

Network Working Group
Internet-Draft
Updates: 7710 (if approved)
Intended status: Standards Track
Expires: September 12, 2019

W. Kumari
Google
E. Kline
Loon
March 11, 2019

Captive-Portal Identification in DHCP / RA
draft-ekwk-capport-rfc7710bis-02

Abstract

In many environments offering short-term or temporary Internet access (such as coffee shops), it is common to start new connections in a captive portal mode. This highly restricts what the customer can do until the customer has authenticated.

This document describes a DHCP option (and a Router Advertisement (RA) extension) to inform clients that they are behind some sort of captive-portal device, and that they will need to authenticate to get Internet access. It is not a full solution to address all of the issues that clients may have with captive portals; it is designed to be used in larger solutions. The method of authenticating to, and interacting with the captive portal is out of scope of this document.

[This document is being collaborated on in Github at: <https://github.com/wkumari/draft-ekwk-capport-rfc7710bis>. The most recent version of the document, open issues, etc should all be available here. The authors (gratefully) accept pull requests. Text in square brackets will be removed before publication.]

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 12, 2019.

Copyright Notice

Copyright (c) 2019 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
1.1. Requirements Notation	3
2. The Captive-Portal Option	3
2.1. IPv4 DHCP Option	4
2.2. IPv6 DHCP Option	4
2.3. The Captive-Portal IPv6 RA Option	5
3. The Captive-Portal Link Relation Type	5
4. Precedence of API URIs	6
5. IANA Considerations	6
5.1. IETF params Registration	6
5.1.1. Registry name: Captive Portal Unrestricted Identifier	6
5.1.2. Registry name: Captive Portal API Link Relation Type	6
6. Security Considerations	7
7. Acknowledgements	8
8. Normative References	8
Appendix A. Changes / Author Notes.	9
Appendix B. Differences from RFC 7710	9
Authors' Addresses	10

1. Introduction

In many environments, users need to connect to a captive-portal device and agree to an Acceptable Use Policy (AUP) and / or provide billing information before they can access the Internet. It is anticipated that the IETF will work on a more fully featured protocol at some point, to ease interaction with Captive Portals. Regardless of how that protocol operates, it is expected that this document will provide needed functionality because the client will need to know when it is behind a captive portal and how to contact it.

In order to present users with the payment or AUP pages, the captive-portal device has to intercept the user's connections and redirect the user to the captive portal, using methods that are very similar to man-in-the-middle (MITM) attacks. As increasing focus is placed on security, and end nodes adopt a more secure stance, these interception techniques will become less effective and/or more intrusive.

This document describes a DHCP ([RFC2131]) option (Captive-Portal) and an IPv6 Router Advertisement (RA) ([RFC4861]) extension that informs clients that they are behind a captive-portal device and how to contact it.

1.1. Requirements Notation

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].

2. The Captive-Portal Option

The Captive Portal DHCP / RA Option informs the client that it is behind a captive portal and provides the URI to access an authentication page. This is primarily intended to improve the user experience by getting them to the captive portal faster; for the foreseeable future, captive portals will still need to implement the interception techniques to serve legacy clients, and clients will need to perform probing to detect captive portals.

In order to support multiple "classes" of clients (e.g. IPv4 only, IPv6 only with DHCPv6 ([RFC3315]), IPv6 only with RA) the captive portal can provide the URI via multiple methods (IPv4 DHCP, IPv6 DHCP, IPv6 RA). The captive portal operator should ensure that the URIs handed out are equivalent to reduce the chance of operational problems. The maximum length of the URI that can be carried in IPv4 DHCP is 255 bytes, so URIs longer than 255 bytes should not be used in IPv6 DHCP or IPv6 RA.

In all variants of this option, the URI SHOULD be that of the captive portal API endpoint, conforming to the recommendations for such URIs [cite:API] (i.e. the URI SHOULD contain a DNS name and SHOULD reference a secure transport, e.g. https). A captive portal MAY do content negotiation ([RFC7231] section 3.4) and attempt to redirect clients querying without an explicit indication of support for the captive portal API content type (i.e. without application/capport+json listed explicitly anywhere within an Accept header vis. [RFC7231] section 5.3). In so doing, the captive portal SHOULD

redirect the client to the value associated with the "user-portal-url" API key.

The URI SHOULD NOT contain an IP address literal.

The URI parameter is not null terminated.

Networks with no captive portals MAY explicitly indicate this condition by using this option with the IANA-assigned URI for this purpose (see Section 5.1.1). Clients observing the URI value "urn:ietf:params:capport-unrestricted" MAY forego time-consuming forms of captive portal detection.

2.1. IPv4 DHCP Option

The format of the IPv4 Captive-Portal DHCP option is shown below.

Code	Len	Data
code	len	URI ...

- o Code: The Captive-Portal DHCPv4 Option (160) (one octet)
- o Len: The length, in octets of the URI.
- o URI: The URI for the captive portal API endpoint to which the user should connect (encoded following the rules in [RFC3986]).

2.2. IPv6 DHCP Option

The format of the IPv6 Captive-Portal DHCP option is shown below.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
option-code										option-len																													
. URI (variable length) .																																							
...																																							

- o option-code: The Captive-Portal DHCPv6Option (103) (two octets)
- o option-len: The length, in octets of the URI.
- o URI: The URI for the captive portal API endpoint to which the user should connect (encoded following the rules in [RFC3986]).

See [RFC7227], Section 5.7 for more examples of DHCP Options with URIs.

