

Network Working Group
Internet Draft
Intended Status: Informational
Expires: February 27, 2017

M. Jenkins
National Security Agency
M. Peck
The MITRE Corporation
K. Burgin
August 26, 2016

AES Encryption with HMAC-SHA2 for Kerberos 5
draft-ietf-kitten-aes-cts-hmac-sha2-11

Abstract

This document specifies two encryption types and two corresponding checksum types for Kerberos 5. The new types use AES in CTS mode (CBC mode with ciphertext stealing) for confidentiality and HMAC with a SHA-2 hash for integrity.

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 February 27, 2017.

Copyright and License 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 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
2. Protocol Key Representation	3
3. Key Derivation Function	3
4. Key Generation from Pass Phrases	5
5. Kerberos Algorithm Protocol Parameters	5
6. Checksum Parameters	8
7. IANA Considerations	8
8. Security Considerations	8
8.1. Random Values in Salt Strings	9
8.2. Algorithm Rationale	9
9. Acknowledgements	10
10. References	10
10.1. Normative References	10
10.2. Informative References	10
Appendix A. Test Vectors	11
Authors' Addresses	18

1. Introduction

This document defines two encryption types and two corresponding checksum types for Kerberos 5 using AES with 128-bit or 256-bit keys.

To avoid ciphertext expansion, we use a variation of the CBC-CS3 mode defined in [SP800-38A+], also referred to as ciphertext stealing or CTS mode. The new types conform to the framework specified in [RFC3961], but do not use the simplified profile, as the simplified profile is not compliant with modern cryptographic best practices such as calculating MACs over ciphertext rather than plaintext.

The encryption and checksum types defined in this document are intended to support environments that desire to use SHA-256 or SHA-384 (defined in [FIPS180]) as the hash algorithm. Differences between the encryption and checksum types defined in this document and the pre-existing Kerberos AES encryption and checksum types specified in [RFC3962] are:

- * The pseudorandom function used by PBKDF2 is HMAC-SHA-256 or HMAC-SHA-384 (HMAC is defined in [RFC2104]).
- * A key derivation function from [SP800-108] using the SHA-256 or SHA-384 hash algorithm is used to produce keys for encryption, integrity protection, and checksum operations.
- * The HMAC is calculated over the cipherstate concatenated with the AES output, instead of being calculated over the confounder and plaintext. This allows the message receiver to verify the integrity of the message before decrypting the message.
- * The HMAC algorithm uses the SHA-256 or SHA-384 hash algorithm for integrity protection and checksum operations.

2. Protocol Key Representation

The AES key space is dense, so we can use random or pseudorandom octet strings directly as keys. The byte representation for the key is described in [FIPS197], where the first bit of the bit string is the high bit of the first byte of the byte string (octet string).

3. Key Derivation Function

We use a key derivation function from Section 5.1 of [SP800-108] which uses the HMAC algorithm as the PRF.

```
function KDF-HMAC-SHA2(key, label, [context,] k):  
    k-truncate(K1)
```

where the value of K1 is computed as below.

key: The source of entropy from which subsequent keys are derived (this is known as Ki in [SP800-108]).

label: An octet string describing the intended usage of the derived key.

context: This parameter is optional. An octet string containing the information related to the derived keying material. This specification does not dictate a specific format for the context field. The context field is only used by the pseudo-random function defined in section 5, where it is set to the pseudo-random function's octet-string input parameter. The content of the octet-string input parameter is defined by the application that uses it.

k: Length in bits of the key to be outputted, expressed in big-endian binary representation in 4 bytes (this is called L in [SP800-108]). Specifically, k=128 is represented as 0x00000080, 192 as 0x000000C0, 256 as 0x00000100, and 384 as 0x00000180.

When the encryption type is aes128-cts-hmac-sha256-128, k must be no greater than 256 bits. When the encryption type is aes256-cts-hmac-sha384-192, k must be no greater than 384 bits.

The k-truncate function is defined in [RFC3961], Section 5.1. It returns the 'k' leftmost bits of the bitstring input.

In all computations in this document, | indicates concatenation.

When the encryption type is aes128-cts-hmac-sha256-128, then K1 is computed as follows:

If the context parameter is not present:

K1 = HMAC-SHA-256(key, 0x00000001 | label | 0x00 | k)

If the context parameter is present:

K1 = HMAC-SHA-256(key, 0x00000001 | label | 0x00 | context | k)

When the encryption type is aes256-cts-hmac-sha384-192, then K1 is computed as follows:

If the context parameter is not present:

K1 = HMAC-SHA-384(key, 0x00000001 | label | 0x00 | k)

If the context parameter is present:

K1 = HMAC-SHA-384(key, 0x00000001 | label | 0x00 | context | k)

In the definitions of K1 above, '0x00000001' is the i parameter (the iteration counter) from Section 5.1 of [SP800-108].

4. Key Generation from Pass Phrases

As defined below, the string-to-key function uses PBKDF2 [RFC2898] and KDF-HMAC-SHA2 to derive the base-key from a passphrase and salt. The string-to-key parameter string is four octets indicating an unsigned number in big-endian order, consistent with [RFC3962], except that the default is decimal 32768 if the parameter is not specified.

To ensure that different long-term base-keys are used with different encyptes, we prepend the enctype name to the salt, separated by a null byte. The enctype-name is "aes128-cts-hmac-sha256-128" or "aes256-cts-hmac-sha384-192" (without the quotes).

The user's long-term base-key is derived as follows:

```
iter_count = string-to-key parameter, default is decimal 32768
saltp = enctype-name | 0x00 | salt
tkey = random-to-key(PBKDF2(passphrase, saltp,
                           iter_count, keylength))
base-key = random-to-key(KDF-HMAC-SHA2(tkey, "kerberos",
                                       keylength))
```

where "kerberos" is the octet-string 0x6B65726265726F73.

where PBKDF2 is the function of that name from RFC 2898, the pseudorandom function used by PBKDF2 is HMAC-SHA-256 when the enctype is "aes128-cts-hmac-sha256-128" and HMAC-SHA-384 when the enctype is "aes256-cts-hmac-sha384-192", the value for keylength is the AES key length (128 or 256 bits), and the algorithm KDF-HMAC-SHA2 is defined in Section 3.

5. Kerberos Algorithm Protocol Parameters

The RFC 3961 cipher state that maintains cryptographic state across different encryption operations using the same key is used as the formal initialization vector (IV) input into CBC-CS3. The plaintext is prepended with a 16-octet random value generated by the message originator, known as a confounder.

The ciphertext is a concatenation of the output of AES in CBC-CS3 mode and the HMAC of the cipher state concatenated with the AES output. The HMAC is computed using either SHA-256 or SHA-384 depending on the encryption type. The output of HMAC-SHA-256 is truncated to 128 bits and the output of HMAC-SHA-384 is truncated to

192 bits. Sample test vectors are given in Appendix A.

Decryption is performed by removing the HMAC, verifying the HMAC against the cipher state concatenated with the ciphertext, and then decrypting the ciphertext if the HMAC is correct. Finally, the first 16 octets of the decryption output (the confounder) is discarded, and the remainder is returned as the plaintext decryption output.

The following parameters apply to the encryption types aes128-cts-hmac-sha256-128 and aes256-cts-hmac-sha384-192.

protocol key format: as defined in Section 2.

specific key structure: three derived keys: { Kc, Ke, Ki }.

Kc: the checksum key, inputted into HMAC to provide the checksum mechanism defined in Section 6.

Ke: the encryption key, inputted into AES encryption and decryption as defined in "encryption function" and "decryption function" below.

Ki: the integrity key, inputted into HMAC to provide authenticated encryption as defined in "encryption function" and "decryption function" below.

required checksum mechanism: as defined in Section 6.

key-generation seed length: key size (128 or 256 bits).

string-to-key function: as defined in Section 4.

default string-to-key parameters: iteration count of decimal 32768.

random-to-key function: identity function.

key-derivation function: KDF-HMAC-SHA2 as defined in Section 3. The key usage number is expressed as four octets in big-endian order.

If the enctype is aes128-cts-hmac-sha256-128:

Kc = KDF-HMAC-SHA2(base-key, usage | 0x99, 128)

Ke = KDF-HMAC-SHA2(base-key, usage | 0xAA, 128)

Ki = KDF-HMAC-SHA2(base-key, usage | 0x55, 128)

If the enctype is aes256-cts-hmac-sha384-192:

Kc = KDF-HMAC-SHA2(base-key, usage | 0x99, 192)

Ke = KDF-HMAC-SHA2(base-key, usage | 0xAA, 256)

Ki = KDF-HMAC-SHA2(base-key, usage | 0x55, 192)

cipher state: a 128-bit CBC initialization vector derived from a previous (if any) ciphertext using the same encryption key, as specified below.

initial cipher state: all bits zero.

encryption function: as follows, where E() is AES encryption in CBC-CS3 mode, and h is the size of truncated HMAC (128 bits or 192 bits as described above).

```
N = random value of length 128 bits (the AES block size)
IV = cipher state
C = E(Ke, N | plaintext, IV)
H = HMAC(Ki, IV | C)
ciphertext = C | H[1..h]
```

Steps to compute the 128-bit cipher state:

```
L = length of C in bits
portion C into 128-bit blocks, placing any remainder
of less than 128 bits into a final block
if L == 128: cipher state = C
else if L mod 128 > 0: cipher state = last full (128-bit)
                                block of C (the
                                next-to-last block)
else if L mod 128 == 0: cipher state = next-to-last block
                                of C
(note that L will never be less than 128 because of the
presence of N in the encryption input)
```

decryption function: as follows, where D() is AES decryption in CBC-CS3 mode, and h is the size of truncated HMAC.

```
(C, H) = ciphertext (Note: H is the last h bits of the ciphertext)
IV = cipher state
if H != HMAC(Ki, IV | C)[1..h]
    stop, report error
(N, P) = D(Ke, C, IV)
Note: N is set to the first block of the decryption output,
P is set to the rest of the output.
cipher state = same as described above in encryption function
```

pseudo-random function:

```
If the enctype is aes128-cts-hmac-sha256-128:
PRF = KDF-HMAC-SHA2(input-key, "prf", octet-string, 256)
```

```
If the enctype is aes256-cts-hmac-sha384-192:
PRF = KDF-HMAC-SHA2(input-key, "prf", octet-string, 384)
```

where "prf" is the octet-string 0x707266

6. Checksum Parameters

The following parameters apply to the checksum types hmac-sha256-128-aes128 and hmac-sha384-192-aes256, which are the associated checksums for aes128-cts-hmac-sha256-128 and aes256-cts-hmac-sha384-192, respectively.

associated cryptosystem: aes128-cts-hmac-sha256-128 or aes256-cts-hmac-sha384-192 as appropriate.

get_mic: HMAC(Kc, message)[1..h].
where h is 128 bits for checksum type hmac-sha256-128-aes128
and 192 bits for checksum type hmac-sha384-192-aes256

verify_mic: get_mic and compare.

7. IANA Considerations

IANA is requested to assign:

Encryption type numbers for aes128-cts-hmac-sha256-128 and aes256-cts-hmac-sha384-192 in the Kerberos Encryption Type Numbers registry.

Etype	Encryption type	Reference
-----	-----	-----
TBD1	aes128-cts-hmac-sha256-128	[this document]
TBD2	aes256-cts-hmac-sha384-192	[this document]

Checksum type numbers for hmac-sha256-128-aes128 and hmac-sha384-192-aes256 in the Kerberos Checksum Type Numbers registry.

Sumtype	Checksum type	Size	Reference
-----	-----	-----	-----
TBD3	hmac-sha256-128-aes128	16	[this document]
TBD4	hmac-sha384-192-aes256	24	[this document]

8. Security Considerations

This specification requires implementations to generate random values. The use of inadequate pseudo-random number generators (PRNGs) can result in little or no security. The generation of quality random numbers is difficult. [RFC4086] offers random number

generation guidance.

This document specifies a mechanism for generating keys from passphrases or passwords. The use of PBKDF2, a salt, and a large iteration count adds some resistance to off-line dictionary attacks by passive eavesdroppers. Salting prevents rainbow table attacks, while large iteration counts slow password guess attempts. Nonetheless, computing power continues to rapidly improve, including the potential for use of graphics processing units (GPUs) in password guess attempts. It is important to choose strong passphrases. Use of Kerberos extensions that protect against off-line dictionary attacks should also be considered, as should the use of public key cryptography for initial Kerberos authentication [RFC4556] to eliminate the use of passwords or passphrases within the Kerberos protocol.

The NIST guidance in section 5.3 of [SP800-38A], requiring that CBC initialization vectors be unpredictable, is satisfied by the use of a random confounder as the first block of plaintext. The confounder fills the cryptographic role typically played by an initialization vector. This approach was chosen to align with other Kerberos cryptosystem approaches.

