

OAuth Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: November 9, 2018

Y. Sheffer
Intuit
D. Hardt
Amazon
M. Jones
Microsoft
May 08, 2018

JSON Web Token Best Current Practices
draft-ietf-oauth-jwt-bcp-03

Abstract

JSON Web Tokens, also known as JWTs, are URL-safe JSON-based security tokens that contain a set of claims that can be signed and/or encrypted. JWTs are being widely used and deployed as a simple security token format in numerous protocols and applications, both in the area of digital identity, and in other application areas. The goal of this Best Current Practices document is to provide actionable guidance leading to secure implementation and deployment of JWTs.

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 November 9, 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 | 3 |
| 1.1. | Target Audience | 3 |
| 1.2. | Conventions used in this document | 4 |
| 2. | Threats and Vulnerabilities | 4 |
| 2.1. | Weak Signatures and Insufficient Signature Validation | 4 |
| 2.2. | Weak symmetric keys | 4 |
| 2.3. | Multiplicity of JSON encodings | 5 |
| 2.4. | Incorrect Composition of Encryption and Signature | 5 |
| 2.5. | Insecure Use of Elliptic Curve Encryption | 5 |
| 2.6. | Substitution Attacks | 5 |
| 2.7. | Cross-JWT Confusion | 5 |
| 3. | Best Practices | 6 |
| 3.1. | Perform Algorithm Verification | 6 |
| 3.2. | Use Appropriate Algorithms | 6 |
| 3.3. | Validate All Cryptographic Operations | 7 |
| 3.4. | Validate Cryptographic Inputs | 7 |
| 3.5. | Ensure Cryptographic Keys have Sufficient Entropy | 7 |
| 3.6. | Avoid Length-Dependent Encryption Inputs | 7 |
| 3.7. | Use UTF-8 | 8 |
| 3.8. | Validate Issuer and Subject | 8 |
| 3.9. | Use and Validate Audience | 8 |
| 3.10. | Do Not Trust Received Claims | 8 |
| 3.11. | Use Explicit Typing | 9 |
| 3.12. | Use Mutually Exclusive Validation Rules for Different Kinds of JWTs | 9 |
| 4. | Security Considerations | 10 |
| 5. | IANA Considerations | 10 |
| 6. | Acknowledgements | 10 |
| 7. | References | 11 |
| 7.1. | Normative References | 11 |
| 7.2. | Informative References | 11 |
| Appendix A. | Document History | 13 |
| A.1. | draft-ietf-oauth-jwt-bcp-03 | 13 |
| A.2. | draft-ietf-oauth-jwt-bcp-02 | 13 |
| A.3. | draft-ietf-oauth-jwt-bcp-01 | 13 |
| A.4. | draft-ietf-oauth-jwt-bcp-00 | 13 |
| A.5. | draft-sheffer-oauth-jwt-bcp-01 | 13 |
| A.6. | draft-sheffer-oauth-jwt-bcp-00 | 13 |
| | Authors' Addresses | 13 |

1. Introduction

JSON Web Tokens, also known as JWTs [[RFC7519](#)], are URL-safe JSON-based security tokens that contain a set of claims that can be signed and/or encrypted. The JWT specification has seen rapid adoption because it encapsulates security-relevant information in one, easy to protect location, and because it is easy to implement using widely-available tools. One application area in which JWTs are commonly used is representing digital identity information, such as OpenID Connect ID Tokens [[OpenID.Core](#)] and OAuth 2.0 [[RFC6749](#)] access tokens and refresh tokens, the details of which are deployment-specific.

Since the JWT specification was published, there have been several widely published attacks on implementations and deployments. Such attacks are the result of under-specified security mechanisms, as well as incomplete implementations and incorrect usage by applications.

The goal of this document is to facilitate secure implementation and deployment of JWTs. Many of the recommendations in this document will actually be about implementation and use of the cryptographic mechanisms underlying JWTs that are defined by JSON Web Signature (JWS) [[RFC7515](#)], JSON Web Encryption (JWE) [[RFC7516](#)], and JSON Web Algorithms (JWA) [[RFC7518](#)]. Others will be about use of the JWT claims themselves.