2.3. The Captive-Portal IPv6 RA Option

This section describes the Captive-Portal Router Advertisement option.




Figure 2: Captive-Portal RA Option Format

Type 37

Length 8-bit unsigned integer. The length of the option (including the Type and Length fields) in units of 8 bytes.

URI The URI for the captive portal API endpoint to which the user should connect. This MUST be padded with NULL (0x00) to make the total option length (including the Type and Length fields) a multiple of 8 bytes.

3. The Captive-Portal Link Relation Type

Some captive portal network deployments may be unable to change, or unwilling to risk changing, the network infrastructure necessary to use any of the above options. In such deployments, when clear text HTTP intercept and redirection are used, a Link relation header ([RFC8288], Section 3.3) MAY be inserted to convey to a HTTP client (user agent) the associated Captive Portal API URI.

HTTP user agents MUST ignore this link relation in any context other than when explicitly probing to detect the presence of a captive portal. Failure to do so could allow an attacker to inject a Captive Portal API URI other than the correct URI for a given network or for networks where there is no captive portal present at all.

4. Precedence of API URIs

A device may learn about Captive Portal API URIs through more than one of (or indeed all of) the above options. It is a network configuration error if the learned URIs are not all identical.

However, if the URIs learned are not in fact all identical the captive device MUST prioritize URIs learned from network provisioning or configuration mechanisms before all other URIs. Specifically, URIs learned via any of the options in Section 2 should take precedence over any URI learned via a mechanism like the one described in Section 3.

If the URIs learned via more than one option described in Section 2 are not all identical, this condition should be logged for the device owner or administrator. URI precedence in this situation is not specified by this document.

5. IANA Considerations

This document requests two new IETF URN protocol parameter ([RFC3553]) entries.

Thanks IANA!

5.1. IETF params Registration

5.1.1. Registry name: Captive Portal Unrestricted Identifier

Registry name: Captive Portal Unrestricted Identifier

URN: urn:ietf:params:capport-unrestricted

Specification: RFC TBD (this document)

Repository: RFC TBD (this document)

Index value: Only one value is defined (see URN above). No hierarchy is defined and therefore no sub-namespace registrations are possible.

5.1.2. Registry name: Captive Portal API Link Relation Type

Registry name: Captive Portal API Link Relation Type

URN: urn:ietf:params:capport-api

Specification: RFC TBD (this document)

Repository: RFC TBD (this document)

Index value: Only one value is defined (see URN above). No hierarchy is defined and therefore no sub-namespace registrations are possible.

6. Security Considerations

An attacker with the ability to inject DHCP messages, RAs, or HTTP headers into cleartext HTTP communications could include an option or link relation from this document and so force users to contact an address of his choosing. As an attacker with this capability could simply list himself as the default gateway (and so intercept all the victim's traffic); this does not provide them with significantly more capabilities, but because this document removes the need for interception, the attacker may have an easier time performing the attack. As the operating systems and application that make use of this information know that they are connecting to a captive-portal device (as opposed to intercepted connections) they can render the page in a sandboxed environment and take other precautions, such as clearly labeling the page as untrusted. The means of sandboxing and user interface presenting this information is not covered in this document - by its nature it is implementation specific and best left to the application and user interface designers.

Devices and systems that automatically connect to an open network could potentially be tracked using the techniques described in this document (forcing the user to continually authenticate, or exposing their browser fingerprint). However, similar tracking can already be performed with the standard captive portal mechanisms, so this technique does not give the attackers more capabilities.

Captive portals are increasingly hijacking TLS connections to force browsers to talk to the portal. Providing the portal's URI via a DHCP or RA option is a cleaner technique, and reduces user expectations of being hijacked - this may improve security by making users more reluctant to accept TLS hijacking, which can be performed from beyond the network associated with the captive portal.

By simplifying the interaction with the captive portal systems, and doing away with the need for interception, we think that users will be less likely to disable useful security safeguards like DNSSEC validation, VPNs, etc. In addition, because the system knows that it is behind a captive portal, it can know not to send cookies, credentials, etc. By handing out a URI using which is protected with TLS, the captive portal operator can attempt to reassure the user that the captive portal is not malicious.

7. Acknowledgements

This document is a -bis of RFC7710. Thanks to all of the original authors (Warren Kumari, Olafur Gudmundsson, Paul Ebersman, Steve Sheng), and original contributors.

Also thanks to the CAPPOT WG for all of the discussion and improvements.

8. 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>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC2939] Droms, R., "Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types", BCP 43, RFC 2939, DOI 10.17487/RFC2939, September 2000, <<https://www.rfc-editor.org/info/rfc2939>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC3553] Mealling, M., Masinter, L., Hardie, T., and G. Klyne, "An IETF URN Sub-namespace for Registered Protocol Parameters", BCP 73, RFC 3553, DOI 10.17487/RFC3553, June 2003, <<https://www.rfc-editor.org/info/rfc3553>>.
- [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>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.

- [RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", BCP 187, RFC 7227, DOI 10.17487/RFC7227, May 2014, <<https://www.rfc-editor.org/info/rfc7227>>.
- [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>>.
- [RFC7710] Kumari, W., Gudmundsson, O., Ebersman, P., and S. Sheng, "Captive-Portal Identification Using DHCP or Router Advertisements (RAs)", RFC 7710, DOI 10.17487/RFC7710, December 2015, <<https://www.rfc-editor.org/info/rfc7710>>.
- [RFC8288] Nottingham, M., "Web Linking", RFC 8288, DOI 10.17487/RFC8288, October 2017, <<https://www.rfc-editor.org/info/rfc8288>>.