8.1. Random Values in Salt Strings

NIST guidance in Section 5.1 of [SP800-132] requires at least 128 bits of the salt to be randomly generated. The string-to-key function as defined in [RFC3961] requires the salt to be valid UTF-8 strings [RFC3629]. Not every 128-bit random string will be valid UTF-8, so a UTF-8 compatible encoding would be needed to encapsulate the random bits. However, using a salt containing a random portion may have the following issues with some implementations:

- * Cross-realm krbtgt keys are typically managed by entering the same password at two KDCs to get the same keys. If each KDC uses a random salt, they won't have the same keys.
- * Random salts may interfere with password history checking.

8.2. Algorithm Rationale

This document has been written to be consistent with common implementations of AES and SHA-2. The encryption and hash algorithm sizes have been chosen to create a consistent level of protection, with consideration to implementation efficiencies. So, for instance, SHA-384, which would normally be matched to AES-192, is instead matched to AES-256 to leverage the fact that there are efficient hardware implementations of AES-256. Note that, as indicated by the

enc-type name "aes256-cts-hmac-sha384-192", the truncation of the HMAC-SHA-384 output to 192-bits results in an overall 192-bit level of security.

9. Acknowledgements

Kelley Burgin was employed at the National Security Agency during much of the work on this document.

10. References

10.1. Normative References

- [RFC2104] Krawczyk, H. et al., "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC2898] Kaliski, B., "PKCS #5: Password-Based Cryptography Specification Version 2.0", RFC 2898, September 2000.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", RFC 3629, November 2003.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, February 2005.
- [RFC3962] Raeburn, K., "Advanced Encryption Standard (AES) Encryption for Kerberos 5", RFC 3962, February 2005.
- [FIPS180] National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-4, August 2015.
- [FIPS197] National Institute of Standards and Technology, "Advanced Encryption Standard (AES)", FIPS PUB 197, November 2001.
- [SP800-38A+] National Institute of Standards and Technology, "Recommendation for Block Cipher Modes of Operation: Three Variants of Ciphertext Stealing for CBC Mode", NIST Special Publication 800-38A Addendum, October 2010.
- [SP800-108] National Institute of Standards and Technology, "Recommendation for Key Derivation Using Pseudorandom Functions", NIST Special Publication 800-108, October 2009.

10.2. Informative References

- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker,

"Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.

[SP800-38A] National Institute of Standards and Technology, "Recommendation for Block Cipher Modes of Operation: Methods and Techniques", NIST Special Publication 800-38A, December 2001.

[SP800-132] National Institute of Standards and Technology, "Recommendation for Password-Based Key Derivation, Part 1: Storage Applications", NIST Special Publication 800-132, June 2010.

Appendix A. Test Vectors

Sample results for string-to-key conversion:

Iteration count = 32768
Pass phrase = "password"

Saltp for creating 128-bit base-key:

61 65 73 31 32 38 2D 63 74 73 2D 68 6D 61 63 2D
73 68 61 32 35 36 2D 31 32 38 00 10 DF 9D D7 83
E5 BC 8A CE A1 73 0E 74 35 5F 61 41 54 48 45 4E
41 2E 4D 49 54 2E 45 44 55 72 61 65 62 75 72 6E

(The saltp is "aes128-cts-hmac-sha256-128" | 0x00 |
random 16 byte valid UTF-8 sequence | "ATHENA.MIT.EDUraeburn")
128-bit base-key:

08 9B CA 48 B1 05 EA 6E A7 7C A5 D2 F3 9D C5 E7

Saltp for creating 256-bit base-key:

61 65 73 32 35 36 2D 63 74 73 2D 68 6D 61 63 2D
73 68 61 33 38 34 2D 31 39 32 00 10 DF 9D D7 83
E5 BC 8A CE A1 73 0E 74 35 5F 61 41 54 48 45 4E
41 2E 4D 49 54 2E 45 44 55 72 61 65 62 75 72 6E

(The saltp is "aes256-cts-hmac-sha384-192" | 0x00 |
random 16 byte valid UTF-8 sequence | "ATHENA.MIT.EDUraeburn")
256-bit base-key:

45 BD 80 6D BF 6A 83 3A 9C FF C1 C9 45 89 A2 22
36 7A 79 BC 21 C4 13 71 89 06 E9 F5 78 A7 84 67

Sample results for key derivation:

enttype aes128-cts-hmac-sha256-128:

128-bit base-key:

37 05 D9 60 80 C1 77 28 A0 E8 00 EA B6 E0 D2 3C

Kc value for key usage 2 (label = 0x0000000299):
B3 1A 01 8A 48 F5 47 76 F4 03 E9 A3 96 32 5D C3
Ke value for key usage 2 (label = 0x00000002AA):
9B 19 7D D1 E8 C5 60 9D 6E 67 C3 E3 7C 62 C7 2E
Ki value for key usage 2 (label = 0x0000000255):
9F DA 0E 56 AB 2D 85 E1 56 9A 68 86 96 C2 6A 6C

enttype aes256-cts-hmac-sha384-192:
256-bit base-key:
6D 40 4D 37 FA F7 9F 9D F0 D3 35 68 D3 20 66 98
00 EB 48 36 47 2E A8 A0 26 D1 6B 71 82 46 0C 52
Kc value for key usage 2 (label = 0x0000000299):
EF 57 18 BE 86 CC 84 96 3D 8B BB 50 31 E9 F5 C4
BA 41 F2 8F AF 69 E7 3D
Ke value for key usage 2 (label = 0x00000002AA):
56 AB 22 BE E6 3D 82 D7 BC 52 27 F6 77 3F 8E A7
A5 EB 1C 82 51 60 C3 83 12 98 0C 44 2E 5C 7E 49
Ki value for key usage 2 (label = 0x0000000255):
69 B1 65 14 E3 CD 8E 56 B8 20 10 D5 C7 30 12 B6
22 C4 D0 0F FC 23 ED 1F

Sample encryptions (all using the default cipher state):

These sample encryptions use the above sample key
derivation results, including use of the same
base-key and key usage values.

The following test vectors are for
enttype aes128-cts-hmac-sha256-128:

Plaintext: (empty)
Confounder:
7E 58 95 EA F2 67 24 35 BA D8 17 F5 45 A3 71 48
128-bit AES key (Ke):
9B 19 7D D1 E8 C5 60 9D 6E 67 C3 E3 7C 62 C7 2E
128-bit HMAC key (Ki):
9F DA 0E 56 AB 2D 85 E1 56 9A 68 86 96 C2 6A 6C
AES Output:
EF 85 FB 89 0B B8 47 2F 4D AB 20 39 4D CA 78 1D
Truncated HMAC Output:
AD 87 7E DA 39 D5 0C 87 0C 0D 5A 0A 8E 48 C7 18
Ciphertext (AES Output | HMAC Output):
EF 85 FB 89 0B B8 47 2F 4D AB 20 39 4D CA 78 1D
AD 87 7E DA 39 D5 0C 87 0C 0D 5A 0A 8E 48 C7 18

Plaintext: (length less than block size)
00 01 02 03 04 05
Confounder:

```
7B CA 28 5E 2F D4 13 0F B5 5B 1A 5C 83 BC 5B 24
128-bit AES key (Ke):
9B 19 7D D1 E8 C5 60 9D 6E 67 C3 E3 7C 62 C7 2E
128-bit HMAC key (Ki):
9F DA 0E 56 AB 2D 85 E1 56 9A 68 86 96 C2 6A 6C
AES Output:
84 D7 F3 07 54 ED 98 7B AB 0B F3 50 6B EB 09 CF
B5 54 02 CE F7 E6
Truncated HMAC Output:
87 7C E9 9E 24 7E 52 D1 6E D4 42 1D FD F8 97 6C
Ciphertext:
84 D7 F3 07 54 ED 98 7B AB 0B F3 50 6B EB 09 CF
B5 54 02 CE F7 E6 87 7C E9 9E 24 7E 52 D1 6E D4
42 1D FD F8 97 6C
```

```
Plaintext: (length equals block size)
00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
Confounder:
56 AB 21 71 3F F6 2C 0A 14 57 20 0F 6F A9 94 8F
128-bit AES key (Ke):
9B 19 7D D1 E8 C5 60 9D 6E 67 C3 E3 7C 62 C7 2E
128-bit HMAC key (Ki):
9F DA 0E 56 AB 2D 85 E1 56 9A 68 86 96 C2 6A 6C
AES Output:
35 17 D6 40 F5 0D DC 8A D3 62 87 22 B3 56 9D 2A
E0 74 93 FA 82 63 25 40 80 EA 65 C1 00 8E 8F C2
Truncated HMAC Output:
95 FB 48 52 E7 D8 3E 1E 7C 48 C3 7E EB E6 B0 D3
Ciphertext:
35 17 D6 40 F5 0D DC 8A D3 62 87 22 B3 56 9D 2A
E0 74 93 FA 82 63 25 40 80 EA 65 C1 00 8E 8F C2
95 FB 48 52 E7 D8 3E 1E 7C 48 C3 7E EB E6 B0 D3
```

```
Plaintext: (length greater than block size)
00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
10 11 12 13 14
Confounder:
A7 A4 E2 9A 47 28 CE 10 66 4F B6 4E 49 AD 3F AC
128-bit AES key (Ke):
9B 19 7D D1 E8 C5 60 9D 6E 67 C3 E3 7C 62 C7 2E
128-bit HMAC key (Ki):
9F DA 0E 56 AB 2D 85 E1 56 9A 68 86 96 C2 6A 6C
AES Output:
72 0F 73 B1 8D 98 59 CD 6C CB 43 46 11 5C D3 36
C7 0F 58 ED C0 C4 43 7C 55 73 54 4C 31 C8 13 BC
E1 E6 D0 72 C1
Truncated HMAC Output:
86 B3 9A 41 3C 2F 92 CA 9B 83 34 A2 87 FF CB FC
```

Ciphertext:

```
72 0F 73 B1 8D 98 59 CD 6C CB 43 46 11 5C D3 36
C7 0F 58 ED C0 C4 43 7C 55 73 54 4C 31 C8 13 BC
E1 E6 D0 72 C1 86 B3 9A 41 3C 2F 92 CA 9B 83 34
A2 87 FF CB FC
```

The following test vectors are for enctype
aes256-cts-hmac-sha384-192:

Plaintext: (empty)

Confounder:

```
F7 64 E9 FA 15 C2 76 47 8B 2C 7D 0C 4E 5F 58 E4
256-bit AES key (Ke):
```

```
56 AB 22 BE E6 3D 82 D7 BC 52 27 F6 77 3F 8E A7
A5 EB 1C 82 51 60 C3 83 12 98 0C 44 2E 5C 7E 49
```

192-bit HMAC key (Ki):

```
69 B1 65 14 E3 CD 8E 56 B8 20 10 D5 C7 30 12 B6
22 C4 D0 0F FC 23 ED 1F
```

AES Output:

```
41 F5 3F A5 BF E7 02 6D 91 FA F9 BE 95 91 95 A0
```

Truncated HMAC Output:

```
58 70 72 73 A9 6A 40 F0 A0 19 60 62 1A C6 12 74
8B 9B BF BE 7E B4 CE 3C
```

Ciphertext:

```
41 F5 3F A5 BF E7 02 6D 91 FA F9 BE 95 91 95 A0
58 70 72 73 A9 6A 40 F0 A0 19 60 62 1A C6 12 74
8B 9B BF BE 7E B4 CE 3C
```

Plaintext: (length less than block size)

```
00 01 02 03 04 05
```

Confounder:

```
B8 0D 32 51 C1 F6 47 14 94 25 6F FE 71 2D 0B 9A
256-bit AES key (Ke):
```

```
56 AB 22 BE E6 3D 82 D7 BC 52 27 F6 77 3F 8E A7
A5 EB 1C 82 51 60 C3 83 12 98 0C 44 2E 5C 7E 49
```

192-bit HMAC key (Ki):

```
69 B1 65 14 E3 CD 8E 56 B8 20 10 D5 C7 30 12 B6
22 C4 D0 0F FC 23 ED 1F
```

AES Output:

```
4E D7 B3 7C 2B CA C8 F7 4F 23 C1 CF 07 E6 2B C7
B7 5F B3 F6 37 B9
```

Truncated HMAC Output:

```
F5 59 C7 F6 64 F6 9E AB 7B 60 92 23 75 26 EA 0D
1F 61 CB 20 D6 9D 10 F2
```