These are intended to be minimum recommendations for the use of JWTs in the vast majority of implementation and deployment scenarios. Other specifications that reference this document can have stricter requirements related to one or more aspects of the format, based on their particular circumstances; when that is the case, implementers are advised to adhere to those stricter requirements. Furthermore, this document provides a floor, not a ceiling, so stronger options are always allowed (e.g., depending on differing evaluations of the importance of cryptographic strength vs. computational load).

Community knowledge about the strength of various algorithms and feasible attacks can change quickly, and experience shows that a Best Current Practice (BCP) document about security is a point-in-time statement. Readers are advised to seek out any errata or updates that apply to this document.

1.1. Target Audience

The targets of this document are:

- Implementers of JWT libraries (and the JWS and JWE libraries used by them),

- Implementers of code that uses such libraries (to the extent that some mechanisms may not be provided by libraries, or until they are), and
- Developers of specifications that rely on JWTs, both inside and outside the IETF.

1.2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

2. Threats and Vulnerabilities

This section lists some known and possible problems with JWT implementations and deployments. Each problem description is followed by references to one or more mitigations to those problems.

2.1. Weak Signatures and Insufficient Signature Validation

Signed JSON Web Tokens carry an explicit indication of the signing algorithm, in the form of the "alg" header parameter, to facilitate cryptographic agility. This, in conjunction with design flaws in some libraries and applications, have led to several attacks:

- The algorithm can be changed to "none" by an attacker, and some libraries would trust this value and "validate" the JWT without checking any signature.
- An "RS256" (RSA, 2048 bit) parameter value can be changed into "HS256" (HMAC, SHA-256), and some libraries would try to validate the signature using HMAC-SHA256 and using the RSA public key as the HMAC shared secret.

For mitigations, see [Section 3.1](#) and [Section 3.2](#).

2.2. Weak symmetric keys

In addition, some applications sign tokens using a weak symmetric key and a keyed MAC algorithm such as "HS256". In most cases, these keys are human memorable passwords that are vulnerable to dictionary attacks [\[Langkemper\]](#).

For mitigations, see [Section 3.5](#).

2.3. Multiplicity of JSON encodings

Previous versions of the JSON format [[RFC8259](#)] allowed several different character encodings: UTF-8, UTF-16 and UTF-32. This is not the case anymore, with the latest standard only allowing UTF-8. However older implementations may result in the JWT being misinterpreted by its recipient.

For mitigations, see [Section 3.7](#).

2.4. Incorrect Composition of Encryption and Signature

Some libraries that decrypt a JWE-encrypted JWT to obtain a JWS-signed object do not always validate the internal signature.

For mitigations, see [Section 3.3](#).

2.5. Insecure Use of Elliptic Curve Encryption

Per [[Sanso](#)], several JOSE libraries fail to validate their inputs correctly when performing elliptic curve key agreement (the "ECDH-ES" algorithm). An attacker that is able to send JWEs of its choosing that use invalid curve points and observe the cleartext outputs resulting from decryption with the invalid curve points can use this vulnerability to recover the recipient's private key.

For mitigations, see [Section 3.4](#).

2.6. Substitution Attacks

There are attacks in which one recipient will have a JWT intended for it and attempt to use it at a different recipient that it was not intended for. If not caught, these attacks can result in the attacker gaining access to resources that it is not entitled to access.

For mitigations, see [Section 3.8](#) and [Section 3.9](#).

2.7. Cross-JWT Confusion

As JWTs are being used by more different protocols in diverse application areas, it becomes increasingly important to prevent cases of JWT tokens that have been issued for one purpose being subverted and used for another. Note that this is a specific type of substitution attack. If the JWT could be used in an application context in which it could be confused with other kinds of JWTs, then mitigations MUST be employed to prevent these substitution attacks.

For mitigations, see [Section 3.8](#), [Section 3.9](#), [Section 3.11](#), and [Section 3.12](#).

3. Best Practices

The best practices listed below should be applied by practitioners to mitigate the threats listed in the preceding section.

3.1. Perform Algorithm Verification

Libraries MUST enable the caller to specify a supported set of algorithms and MUST NOT use any other algorithms when performing cryptographic operations. The library MUST ensure that the "alg" or "enc" header specifies the same algorithm that is used for the cryptographic operation. Moreover, each key MUST be used with exactly one algorithm, and this MUST be checked when the cryptographic operation is performed.