Appendix A. Changes / Author Notes.

[RFC Editor: Please remove this section before publication]

From initial to -00.

- o Import of RFC7710.

Appendix B. Differences from RFC 7710

This document incorporates the following differences from [RFC7710].

- o Clarify that IP string literals are NOT RECOMMENDED.
- o Clarify that the option URI SHOULD be that of the captive portal API endpoint.
- o Clarify that captive portals MAY do content negotiation.
- o Added text about Captive Portal API URI precedence in the event of a network configuration error.
- o Added urn:ietf:params:capport-unrestricted URN.
- o Added urn:ietf:params:capport-api URN.

Authors' Addresses

Warren Kumari
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: warren@kumari.net

Erik Kline
Loon
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: ek@google.com

Captive Portal Interaction
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

T. Pauly, Ed.
Apple Inc.
D. Thakore, Ed.
CableLabs
March 11, 2019

Captive Portal API
draft-ietf-capport-api-02

Abstract

This document describes an HTTP API that allows clients 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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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. Terminology	2
3. Workflow	3
4. API Details	3
4.1. URI of Captive Portal API endpoint	3
4.1.1. Server Authentication	4
4.2. JSON Keys	4
4.3. An Example Interaction	5
5. Security Considerations	6
5.1. Privacy Considerations	6
6. IANA Considerations	6
7. Acknowledgments	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	8

1. Introduction

This document describes a HyperText Transfer Protocol (HTTP) Application Program Interface (API) that allows clients 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 client has access to the Internet)
- o A URI that a client browser can present to a user to get out of captivity
- o An encrypted connection (TLS for both the API and portal URI)

2. Terminology

This document leverages the terminology and components described in [I-D.ietf-capport-architecture] and additionally uses the following association:

- o Captive Portal Client: The client that interacts with the captive portal API is typically some application running on the User Equipment that is connected to the Captive Network. This is also referred to as the "client" in this document.

- o Captive Portal API Server: The server exposing the API's defined in this document to the client. This is also referred to as the "API server" in this document.

3. Workflow

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

1. Provisioning, in which a client discovers that a network has a captive portal, and learns the URI of the API server
2. API Server interaction, in which a client 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 clients 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 client as part of the first step. The mechanism for discovering the API Server endpoint is not covered by this document.

4. API Details

4.1. 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 client 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 might be unique for each client host and between sessions for the same client 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"

4.1.1.1. Server Authentication

The purpose of accessing the Captive Portal API over an HTTPS connection is twofold: first, the encrypted connection protects the integrity and confidentiality of the API exchange from other parties on the local network; and second, it provides the client of the API an opportunity to authenticate the server that is hosting the API. This authentication is aimed at allowing a user to be reasonably confident that the entity providing the Captive Portal API has a valid certificate for the hostname in the URI (such as "example.com"). The hostname of the API SHOULD be displayed to the user in order to indicate the entity which is providing the API service.

Clients performing revocation checking will need some means of accessing revocation information for certificates presented by the API server. Online Certificate Status Protocol [RFC6960] (OCSP) stapling, using the TLS Certificate Status Request extension [RFC6066] SHOULD be used. OCSP stapling allows a client to perform revocation checks without initiating new connections. To allow for other forms of revocation checking, a captive network could permit connections to OCSP responders or Certificate Revocation Lists (CRLs) that are referenced by certificates provided by the API server. In addition to connections to OCSP responders and CRLs, a captive network SHOULD also permit connections to Network Time Protocol (NTP) [RFC5905] servers or other time-sync mechanisms to allow clients to accurately validate certificates.

Certificates with missing intermediate certificates that rely on clients validating the certificate chain using the URI specified in the Authority Information Access (AIA) extension [RFC5280] SHOULD NOT be used by the Captive Portal API server. If the certificates do require the use of AIA, the captive network will need to allow client access to the host specified in the URI.

If the client is unable to validate the certificate presented by the API server, it MUST NOT proceed with any of the behavior for API interaction described in this document. The client will proceed to interact with the captive network as if the API capabilities were not present. It may still be possible for the user to access the network by being redirected to a web portal.

4.2. JSON Keys

The Captive Portal API data structures are specified in JavaScript Object Notation (JSON) [RFC7159]. Requests and responses for the Captive Portal API use the "application/captive+json" media type. Clients SHOULD include this media type as an Accept header in their

GET requests, and servers MUST mark this media type as their Content-Type header in responses.

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

- o "captive" (required, boolean): indicates whether the client is in a state of captivity, i.e it has not satisfied the conditions to access the external network. If the client is captive (i.e. captive=true), it can still be allowed enough access for it to perform server authentication Section 4.1.1.
- o "user-portal-url" (required, string): provides the URL of a web portal with which a user can interact.
- o "vendor-info-url" (optional, string): provides the URL of a webpage or site on which the operator of the network has information that it wishes to share with the user (e.g. store info, maps, flight status, or entertainment).
- o "expire-date" (optional, string formatted as [RFC3339] datetime): indicates the date and time after which the client will be in a captive state. The API server SHOULD include this value if the client is not captive (i.e. captive=false) and SHOULD omit this value for captive clients.
- o "bytes-remaining" (optional, integer): indicates the number of bytes remaining, after which the client will be in placed into a captive state.