Ciphertext:

```
4E D7 B3 7C 2B CA C8 F7 4F 23 C1 CF 07 E6 2B C7
B7 5F B3 F6 37 B9 F5 59 C7 F6 64 F6 9E AB 7B 60
92 23 75 26 EA 0D 1F 61 CB 20 D6 9D 10 F2
```

```
Plaintext: (length equals block size)
 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
Confounder:
 53 BF 8A 0D 10 52 65 D4 E2 76 42 86 24 CE 5E 63
256-bit AES key (Ke):
 56 AB 22 BE E6 3D 82 D7 BC 52 27 F6 77 3F 8E A7
 A5 EB 1C 82 51 60 C3 83 12 98 0C 44 2E 5C 7E 49
192-bit HMAC key (Ki):
 69 B1 65 14 E3 CD 8E 56 B8 20 10 D5 C7 30 12 B6
 22 C4 D0 0F FC 23 ED 1F
AES Output:
 BC 47 FF EC 79 98 EB 91 E8 11 5C F8 D1 9D AC 4B
 BB E2 E1 63 E8 7D D3 7F 49 BE CA 92 02 77 64 F6
Truncated HMAC Output:
 8C F5 1F 14 D7 98 C2 27 3F 35 DF 57 4D 1F 93 2E
 40 C4 FF 25 5B 36 A2 66
Ciphertext:
 BC 47 FF EC 79 98 EB 91 E8 11 5C F8 D1 9D AC 4B
 BB E2 E1 63 E8 7D D3 7F 49 BE CA 92 02 77 64 F6
 8C F5 1F 14 D7 98 C2 27 3F 35 DF 57 4D 1F 93 2E
 40 C4 FF 25 5B 36 A2 66

Plaintext: (length greater than block size)
 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
 10 11 12 13 14
Confounder:
 76 3E 65 36 7E 86 4F 02 F5 51 53 C7 E3 B5 8A F1
256-bit AES key (Ke):
 56 AB 22 BE E6 3D 82 D7 BC 52 27 F6 77 3F 8E A7
 A5 EB 1C 82 51 60 C3 83 12 98 0C 44 2E 5C 7E 49
192-bit HMAC key (Ki):
 69 B1 65 14 E3 CD 8E 56 B8 20 10 D5 C7 30 12 B6
 22 C4 D0 0F FC 23 ED 1F
AES Output:
 40 01 3E 2D F5 8E 87 51 95 7D 28 78 BC D2 D6 FE
 10 1C CF D5 56 CB 1E AE 79 DB 3C 3E E8 64 29 F2
 B2 A6 02 AC 86
Truncated HMAC Output:
 FE F6 EC B6 47 D6 29 5F AE 07 7A 1F EB 51 75 08
 D2 C1 6B 41 92 E0 1F 62
Ciphertext:
 40 01 3E 2D F5 8E 87 51 95 7D 28 78 BC D2 D6 FE
 10 1C CF D5 56 CB 1E AE 79 DB 3C 3E E8 64 29 F2
 B2 A6 02 AC 86 FE F6 EC B6 47 D6 29 5F AE 07 7A
 1F EB 51 75 08 D2 C1 6B 41 92 E0 1F 62
```

Sample checksums:

These sample checksums use the above sample key derivation results, including use of the same base-key and key usage values.

Checksum type: hmac-sha256-128-aes128

128-bit HMAC key (Kc):

B3 1A 01 8A 48 F5 47 76 F4 03 E9 A3 96 32 5D C3

Plaintext:

00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F

10 11 12 13 14

Checksum:

D7 83 67 18 66 43 D6 7B 41 1C BA 91 39 FC 1D EE

Checksum type: hmac-sha384-192-aes256

192-bit HMAC key (Kc):

EF 57 18 BE 86 CC 84 96 3D 8B BB 50 31 E9 F5 C4

BA 41 F2 8F AF 69 E7 3D

Plaintext:

00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F

10 11 12 13 14

Checksum:

45 EE 79 15 67 EE FC A3 7F 4A C1 E0 22 2D E8 0D

43 C3 BF A0 66 99 67 2A

Sample pseudorandom function (PRF) invocations:

PRF input octet-string: "test" (0x74657374)

enttype aes128-cts-hmac-sha256-128:

input-key value / HMAC-SHA-256 key:

37 05 D9 60 80 C1 77 28 A0 E8 00 EA B6 E0 D2 3C

HMAC-SHA-256 input message:

00 00 00 01 70 72 66 00 74 65 73 74 00 00 01 00

PRF output:

9D 18 86 16 F6 38 52 FE 86 91 5B B8 40 B4 A8 86

FF 3E 6B B0 F8 19 B4 9B 89 33 93 D3 93 85 42 95

enttype aes256-cts-hmac-sha384-192:

input-key value / HMAC-SHA-384 key:

6D 40 4D 37 FA F7 9F 9D F0 D3 35 68 D3 20 66 98

00 EB 48 36 47 2E A8 A0 26 D1 6B 71 82 46 0C 52

HMAC-SHA-384 input message:

00 00 00 01 70 72 66 00 74 65 73 74 00 00 01 80

PRF output:

98 01 F6 9A 36 8C 2B F6 75 E5 95 21 E1 77 D9 A0

7F 67 EF E1 CF DE 8D 3C 8D 6F 6A 02 56 E3 B1 7D

B3 C1 B6 2A D1 B8 55 33 60 D1 73 67 EB 15 14 D2

Authors' Addresses

Michael J. Jenkins
National Security Agency

EMail: mjjenki@tycho.ncsc.mil

Michael A. Peck
The MITRE Corporation

EMail: mpeck@mitre.org

Kelley W. Burgin

Email: kelley.burgin@gmail.com

KITTEN
Internet-Draft
Updates: 2743, 2744 (if approved)
Intended status: Standards Track
Expires: August 9, 2019

R. Harwood
Red Hat
N. Williams
Cryptonector
February 5, 2019

Channel Binding Signalling for the Generic Security Services Application
Programming Interface
draft-ietf-kitten-channel-bound-flag-04

Abstract

Channel binding is a technique that allows applications to use a secure channel at a lower layer without having to use authentication at that lower layer. The concept of channel binding comes from the Generic Security Services Application Programming Interface (GSS-API). It turns out that the semantics commonly implemented are different than those specified in the base GSS-API RFC (RFC2743), and that that specification has a serious bug. This document addresses both the inconsistency as-implemented and the specification bug.

This Internet-Draft proposes the addition of a "channel bound" return flag for the `GSS_Init_sec_context()` and `GSS_Accept_sec_context()` functions. Two behaviors are specified: a default, safe behavior reflecting existing implementation deployments, and a behavior that is only safe when the application specifically tells the GSS-API that it (the application) supports the new behavior. Additional API elements related to this are also added, including a new security context establishment API.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 9, 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
1.1. Design and Future directions	3
1.2. Conventions used in this document	3
2. Channel Binding State Extension	4
2.1. GSS_Create_sec_context()	4
2.1.1. C-Bindings	4
2.2. GSS_Set_context_flags()	5
2.2.1. C-Bindings	5
2.3. Return Flag for Channel Binding State Signalling	6
2.3.1. C-Bindings	6
2.4. New Mechanism Attribute	6
2.5. Request Flag for Acceptor Confirmation of Channel Binding	6
2.5.1. C-Bindings	6
2.6. Handling Empty Contexts in Other GSS-API Functions	6
3. Modified Channel Binding Semantics	7
4. Security Considerations	8
5. IANA Considerations	8
6. References	8
6.1. Normative References	8
6.2. Informative References	8
Authors' Addresses	9

1. Introduction

The GSS-API [RFC2743] supports "channel binding" [RFC5056], a technique for detection of man-in-the-middle (MITM) attacks in secure channels at lower network layers. This facility is meant to be all-or-nothing: either both the initiator and acceptor use it and it succeeds, or both must not use it. This has created a negotiation problem when retrofitting the use of channel binding into existing application protocols.

However, GSS-APIv2u1 [RFC2743] does not specify channel binding behavior when only one party provides provides none. In practice, some mechanisms (such as Kerberos [RFC4121]) ignore channel bindings when the acceptor provides none, but not when the initiator provides none. Note that it would be useless to allow security context establishment to succeed when the initiator does not provide channel bindings but the acceptor does, at least as long as there's no outward indication of whether channel binding was used! Since the GSS-APIv2u1 does not provide any such indication, this document corrects that flaw.

Allowing the connection to succeed when an initiator provides bindings but an acceptor does not has helped deployment of channel binding in existing applications: first fix all the initiators, then fix all the acceptors. But even this technique is insufficient when there are many clients to fix, such that fixing them all will take a long time. Additionally, it limits the usefulness of channel bindings, while allowing the acceptor to provide but not enforce would protect against man in the middle attacks (for channel binding aware initiators).

This document proposes a new method for deployment of channel binding that allows the feature to be enabled on the acceptor side before fixing all initiators. If the GSS-API had always had a return flag by which to indicate channel binding state then we could have had a simpler method of deploying channel binding: applications check that return flag and act accordingly (e.g., fail when channel binding is required). We cannot safely introduce this behavior now without an indication of support by the application.

Additionally, there may be applications where it is important for initiators to know that acceptors did use channel binding, and even to know whether a mechanism is capable of indicating as much. We add a request flag and a mechanism attribute for such applications.

1.1. Design and Future directions

The design for signalling application flag support and empty contexts is based on the Java Bindings of the GSS-API [RFC5653]. This document begins introduction of additional context inquiry and mutation functions, which eventually will also allow for simplified stepping to replace the GSS_Init/Accept_sec_context() loop.

1.2. Conventions used in this document

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. Channel Binding State Extension

We propose a new return flag for `GSS_Init_sec_context()` and `GSS_Accept_sec_context()`, as well as a pair of functions for a) creating "empty" security context handles, b) requesting flags and indicating which flags the application understands. We also add a new mechanism attribute for supporting channel binding confirmation.

C bindings of these extensions are provided along the lines of [RFC2744] and [RFC5587].

In the future we might move more of the many input (and output) arguments to `GSS_Init_sec_context()` and `GSS_Accept_sec_context()` into mutators on empty security context handles.

2.1. `GSS_Create_sec_context()`

Inputs:

- o <none>

Outputs:

- o `major_status` INTEGER
- o `minor_status` INTEGER -- note: mostly useless, but we should keep it
- o `context` SECURITY CONTEXT -- "empty" security context

Return major status codes:

- o `GSS_S_COMPLETE` indicates success.
- o `GSS_S_UNAVAILABLE` indicates that memory is not available, for example.
- o `GSS_S_FAILURE` indicates a general failure.

This function creates an "empty" security context handle that can be passed to `GSS_Init_sec_context()` or `GSS_Accept_sec_context()` where they expect `GSS_C_NO_CONTEXT`.

2.1.1. C-Bindings

```
OM_uint32
gss_create_sec_context(OM_uint32 *minor_status,
                      gss_ctx_id_t *context);
```

2.2. GSS_Set_context_flags()

Inputs:

context CONTEXT HANDLE

req_flags FLAGS Requested flags. Applicable to acceptors and initiators.

ret_flags_understood FLAGS The set of return flags understood by the caller.

Outputs:

o major_status INTEGER

o minor_status INTEGER

Return major status codes:

o GSS_S_COMPLETE indicates success.

o GSS_S_FAILURE indicates a general failure.

This function tells the mechanism (when one is eventually chosen and invoked) that the application requests the given req_flags and is prepared to check the flags in the given ret_flags_understood. Mechanisms SHOULD NOT limit flags returned to those in ret_flags_understood, but MAY alter behavior accordingly. Initiators can override the req_flags in their GSS_Init_sec_context() call, but if no flags are requested there then the req_flags set on the empty context will be used. GSS_Accept_sec_context() is not required to perform any action based on req_flags at this time.

NOTE: The abstract GSS-API [RFC2743] uses individual elements--one per-flag--instead of a "FLAGS" type. This is unwieldy, therefore we introduce an abstract type named "FLAGS" to act as a set of all the request/return flags defined for the abstract GSS-API.

2.2.1. C-Bindings

```
OM_uint32
gss_set_context_flags(OM_uint32 *minor_status,
                     gss_ctx_id_t context,
                     uint64_t req_flags,
                     uint64_t ret_flags_understood);
```

2.3. Return Flag for Channel Binding State Signalling

Whenever both the initiator and the acceptor provide matching channel bindings to `GSS_Init_sec_context()` and `GSS_Accept_sec_context()`, respectively, then the mechanism SHALL indicate that the context is channel bound via an output flag, `ret_channel_bound_flag`, for the established context. Note that some mechanisms have no way for the acceptor to signal CB success to the initiator, in which case `GSS_Init_sec_context()` MUST NOT output the `ret_channel_bound_flag`.

2.3.1. C-Bindings

```
#define GSS_C_CHANNEL_BOUND_FLAG 2048 /* 0x00000800 */
```

2.4. New Mechanism Attribute

- o We add a new mechanism attribute, `GSS_C_MA_CBINDING_CONFIRM`, to indicate that the initiator can and always does learn whether the acceptor application supplied channel bindings (assuming mutual auth has been requested). Mechanisms advertising this attribute MUST always indicate acceptor channel bound state to the initiator.

OID assignments TBD.

2.5. Request Flag for Acceptor Confirmation of Channel Binding

We add a new request flag for `GSS_Init_sec_context()`, `req_cb_confirmation_flag`, to be used by initiators that insist on acceptors providing channel bindings. If set, the mechanism MUST prefer establishment of contexts which provide channel binding confirmation. It SHOULD NOT fail to negotiate just because it cannot provide the `GSS_C_MA_CBINDING_CONFIRM` attribute.

