

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 7, 2018

T. Pauly, Ed.
Apple Inc.
D. Thakore, Ed.
CableLabs
February 03, 2018

Captive Portal API
draft-ietf-capport-api-00

Abstract

This document describes an HTTP API that allows hosts to interact with a Captive Portal system.

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 August 7, 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.	Workflow	2
3.	API Details	3
3.1.	URI of Captive Portal API endpoint	3
3.2.	JSON Keys	3
3.3.	Example Exchange	4
4.	Security Considerations	4
4.1.	Privacy Considerations	4
5.	IANA Considerations	5
6.	Acknowledgments	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	5
	Authors' Addresses	5

[1.](#) Introduction

This document describes a HyperText Transfer Protocol (HTTP) Application Program Interface (API) that allows hosts to interact with a Captive Portal system. The API defined in this document has been designed to meet the requirements in the Captive Portal Architecture [[I-D.ietf-capport-architecture](#)]. Specifically, the API provides:

- o The state of captivity (whether or not the host has access to the Internet)
- o A URI that a host's browser can present to a user to get out of captivity
- o An encrypted connection (TLS for both the API and portal URI)

[2.](#) Workflow

The Captive Portal Architecture defines three steps of interaction between hosts and a Captive Portal service:

1. Provisioning, in which a host discovers that a network has a captive portal, and learns the URI of the API server
2. API Server interaction, in which a host queries the state of the captive portal and retrieves the necessary information to get out of captivity

3. Enforcement, in which the enforcement device in the network blocks disallowed traffic, and sends ICMP messages to let hosts know they are blocked by the captive portal

This document is focused on the second step. It is assumed that the location of the Captive Portal API server has been discovered by the host as part of the first step. The mechanism for discovering the API Server endpoint is not covered by this document.

3. API Details

3.1. URI of Captive Portal API endpoint

The URI of the API endpoint MUST be accessed using HTTP over TLS (HTTPS) and SHOULD be served on port 443 [[RFC2818](#)]. The host SHOULD NOT assume that the URI for a given network attachment will stay the same, and SHOULD rely on the discovery or provisioning process each time it joins the network. Depending on how the Captive Portal system is configured, the URI may be unique for each host and between sessions for the same host.

For example, if the Captive Portal API server is hosted at example.org, the URI's of the API could be:

- o "https://example.org/captive-portal/api"
- o "https://example.org/captive-portal/api/X54PD"

3.2. JSON Keys

The Captive Portal API data structures are specified in JavaScript Object Notation (JSON) [[RFC7159](#)].

The following keys are defined at the top-level of the JSON structure returned by the API server:

- o "permitted" (required, boolean): indicates whether or not the Captive Portal is open to the requesting host
- o "hmac-key" (required, string): provides a per-host key that can be used to authenticate messages from the Captive Portal enforcement server
- o "user-portal-url" (required, string): provides the URL of a web portal that can be presented to a user to interact with

- o "expire-date" (optional, string formatted as [[RFC3339](#)] datetime): indicates the date and time after which the host will be in a captive state
- o "bytes-remaining" (optional, integer): indicates the number of bytes left, after which the host will be in a captive state

Note that the use of the hmac-key is not defined in this document, but is intended for use in the enforcement step of the Captive Portal Architecture.

[3.3.](#) Example Exchange

To request the Captive Portal JSON content, a host sends an HTTP GET request:

```
GET /captive-portal/api/X54PD
Host: example.org
Accept: application/json
```

The server then responds with the JSON content for that client:

```
HTTP/1.1 200 OK
Cache-Control: private
Date: Mon, 04 Dec 2013 05:07:35 GMT
Content-Type: application/json

{
  "permitted": false,
  "hmac-key":
    "7cec81acce3176b262a46363666a01881b0e3bf60d97a98b5409b71cc60a1ac0"
  "user-portal-url": "https://example.org/portal.html"
  "expire-date": "2014-01-01T23:28:56.782Z"
}
```

[4.](#) Security Considerations

TBD: Provide complete security requirements and analysis.

[4.1.](#) Privacy Considerations

Information passed in this protocol may include a user's personal information, such as a full name and credit card details. Therefore, it is important that Captive Portal API Servers do not allow access to the Captive Portal API over unencrypted sessions.

5. IANA Considerations

TBD: None

6. Acknowledgments

This work in this document was started by Mark Donnelly and Margaret Cullen. Thanks to everyone in the CAPPOR Working Group who has given input.

7. References

7.1. Normative References

- [RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), DOI 10.17487/RFC2818, May 2000, <<https://www.rfc-editor.org/info/rfc2818>>.
- [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>>.
- [RFC5785] 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>>.
- [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>>.

7.2. Informative References

- [I-D.ietf-cappor-architecture]
Larose, K. and D. Dolson, "CAPPOR Architecture", [draft-ietf-cappor-architecture-00](#) (work in progress), September 2017.

Authors' Addresses

Tommy Pauly (editor)
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
United States of America

Email: tpauly@apple.com

Darshak Thakore (editor)

CableLabs

858 Coal Creek Circle

Louisville, CO 80027

United States of America

Email: d.thakore@cablelabs.com