4.3. An Example Interaction

A client connected to a captive network upon discovering the URI of the API server will query the API server to retrieve information about its captive state and conditions to escape captivity. To request the Captive Portal JSON content, a client sends an HTTP GET request:

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

```
{
  "captive": true,
  "user-portal-url": "https://example.org/portal.html",
  "vendor-info-url": "https://flight.example.com/entertainment",
  "expire-date": "2014-01-01T23:28:56.782Z"
}
```

Upon receiving this information the client will provide this information to the user so that they may navigate the web portal (as specified by the user-portal-url value) to enable access to the external network. Once the user satisfies the requirements for external network access, the client SHOULD query the API server again to verify that it is no longer captive.

5. Security Considerations

TBD: Provide complete security requirements and analysis.

5.1. 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.

6. IANA Considerations

This document registers the media type for Captive Portal API JSON text, "application/captive+json".

Type name: application

Subtype name: captive+json

Required parameters: None

Optional parameters: None

Encoding considerations: Encoding considerations are identical to those specified for the "application/json" media type.

Security considerations: See Section 5

Interoperability considerations: This document specifies format of conforming messages and the interpretation thereof.

Published specification: This document

Applications that use this media type: This media type is intended to be used by servers presenting the Captive Portal API, and clients connecting to such captive networks.

Additional information: None

Person & email address to contact for further information: See Authors' Addresses section.

Intended usage: COMMON

Restrictions on usage: None

Author: CAPPORT IETF WG

Change controller: IETF

7. Acknowledgments

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

8. References

8.1. 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>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.

- [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>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI 10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/info/rfc6066>>.
- [RFC6960] Santesson, S., Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 6960, DOI 10.17487/RFC6960, June 2013, <<https://www.rfc-editor.org/info/rfc6960>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-capport-architecture]
Larose, K. and D. Dolson, "CAPPORT Architecture", draft-ietf-capport-architecture-03 (work in progress), December 2018.

Authors' Addresses

Tommy Pauly (editor)
Apple Inc.
One Apple Park Way
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

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: June 30, 2019

K. Larose
Agilicus
D. Dolson
December 27, 2018

CAPPORT Architecture
draft-ietf-capport-architecture-03

Abstract

This document aims to document consensus on the CAPPORT architecture. DHCP or Router Advertisements, an optional signaling protocol, and an HTTP API are used to provide the solution. The role of Provisioning Domains (PvDs) is described.

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 June 30, 2019.

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
1.1. Requirements Language	4
1.2. Terminology	4
2. Components	5
2.1. User Equipment	5
2.2. Provisioning Service	6
2.2.1. DHCP or Router Advertisements	6
2.2.2. Provisioning Domains	6
2.3. Captive Portal API Server	7
2.4. Captive Portal Enforcement	7
2.5. Captive Portal Signal	8
2.6. Component Diagram	9
3. User Equipment Identity	10
3.1. Identifiers	10
3.2. Recommended Properties	11
3.2.1. Uniquely Identify User Equipment	11
3.2.2. Hard to Spoof	11
3.2.3. Visible to the API	12
3.2.4. Visible to the Enforcement Device	12
3.3. Evaluating an Identifier	12
3.4. Examples of an Identifier	12
3.4.1. Physical Interface	12
3.4.2. IP Address	13
4. Solution Workflow	14
4.1. Initial Connection	14
4.2. Conditions About to Expire	14
5. Acknowledgments	15
6. IANA Considerations	15
7. Security Considerations	15
7.1. Trusting the Network	15
7.2. Authenticated APIs	16
7.3. Secure APIs	16
7.4. Risk of Nuisance Captive Portal	16
7.5. User Options	16
8. References	17
8.1. Normative References	17
8.2. Informative References	17
Authors' Addresses	17

1. Introduction

In this document, "Captive Portal" is used to describe a network to which a device may be voluntarily attached, such that network access is limited until some requirements have been fulfilled. Typically a user is required to use a web browser to fulfill requirements imposed

by the network operator, such as reading advertisements, accepting an acceptable-use policy, or providing some form of credentials.

Implementations generally require a web server, some method to allow/block traffic, and some method to alert the user. Common methods of alerting the user involve modifying HTTP or DNS traffic.

Problems with captive portal implementations have been described in [I-D.nottingham-capport-problem]. [If that document cannot be published, consider putting its best parts into an appendix of this document.]

This document standardizes an architecture for implementing captive portals that provides tools for addressing most of those problems. We are guided by these principles:

- o Solutions SHOULD NOT require the forging of responses from DNS or HTTP servers, or any other protocol. In particular, solutions SHOULD NOT require man-in-the-middle proxy of TLS traffic.
- o Solutions MUST operate at the layer of Internet Protocol (IP) or above, not being specific to any particular access technology such as Cable, WiFi or 3GPP.
- o Solutions MAY allow a device to be alerted that it is in a captive network when attempting to use any application on the network.
- o Solutions SHOULD allow a device to learn that it is in a captive network before any application attempts to use the network.
- o The state of captivity SHOULD be explicitly available to devices (in contrast to modification of DNS or HTTP, which is only indirectly machine-detectable by the client--by comparing responses to well-known queries with expected responses).
- o The architecture MUST provide a path of incremental migration, acknowledging a huge variety of portals and end-user device implementations and software versions.

A side-benefit of the architecture described in this document is that devices without user interfaces are able to identify parameters of captivity. However, this document does not yet describe a mechanism for such devices to escape captivity.

The architecture uses the following mechanisms:

- o Network provisioning protocols provide end-user devices with a URI for the API that end-user devices query for information about what

is required to escape captivity. DHCP, DHCPv6, and Router-Advertisement options for this purpose are available in [RFC7710]. Other protocols (such as RADIUS), Provisioning Domains [I-D.ietf-intarea-provisioning-domains], or static configuration may also be used. A device MAY query this API at any time to determine whether the network is holding the device in a captive state.

