)
- [4.](#) The Details Object [5](#)
- [5.](#) Example Output [7](#)
- [6.](#) Security Considerations [9](#)
- [7.](#) IANA Considerations [9](#)
- [8.](#) Acknowledgements [10](#)
- [9.](#) Creating and Serving Health Responses [10](#)
- [10.](#) Consuming Health Check Responses [11](#)
- [11.](#) References [11](#)
 - [11.1.](#) Normative References [11](#)
 - [11.2.](#) Informative References [12](#)
 - [11.3.](#) URIs [12](#)
- Author's Address [12](#)

1. Introduction

The vast majority of modern APIs driving data to web and mobile applications use HTTP [[RFC7230](#)] as their protocol. The health and uptime of these APIs determine availability of the applications themselves. In distributed systems built with a number of APIs, understanding the health status of the APIs and making corresponding decisions, for failover or circuit-breaking, are essential for providing highly available solutions.

There exists a wide variety of operational software that relies on the ability to read health check response of APIs. There is currently no standard for the health check output response, however, so most applications either rely on the basic level of information included in HTTP status codes [[RFC7231](#)] or use task-specific formats.

Usage of task-specific or application-specific formats creates significant challenges, disallowing any meaningful interoperability across different implementations and between different tooling.

Standardizing a format for health checks can provide any of a number of benefits, including:

- o Flexible deployment - since operational tooling and API clients can rely on rich, uniform format, they can be safely combined and substituted as needed.
- o Evolvability - new APIs, conforming to the standard, can safely be introduced in any environment and ecosystem that also conforms to the same standard, without costly coordination and testing requirements.

This document defines a "health check" format using the JSON format [[RFC8259](#)] for APIs to use as a standard point for the health information they offer. Having a well-defined format for this purpose promotes good practice and tooling.

2. Notational Conventions

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. API Health Response

The API Health Response Format (or, interchangeably, "health check response format") uses the JSON format described in [[RFC8259](#)] and has the media type "application/health+json".

Its content consists of a single mandatory root field ("status") and several optional fields:

- o status: (required) indicates whether the service status is acceptable or not. API publishers SHOULD use following values for the field:
 - * "pass": healthy,
 - * "fail": unhealthy, and
 - * "warn": healthy, with some concerns.

The value of the status field is tightly related with the HTTP response code returned by the health endpoint. For "pass" and

"warn" statuses HTTP response code in the 2xx-3xx range MUST be used. For "fail" status HTTP response code in the 4xx-5xx range MUST be used. In case of the "warn" status, endpoint SHOULD return HTTP status in the 2xx-3xx range and additional information SHOULD be provided, utilizing optional fields of the response.

A health endpoint is only meaningful in the context of the component it indicates the health of. It has no other meaning or purpose. As such, its health is a conduit to the health of the component. Clients SHOULD assume that the HTTP response code returned by the health endpoint is applicable to the entire component (e.g. a larger API or a microservice). This is compatible with the behavior that current infrastructural tooling expects: load-balancers, service discoveries and others, utilizing health-checks.

- o version: (optional) public version of the service.
- o releaseID: (optional) in well-designed APIs, backwards-compatible changes in the service should not update a version number. APIs usually change their version number as infrequently as possible, to preserve stable interface. However implementation of an API may change much more frequently, which leads to the importance of having separate "release number" or "releaseID" that is different from the public version of the API.
- o notes: (optional) array of notes relevant to current state of health
- o output: (optional) raw error output, in case of "fail" or "warn" states. This field SHOULD be omitted for "pass" state.
- o details: (optional) an object representing status of sub-components of the service in question. Please refer to the "The Details Object" section for more information.
- o links: (optional) an array of objects containing link relations and URIs [[RFC3986](#)] for external links that MAY contain more information about the health of the endpoint. Per web-linking standards [[RFC5988](#)] a link relationship SHOULD either be a common/registered one or be indicated as a URI, to avoid name clashes. If a "self" link is provided, it MAY be used by clients to check health via HTTP response code, as mentioned above.
- o serviceID: (optional) unique identifier of the service, in the application scope.
- o description: (optional) human-friendly description of the service.

4. The Details Object

The "details" object MAY have a number of unique keys, one for each logical sub-components. Since each sub-component may be backed by several nodes with varying health statuses, the key points to an array of objects. In case of a single-node sub-component (or if presence of nodes is not relevant), a single-element array should be used as the value, for consistency.