2.5.1. C-Bindings

Because `GSS_C_CHANNEL_BOUND_FLAG` is a return flag only, and this flag is a request flag only, and to save on precious flag bits, we use the same flag bit assignment for both flags:

```
#define GSS_C_CB_CONFIRM_FLAG 2048 /* 0x00000800 */
```

2.6. Handling Empty Contexts in Other GSS-API Functions

`GSS_Init_sec_context()` and `GSS_Accept_sec_context()` operate on empty security contexts as specified above (i.e., examining flags).

All other GSS-API functions MUST treat empty contexts as they would GSS_C_NO_CONTEXT as well. For most functions, this will result in returning GSS status code GSS_S_NO_CONTEXT.

GSS_Delete_sec_context() MUST NOT output a context deletion token when applied to empty security contexts.

3. Modified Channel Binding Semantics

The channel binding semantics of the base GSS-API are modified as follows:

- o Whenever both the initiator and acceptor have provided input_channel_bindings to GSS_Init/Accept_sec_context() and the channel bindings do not match, then the mechanism MUST fail to establish a security context token. (This is a restatement of an existing requirement in the base specification.)
- o Whenever the acceptor application has a) provided channel bindings to GSS_Accept_sec_context(), and b) not indicated support for the ret_channel_bound_flag flag, then the mechanism MUST fail to establish a security context if the initiator did not provide channel bindings data. This requirement is critical for security purposes, to make applications predating this document secure, and this requirement reflects actual implementations as deployed.
- o Whenever the initiator application has a) provided channel bindings to GSS_Init_sec_context(), and b) not indicated support for the ret_channel_bound_flag flag, then the mechanism SHOULD NOT fail to establish a security context just because the acceptor failed to provide channel bindings data. This strong suggestion is for interoperability purposes, and reflects actual implementations that have been deployed.
- o Whenever the application has a) provided channel bindings to GSS_Init_sec_context() or GSS_Accept_sec_context(), and b) indicated support for the ret_channel_bound_flag flag, then the mechanism SHOULD NOT fail to establish a security context just because the peer did not provide channel bindings data. The mechanism MUST output the ret_channel_bound_flag if both peers provided the same input_channel_bindings to GSS_Init_sec_context() and GSS_Accept_sec_context(). The mechanism MUST NOT output the ret_channel_bound_flag if either (or both) peer did not provide input_channel_bindings to GSS_Init/Accept_sec_context(). This requirement restores the original base GSS-API specified behavior, with the addition of the ret_channel_bound_flag flag.

4. Security Considerations

This document deals with security. There are no security considerations that should be documented separately in this section. To recap, this document fixes a significant flaw in the base GSS-API [RFC2743] specification that fortunately has not been implemented, and it adds a feature (that should have been in the base specification) for improved negotiation of use of channel binding [RFC5056].

5. IANA Considerations

The GSS-API mechanism attribute is to be added to the "SMI Security for Mechanism gsscma Codes" registry established by RFC5587 [RFC5587]. See Section 2.4.

6. References

6.1. 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>>.
- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, DOI 10.17487/RFC2743, January 2000, <<https://www.rfc-editor.org/info/rfc2743>>.
- [RFC2744] Wray, J., "Generic Security Service API Version 2 : C-bindings", RFC 2744, DOI 10.17487/RFC2744, January 2000, <<https://www.rfc-editor.org/info/rfc2744>>.
- [RFC5056] Williams, N., "On the Use of Channel Bindings to Secure Channels", RFC 5056, DOI 10.17487/RFC5056, November 2007, <<https://www.rfc-editor.org/info/rfc5056>>.
- [RFC5587] Williams, N., "Extended Generic Security Service Mechanism Inquiry APIs", RFC 5587, DOI 10.17487/RFC5587, July 2009, <<https://www.rfc-editor.org/info/rfc5587>>.

6.2. Informative References

- [RFC4121] Zhu, L., Jaganathan, K., and S. Hartman, "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2", RFC 4121, DOI 10.17487/RFC4121, July 2005, <<https://www.rfc-editor.org/info/rfc4121>>.
- [RFC5653] Upadhyay, M. and S. Malkani, "Generic Security Service API Version 2: Java Bindings Update", RFC 5653, DOI 10.17487/RFC5653, August 2009, <<https://www.rfc-editor.org/info/rfc5653>>.

Authors' Addresses

Robbie Harwood
Red Hat, Inc.

Email: rharwood@redhat.com

Nicolas Williams
Cryptonector, LLC

Email: nico@cryptonector.com

NETWORK WORKING GROUP
Internet-Draft
Intended status: Standards Track
Expires: October 1, 2017

N. Williams
Cryptonector LLC
A. Melnikov
Isode Ltd
March 30, 2017

Namespace Considerations and Registries for GSS-API Extensions
draft-ietf-kitten-gssapi-extensions-iana-11.txt

Abstract

This document describes the ways in which the GSS-API may be extended and directs the creation of an IANA registry for various GSS-API namespaces.

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 October 1, 2017.

Copyright Notice

Copyright (c) 2017 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. Conventions used in this document	2
2. Introduction	2
3. Extensions to the GSS-API	2
4. Generic GSS-API Namespaces	3
5. Language Binding-Specific GSS-API Namespaces	3
6. Extension-Specific GSS-API Namespaces	4
7. Registration Form	4
8. IANA Considerations	6
8.1. Initial Namespace Registrations	7
8.1.1. Example registrations	7
8.2. Registration Maintenance Guidelines	9
8.2.1. Sub-Namespace Symbol Pattern Matching	10
8.2.2. Expert Reviews of Individual Submissions	10
8.2.3. Change Control	11
9. Security Considerations	12
10. References	12
10.1. Normative References	12
10.2. Informative References	12
Authors' Addresses	13

1. Conventions used in this document

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. Introduction

There is a need for private-use and mechanism-specific extensions to the Generic Security Services Application Programming Interface (GSS-API). As such extensions are designed and standardized (or not), both at the IETF and elsewhere, there is a non-trivial risk of namespace pollution and conflicts. To avoid this we set out guidelines for extending the GSS-API and direct the creation of an IANA registry for GSS-API namespaces.

Registrations of individual items and sub-namespaces are allowed. Each sub-namespace may provide different rules for registration, e.g., for mechanism-specific and private-use extensions.

3. Extensions to the GSS-API

Extensions to the GSS-API can be categorized as follows:

- o Abstract API extensions

- o Implementation-specific
- o Mechanism-specific
- o Language binding-specific

Extensions to the GSS-API may be purely semantic, without effect on the GSS-API's namespaces. Or they may introduce new functions, constants, types, etc...; these clearly affect the GSS-API namespaces.

Extensions that affect the GSS-API namespaces should be registered with the IANA as described herein.

4. Generic GSS-API Namespaces

The abstract API namespaces for the GSS-API are:

- o Type names
- o Function names
- o Constant names for various types
- o Constant values for various types
- o Name types (OID, type name and syntaxes)

Additionally we have namespaces associates with the OBJECT IDENTIFIER (OID) type. The IANA already maintains a registry of such OIDs:

- o Mechanism OIDs
- o Name Type OIDs

5. Language Binding-Specific GSS-API Namespaces

Language binding specific namespaces include, among others:

- o Header/interface module names
- o Object classes and/or types
- o Methods and/or functions
- o Constant names
- o Constant values

6. Extension-Specific GSS-API Namespaces

Extensions to the GSS-API may create additional namespaces. See Section 8.2.

7. Registration Form

Registrations for GSS-API namespaces SHALL take the following form:

Registration Field	Possible Values	Description
Bindings	'Generic', 'C-bindings', 'Java', 'C#', <programming language name>	Indicates the name of the programming language that this registration involves, or, if 'Generic', that this is an entry for the generic abstract GSS-API (i.e., not specific to any programming language).
Registration type	'Instance', 'Sub-Namespace'	Indicates whether this entry reserves a given symbol name (and possibly, constant value), or whether it reserves an entire sub-namespace (the name is a pattern) or constant value range.
Object Type	<Symbol> defined by the binding language (for example 'Data-Type', 'Function', 'Method', 'Integer', 'String', 'OID', 'Context-Flag', 'Name-Type', 'Macro', 'Header-File-Name', 'Module-Name', 'Class')	Indicates the type of the object whose symbolic name or constant value this entry registers. The possible values of this field depend on the programming language in question, therefore they are not all specified here.
Symbol Name/Prefix	<Symbol name or name pattern>	The name of a symbol or symbol sub-namespace being

		registered. See Section 8.2.1
Binding of	<Name of abstract API element of which this object is a binding>	If the registration is for a specific language binding of the GSS-API, then this names the abstract API element of which it is a binding (OPTIONAL).
Constant Value/Range	<Constant value> or <constant value range>	The value of the constant named by the <Symbol Name/Prefix>. This field is present only for Instance and Sub-namespace registrations of Constant object types.
Description	<Text>	Description of the registration. Multiple instances of this field may result (see Section 8.2.3).
Registration Rules	<Reference> to an IANA registration Policy defined in [RFC5226] (or an RFC that updates it), for instance 'IESG Approval', 'Expert Review', 'First Come First Served', 'Private Use'.	Describes the rules for allocation of items that fall in this sub-namespace, for entries with Registration Type of Sub-namespace (OPTIONAL). For private use sub-namespaces the submitter MUST provide the e-mail address of a responsible contact. If this field is not specified for a sub-namespace, the default registration rules specified in Section 8.2 apply.
Reference	<Reference>	Reference to a document that describes the registration, if any (OPTIONAL). Multiple instances of this field are allowed, with one reference each.
Expert Reviewer	<Name of expert reviewers, possibly	OPTIONAL, see Section 8.2.2. Multiple instances of this

	WG names>	field are allowed, with one expert reviewer per-instance. Leave this field blank when requesting a registration. It will be filled in by the Expert who reviews the registration.
Expert Review Notes	<Notes from the expert review>	Expert reviewers may request that some comments be included with the registration, e.g., regarding security considerations of the registered extension.
Status	'Registered' or 'Obsoleted'	Status of the registration.
Obsoleting Reference	<Reference>	Reference to a document, if any, that obsoletes this registration. Multiple instances of this field are allowed, with one reference each. (OPTIONAL)

The IANA should create a single GSS-API namespace registry, or multiple registries, one for symbolic names and one for constant values, and/or it may create a registry per-programming language, at its convenience.

Entries in these registries should consist of all the fields from their corresponding registration entries.

Entries should be sorted by: programming language, registration type, object type, and symbol name/pattern.

8. IANA Considerations

This document deals with IANA considerations throughout. Specifically it creates a single registry of various kinds of things, though the IANA may instead create multiple registries, each for one of those kinds of things. Of particular interest may be that IANA will now be the registration authority for the GSS-API name type OID space.

8.1. Initial Namespace Registrations

Initial registry content corresponding to the items defined in [RFC2743], [RFC2744], [RFC2853], [RFC1964] and [RFC4121] and others will be supplied during the IANA review portion of the RFC publishing process. [[Note to RFC Editor: Delete the following sentence before publication:]] The KITTEN WG chairs MUST indicate that such content has been reviewed by the WG and that there is WG consensus that the entries are in agreement with those RFCs.

8.1.1. Example registrations

In order to sanity check recommended IANA registration templates, this section registers several entries.

Registration Field	Possible Values
Bindings	C-bindings
Registration type	Instance
Object Type	Function
Symbol Name	gss_init_sec_context
Binding of	GSS_Init_sec_context
Constant Value/Range	N/A
Description	Create a security context by initiator
Registration Rules	N/A
Reference	RFC 2744
Expert Reviewer	Kitten WG
Expert Review Notes	
Status	Registered
Obsoleting Reference	N/A

Registration Field	Possible Values
Bindings	C-bindings
Registration type	Instance
Object Type	Function
Symbol Name	gss_accept_sec_context
Binding of	GSS_Accept_sec_context
Constant Value/Range	N/A
Description	Accept a security context from initiator
Registration Rules	N/A
Reference	RFC 2744
Expert Reviewer	Kitten WG
Expert Review Notes	
Status	Registered
Obsoleting Reference	N/A

Registration Field	Possible Values
Bindings	C-bindings
Registration type	Instance
Object Type	Context-Flag
Symbol Name	GSS_C_DELEG_FLAG
Binding of	deleg_state or deleg_req_flag
Constant Value/Range	1
Description	On output (if set): Delegated credentials are available via the delegated_cred_handle parameter of GSS_Accept_sec_context. On input (if set): With the call to GSS_Init_sec_context, delegate credentials to the acceptor.
Registration Rules	N/A
Reference	RFC 2744
Expert Reviewer	Kitten WG
Expert Review Notes	
Status	Registered
Obsoleting Reference	N/A

8.2. Registration Maintenance Guidelines

Standards-Track RFCs can create new items with any non-conflicting Symbol Name/Prefix value for this registry by virtue of IESG approval to publish as a Standards-Track RFC -- that is, without additional expert review.