- o End-user devices can be notified of captivity with Captive Portal Signals in response to traffic. This notification should work with any Internet protocol, not just clear-text HTTP. This notification does not carry the portal URI; rather it provides a notification to the User Equipment that it is in a captive state. This document will specify requirements for a signaling protocol which could generate Captive Portal Signals.
- o Receipt of a Captive Portal Signal informs an end-user device that it could be captive. In response, the device MAY query the provisioned API to obtain information about the network state. The device MAY take immediate action to satisfy the portal (according to its configuration/policy).

The architecture attempts to provide privacy, authentication, and safety mechanisms to the extent possible.

1.1. Requirements Language

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].

1.2. Terminology

Captive Network: A network which limits communication of attached devices to restricted hosts until the user has satisfied Captive Portal Conditions, after which access is permitted to a wider set of hosts (typically the internet).

Captive Portal Conditions: site-specific requirements that a user or device must satisfy in order to gain access to the wider network.

Captive Portal Enforcement: The network equipment which enforces the traffic restriction.

Captive Portal User Equipment: Also known as User Equipment. A device which has voluntarily joined a network for purposes of communicating beyond the constraints of the captive network.

Captive Portal Server: The web server providing a user interface for assisting the user in satisfying the conditions to escape captivity.

Captive Portal Signal: A notification from the network used to inform the User Equipment that the state of its captivity could have changed.

Captive Portal Signaling Protocol: Also known as Signaling Protocol. The protocol for communicating Captive Portal Signals.

2. Components

2.1. User Equipment

The User Equipment is the device that a user desires to be attached to a network with full access to all hosts on the network (e.g., to have Internet access). The User Equipment communication is typically restricted by the Captive Portal Enforcement, described in Section 2.4, until site-specific requirements have been met.

At this time we consider only devices with web browsers, with web applications being the means of satisfying Captive Portal Conditions.

- o An example interactive User Equipment is a smart phone.
- o SHOULD support provisioning of the URI for the Captive Portal API (e.g., by DHCP)
- o SHOULD distinguish Captive Portal API access per network interface, in the manner of Provisioning Domain Architecture [RFC7556].
- o SHOULD have a mechanism for notifying the user of the Captive Portal
- o SHOULD have a web browser so that the user may navigate the Captive Portal user interface.
- o MAY restrict application access to networks not granting full network access. E.g., a device connected to a mobile network may be connecting to a WiFi network; the operating system MAY avoid updating the default route until network access restrictions have been lifted (excepting access to the Captive Portal server). This has been termed "make before break".

None of the above requirements are mandatory because (a) we do not wish to say users or devices must seek access beyond the captive network, (b) the requirements may be fulfilled by manually visiting

the captive portal web application, and (c) legacy devices must continue to be supported.

2.2. Provisioning Service

Here we discuss candidate mechanisms for provisioning the User Equipment with the URI of the API to query captive portal state and navigate the portal.

2.2.1. DHCP or Router Advertisements

A standard for providing a portal URI using DHCP or Router Advertisements is described in [RFC7710]. The CAPPORT architecture expects this URI to indicate the API described in Section 2.3.

Although it is not clear from RFC7710 what protocol should be executed at the specified URI, some readers might have assumed it to be an HTML page, and hence there might be User Equipment assuming a browser should open this URI. For backwards compatibility, it is RECOMMENDED that the server check the "Accept" field when serving the URI, and serve HTML pages for "text/html" and serve the API for "application/json". [REVISIT: are these details appropriate?]

2.2.2. Provisioning Domains

Although still a work in progress, [I-D.ietf-intarea-provisioning-domains] proposes a mechanism for User Equipment to be provided with PvD Bootstrap Information containing the URI for a JSON file containing key-value pairs to be downloaded over HTTPS. This JSON file would fill the role of the Captive Portal API described in Section 2.3.

The PvD security model provides secure binding between the information provided by the trusted Router Advertisement and the HTTPS server.

One key-value pair can be used to indicate the network has restricted access, requiring captive portal navigation by a user. E.g., key="captivePortal" and value=<URI of portal>. The key-value pair should provide a different result when access is not restricted. E.g., key="captivePortal" and value="".

This JSON file is extensible, allowing new key-value pairs to indicate such things as network access expiry time, URI for API access by IOT devices, etc.

The PvD server MUST support multiple (repeated) queries from each User Equipment, always returning the current captive portal

information. The User Equipment is expected to make this query upon receiving (and validating) a Captive Portal Signal (see Section 2.5).

2.3. Captive Portal API Server

The purpose of a Captive Portal API is to permit a query of Captive Portal state without interrupting the user. This API thereby removes the need for a device to perform clear-text "canary" HTTP queries to check for response tampering.

The URI of this API will have been provisioned to the User Equipment. (Refer to Section 2.2).

This architecture expects the User Equipment to query the API when the User Equipment attaches to the network and multiple times thereafter. Therefore the API MUST support multiple repeated queries from the same User Equipment, returning the current state of captivity for the equipment.

At minimum the API MUST provide: (1) the state of captivity and (2) a URI for a browser to present the portal application to the user. The API SHOULD provide evidence to the caller that it supports the present architecture.

When user equipment receives Captive Portal Signals, the user equipment MAY query the API to check the state. The User Equipment SHOULD rate-limit these API queries in the event of the signal being flooded. (See Section 7.)