The key identifying an element in the object should be a unique string within the details section. It MAY have two parts: "{componentName}:{metricName}", in which case the meaning of the parts SHOULD be as follows:

- o componentName: (optional) human-readable name for the component. MUST not contain a colon, in the name, since colon is used as a separator.
- o metricName: (optional) name of the metrics that the status is reported for. MUST not contain a colon, in the name, since colon is used as a separator and can be one of:
 - * Pre-defined value from this spec. Pre-defined values include:
 - + utilization
 - + responseTime
 - + connections
 - + uptime
 - * A common and standard term from a well-known source such as schema.org, IANA or microformats.
 - * A URI that indicates extra semantics and processing rules that MAY be provided by a resource at the other end of the URI. URIs do not have to be dereferenceable, however. They are just a namespace, and the meaning of a namespace CAN be provided by any convenient means (e.g. publishing an RFC, Swagger document or a nicely printed book).

On the value eside of the equation, each "component details" object in the array MAY have one of the following object keys:

- o componentId: (optional) unique identifier of an instance of a specific sub-component/dependency of a service. Multiple objects

with the same componentID MAY appear in the details, if they are from different nodes.

- o componentType: (optional) SHOULD be present if componentName is present. Type of the component. Could be one of:
 - * Pre-defined value from this spec. Pre-defined values include:
 - + component
 - + datastore
 - + system
 - * A common and standard term from a well-known source such as schema.org, IANA or microformats.
 - * A URI that indicates extra semantics and processing rules that MAY be provided by a resource at the other end of the URI. URIs do not have to be dereferenceable, however. They are just a namespace, and the meaning of a namespace CAN be provided by any convenient means (e.g. publishing an RFC, Swagger document or a nicely printed book).
- o metricValue: (optional) could be any valid JSON value, such as: string, number, object, array or literal.
- o metricUnit: (optional) SHOULD be present if metricValue is present. Could be one of:
 - * A common and standard term from a well-known source such as schema.org, IANA, microformats, or a standards document such as [[RFC3339](#)].
 - * A URI that indicates extra semantics and processing rules that MAY be provided by a resource at the other end of the URI. URIs do not have to be dereferenceable, however. They are just a namespace, and the meaning of a namespace CAN be provided by any convenient means (e.g. publishing an RFC, Swagger document or a nicely printed book).
- o time: the date-time, in ISO8601 format, at which the reading of the metricValue was recorded. This assumes that the value can be cached and the reading typically doesn't happen in real time, for performance and scalability purposes.
- o output: (optional) has the exact same meaning as the top-level "output" element, but for the sub-component.

- o links: (optional) has the exact same meaning as the top-level "output" element, but for the sub-component.

5. Example Output

```
GET /health HTTP/1.1
Host: example.org
Accept: application/health+json

HTTP/1.1 200 OK
Content-Type: application/health+json
Cache-Control: max-age=3600
Connection: close
```

```
{
  "status": "pass",
  "version": "1",
  "releaseID": "1.2.2",
  "notes": [""],
  "output": "",
  "serviceID": "f03e522f-1f44-4062-9b55-9587f91c9c41",
  "description": "health of authz service",
  "details": {
    "cassandra:responseTime": [
      {
        "componentId": "dfd6cf2b-1b6e-4412-a0b8-f6f7797a60d2",
        "componentType": "datastore",
        "metricValue": 250,
        "metricUnit": "ms",
        "status": "pass",
        "time": "2018-01-17T03:36:48Z",
        "output": ""
      }
    ],
    "cassandra:connections": [
      {
        "componentId": "dfd6cf2b-1b6e-4412-a0b8-f6f7797a60d2",
        "type": "datastore",
        "metricValue": 75,
        "status": "warn",
        "time": "2018-01-17T03:36:48Z",
        "output": "",
        "links": {
          "self": "http://api.example.com/dbnode/dfd6cf2b/health"
        }
      }
    ],
    "uptime": [
```



```
{
  "componentType": "system",
  "metricValue": 1209600.245,
  "metricUnit": "s",
  "status": "pass",
  "time": "2018-01-17T03:36:48Z"
}
],
"cpu:utilization": [
  {
    "componentId": "6fd416e0-8920-410f-9c7b-c479000f7227",
    "node": 1,
    "componentType": "system",
    "metricValue": 85,
    "metricUnit": "percent",
    "status": "warn",
    "time": "2018-01-17T03:36:48Z",
    "output": ""
  },
  {
    "componentId": "6fd416e0-8920-410f-9c7b-c479000f7227",
    "node": 2,
    "componentType": "system",
    "metricValue": 85,
    "metricUnit": "percent",
    "status": "warn",
    "time": "2018-01-17T03:36:48Z",
    "output": ""
  }
],
"memory:utilization": [
  {
    "componentId": "6fd416e0-8920-410f-9c7b-c479000f7227",
    "node": 1,
    "componentType": "system",
    "metricValue": 8.5,
    "metricUnit": "GiB",
    "status": "warn",
    "time": "2018-01-17T03:36:48Z",
    "output": ""
  },
  {
    "componentId": "6fd416e0-8920-410f-9c7b-c479000f7227",
    "node": 2,
    "componentType": "system",
    "metricValue": 5500,
    "metricUnit": "MiB",
    "status": "pass",
```



```
        "time": "2018-01-17T03:36:48Z",
        "output": ""
    }
]
},
"links": {
  "about": "http://api.example.com/about/authz",
  "http://api.x.io/rel/thresholds":
    "http://api.x.io/about/authz/thresholds"
}
}
```