Standards-Track RFCs can mark existing entries as obsolete, and can even create conflicting entries if explicitly stated (the IESG, of course, should review conflicts carefully, and may reject them).

IANA shall also consider submissions from individuals, and via Informational and Experimental RFCs, subject to Expert Review. IANA SHALL allow such registrations if a) they are not conflicting, b) provided that the registration is for object types other than Context-Flags, and c) subject to expert review. Guidelines for expert reviews are given below.

8.2.1. Sub-Namespace Symbol Pattern Matching

Sub-namespace registrations must provide a pattern for matching symbols for which the sub-namespace's registration rules apply. The pattern consists of a string with the following special tokens:

- o '*', meaning "match any string."
- o "%m", meaning "match any mechanism family short-hand name."
- o "%i", meaning "match any implementor vanity short-hand name."

For example, "GSS_%m*" matches "GSS_krb5_foo" since "krb5" is a common short-hand for the Kerberos V GSS-API mechanism [RFC1964]. But "GSS_%m*" does not match "GSS_foo_bar" unless "foo" is asserted to be a short-hand for some mechanism.

8.2.2. Expert Reviews of Individual Submissions

[[The following paragraph should be deleted from the document before publication, as it will not age well. It should be moved to the shepherding write-up.]]

Expert review selection SHALL be done as follows. If, at the time that the IANA receives an individual submission for registration in this registry, there are any IETF Working Groups chartered to produce GSS-API-related documents, then the IANA SHALL ask the chairs of such WGs to be expert reviewers or to name one. If there are no such WGs at that time, then the IANA SHALL ask past chairs of the KITTEN WG and the author/editor of this RFC to act as expert reviewers or name an alternate.

Expert reviewers of individual registration submissions with Registration Type == Sub-namespace should check that the registration request has a suitable description (which doesn't need to be sufficiently detailed for others to implement) and that the Symbol Name/Prefix is sufficiently descriptive of the purpose of the sub-namespace or reflective of the name of the submitter or associated company.

Expert reviewers of individual registration submissions with

Registration Type == Instance should check that the Symbol Name falls under a sub-namespace controlled by the submitter. Registration of such entries which do not fall under such a sub-namespace may be allowed provided that they correspond to long existing non-standard extensions to the GSS-API and this can be easily checked or demonstrated, otherwise IESG Protocol Action is REQUIRED (see previous section). Also, reviewers should check that any registration of constant values have a detailed description that is suitable for other implementors to reproduce, and that they don't conflict with other usages or are otherwise dangerous in the reviewers estimation.

Expert reviewers should review impact on mechanisms, security and interoperability, and may reject or annotate registrations which can have mechanism impact that requires IESG protocol action. Consider, for example, new versions of `GSS_Init_sec_context()` and/or `GSS_Accept_sec_context` which have new input and/or output parameters which imply changes on the wire or in behaviour that may result in interoperability issues. A reviewer could choose to add notes to the registration describing such issues, or the reviewer might conclude that the danger to Internet interoperability is sufficient to warrant rejecting the registration.

8.2.3. Change Control

Registered entries may be marked obsoleted using the same expert review process as for registering new entries. Obsoleted entries are not, however, to be deleted, but merely marked having Obsoleted Status. Note that entries may be created as obsoleted to record the fact that the given symbol(s) have been used before, even though continued use of them is discouraged.

Registered entries may also be updated in two other ways: additional references, obsoleting references, and descriptions may be added.

All changes are subject to expert review, except for changes to registrations in a sub-namespace which are subject to the rules of the relevant sub-namespace. The submitter of a change request need not be the same as the original submitter.

Registrations may be modified by addition, but under no circumstance may any fields be modified except for the Status field or Contact Address, or to correct for transcription errors in filing or processing registration requests.

The IANA SHALL add a field describing the date that a an addition or modification was made, and a description of the change.

9. Security Considerations

General security considerations relating to IANA registration services apply; see [RFC5226].

Also, expert reviewers should look for and may document security related issues with submitters' GSS-API extensions, to the best of the reviewers' ability given the information furnished by the submitter. Reviewers may add comments regarding their limited ability to review a submission for security problems if the submitter is unwilling to provide sufficient documentation.

10. References

10.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, DOI 10.17487/RFC2743, January 2000, <<http://www.rfc-editor.org/info/rfc2743>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

10.2. Informative References

- [RFC1964] Linn, J., "The Kerberos Version 5 GSS-API Mechanism", RFC 1964, DOI 10.17487/RFC1964, June 1996, <<http://www.rfc-editor.org/info/rfc1964>>.
- [RFC2744] Wray, J., "Generic Security Service API Version 2 : C-bindings", RFC 2744, DOI 10.17487/RFC2744, January 2000, <<http://www.rfc-editor.org/info/rfc2744>>.
- [RFC2853] Kabat, J. and M. Upadhyay, "Generic Security Service API Version 2 : Java Bindings", RFC 2853, DOI 10.17487/RFC2853, June 2000, <<http://www.rfc-editor.org/info/rfc2853>>.

[RFC4121] Zhu, L., Jaganathan, K., and S. Hartman, "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2", RFC 4121, DOI 10.17487/RFC4121, July 2005, <<http://www.rfc-editor.org/info/rfc4121>>.

Authors' Addresses

Nicolas Williams
Cryptonector LLC

Email: nico@cryptonector.com

Alexey Melnikov
Isode Ltd
5 Castle Business Village
36 Station Road
Hampton, Middlesex TW12 2BX
UK

Email: Alexey.Melnikov@isode.com

NETWORK WORKING GROUP
Internet-Draft
Updates: 4120,4121 (if approved)
Intended status: Standards Track
Expires: October 1, 2017

B. Kaduk, Ed.
Akamai
J. Schaad, Ed.
Soaring Hawk Consulting
L. Zhu
Microsoft Corporation
J. Altman
Secure Endpoints
March 30, 2017

Initial and Pass Through Authentication Using Kerberos V5 and the GSS-
API (IAKERB)
draft-ietf-kitten-iakerb-03

Abstract

This document defines extensions to the Kerberos protocol and the GSS-API Kerberos mechanism that enable a GSS-API Kerberos client to exchange messages with the KDC by using the GSS-API acceptor as a proxy, encapsulating the Kerberos messages inside GSS-API tokens. With these extensions a client can obtain Kerberos tickets for services where the KDC is not accessible to the client, but is accessible to 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 October 1, 2017.

Copyright Notice

Copyright (c) 2017 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
2. Conventions Used in This Document	3
3. GSS-API Encapsulation	3
3.1. Enterprise principal names	6
4. Finish Message	7
5. Addresses in Tickets	8
6. Security Considerations	8
7. Acknowledgements	9
8. Assigned Numbers	10
9. IANA Considerations	10
10. References	10
10.1. Normative References	10
10.2. Informative references	11
Appendix A. Interoperate with Previous MIT version	11
Authors' Addresses	12

1. Introduction

When authenticating using Kerberos V5, clients obtain tickets from a KDC and present them to services. This model of operation cannot work if the client does not have access to the KDC. For example, in remote access scenarios, the client must initially authenticate to an access point in order to gain full access to the network. Here the client may be unable to directly contact the KDC either because it does not have an IP address, or the access point packet filter does not allow the client to send packets to the Internet before it authenticates to the access point. The Initial and Pass Through Authentication Using Kerberos (IAKERB) mechanism allows for the use of Kerberos in such scenarios where the client is unable to directly contact the KDC, by using the service to pass messages between the client and the KDC. This allows the client to obtain tickets from the KDC and present them to the service, as in normal Kerberos operation.

Recent advancements in extending Kerberos permit Kerberos authentication to complete with the assistance of a proxy. The

Kerberos [RFC4120] pre-authentication framework [RFC6113] prevents the exposure of weak client keys over the open network. The Kerberos support of anonymity [RFC6112] provides for privacy and further complicates traffic analysis. The kdc-referrals option defined in [RFC6113] may reduce the number of messages exchanged while obtaining a ticket to exactly two even in cross-realm authentications.

Building upon these Kerberos extensions, this document extends [RFC4120] and [RFC4121] such that the client can communicate with the KDC using a Generic Security Service Application Program Interface (GSS-API) [RFC2743] acceptor as a message-passing proxy. (This is completely unrelated to the type of proxy specified in [RFC4120].) The client acts as a GSS-API initiator, and the GSS-API acceptor relays the KDC request and reply messages between the client and the KDC, transitioning to normal [RFC4121] GSS-krb5 messages once the client has obtained the necessary credentials. Consequently, IAKERB as defined in this document requires the use of the GSS-API.

The GSS-API acceptor, when relaying these Kerberos messages, is called an IAKERB proxy.

2. Conventions Used in This Document

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

3. GSS-API Encapsulation


The GSS-API mechanism Objection Identifier (OID) for IAKERB is id-kerberos-iakerb:

```
id-kerberos-iakerb ::=
{ iso(1) org(3) dod(6) internet(1) security(5) kerberosV5(2)
  iakerb(5) }
```

All context establishment tokens of IAKERB MUST have the token framing described in section 4.1 of [RFC4121] with the mechanism OID being id-kerberos-iakerb. MIT implemented an earlier draft of this specification; details on how to interoperate with that implementation can be found in Appendix A.

The client starts by constructing a ticket request, as if it is being made directly to the KDC. Instead of contacting the KDC directly, the client encapsulates the request message into the output token of the GSS_Init_security_context() call and returns GSS_S_CONTINUE_NEEDED [RFC2743], indicating that at least one more token is required in order to establish the context. The output

token is then passed over the application protocol for use as the input token to the `GSS_Accept_sec_context()` call in accordance with GSS-API. The GSS-API acceptor extracts the Kerberos request from the input token, locates the target KDC, and sends the request on behalf of the client. After receiving the KDC reply, the GSS-API acceptor then encapsulates the reply message into the output token of `GSS_Accept_sec_context()`. The GSS-API acceptor returns `GSS_S_CONTINUE_NEEDED` [RFC2743] indicating that at least one more token is required in order to establish the context. The output token is passed to the initiator over the application protocol in accordance with GSS-API.



For all context tokens generated by the IAKERB mechanism, the innerToken described in section 4.1 of [RFC4121] has the following format: it starts with a two-octet token-identifier (`TOK_ID`), which is followed by an IAKERB message or a Kerberos message.

Only one IAKERB specific message, namely the `IAKERB_PROXY` message, is defined in this document. The `TOK_ID` values for Kerberos messages are the same as defined in [RFC4121].

Token	TOK_ID Value in Hex
IAKERB_PROXY	05 01

The content of the `IAKERB_PROXY` message is defined as an `IAKERB-HEADER` structure immediately followed by a Kerberos message, which is optional. The Kerberos message can be an `AS-REQ`, an `AS-REP`, a `TGS-REQ`, a `TGS-REP`, or a `KRB-ERROR` as defined in [RFC4120].

```

IAKERB-HEADER ::= SEQUENCE {
    -- Note that the tag numbers start at 1, not 0, which would
    -- be more conventional for Kerberos.
    target-realm      [1] UTF8String,
    -- The name of the target realm.
    cookie            [2] OCTET STRING OPTIONAL,
    -- Opaque data, if sent by the server,
    -- MUST be copied by the client verbatim into
    -- the next IAKRB_PROXY message.
    ...
}
  
```

The IAKERB-HEADER structure and all the Kerberos messages MUST be encoded using Abstract Syntax Notation One (ASN.1) Distinguished Encoding Rules (DER) [CCITT.X680.2002] [CCITT.X690.2002].

The client fills out the IAKERB-HEADER structure as follows: the target-realm contains the realm name the ticket request is addressed to. In the initial message from the client, the cookie field is absent. The client MAY send a completely empty IAKERB_PROXY message (consisting solely of the octets 05 01 and an IAKERB_HEADER with zero-length target-realm) in order to query the Kerberos realm of the acceptor, see Section 3.1. In all other cases, the client MUST specify a target-realm. This can be the realm of the client's host, if no other realm information is available. client's host.

Upon receipt of the IAKERB_PROXY message, the GSS-API acceptor inspects the target-realm field in the IAKERB_HEADER, locates a KDC for that realm, and sends the ticket request to that KDC. The IAKERB proxy MAY engage in fallback behavior, retransmitting packets to a given KDC and/or sending the request to other KDCs in that realm if the initial transmission does not receive a reply, as would be done if the proxy was making requests on its own behalf.

The GSS-API acceptor encapsulates the KDC reply message in the returned IAKERB message. It fills out the target realm using the realm sent by the client and the KDC reply message is included immediately following the IAKERB-HEADER header.