The API MUST be extensible to support future use-cases by allowing extensible information elements.

The API MUST use TLS for privacy and server authentication. The implementation of the API MUST ensure both privacy and integrity of any information provided by or required by it.

This document does not specify the details of the API.

2.4. Captive Portal Enforcement

The Captive Portal Enforcement component restricts network access to User Equipment according to site-specific policy. Typically User Equipment is permitted access to a small number of services and is denied general network access until it has performed some action.

The Captive Portal Enforcement component:

- o Allows traffic through for allowed User Equipment.

- o Blocks (discards) traffic for disallowed User Equipment.
- o May signal User Equipment using the Captive Portal Signaling protocol if traffic is blocked.
- o Permits disallowed User Equipment to access necessary APIs and web pages to fulfill requirements of exiting captivity.
- o Updates allow/block rules per User Equipment in response to operations from the Captive Portal back-end.

2.5. Captive Portal Signal

User Equipment may send traffic outside the captive network prior to the Enforcement device granting it access. The Enforcement Device rightly blocks or resets these requests. However, lacking a signal from the Enforcement Device or interaction with the API server, the User Equipment can only guess at whether it is captive. Consequently, allowing the Enforcement Device to signal to the User Equipment that there is a problem with its connectivity may improve the user's experience.

An Enforcement Device may also want to inform the User Equipment of a pending expiry of its access to the external network, so providing the Enforcement Device the ability to preemptively signal may be desirable.

A specific Captive Portal Signaling Protocol is out of scope for this document. However, in order to ensure that future protocols fit into the architecture, requirements for a Captive Portal Signaling Protocol follow:

1. The notification SHOULD NOT be easy to spoof. If an attacker can send spoofed notifications to the User Equipment, they can cause the User Equipment to unnecessarily access the API. Rather than relying solely on rate limits to prevent problems, a good protocol will strive to limit the feasibility of such attacks.
2. It SHOULD be possible to send the notification before the captive portal closes. This will help ensure seamless connectivity for the user, as the User Equipment will not need to wait for a network failure to refresh its login. On receipt of preemptive notification, the User Equipment can prompt the user to refresh.
3. The signal SHOULD NOT include any information other than an indication that traffic is restricted, which can be used as a prompt to contact the API.

The Captive Portal Signaling Protocol does not provide any means of indicating that the network prevents access to some destinations. The intent is to rely on the Captive Portal API and the web portal to which it points to communicate local network policies.

The Captive Portal Enforcement function MAY send Captive Portal Signals when disallowed User Equipment attempts to send to the network. These signals MUST be rate-limited to a configurable rate.

The signals MUST NOT be sent to the Internet devices. The indications are only sent to the User Equipment.

2.6. Component Diagram

The following diagram shows the communication between each component.

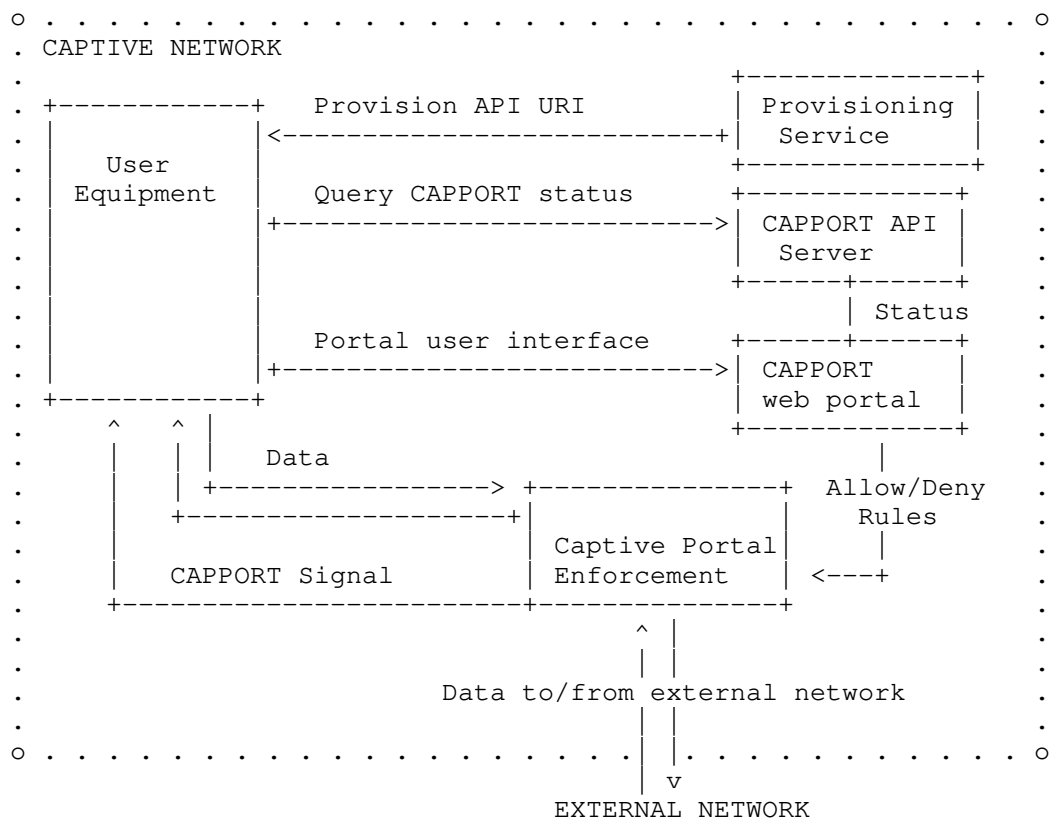


Figure 1: Captive Portal Architecture Component Diagram

In the diagram:

- o During provisioning (e.g., DHCP), the User Equipment acquires the URI for the CAPPORT API.
- o The User Equipment queries the API to learn of its state of captivity. If captive, the User Equipment presents the portal user interface to the user.
- o The User Equipment attempts to communicate to the external network through the Captive Portal enforcement device.
- o The Captive Portal Enforcement device either allows the User Equipment's packets to the external network, or if a signal has been implemented, responds with a Captive Portal Signal.
- o The CAPPORT web portal server directs the Captive Portal Enforcement device to either allow or deny external network access for the User Equipment.

