

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 8, 2016

M. Thomson
Mozilla
P. Beverloo
Google
January 5, 2016

Voluntary Application Server Identification for Web Push
draft-thomson-webpush-vapid-01

Abstract

An application server can voluntarily identify itself to a push service using the described technique. This identification information can be used by the push service to attribute requests that are made by the same application server to a single entity, and reduce the need for endpoint secrecy by being able to associate subscriptions with an application server. An application server is further able include additional information the operator of a push service can use to contact the operator of the application server.

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 July 8, 2016.

Copyright Notice

Copyright (c) 2016 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

Internet-Draft

Self Identification

January 2016

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.	Notational Conventions	3
2.	Self-Identification Alternatives	4
2.1.	Certificates	4
2.2.	Server-Vended Tokens	4
2.3.	Client-Vended Tokens	5
2.4.	Contact Information Header Field	6
2.5.	Request Signing	6
2.6.	Token Binding	6
3.	Subscription Association	7
3.1.	Amendments to Subscribing for Push Messages	7
3.2.	Amendments to Requesting Push Message Delivery	7
4.	IANA Considerations	7
5.	Security Considerations	7
6.	Acknowledgements	8
7.	References	8
7.1.	Normative References	8
7.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

The Web Push protocol [[I-D.ietf-webpush-protocol](#)] describes how an application server is able to request that a push service deliver a push message to a user agent.

As a consequence of the expected deployment architecture, there is no basis for an application server to be known to a push service prior to requesting delivery of a push message. By the same measure, requesting the creation of a subscription for push message receipts has no prior authentication. Requiring that the push service be able to authenticate application servers places an unwanted constraint on the interactions between user agents and application servers, who are the ultimate users of a push service. That constraint would also degrade the privacy-preserving properties the protocol provides. For these reasons, [[I-D.ietf-webpush-protocol](#)] does not define a

mandatory system for authentication of application servers.

An unfortunate consequence of this design is that a push service is exposed to a greater risk of denial of service attack. While requests from application servers can be indirectly attributed to

user agents, this is not always efficient or even sufficient. Providing more information about the application server more directly to a push service allows the push service to better distinguish between legitimate and bogus requests.

Additionally, this design also relies on endpoint secrecy as any application server in possession of the endpoint is able to send messages, albeit without payloads. In situations where usage of a subscription can be limited to a single application server, the ability to associate said subscription with the application server provides could reduce the impact of a data leak.

This document describes a system whereby an application server can volunteer information about itself to a push service. At a minimum, this provides a stable identity for the application server, though this could also include contact information, such as an email address.

A consistent identity can be used by a push service to establish behavioral expectations for an application server. Significant deviations from an established norm can then be used to trigger exception handling procedures.

Voluntarily-provided contact information can be used to contact an application server operator in the case of exceptional situations.

Experience with push service deployment has shown that software errors or unusual circumstances can cause large increases in push message volume. Contacting the operator of the application server has proven to be valuable.

Even in the absence of usable contact information, an application server that has a well-established reputation might be given preference over an unidentified application server when choosing whether to discard a push message.

1.1. Notational Conventions

The words "MUST", "MUST NOT", "SHOULD", and "MAY" are used in this document. It's not shouting, when they are capitalized, they have the special meaning described in [[RFC2119](#)].

The terms "push message", "push service", "application server", and "user agent" are used as defined in [[I-D.ietf-webpush-protocol](#)]

2. Self-Identification Alternatives

Several options have been proposed, here are some of the more concrete options.

2.1. Certificates

A push service that supports application server self-identification requests a client certificate from application servers. The client certificate is requested during the TLS [[RFC5246](#)] handshake.

An application server that does not have a client certificate offers no certificate in response; an application server that wishes to self-identify includes a certificate.

The certificate that the application server offers SHOULD be self-signed (see [Section 3.2 of \[RFC5280\]](#)). The certificate MAY contain an alternative name extension ([Section 4.2.1.6 of \[RFC5280\]](#)) that includes contact information. Of the available options, an email address using the "rfc822Name" form or an HTTP [[RFC7230](#)] (or HTTPS [[RFC2818](#)]) "uniformResourceIdentifier" SHOULD be included, though including other options are not prohibited.

The offered end-entity certificate or the public key it contains becomes an identifier for the application server. Push services are able to reduce the data they retain for an application server, either by extracting important information like the subject public key information (SPKI), by hashing, or a combination. Of course, a push service might choose to ignore the provided information.

To avoid requesting certificates from user agents, it might be necessary to serve requests from user agents and requests from application servers on different hostnames or port numbers.

[2.2.](#) Server-Vended Tokens

In this model, the push service vends a token to each application server that requests it. That token is kept secret and used as the basis for authentication.

This doesn't address the need for contact information, but it addresses the need for a consistent application server identifier.