When the GSS-API acceptor is unable to obtain an IP address for a KDC in the client's realm, it sends a KRB_ERROR message with the code KRB_AP_ERR_IAKERB_KDC_NOT_FOUND to the client in place of an actual reply from the KDC, and the context fails to establish. There is no accompanying error data defined in this document for this error code.

```
KRB_AP_ERR_IAKERB_KDC_NOT_FOUND      85
-- The IAKERB proxy could not find a KDC.
```

When the GSS-API acceptor has an IP address for at least one KDC in the target realm, but does not receive a response from any KDC in the realm (including in response to retries), it sends a KRB_ERROR message with the code KRB_AP_ERR_IAKERB_KDC_NO_RESPONSE to the client and the context fails to establish. There is no accompanying error data defined in this document for this error code.

```
KRB_AP_ERR_IAKERB_KDC_NO_RESPONSE    86
-- The KDC did not respond to the IAKERB proxy.
```

The IAKERB proxy can send opaque data in the cookie field of the IAKERB-HEADER structure in the server reply to the client, in order

to, for example, minimize the amount of state information kept by the GSS-API acceptor. The content and the encoding of the cookie field is a local matter of the IAKERB proxy. Whenever the cookie is present in a token received by the initiator, the initiator **MUST** copy the cookie verbatim into its subsequent response tokens which contain IAKERB_PROXY messages.

The client and the server can repeat the sequence of sending and receiving the IAKERB messages as described above for an arbitrary number of message exchanges, in order to allow the client to interact with the KDC through the IAKERB proxy, and to obtain Kerberos tickets as needed to authenticate to the acceptor.

Once the client has obtained the service ticket needed to authenticate to the acceptor, subsequent GSS-API context tokens are of type KRB_AP_REQ, not IAKERB_PROXY, and the client performs the client-server application exchange as defined in [RFC4120] and [RFC4121].

For implementations conforming to this specification, both the authenticator subkey and the GSS_EXTS_FINISHED extension as defined in Section 4 **MUST** be present in the AP-REQ authenticator. This checksum provides integrity protection for the IAKERB messages previously exchanged, including the unauthenticated clear texts in the IAKERB-HEADER structure.

If the pre-authentication data is encrypted in the long-term password-based key of the principal, the risk of security exposures is significant. Implementations **SHOULD** utilize the AS_REQ armoring as defined in [RFC6113] unless an alternative protection is deployed. In addition, the anonymous Kerberos FAST option is **RECOMMENDED** for the client to complicate traffic analysis.

3.1. Enterprise principal names

The introduction of principal name canonicalization by [RFC6806] created the possibility for a client to have a principal name (of type NT-ENTERPRISE) for which it is trying to obtain credentials, but no information about what realm's KDC to contact to obtain those credentials. A Kerberos client not using IAKERB would typically resolve the NT-ENTERPRISE name to a principal name by starting from the realm of the client's host and finding out the true realm of the enterprise principal based on referrals [RFC6806].

A client using IAKERB may not have any realm information, even for the realm of the client's host, or may know that the client host's realm is not appropriate for a given enterprise principal name. In such cases, the client can retrieve the realm of the GSS-API acceptor

as follows: the client returns GSS_S_CONTINUE_NEEDED with the output token containing an IAKERB message with an empty target-realm in the IAKERB-HEADER and no Kerberos message following the IAKERB-HEADER structure. Upon receipt of the realm request, the GSS-API acceptor fills out an IAKERB_PROXY response message, filling the target-realm field with the realm of the acceptor, and returns GSS_S_CONTINUE_NEEDED with the output token containing the IAKERB message with the server's realm and no Kerberos message following the IAKERB-HEADER header. The GSS-API initiator can then use the returned realm in subsequent IAKERB messages to resolve the NT-ENTERPRISE name type. Since the GSS-API acceptor can act as a Kerberos acceptor, it always has an associated Kerberos realm.

4. Finish Message

For implementations conforming to this specification, the authenticator subkey in the AP-REQ MUST always be present, and the Exts field in the GSS-API authenticator [RFC6542] MUST contain an extension of type GSS_EXTS_FINISHED with extension data containing the ASN.1 DER encoding of the structure KRB-FINISHED.

```
GSS_EXTS_FINISHED          2
    --- Data type for the IAKERB checksum.

KRB-FINISHED ::= {
    -- Note that the tag numbers start at 1, not 0, which would be
    -- more conventional for Kerberos.
    gss-mic [1] Checksum,
        -- Contains the checksum [RFC3961] of the GSS-API tokens
        -- exchanged between the initiator and the acceptor,
        -- and prior to the containing AP-REQ GSS-API token.
        -- The checksum is performed over the GSS-API tokens
        -- exactly as they were transmitted and received,
        -- in the order that the tokens were sent.
    ...
}
```

The gss-mic field in the KRB-FINISHED structure contains a Kerberos checksum [RFC3961] of all the preceding context tokens of this GSS-API context (including the generic token framing of the GSSAPI-Token type from [RFC4121]), concatenated in chronological order (note that GSS-API context token exchanges are synchronous). The checksum type is the required checksum type of the enctype of the subkey in the authenticator, the protocol key for the checksum operation is the authenticator subkey, and the key usage number is KEY_USAGE_FINISHED.

```
KEY_USAGE_FINISHED          41
```

The GSS-API acceptor MUST then verify the checksum contained in the GSS_EXTS_FINISHED extension. This checksum provides integrity protection for the messages exchanged including the unauthenticated clear texts in the IAKERB-HEADER structure.

5. Addresses in Tickets

In IAKERB, the machine sending requests to the KDC is the GSS-API acceptor and not the client. As a result, the client should not include its addresses in any KDC requests for two reasons. First, the KDC may reject the forwarded request as being from the wrong client. Second, in the case of initial authentication for a dial-up client, the client machine may not yet possess a network address. Hence, as allowed by [RFC4120], the addresses field of the AS-REQ and TGS-REQ requests SHOULD be blank.

6. Security Considerations

The IAKERB proxy is a man-in-the-middle for the client's Kerberos exchanges. The Kerberos protocol is designed to be used over an untrusted network, so this is not a critical flaw, but it does expose to the IAKERB proxy all information sent in cleartext over those exchanges, such as the principal names in requests. Since the typical usage involves the client obtaining a service ticket for the service operating the proxy, which will receive the client principal as part of normal authentication, this is also not a serious concern. However, an IAKERB client not using an armored FAST channel [RFC6113] sends an AS_REQ with pre-authentication data encrypted in the long-term keys of the user, even before the acceptor is authenticated. This subjects the user's long-term key to an offline attack by the proxy. To mitigate this threat, the client SHOULD use FAST [RFC6113] and its KDC authentication facility to protect the user's credentials.

Similarly, the client principal name is in cleartext in the AS and TGS exchanges, whereas in the AP exchanges embedded in GSS context tokens for the regular krb5 mechanism, the client principal name is present only in encrypted form. Thus, more information is exposed over the network path between initiator and acceptor when IAKERB is used than when the krb5 mechanism is used, unless FAST armor is employed. (This information would be exposed in other traffic from the initiator when the krb5 mech is used.) As such, to complicate traffic analysis and provide privacy for the client, the client SHOULD request the anonymous Kerberos FAST option [RFC6113].

Similar to other network access protocols, IAKERB allows an unauthenticated client (possibly outside the security perimeter of an organization) to send messages that are proxied to servers inside the

perimeter. To reduce the attack surface, firewall filters can be applied to restrict from which hosts client requests can be proxied, and the proxy can further restrict the set of realms to which requests can be proxied.

In the intended use scenario, the client uses the proxy to obtain a TGT and then a service ticket for the service it is authenticating to (possibly preceeded by exchanges to produce FAST armor). However, the protocol allows arbitrary KDC-REQs to be passed through, and there is no limit to the number of exchanges that may be proxied. The client can send KDC-REQs unrelated to the current authentication, and obtain service tickets for other service principals in the database of the KDC being contacted.

In a scenario where DNS SRV RR's are being used to locate the KDC, IAKERB is being used, and an external attacker can modify DNS responses to the IAKERB proxy, there are several countermeasures to prevent arbitrary messages from being sent to internal servers:

1. KDC port numbers can be statically configured on the IAKERB proxy. In this case, the messages will always be sent to KDC's. For an organization that runs KDC's on a static port (usually port 88) and does not run any other servers on the same port, this countermeasure would be easy to administer and should be effective.
2. The proxy can do application level sanity checking and filtering. This countermeasure should eliminate many of the above attacks.
3. DNS security can be deployed. This countermeasure is probably overkill for this particular problem, but if an organization has already deployed DNS security for other reasons, then it might make sense to leverage it here. Note that Kerberos could be used to protect the DNS exchanges. The initial DNS SRV KDC lookup by the proxy will be unprotected, but an attack here is at most a denial of service (the initial lookup will be for the proxy's KDC to facilitate Kerberos protection of subsequent DNS exchanges between itself and the DNS server).

7. Acknowledgements

Jonathan Trostle, Michael Swift, Bernard Aboba and Glen Zorn wrote earlier revision of this document.

The hallway conversations between Larry Zhu and Nicolas Williams formed the basis of this document.

8. Assigned Numbers

The value for the error code KRB_AP_ERR_IAKERB_KDC_NOT_FOUND is 85.

The value for the error code KRB_AP_ERR_IAKERB_KDC_NO_RESPONSE is 86.

The key usage number KEY_USAGE_FINISHED is 41.

The key usage number KEY_USAGE_IAKERB_FINISHED is 42.

9. IANA Considerations

IANA is requested to make a modification in the "Kerberos GSS-API Token Type Identifiers" registry.

The following data to the table:

ID	Description	Reference
05 01	IAKERB_PROXY	[THIS RFC]

10. References

10.1. Normative References

- [CCITT.X680.2002]
International Telephone and Telegraph Consultative Committee, "Abstract Syntax Notation One (ASN.1): Specification of basic notation", CCITT Recommendation X.680, July 2002.
- [CCITT.X690.2002]
International Telephone and Telegraph Consultative Committee, "ASN.1 encoding rules: Specification of basic encoding Rules (BER), Canonical encoding rules (CER) and Distinguished encoding rules (DER)", CCITT Recommendation X.690, July 2002.
- [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>>.

- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, DOI 10.17487/RFC2743, January 2000, <<http://www.rfc-editor.org/info/rfc2743>>.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, DOI 10.17487/RFC3961, February 2005, <<http://www.rfc-editor.org/info/rfc3961>>.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, DOI 10.17487/RFC4120, July 2005, <<http://www.rfc-editor.org/info/rfc4120>>.
- [RFC4121] Zhu, L., Jaganathan, K., and S. Hartman, "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2", RFC 4121, DOI 10.17487/RFC4121, July 2005, <<http://www.rfc-editor.org/info/rfc4121>>.
- [RFC6542] Emery, S., "Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Channel Binding Hash Agility", RFC 6542, DOI 10.17487/RFC6542, March 2012, <<http://www.rfc-editor.org/info/rfc6542>>.

10.2. Informative references

- [RFC6112] Zhu, L., Leach, P., and S. Hartman, "Anonymity Support for Kerberos", RFC 6112, DOI 10.17487/RFC6112, April 2011, <<http://www.rfc-editor.org/info/rfc6112>>.
- [RFC6113] Hartman, S. and L. Zhu, "A Generalized Framework for Kerberos Pre-Authentication", RFC 6113, DOI 10.17487/RFC6113, April 2011, <<http://www.rfc-editor.org/info/rfc6113>>.
- [RFC6806] Hartman, S., Ed., Raeburn, K., and L. Zhu, "Kerberos Principal Name Canonicalization and Cross-Realm Referrals", RFC 6806, DOI 10.17487/RFC6806, November 2012, <<http://www.rfc-editor.org/info/rfc6806>>.

Appendix A. Interoperate with Previous MIT version

MIT implemented an early draft version of this document. This section gives a method for detecting and interoperating with that version.

Initiators behave as follows:

- o If the first acceptor token begins with generic token framing as described in section 3.1 of [RFC2743], then use the protocol as defined in this document.
- o If the first acceptor token is missing the generic token framing (i.e., the token begins with the two-byte token ID 05 01), then
 - * When creating the finish message, the value of one (1) should be used in place of GSS_EXTS_FINISHED.
 - * When computing the checksum, the value of KEY_USAGE_IAKERB_FINISHED should be used in place of KEY_USAGE_FINISHED.

KEY_USAGE_IAKERB_FINISHED

42

Acceptors behave as follows:

- o After the first initiator token, allow initiator tokens to omit generic token framing. This allowance is required only for IAKERB_PROXY messages (those using token ID 05 01), not for tokens defined in [RFC4121].
- o If the AP-REQ authenticator contains an extension of type 1 containing a KRB-FINISHED message, then process the extension as if it were of type GSS_EXTS_FINISHED, except with a key usage of KEY_USAGE_IAKERB_FINISHED (42) instead of KEY_USAGE_FINISHED (41).