Although the provisioning, API, and web portal functions are shown as discrete blocks, they could of course be combined into a single element.

3. User Equipment Identity

Multiple components in the architecture interact with both the User Equipment and each other. Since the User Equipment is the focus of these interactions, the components must be able to both identify the user equipment from their interactions with it, and be able to agree on the identity of the user equipment when interacting with each other.

The methods by which the components interact restrict the type of information that may be used as an identifying characteristics. This section discusses the identifying characteristics.

3.1. Identifiers

An Identifier is a characteristic of the User Equipment used by the components of a Captive Portal to uniquely determine which specific User Equipment is interacting with them. An Identifier MAY be a field contained in packets sent by the User Equipment to the External Network. Or, an Identifier MAY be an ephemeral property not contained in packets destined for the External Network, but instead correlated with such information through knowledge available to the different components.

3.2. Recommended Properties

The set of possible identifiers is quite large. However, in order to be considered a good identifier, an identifier SHOULD meet the following criteria. Note that the optimal identifier will likely change depending on the position of the components in the network as well as the information available to them. An identifier SHOULD:

- o Uniquely Identify the User Equipment
- o Be Hard to Spoof
- o Be Visible to the API
- o Be Visible to the Enforcement Device

An identifier might only apply to the current point of network attachment. If the device moves to a different network location its identity could change.

3.2.1. Uniquely Identify User Equipment

In order to uniquely identify the User Equipment, at most one user equipment interacting with the other components of the Captive Portal MUST have a given value of the identifier.

Over time, the user equipment identified by the value MAY change. Allowing the identified device to change over time ensures that the space of possible identifying values need not be overly large.

Independent Captive Portals MAY use the same identifying value to identify different User Equipment. Allowing independent captive portals to reuse identifying values allows the identifier to be a property of the local network, expanding the space of possible identifiers.

3.2.2. Hard to Spoof

A good identifier does not lend itself to being easily spoofed. At no time should it be simple or straightforward for one User Equipment to pretend to be another User Equipment, regardless of whether both are active at the same time. This property is particularly important when the user equipment is extended externally to devices such as billing systems, or where the identity of the User Equipment could imply liability.

3.2.3. Visible to the API

Since the API will need to perform operations which rely on the identity of the user equipment, such as query whether it is captive, the API needs to be able to relate requests to the User Equipment making the request.

3.2.4. Visible to the Enforcement Device

The Enforcement Device will decide on a per packet basis whether it should be permitted to communicate with the external network. Since this decision depends on which User Equipment sent the packet, the Enforcement Device requires that it be able to map the packet to its concept of the User Equipment.

3.3. Evaluating an Identifier

To evaluate whether an identifier is appropriate, one should consider every recommended property from the perspective of interactions among the components in the architecture. When comparing identifiers, choose the one which best satisfies all of the recommended properties. The architecture does not provide an exact measure of how well an identifier satisfies a given property; care should be taken in performing the evaluation.

3.4. Examples of an Identifier

This section provides some examples of identifiers, along with some evaluation of whether they are good identifiers. The list of identifiers is not exhaustive. Other identifiers may be used. An important point to note is that whether the identifiers are good depends heavily on the capabilities of the components and where in the network the components exist.

3.4.1. Physical Interface

The physical interface by which the User Equipment is attached to the network can be used to identify the User Equipment. This identifier has the property of being extremely difficult to spoof: the User Equipment is unaware of the property; one User Equipment cannot manipulate its interactions to appear as though it is another.

Further, if only a single User Equipment is attached to a given physical interface, then the identifier will be unique. If multiple User Equipment is attached to the network on the same physical interface, then this property is not appropriate.

Another consideration related to uniqueness of the User Equipment is that if the attached User Equipment changes, both the API server and the Enforcement Device must invalidate their state related to the User Equipment.

The Enforcement Device needs to be aware of the physical interface, which constrains the environment: it must either be part of the device providing physical access (e.g., implemented in firmware), or packets traversing the network must be extended to include information about the source physical interface (e.g. a tunnel).

The API server faces a similar problem, implying that it should co-exist with the Enforcement Device, or that the enforcement device should extend requests to it with the identifying information.

3.4.2. IP Address

A natural identifier to consider is the IP address of the User Equipment. At any given time, no device on the network can have the same IP address without causing the network to malfunction, so it is appropriate from the perspective of uniqueness.

However, it may be possible to spoof the IP address, particularly for malicious reasons where proper functioning of the network is not necessary for the malicious actor. Consequently, any solution using the IP address should proactively try to prevent spoofing of the IP address. Similarly, if the mapping of IP address to User Equipment is changed, the components of the architecture must remove or update their mapping to prevent spoofing. Demonstrations of return routeability, such as that required for TCP connection establishment, might be sufficient defense against spoofing, though this might not be sufficient in networks that use broadcast media (such as some wireless networks).