3.2. Use Appropriate Algorithms

As [Section 5.2 of \[RFC7515\]](#) says, "it is an application decision which algorithms may be used in a given context. Even if a JWS can be successfully validated, unless the algorithm(s) used in the JWS are acceptable to the application, it SHOULD consider the JWS to be invalid."

Therefore, applications MUST only allow the use of cryptographically current algorithms that meet the security requirements of the application. This set will vary over time as new algorithms are introduced and existing algorithms are deprecated due to discovered cryptographic weaknesses. Applications must therefore be designed to enable cryptographic agility.

That said, if a JWT is cryptographically protected by a transport layer, such as TLS using cryptographically current algorithms, there may be no need to apply another layer of cryptographic protections to the JWT. In such cases, the use of the "none" algorithm can be perfectly acceptable. JWTs using "none" are often used in application contexts in which the content is optionally signed; then the URL-safe claims representation and processing can be the same in both the signed and unsigned cases.

Applications SHOULD follow these algorithm-specific recommendations:

- Avoid all RSA-PKCS1 v1.5 encryption algorithms, preferring RSA-OAEP.

- ECDSA signatures require a unique random value for every message that is signed. If even just a few bits of the random value are predictable across multiple messages then the security of the signature scheme may be compromised. In the worst case, the private key may be recoverable by an attacker. To counter these attacks, JWT libraries SHOULD implement ECDSA using the deterministic approach defined in [\[RFC6979\]](#). This approach is completely compatible with existing ECDSA verifiers and so can be implemented without new algorithm identifiers being required.

[3.3.](#) Validate All Cryptographic Operations

All cryptographic operations used in the JWT MUST be validated and the entire JWT MUST be rejected if any of them fail to validate. This is true not only of JWTs with a single set of Header Parameters but also for Nested JWTs, in which both outer and inner operations MUST be validated using the keys and algorithms supplied by the application.

[3.4.](#) Validate Cryptographic Inputs

Some cryptographic operations, such as Elliptic Curve Diffie-Hellman key agreement ("ECDH-ES") take inputs that may contain invalid values, such as points not on the specified elliptic curve or other invalid points (see e.g. [\[Valenta\]](#), Sec. 7.1). Either the JWS/JWE library itself must validate these inputs before using them or it must use underlying cryptographic libraries that do so (or both!).

ECDH-ES ephemeral public key (epk) inputs should be validated according to the recipient's chosen elliptic curve. For the NIST prime-order curves P-256, P-384 and P-521, validation MUST be performed according to [Section 5.6.2.3.4](#) "ECC Partial Public-Key Validation Routine" of NIST Special Publication 800-56A revision 3 [\[nist-sp-800-56a-r3\]](#).

[3.5.](#) Ensure Cryptographic Keys have Sufficient Entropy

The Key Entropy and Random Values advice in [Section 10.1 of \[RFC7515\]](#) and the Password Considerations in [Section 8.8 of \[RFC7518\]](#) MUST be followed. In particular, human-memorizable passwords MUST NOT be directly used as the key to a keyed-MAC algorithm such as "HS256".

[3.6.](#) Avoid Length-Dependent Encryption Inputs

Many encryption algorithms leak information about the length of the plaintext, with a varying amount of leakage depending on the algorithm and mode of operation. Sensitive information, such as passwords, SHOULD be padded before being encrypted. It is

RECOMMENDED to avoid any compression of data before encryption since such compression often reveals information about the plaintext.

3.7. Use UTF-8

[[RFC7515](#)], [[RFC7516](#)], and [[RFC7519](#)] all specify that UTF-8 be used for encoding and decoding JSON used in Header Parameters and JWT Claims Sets. This is also in line with the latest JSON specification [[RFC8259](#)]. Implementations and applications MUST do this, and not use or admit the use of other Unicode encodings for these purposes.

3.8. Validate Issuer and Subject

When a JWT contains an "iss" (issuer) claim, the application MUST validate that the cryptographic keys used for the cryptographic operations in the JWT belong to the issuer. If they do not, the application MUST reject the JWT.