Authors' Addresses

Benjamin Kaduk (editor)
Akamai

Email: kaduk@mit.edu

Jim Schaad (editor)
Soaring Hawk Consulting

Email: ietf@augustcellars.com

Larry Zhu
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
US

Email: lzhu@microsoft.com

Jeffery Altman
Secure Endpoints
255 W 94th St
New York, NY 10025
US

Email: jaltman@secure-endpoints.com

Network Working Group
Internet-Draft
Updates: rfc4120 (if approved)
Intended status: Standards Track
Expires: October 1, 2017

T. Yu
MIT Kerberos Consortium
March 30, 2017

Move Kerberos protocol parameter registries to IANA
draft-ietf-kitten-kerberos-iana-registries-04

Abstract

The Keberos 5 network authentication protocol has several numeric protocol parameters. Most of these parameters are not currently under IANA maintenance. This document requests that IANA take over the maintenance of the remainder of these Kerberos parameters.

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 October 1, 2017.

Copyright Notice

Copyright (c) 2017 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.

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. Introduction

The Kerberos 5 network authentication protocol [RFC4120] [RFC1510] has several numeric protocol parameters. This document requests that IANA take over the maintenance of the Kerberos protocol parameters that are not currently under IANA maintenance. Several instances of number conflicts in Kerberos implementations could have been prevented by having IANA registries for those numbers. This document updates [RFC4120].

3. General registry format

Unless otherwise specified, each Kerberos protocol number registry will have the following fields: "number", "name", "reference", and "comments".

The name must begin with a lowercase letter, and must consist of ASCII letters, digits, and hyphens. Two or more hyphens must not appear directly adjacent to each other. A hyphen must not appear at the end of a name. It is preferred that words in a name be separated by hyphens, and that all of the letters be lowercase.

(These rules are consistent with the lexical rules for an ASN.1 valuereference or identifier. Where the constraints are stricter than the ASN.1 lexical rules, they make it easier to systematically transform the names for use in implementation languages.)

Names for numeric parameter values have no inherent meaning in the Kerberos protocol, but they can guide choices for internal implementation symbol names and for user-visible non-numeric representations. When written in English prose in specifications, or when used as symbolic constants in implementation languages (e.g., C preprocessor macros), it is common to transform the name into all uppercase letters, and possibly to replace hyphens with underscores.

4. General registration procedure

This document requests that the IESG establish a pool of Kerberos experts who will manage the Kerberos registries using these guidelines. The IESG may wish to consider including the set of designated IANA experts for existing Kerberos IANA registries as candidates for this pool.

IANA will select an expert from this pool for each registration request. The expert will review the registration request and may approve the registration, decline the registration with comments, or recommend that the registration request should follow a specific alternative process. The alternative processes that the expert may recommend are the IETF review process and the standards action process.

Initially, the expert reviewers will use a permissive process, generally approving registrations that are architecturally consistent with Kerberos and the protocol parameter in question. Over time, with input from the community, the experts may refine the requirements that registrations are expected to meet. The experts will maintain a current version of these guidelines in a manner that is generally accessible to the entire community. As the guidelines evolve, experts may consider the technical quality of specifications, security impacts of the registrations, architectural consistency, and interoperability impact. Experts may require a publicly available specification in order to make certain registrations.

[For the individual registries, include "Registrations in this registry are managed by the expert review process [RFC5226] or in exceptional cases by IESG approval. See section x for guidelines for the experts to be used with this registry."]

5. Integer assignments

Names for integer assignments must be unique across all Kerberos integer parameter registries. This is normally accomplished by including a name prefix that identifies the registry.

Assignments for integers parameters will follow the general registration procedure outlined above, except as otherwise noted in the section that contains the description of the parameter. Kerberos integer parameters take on signed 32-bit values (-2147483648 to 2147483647). Negative values are for private or local use.

5.1. Address types

Registry name: Address types

Assignment policy: General registration procedure

Valid values: Signed 32-bit integers

Address types historically align with numeric constants used in the Berkeley sockets API. Future address type assignments should conform

to this historical practice when possible. The name prefix for address types is "addrtype-".

5.2. Authorization data types

Registry name: Authorization data types

Assignment policy: General registration procedure

Valid values: Signed 32-bit integers

The name prefix for authorization data types is "ad-".

5.3. Error codes

Registry name: Error codes

Assignment policy: Standards action

Valid values: Signed 32-bit integers

Assignments for error codes require standards action due to their scarcity: assigning error codes greater than 127 could require significant changes to certain implementations. The name prefixes for error codes are "kdc-err-", "krb-err-", and "krb-ap-err-".

5.4. Key usages

Registry name: Key usages

Assignment policy: General registration procedure

Valid values: Unsigned 32-bit integers

Key usages are unsigned 32-bit integers (0 to 4294967295). Zero is reserved and may not be assigned.

The name prefix for key usages is "ku-".

5.5. Name types

Registry name: Name types

Assignment policy: General registration procedure

Valid values: Signed 32-bit integers

The name prefix for name types is "nt-".

number	name	reference	comment
0	nt-unknown	RFC4120	Name type not known
1	nt-principal	RFC4120	Just the name of the principal as in DCE, or for users
2	nt-srv-inst	RFC4120	Service and other unique instance (krbtgt)
3	nt-srv-hst	RFC4120	Service with host name as instance (telnet, rcommands)
4	nt-srv-xhst	RFC4120	Service with host as remaining components
5	nt-uid	RFC4120	Unique ID
6	nt-x500-principal	RFC4120	Encoded X.509 Distinguished name [RFC2253]
7	nt-smtp-name	RFC4120	Name in form of SMTP email name (e.g., user@example.com)
10	nt-enterprise	RFC4120	Enterprise name - may be mapped to principal name
11	nt-wellknown	RFC6111	Well-known principal name
12	nt-srv-hst-domain	RFC5179	Domain-based names

5.6. Pre-authentication and typed data

Registry name: Pre-authentication and typed data

Assignment policy: General registration procedure

Valid values: Signed 32-bit integers

This document requests that IANA modify the existing Kerberos Pre-authentication and typed data registry to be consistent with the procedures in this document.

The name prefix for pre-authentication type numbers is "pa-". The name prefix for typed data numbers is "td-". Pre-authentication and typed data numbers are in the same registry, but a pre-authentication number may be also be assigned to a related typed data number.

6. Named bit assignments

Assignments for named bits require standards action, due to their scarcity: assigning bit numbers greater than 31 could require significant changes to implementations. Names for named bit assignments must be unique within a given named bit registry, and typically do not have name prefixes that identify which registry they belong to.

6.1. AP-REQ options

Registry name: AP-REQ options
Assignment policy: Standards action
Valid values: ASN.1 bit numbers 0 through 31

6.2. KDC-REQ options

Registry name: KDC-REQ options
Assignment policy: Standards action
Valid values: ASN.1 bit numbers 0 through 31

6.3. Ticket flags

Registry name: Ticket flags
Assignment policy: Standards action
Valid values: ASN.1 bit numbers 0 through 31

7. Numbers that will not be registered

ASN.1 application tag numbers (which are always equal to the "msg-type" field in Kerberos messages where they appear) will not be registered. Any Kerberos protocol change that requires a new application tag number will be a sufficiently major change that the specification of the change MUST define a new ASN.1 module and MUST be Standards Track.

Transited encoding values will not be registered. There is only one transited encoding type for the Kerberos protocol. The interoperability concerns inherent to the cross-realm operation of Kerberos mean that specifications of new transited encoding types are very unlikely. Any specification of new transited encoding types MUST be Standards Action.

Protocol version number (pvno) values will not be registered. The location of the "pvno" value in Kerberos messages is not in a place that implementations can meaningfully use to distinguish among different variants of the Kerberos protocol.

8. Contributors

Sam Hartman proposed the text of the expert review guidelines. Love Hornquist Astrand wrote a previous document (draft-lha-krb-wg-some-numbers-to-iana-00) with the same goals as this document.

9. Acknowledgments

Thanks to Tom Petch for providing useful feedback on previous versions of this document.

10. Security Considerations

Assignments of new Kerberos protocol parameter values can have security implications. In cases where the assignment policy calls for expert review, the reviewer is responsible for evaluating whether adequate documentation exists concerning the security considerations for the requested assignment. For assignments that require IETF review or standards action, the normal IETF processes ensure adequate treatment of security considerations.

11. IANA Considerations

This document requests that IANA create several registries for Kerberos protocol parameters:

- o Address types
- o Authorization data types
- o Error codes
- o Key usages
- o Name types
- o AP-REQ options
- o KDC-REQ options
- o Ticket flags

This document requests that IANA modify the existing "Pre-authentication data and typed data" registry to contain an additional reference to this document, and to transform existing names in that registry to the lowercase-and-hyphens style.

12. Open issues

Do we make a registry for application tag numbers (equal to message type numbers)? We've said that we would replace the entire ASN.1 module in that case, but Nico's recent proposal doesn't do that, and if we want to accommodate that sort of proposal, it would probably be best to establish a registry. (It should require standards action for registrations.)

Do transited encodings need a registry? They would probably require standards action, even if there were a registry.

13. References

13.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, DOI 10.17487/RFC3961, February 2005, <<http://www.rfc-editor.org/info/rfc3961>>.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, DOI 10.17487/RFC4120, July 2005, <<http://www.rfc-editor.org/info/rfc4120>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

13.2. Informative References

- [RFC1510] Kohl, J. and C. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, DOI 10.17487/RFC1510, September 1993, <<http://www.rfc-editor.org/info/rfc1510>>.

Author's Address

Tom Yu
MIT Kerberos Consortium
77 Massachusetts Ave
Cambridge, Massachusetts
USA

Email: tlyu@mit.edu

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 11, 2021

S. Cantor
Shibboleth Consortium
M. Cullen
Painless Security
S. Josefsson
SJD AB
May 10, 2021

SAML Enhanced Client SASL and GSS-API Mechanisms
draft-ietf-kitten-sasl-saml-ec-20

Abstract

Security Assertion Markup Language (SAML) 2.0 is a generalized framework for the exchange of security-related information between asserting and relying parties. Simple Authentication and Security Layer (SASL) and the Generic Security Service Application Program Interface (GSS-API) are application frameworks that facilitate an extensible authentication model, among other things. This document specifies a SASL and GSS-API mechanism for SAML 2.0 that leverages the capabilities of a SAML-aware "enhanced client" to address significant barriers to federated authentication in a manner that encourages reuse of existing SAML bindings and profiles designed for non-browser scenarios.

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 11, 2021.

Copyright Notice

Copyright (c) 2021 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
2. Terminology	5
3. Applicability for Non-HTTP Use Cases	6
4. SAML Enhanced Client SASL Mechanism Specification	8
4.1. Advertisement	8
4.2. Initiation	9
4.3. Server Response	9
4.4. User Authentication with Identity Provider	10
4.5. Client Response	10
4.6. Outcome	10
4.7. Additional Notes	10
5. SAML EC GSS-API Mechanism Specification	11
5.1. GSS-API Credential Delegation	12
5.2. GSS-API Channel Binding	13
5.3. Session Key Derivation	13
5.3.1. Generated by Identity Provider	14
5.3.2. Alternate Key Derivation Mechanisms	15
5.4. Per-Message Tokens	15
5.5. Pseudo-Random Function (PRF)	15
5.6. GSS-API Principal Name Types for SAML EC	16
5.6.1. User Naming Considerations	16
5.6.2. Service Naming Considerations	17
6. Example	17
7. Security Considerations	25
7.1. Risks Left Unaddressed	26
7.2. User Privacy	26
7.3. Collusion between RPs	27
8. IANA Considerations	27
8.1. GSS-API and SASL Mechanism Registration	27
8.2. XML Namespace Name for SAML-EC	27
9. References	28
9.1. Normative References	28
9.2. Informative References	30
Appendix A. XML Schema	31
Appendix B. Acknowledgments	33
Appendix C. Changes	33

Authors' Addresses	34
--------------------	----

1. Introduction

Security Assertion Markup Language (SAML) 2.0

[OASIS.saml-core-2.0-os] is a modular specification that provides various means for a user to be identified to a relying party (RP) through the exchange of (typically signed) assertions issued by an identity provider (IdP).

Simple Authentication and Security Layer (SASL) [RFC4422] is a generalized mechanism for identifying and authenticating a user and for optionally negotiating a security layer for subsequent protocol interactions. SASL is used by application protocols like IMAP [RFC3501], the Post Office Protocol (POP [RFC1939]) and XMPP [RFC6120]. The effect of SASL is to make authentication modular, so that newer authentication mechanisms can be added as needed.

There are related protocols, protocol bindings [OASIS.saml-bindings-2.0-os], and interoperability profiles [OASIS.saml-profiles-2.0-os] designed for different use cases. Additional profiles and extensions are also routinely developed and published.

The Generic Security Service Application Program Interface (GSS-API) [RFC2743] provides a framework for applications to support multiple authentication mechanisms through a unified programming interface, as well as additional optional cryptographic functionality. This document defines a pure SASL mechanism for SAML, but it conforms to the bridge between SASL and GSS-API called GS2 [RFC5801]. This means that this document defines both a SASL mechanism and a GSS-API mechanism. The GSS-API interface is optional for SASL implementers, and the GSS-API considerations can be avoided in environments that use SASL directly without GSS-API.