Since the IP address may traverse multiple segments of the network, more flexibility is afforded to the Enforcement Device and the API server: they simply must exist on a segment of the network where the IP address is still unique. However, consider that a NAT may be deployed between the User Equipment and the Enforcement Device. In such cases, it is possible for the components to still uniquely identify the device if they are aware of the port mapping.

In some situations, the User Equipment may have multiple IP addresses, while still satisfying all of the recommended properties. This raises some challenges to the components of the network. For example, if the user equipment tries to access the network with multiple IP addresses, should the enforcement device and API server treat each IP address as a unique User Equipment, or should it tie

the multiple addresses together into one view of the subscriber? An implementation MAY do either. Attention should be paid to IPv6 and the fact that it is expected for a device to have multiple IPv6 addresses on a single link. In such cases, identification could be performed by subnet, such as the /64 to which the IP belongs.

4. Solution Workflow

This section aims to improve understanding by describing a possible workflow of solutions adhering to the architecture.

4.1. Initial Connection

This section describes a possible work-flow when User Equipment initially joins a Captive Network.

1. The User Equipment joins the Captive Network by acquiring a DHCP lease, RA, or similar, acquiring provisioning information.
2. The User Equipment learns the URI for the Captive Portal API from the provisioning information (e.g., [RFC7710]).
3. The User Equipment accesses the CAPPORT API to receive parameters of the Captive Network, including web-portal URI. (This step replaces the clear-text query to a canary URL.)
4. If necessary, the User navigates the web portal to gain access to the external network.
5. The Captive Portal API server indicates to the Captive Portal Enforcement device that the User Equipment is allowed to access the external network.
6. The User Equipment attempts a connection outside the captive network
7. If the requirements have been satisfied, the access is permitted; otherwise the "Expired" behavior occurs.
8. The User Equipment accesses the network until conditions Expire.

4.2. Conditions About to Expire

This section describes a possible work-flow when access is about to expire.

1. Precondition: the API server has provided the User Equipment with a duration over which its access is valid

2. The User Equipment is communicating with the outside network
 3. The User Equipment's UI indicates that the length of time left for its access has fallen below a threshold
 4. The User Equipment visits the API again to validate the expiry time
 5. If expiry is still imminent, the User Equipment prompts the user to access the web-portal URI again
 6. The User extends their access through the web-portal
 7. The User Equipment's access to the outside network continues uninterrupted
5. Acknowledgments

The authors thank Lorenzo Colitti for providing the majority of the content for the Captive Portal Signal requirements.

The authors thank various individuals for their feedback on the mailing list and during the IETF98 hackathon: David Bird, Erik Kline, Alexis La Goulette, Alex Roscoe, Darshak Thakore, and Vincent van Dam.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

7.1. Trusting the Network

When joining a network, some trust is placed in the network operator. This is usually considered to be a decision by a user on the basis of the reputation of an organization. However, once a user makes such a decision, protocols can support authenticating a network is operated by who claims to be operating it. The Provisioning Domain Architecture [RFC7556] provides some discussion on authenticating an operator.

Given that a user chooses to visit a Captive Portal URI, the URI location SHOULD be securely provided to the user's device. E.g., the DHCPv6 AUTH option can sign this information.

If a user decides to incorrectly trust an attacking network, they might be convinced to visit an attacking web page and unwittingly

provide credentials to an attacker. Browsers can authenticate servers but cannot detect cleverly misspelled domains, for example.

7.2. Authenticated APIs

The solution described here assumes that when the User Equipment needs to trust the API server, server authentication will be performed using TLS mechanisms.

7.3. Secure APIs

The solution described here requires that the API be secured using TLS. This is required to allow the user equipment and API server to exchange secrets which can be used to validate future interactions. The API must ensure the integrity of this information, as well as its confidentiality.

7.4. Risk of Nuisance Captive Portal

If a Signaling Protocol is implemented, it may be possible for any user on the Internet to send signals in attempt to cause the receiving equipment to communicate with the Captive Portal API. This has been considered, and implementations may address it in the following ways:

- o The signal only informs the User Equipment to query the API. It does not carry any information which may mislead or misdirect the User Equipment.
- o Even when responding to the signal, the User Equipment securely authenticates with API servers.
- o Accesses to the API server are rate limited, limiting the impact of a repeated attack.

7.5. User Options

The Signal could inform the User Equipment that it is being held captive. There is no requirement that the User Equipment do something about this. Devices MAY permit users to disable automatic reaction to captive-portal indications for privacy reasons. However, there is the trade-off that the user doesn't get notified when network access is restricted. Hence, end-user devices MAY allow users to manually control captive portal interactions, possibly on the granularity of Provisioning Domains.

8. References

8.1. 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>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.
- [RFC7710] Kumari, W., Gudmundsson, O., Ebersman, P., and S. Sheng, "Captive-Portal Identification Using DHCP or Router Advertisements (RAs)", RFC 7710, DOI 10.17487/RFC7710, December 2015, <<https://www.rfc-editor.org/info/rfc7710>>.

8.2. Informative References

- [I-D.ietf-intarea-provisioning-domains] Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", draft-ietf-intarea-provisioning-domains-03 (work in progress), October 2018.
- [I-D.nottingham-capport-problem] Nottingham, M., "Captive Portals Problem Statement", draft-nottingham-capport-problem-01 (work in progress), April 2016.

Authors' Addresses

Kyle Larose
Agilicus

Email: kyle@agilicus.com

David Dolson

Email: ddolson@acm.org