The means of determining the keys owned by an issuer is application-specific. As one example, OpenID Connect [[OpenID.Core](#)] issuer values are "https" URLs that reference a JSON metadata document that contains a "jwks_uri" value that is an "https" URL from which the issuer's keys are retrieved as a JWK Set [[RFC7517](#)]. This same mechanism is used by [[I-D.ietf-oauth-discovery](#)]. Other applications may use different means of binding keys to issuers.

Similarly, when the JWT contains a "sub" (subject) claim, the application MUST validate that the subject value corresponds to a valid subject and/or issuer/subject pair at the application. This may include confirming that the issuer is trusted by the application. If the issuer, subject, or the pair are invalid, the application MUST reject the JWT.

3.9. Use and Validate Audience

If the same issuer can issue JWTs that are intended for use by more than one relying party or application, the JWT MUST contain an "aud" (audience) claim that can be used to determine whether the JWT is being used by an intended party or was substituted by an attacker at an unintended party. Furthermore, the relying party or application MUST validate the audience value and if the audience value is not present or not associated with the recipient, it MUST reject the JWT.

3.10. Do Not Trust Received Claims

The "kid" (key ID) header is used by the relying application to perform key lookup. Applications should ensure that this does not create SQL or LDAP injection vulnerabilities.

Similarly, blindly following a "jku" (JWK set URL) header, which may contain an arbitrary URL, could result in server-side request forgery (SSRF) attacks.

3.11. Use Explicit Typing

Confusion of one kind of JWT for another can be prevented by having all the kinds of JWTs that could otherwise potentially be confused include an explicit JWT type value and include checking the type value in their validation rules. Explicit JWT typing is accomplished by using the "typ" header parameter. For instance, the [\[I-D.ietf-secevent-token\]](#) specification uses the "application/secevent+jwt" media type to perform explicit typing of Security Event Tokens (SETs).

Per the definition of "typ" in [Section 4.1.9 of \[RFC7515\]](#), it is RECOMMENDED that the "application/" prefix be omitted from the "typ" value. Therefore, for example, the "typ" value used to explicitly include a type for a SET SHOULD be "secevent+jwt". When explicit typing is employed for a JWT, it is RECOMMENDED that a media type name of the format "application/example+jwt" be used, where "example" is replaced by the identifier for the specific kind of JWT.

When applying explicit typing to a Nested JWT, the "typ" header parameter containing the explicit type value MUST be present in the inner JWT of the Nested JWT (the JWT whose payload is the JWT Claims Set). The same "typ" header parameter value MAY be present in the outer JWT as well, to explicitly type the entire Nested JWT.

Note that the use of explicit typing may not achieve disambiguation from existing kinds of JWTs, as the validation rules for existing kinds of JWTs often do not use the "typ" header parameter value. Explicit typing is RECOMMENDED for new uses of JWTs.

3.12. Use Mutually Exclusive Validation Rules for Different Kinds of JWTs

Each application of JWTs defines a profile specifying the required and optional JWT claims and the validation rules associated with them. If more than one kind of JWT can be issued by the same issuer, the validation rules for those JWTs MUST be written such that they are mutually exclusive, rejecting JWTs of the wrong kind. To prevent substitution of JWTs from one context into another, a number of strategies may be employed:

- Use explicit typing for different kinds of JWTs. Then the distinct "typ" values can be used to differentiate between the different kinds of JWTs.

- Use different sets of required claims or different required claim values. Then the validation rules for one kind of JWT will reject those with different claims or values.
- Use different sets of required header parameters or different required header parameter values. Then the validation rules for one kind of JWT will reject those with different header parameters or values.
- Use different keys for different kinds of JWTs. Then the keys used to validate one kind of JWT will fail to validate other kinds of JWTs.
- Use different "aud" values for different uses of JWTs from the same issuer. Then audience validation will reject JWTs substituted into inappropriate contexts.
- Use different issuers for different kinds of JWTs. Then the distinct "iss" values can be used to segregate the different kinds of JWTs.

Given the broad diversity of JWT usage and applications, the best combination of types, required claims, values, header parameters, key usages, and issuers to differentiate among different kinds of JWTs will, in general, be application specific.

4. Security Considerations

This entire document is about security considerations when implementing and deploying JSON Web Tokens.