A Cookie [[RFC6265](#)] is a token of this nature. All the considerations regarding the use (and misuse) of HTTP cookies apply should this option be chosen. A mechanism that makes token theft more difficult, such as [[I-D.ietf-tokbind-https](#)] might be needed. However that suggests a separate option (see [Section 2.6](#)).

This information would be repeated with each request, but that overhead is greatly reduced by header compression [[RFC7541](#)] in HTTP/2 [[RFC7540](#)].

[2.3.](#) Client-Vended Tokens

In this model, clients generates a token that it uses to prove ownership over a private key. Use of the same key over time establishes a continuous identity.

Push message requests can be accompanied by a JSON Web Token (JWT) [[RFC7519](#)]. An "aud" (Audience) claim in the token MUST include the push resource URL. This binds the token to a specific push subscription, but not a specific push message. As a result, the token is reusable, limited by the value of the "exp" (Expiry) claim in the token. An "exp" claim MUST be included.

The JWT is included in an Authorization header field, using an auth-scheme of "WebPush".

Editor's Note: The definition of the "Bearer" auth-scheme in [[RFC6750](#)] is almost perfect, but context would seem to indicate

that this is only valid for use with OAuth.

The corresponding public key is included in a JSON Web Key (JWK) [[RFC7517](#)]. This would be included in either a newly-defined "jwk" parameter of the Crypto-Key header field [[I-D.ietf-httpbis-encryption-encoding](#)]; alternatively, the "p256ecdsa" parameter defined in [[I-D.thomson-http-content-signature](#)] could be used to transport a raw key.

For voluntarily-provided contact details, a separate header field could be used (as in [Section 2.4](#)) or the JWT could include claims about identity. For the latter, the "sub" (subject) claim could include a contact URI for the application server.

The JWT MUST use a JSON Web Signature (JWS) [[RFC7515](#)]. Both the JWS and JWK MUST use an elliptic curve digital signature (ECDSA) key on the NIST P-256 curve [[FIPS186](#)].

A JWT also offers the option of including contact information as an additional claim. An "iss" (Issuer) claim can include a contact URI: either a "mailto:" (email) [[RFC6068](#)] or an "https:" [[RFC2818](#)] SHOULD be used.

[2.4.](#) Contact Information Header Field

Contact information for an application server could be included in a header field, either directly (e.g., a From header field, [Section 5.5.1 of \[RFC7231\]](#)), or by reference (e.g., a new "contact" link relation [[RFC5988](#)] that identified a vCard [[RFC6350](#)]). Note that a From header field is limited to email addresses.

Like an server- or client-vended token ([Section 2.2](#), [Section 2.3](#)), contact information would need to be repeated, though that cost is reduced with HTTP/2.

[2.5.](#) Request Signing

Signing of push message requests would allow the push service to

attribute requests to an application server based on an asymmetric key. This could be done in any number of ways JWS [[RFC7515](#)] and HTTP signatures [[I-D.cavage-http-signatures](#)] being the most likely options. Note that the latter does not provide a means of conveying the signing key, which would be necessary for this application.

Request signing shares much in common with client-vended tokens [Section 2.3](#), but it removes any possibility of token reuse in the interests of security.

Request signing has a few shortcomings:

- o Deciding what to sign is challenging. Signing only the body of a message is not sufficient to prevent message replay attacks.
- o Every message contains a signature, which can increase the load on a server significantly. This is especially bad if a signature validation result is input to denial of service mitigation decision making.

[2.6.](#) Token Binding

The mechanism proposed in [[I-D.ietf-tokbind-https](#)] can be used to provide a stable identifier for application servers. This includes a signature over material that is exported from the underlying TLS connection. Importantly, this does not require a new signature for each request: the same signature is repeated for every request, HTTP/2 is again used to reduce the cost of the repeated information.

Token binding could be used independently of cookies. Consequently, an application server would not be required to accept and store cookies, though the push service would not be able to offload any state as a result.

[3.](#) Subscription Association

A stable identifier for an application server - such as a public key or a token - could also be used to associate a subscription with the application server.

Subscription association reduces the reliance on endpoint secrecy by requiring proof of possession to be demonstrated by an application

server when requesting delivery of a push message. This provides an additional level of protection against leaking of the details of the push subscription.

[3.1.](#) Amendments to Subscribing for Push Messages

The user agent includes the public key or token of the application server when requesting the creation of a subscription. For example, the Web Push API [[API](#)] could allow an application to provide a public key as part of a new field on the "PushSubscriptionOptions" dictionary.

This token might then be added to the request to create a push subscription.

Allowing the inclusion of multiple keys when creating a subscription would allow a subscription to be associated with multiple application servers or application server instances. This would require more state to be maintained by the push service for each subscription.

[3.2.](#) Amendments to Requesting Push Message Delivery

When a subscription has been associated with an application server, the request for push message delivery MUST include proof of possession for the associated private key or token that was used when creating the subscription. Requests that do not contain proof of possession are rejected with a 401 (Unauthorized) status code.