6. Security Considerations

Clients need to exercise care when reporting health information. Malicious actors could use this information for orchestrating attacks. In some cases the health check endpoints may need to be authenticated and institute role-based access control.

7. IANA Considerations

The media type for health check response is `application/health+json`.

- o Media type name: `application`
- o Media subtype name: `health+json`
- o Required parameters: n/a
- o Optional parameters: n/a
- o Encoding considerations: `binary`
- o Security considerations: Health+JSON shares security issues common to all JSON content types. See [RFC 8259](#) Section #12 for additional information.

Health+JSON allows utilization of Uniform Resource Identifiers (URIs) and as such shares security issues common to URI usage. See [RFC 3986](#) Section #7 for additional information.

Since Hyper+JSON can carry wide variety of data, some data may require privacy or integrity services. This specification does not prescribe any specific solution and assumes that concrete implementations will utilize common, trusted approaches such as TLS/HTTPS, OAuth2 etc.

- o Interoperability considerations: None

- o Published specification: this RFC draft
- o Applications which use this media: Various
- o Fragment identifier considerations: Health+JSON follows [RFC6901](#) for implementing URI Fragment Identification standard to JSON content types.
- o Restrictions on usage: None
- o Additional information:
 1. Deprecated alias names for this type: n/a
 2. Magic number(s): n/a
 3. File extension(s): .json
 4. Macintosh file type code: TEXT
 5. Object Identifiers: n/a
- o General Comments:
- o Person to contact for further information:
 1. Name: Irakli Nadareishvili
 2. Email: irakli@gmail.com
- o Intended usage: Common
- o Author/Change controller: Irakli Nadareishvili

8. Acknowledgements

Thanks to Mike Amundsen, Erik Wilde, Justin Bachorik and Randall Randall for their suggestions and feedback. And to Mark Nottingham for blueprint for authoring RFCs easily.

9. Creating and Serving Health Responses

When making an health check endpoint available, there are a few things to keep in mind:

- o A health response endpoint is best located at a memorable and commonly-used URI, such as "health" because it will help self-discoverability by clients.

- o Health check responses can be personalized. For example, you could advertise different URIs, and/or different kinds of link relations, to afford different clients access to additional health check information.
- o Health check responses must be assigned a freshness lifetime (e.g., "Cache-Control: max-age=3600") so that clients can determine how long they could cache them, to avoid overly frequent fetching and unintended DDOS-ing of the service.
- o Custom link relation types, as well as the URIs for variables, should lead to documentation for those constructs.

10. Consuming Health Check Responses

Clients might use health check responses in a variety of ways.

Note that the health check response is a "living" document; links from the health check response MUST NOT be assumed to be valid beyond the freshness lifetime of the health check response, as per HTTP's caching model [[RFC7234](#)].

As a result, clients ought to cache the health check response (as per [[RFC7234](#)]), to avoid fetching it before every interaction (which would otherwise be required).

Likewise, a client encountering a 404 (Not Found) on a link is encouraged to obtain a fresh copy of the health check response, to assure that it is up-to-date.

11. References

11.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#), DOI 10.17487/RFC5988, October 2010, <<https://www.rfc-editor.org/info/rfc5988>>.

- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", [RFC 7234](#), DOI 10.17487/RFC7234, June 2014, <<https://www.rfc-editor.org/info/rfc7234>>.
- [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>>.

11.2. Informative References

- [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>>.
- [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>>.
- [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>>.

11.3. URIs

- [1] <https://github.com/inadarei/rfc-healthcheck/issues>
- [2] <https://inadarei.github.io/rfc-healthcheck/>
- [3] <https://github.com/inadarei/rfc-healthcheck/commits/master>
- [4] <https://datatracker.ietf.org/doc/draft-inadarei-api-health-check/>

Author's Address

Irakli Nadareishvili
114 5th Avenue
New York
United States

Email: irakli@gmail.com

URI: <http://www.freshblurbs.com>