5. IANA Considerations

This document requires no IANA actions.

6. Acknowledgements

Thanks to Antonio Sanso for bringing the "ECDH-ES" invalid point attack to the attention of JWE and JWT implementers. Tim McLean published the RSA/HMAC confusion attack. Thanks to Nat Sakimura for advocating the use of explicit typing. Thanks to Neil Madden for his numerous comments, and to Carsten Bormann and Brian Campbell for their reviews.

7. References

7.1. 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>>.
- [RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", [RFC 6979](#), DOI 10.17487/RFC6979, August 2013, <<https://www.rfc-editor.org/info/rfc6979>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", [RFC 7515](#), DOI 10.17487/RFC7515, May 2015, <<https://www.rfc-editor.org/info/rfc7515>>.
- [RFC7516] Jones, M. and J. Hildebrand, "JSON Web Encryption (JWE)", [RFC 7516](#), DOI 10.17487/RFC7516, May 2015, <<https://www.rfc-editor.org/info/rfc7516>>.
- [RFC7518] Jones, M., "JSON Web Algorithms (JWA)", [RFC 7518](#), DOI 10.17487/RFC7518, May 2015, <<https://www.rfc-editor.org/info/rfc7518>>.
- [RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", [RFC 7519](#), DOI 10.17487/RFC7519, May 2015, <<https://www.rfc-editor.org/info/rfc7519>>.
- [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>>.

7.2. Informative References

- [I-D.ietf-oauth-discovery]
Jones, M., Sakimura, N., and J. Bradley, "OAuth 2.0 Authorization Server Metadata", [draft-ietf-oauth-discovery-10](#) (work in progress), March 2018.
- [I-D.ietf-secevent-token]
Hunt, P., Jones, M., Denniss, W., and M. Ansari, "Security Event Token (SET)", [draft-ietf-secevent-token-10](#) (work in progress), May 2018.

[Langkemper]

Langkemper, S., "Attacking JWT Authentication", September 2016, <<https://www.sjoerdlangkemper.nl/2016/09/28/attacking-jwt-authentication/>>.

[nist-sp-800-56a-r3]

Barker, E., Chen, L., Keller, S., Roginsky, A., Vassilev, A., and R. Davis, "Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography, Draft NIST Special Publication 800-56A Revision 3", August 2017, <<https://csrc.nist.gov/CSRC/media/Publications/sp/800-56a/rev-3/draft/documents/sp800-56ar3-draft.pdf>>.

[OpenID.Core]

Sakimura, N., Bradley, J., Jones, M., Medeiros, B., and C. Mortimore, "OpenID Connect Core 1.0", November 2014, <http://openid.net/specs/openid-connect-core-1_0.html>.

[RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", [RFC 6749](#), DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.

[RFC7517] Jones, M., "JSON Web Key (JWK)", [RFC 7517](#), DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.

[Sanso] Sanso, A., "Critical Vulnerability Uncovered in JSON Encryption", March 2017, <<https://blogs.adobe.com/security/2017/03/critical-vulnerability-uncovered-in-json-encryption.html>>.

[Valenta] Valenta, L., Sullivan, N., Sanso, A., and N. Heninger, "In search of CurveSwap: Measuring elliptic curve implementations in the wild", March 2018, <<https://ia.cr/2018/298>>.

Appendix A. Document History

[[to be removed by the RFC editor before publication as an RFC]]

A.1. draft-ietf-oauth-jwt-bcp-03

- Acknowledgements.

A.2. draft-ietf-oauth-jwt-bcp-02

- Implemented WGLC feedback.

A.3. draft-ietf-oauth-jwt-bcp-01

- Feedback from Brian Campbell.

A.4. draft-ietf-oauth-jwt-bcp-00

- Initial WG draft. No change from the latest individual version.

A.5. draft-sheffer-oauth-jwt-bcp-01

- Added explicit typing.

A.6. draft-sheffer-oauth-jwt-bcp-00

- Initial version.

Authors' Addresses

Yaron Sheffer
Intuit

EMail: yaronf.ietf@gmail.com

Dick Hardt
Amazon

EMail: dick@amazon.com

Michael B. Jones
Microsoft

EMail: mbj@microsoft.com

URI: <http://self-issued.info/>