[4.](#) IANA Considerations

This document has no IANA actions (yet).

[5.](#) Security Considerations

TLS 1.2 [[RFC5246](#)] does not provide any confidentiality protections for client certificates. A network attacker can therefore see the identification information that is provided by the application server. A push service MAY choose to offer confidentiality protection for application server identity by initiating TLS renegotiation immediately after establishing the TLS connection at

the cost of some additional latency. Using TLS 1.3

[[I-D.ietf-tls-tls13](#)] provides confidentiality protection for this information without additional latency.

An application server might offer falsified contact information. A push service operator therefore cannot use the presence of unvalidated contact information as input to any security-critical decision-making process.

Many of the alternative solutions are vulnerable to replay attacks. Applying narrow limits to the period over which a replayable token can be reused limits the potential value of a stolen token to an attacker and can increase the difficulty of stealing a token. The token binding solution, which binds tokens to a single TLS connection can make tokens less reusable.

[6.](#) Acknowledgements

This document would have been much worse than it currently is if not for the contributions of Benjamin Bangert, Chris Karlof, Costin Manolache, and others.

[7.](#) References

[7.1.](#) Normative References

[FIPS186] National Institute of Standards and Technology (NIST), "Digital Signature Standard (DSS)", NIST PUB 186-4 , July 2013.

[[I-D.ietf-webpush-protocol](#)]
Thomson, M., Damaggio, E., and B. Raymor, "Generic Event Delivery Using HTTP Push", [draft-ietf-webpush-protocol-02](#) (work in progress), November 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

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

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.

- [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, <<http://www.rfc-editor.org/info/rfc5280>>.
- [RFC6068] Duerst, M., Masinter, L., and J. Zawinski, "The 'mailto' URI Scheme", [RFC 6068](#), DOI 10.17487/RFC6068, October 2010, <<http://www.rfc-editor.org/info/rfc6068>>.
- [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, <<http://www.rfc-editor.org/info/rfc7230>>.

[7.2.](#) Informative References

- [API] van Ouwerkerk, M. and M. Thomson, "Web Push API", 2015, <<https://w3c.github.io/push-api/>>.
- [I-D.cavage-http-signatures]
Cavage, M. and M. Sporny, "Signing HTTP Messages", [draft-cavage-http-signatures-05](#) (work in progress), October 2015.
- [I-D.ietf-httpbis-encryption-encoding]
Thomson, M., "Encrypted Content-Encoding for HTTP", [draft-ietf-httpbis-encryption-encoding-00](#) (work in progress), December 2015.
- [I-D.ietf-tls-tls13]
Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", [draft-ietf-tls-tls13-11](#) (work in progress), December 2015.
- [I-D.ietf-tokbind-https]
Popov, A., Nystrom, M., Balfanz, D., and A. Langley, "Token Binding over HTTP", [draft-ietf-tokbind-https-02](#) (work in progress), October 2015.
- [I-D.thomson-http-content-signature]
Thomson, M., "Content-Signature Header Field for HTTP", [draft-thomson-http-content-signature-00](#) (work in progress), July 2015.
- [RFC5988] Nottingham, M., "Web Linking", [RFC 5988](#),

Internet-Draft

Self Identification

January 2016

- [RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), DOI 10.17487/RFC6265, April 2011, <<http://www.rfc-editor.org/info/rfc6265>>.
- [RFC6350] Perreault, S., "vCard Format Specification", [RFC 6350](#), DOI 10.17487/RFC6350, August 2011, <<http://www.rfc-editor.org/info/rfc6350>>.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", [RFC 6750](#), DOI 10.17487/RFC6750, October 2012, <<http://www.rfc-editor.org/info/rfc6750>>.
- [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, <<http://www.rfc-editor.org/info/rfc7231>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", [RFC 7515](#), DOI 10.17487/RFC7515, May 2015, <<http://www.rfc-editor.org/info/rfc7515>>.
- [RFC7517] Jones, M., "JSON Web Key (JWK)", [RFC 7517](#), DOI 10.17487/RFC7517, May 2015, <<http://www.rfc-editor.org/info/rfc7517>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", [RFC 7519](#), DOI 10.17487/RFC7519, May 2015, <<http://www.rfc-editor.org/info/rfc7519>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", [RFC 7540](#), DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.
- [RFC7541] Peon, R. and H. Ruellan, "HPACK: Header Compression for HTTP/2", [RFC 7541](#), DOI 10.17487/RFC7541, May 2015, <<http://www.rfc-editor.org/info/rfc7541>>.

Authors' Addresses

Martin Thomson
Mozilla

Email: martin.thomson@gmail.com

Thomson & Beverloo

Expires July 8, 2016

[Page 10]

Internet-Draft

Self Identification

January 2016

Peter Beverloo
Google

Email: beverloo@google.com