The mechanisms specified in this document allow a SASL- or GSS-API-enabled server to act as a SAML relying party, or service provider (SP), by advertising this mechanism as an option for SASL or GSS-API clients that support the use of SAML to communicate identity and attribute information. Clients supporting this mechanism are termed "enhanced clients" in SAML terminology because they understand the federated authentication model and have specific knowledge of the IdP(s) associated with the user. This knowledge, and the ability to act on it, addresses a significant problem with browser-based SAML profiles known as the "discovery", or "where are you from?" (WAYF) problem. In a "dumb" client such as a web browser, various intrusive user interface techniques are used to determine the appropriate IdP to use because the request to the IdP is generated as an HTTP

redirect by the RP, which does not generally have prior knowledge of the IdP to use. Obviating the need for the RP to interact with the client to determine the right IdP (and its network location) is both a user interface and security improvement.

The SAML mechanism described in this document is an adaptation of an existing SAML profile, the Enhanced Client or Proxy (ECP) Profile (V2.0) [SAMLECP20].

Figure 1 describes the interworking between SAML and SASL: this document requires enhancements to the RP and to the client (as the two SASL communication endpoints) but no changes to the SAML IdP are assumed apart from its support for the applicable SAML profile. To accomplish this, a SAML protocol exchange between the RP and the IdP, brokered by the client, is tunneled within SASL. There is no assumed communication between the RP and the IdP, but such communication may occur in conjunction with additional SAML-related profiles not in scope for this document.




Figure 1: Interworking Architecture




Figure 2: Communication Flow

2. Terminology

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 BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

The reader is also assumed to be familiar with the terms used in the SAML 2.0 specification, and an understanding of the Enhanced Client or Proxy (ECP) Profile (V2.0) [SAMLECP20] is necessary, as part of this mechanism explicitly reuses and references it.

This document can be implemented without knowledge of GSS-API since the normative aspects of the GS2 protocol syntax have been duplicated in this document. The document may also be implemented to provide a GSS-API mechanism, and then knowledge of GSS-API is essential.

3. Applicability for Non-HTTP Use Cases

While SAML is designed to support a variety of application scenarios, the profiles for authentication defined in the original standard are designed around HTTP [RFC7230] applications. They are not, however, limited to browsers, because browsers do not always meet the needs of more security-sensitive applications. Specifically, the notion of an "Enhanced Client" (or a proxy acting as one on behalf of a browser, thus the term "ECP") was specified for a software component that acts somewhat like a browser from an application perspective, but includes limited, but sufficient, awareness of SAML to play a more conscious role in the authentication exchange between the RP and the IdP. What follows is an outline of the Enhanced Client or Proxy (ECP) Profile (V2.0) [SAMLECP20], as applied to the web/HTTP service use case:

1. The Enhanced Client requests a resource of a Relying Party (RP) (via an HTTP request). In doing so, it advertises its "enhanced" capability using HTTP headers.
2. The RP, desiring SAML authentication and noting the client's capabilities, responds not with an HTTP redirect or form, but with a SOAP [W3C.soap11] envelope containing a SAML <AuthnRequest> along with some supporting headers. This request identifies the RP (and may be signed), and may provide hints to the client as to what IdPs the RP finds acceptable, but the choice of IdP is generally left to the client.
3. The client is then responsible for delivering the body of the SOAP message to the IdP it is instructed to use (often via out-of-band configuration). The user authenticates to the IdP ahead of, during, or after the delivery of this message, and perhaps explicitly authorizes the response to the RP.
4. Whether authentication succeeds or fails, the IdP responds with its own SOAP envelope, generally containing a SAML <Response> message for delivery to the RP. In a successful case, the message will include one or more SAML <Assertion> elements containing authentication, and possibly attribute, statements about the subject. Either the response or each assertion is signed, and the assertion(s) may be encrypted to a key negotiated with or known to belong to the RP.
5. The client then delivers the SOAP envelope containing the <Response> to the RP at a location the IdP directs (which acts as an additional, though limited, defense against MITM attacks). This completes the SAML exchange.

6. The RP now has sufficient identity information to approve the original HTTP request or not, and acts accordingly. Everything between the original request and this response can be thought of as an "interruption" of the original HTTP exchange.

When considering this flow in the context of an arbitrary application protocol and SASL, the RP and the client both must change their code to implement this SASL mechanism, but the IdP can remain unmodified. The existing RP/client exchange that is tunneled through HTTP maps well to the tunneling of that same exchange in SASL. In the parlance of SASL [RFC4422], this mechanism is "client-first" for consistency with GS2. The steps are shown below:

1. The server MAY advertise the SAML20EC and/or SAML20EC-PLUS mechanisms.
2. The client initiates a SASL authentication with SAML20EC or SAML20EC-PLUS.
3. The server sends the client a challenge consisting of a SOAP envelope containing its SAML <AuthnRequest>.
4. The SASL client unpacks the SOAP message and communicates with its chosen IdP to relay the SAML <AuthnRequest> to it. This communication, and the authentication with the IdP, proceeds separately from the SASL process.
5. Upon completion of the exchange with the IdP, the client responds to the SASL server with a SOAP envelope containing the SAML <Response> it obtained, or a SOAP fault, as warranted.
6. The SASL Server indicates success or failure.

Note: The details of the SAML processing, which are consistent with the Enhanced Client or Proxy (ECP) Profile (V2.0) [SAMLECP20], are such that the client MUST interact with the IdP in order to complete any SASL exchange with the RP. The assertions issued by the IdP for the purposes of the profile, and by extension this SASL mechanism, are short lived, and therefore cannot be cached by the client for later use.

Encompassed in step four is the client-driven selection of the IdP, authentication to it, and the acquisition of a response to provide to the SASL server. These processes are all external to SASL.

Note also that unlike an HTTP-based profile, the IdP cannot participate in the selection of, or evaluation of, the location to which the SASL Client Response will be delivered by the client. The

use of GSS-API Channel Binding is an important mitigation of the risk of a "Man in the Middle" attack between the client and RP, as is the use of a negotiated or derived session key in whatever protocol is secured by this mechanism.

With all of this in mind, the typical flow appears as follows:




Figure 3: Authentication flow

4. SAML Enhanced Client SASL Mechanism Specification

Based on the previous figures, the following operations are defined by the SAML SASL mechanism:

4.1. Advertisement

To advertise that a server supports this mechanism, during application session initiation, it displays the name "SAML20EC" and/or "SAML20EC-PLUS" in the list of supported SASL mechanisms.

In accordance with [RFC5801] the "-PLUS" variant indicates that the server supports channel binding and would be selected by a client with that capability.

4.2. Initiation

A client initiates "SAML20EC" or "SAML20EC-PLUS" authentication. If supported by the application protocol, the client MAY include an initial response, otherwise it waits until the server has issued an empty challenge (because the mechanism is client-first).

The format of the initial client response ("initresp") is as follows:

```
hok = "urn:oasis:names:tc:SAML:2.0:cm:holder-of-key"
```

```
mut = "urn:oasis:names:tc:SAML:2.0:profiles:SSO:ecp:2.0:" \
      "WantAuthnRequestsSigned"
```

```
del = "urn:oasis:names:tc:SAML:2.0:conditions:delegation"
```

```
initresp = gs2-cb-flag "," [gs2-authzid] "," [hok] "," [mut] "," [del]
```

The gs2-cb-flag flag MUST be set as defined in [RFC5801] to indicate whether the client supports channel binding. This takes the place of the PAOS HTTP header extension used in [SAMLECP20] to indicate channel binding support.

The optional "gs2-authzid" field holds the authorization identity, as requested by the client.

The optional "hok" field is a constant that signals the client's support for stronger security by means of a locally held key. This takes the place of the PAOS HTTP header extension used in [SAMLECP20] to indicate "holder of key" support.

The optional "mut" field is a constant that signals the client's desire for mutual authentication. If set, the SASL server MUST digitally sign its SAML <AuthnRequest> message. The URN constant above is a single string; the linefeed is shown for RFC formatting reasons.

The optional "del" field is a constant that signals the client's desire for the acceptor to request an assertion usable for delegation of the client's identity to the acceptor.

4.3. Server Response

The SASL server responds with a SOAP envelope constructed in accordance with section 2.3.2 of [SAMLECP20]. This includes adhering to the SOAP header requirements of the SAML PAOS Binding [OASIS.saml-bindings-2.0-os], for compatibility with the existing

profile. Various SOAP headers are also consumed by the client in exactly the same manner prescribed by that section.

4.4. User Authentication with Identity Provider

Upon receipt of the Server Response (Section 4.3), the steps described in sections 2.3.3 through 2.3.6 of [SAMLECP20] are performed between the client and the chosen IdP. The means by which the client determines the IdP to use, and where it is located, are out of scope of this mechanism.

The exact means of authentication to the IdP are also out of scope, but clients supporting this mechanism MUST support HTTP Basic Authentication as defined in [RFC7617] and TLS 1.3 client authentication as defined in [RFC8446].

4.5. Client Response

Assuming a response is obtained from the IdP, the client responds to the SASL server with a SOAP envelope constructed in accordance with section 2.3.7 of [SAMLECP20]. This includes adhering to the SOAP header requirements of the SAML PAOS Binding [OASIS.saml-bindings-2.0-os], for compatibility with the existing profile. If the client is unable to obtain a response from the IdP, or must otherwise signal failure, it responds to the SASL server with a SOAP envelope containing a SOAP fault.

4.6. Outcome

The SAML protocol exchange having completed, the SASL server will transmit the outcome to the client depending on local validation of the client responses (including the assertion conveyed from the chosen IDP). This outcome is transmitted in accordance with the application protocol in use.

4.7. Additional Notes

Because this mechanism is an adaptation of an HTTP-based profile, there are a few requirements outlined in [SAMLECP20] that make reference to a response URL that is normally used to regulate where the client returns information to the RP. There are also security-related checks built into the profile that involve this location.

For compatibility with existing IdP and profile behavior, and to provide for mutual authentication, the SASL server MUST populate the responseConsumerURL and AssertionConsumerServiceURL attributes with its service name. As discussed in Section 5.6.2, most SASL profiles rely on a service name format of "service@host", but regardless of

the form, the service name is used directly rather than transformed into an absolute URI if it is not already one, and MUST be percent-encoded per [RFC3986].

The IdP MUST securely associate the service name with the SAML entityID claimed by the SASL server, such as through the use of SAML metadata [OASIS.saml-metadata-2.0-os]. If metadata is used, a SASL service's <SPSSODescriptor> role MUST contain a corresponding <AssertionConsumerService> whose Location attribute contains the appropriate service name, as described above. The Binding attribute MUST be one of "urn:ietf:params:xml:ns:saml" (RECOMMENDED) or "urn:oasis:names:tc:SAML:2.0:bindings:PAOS" (for compatibility with older implementations of the ECP profile in existing IdP software).

Finally, note that the use of HTTP status signaling between the RP and client mandated by [SAML ECP 20] may not be applicable.

5. SAML EC GSS-API Mechanism Specification

This section and its sub-sections and all normative references of it not referenced elsewhere in this document are INFORMATIONAL for SASL implementors, but they are NORMATIVE for GSS-API implementors.

The SAML Enhanced Client SASL mechanism is also a GSS-API mechanism. The messages are the same, but a) the GS2 [RFC5801] header on the client's first authentication message is excluded when SAML EC is used as a GSS-API mechanism, and b) the [RFC2743] section 3.1 initial context token header is used for the client's first authentication message (context token) instead, with the body of the message being the same as for the SASL mechanism case.

The GSS-API mechanism OID for SAML EC is OID-TBD (IANA to assign: see IANA considerations). The DER encoding of the OID is TBD.

The mutual_state request flag (GSS_C_MUTUAL_FLAG) MAY be set to TRUE, resulting in the "mut" option set in the initial client response. The security context mutual_state flag is set to TRUE only if the server digitally signs its SAML <AuthnRequest> message and the signature and signing credential are appropriately verified by the IdP. The IdP signals this to the client in an <ecp:RequestAuthenticated> SOAP header block.

The lifetime of a security context established with this mechanism SHOULD be limited by the value of a SessionNotOnOrAfter attribute, if any, in the <AuthnStatement> element(s) of the SAML assertion(s) received by the RP. By convention, in the rare case that multiple valid/confirmed assertions containing <AuthnStatement> elements are

received, the most restrictive SessionNotOnOrAfter is generally applied.

5.1. GSS-API Credential Delegation

This mechanism can support credential delegation through the issuance of SAML assertions that an IdP will accept as proof of authentication by a service on behalf of a subject. An initiator may request delegation of its credentials by setting the "del" option field in the initial client response to "urn:oasis:names:tc:SAML:2.0:conditions:delegation".

An acceptor, upon receipt of this constant, requests a delegated assertion by including in its <AuthnRequest> message a <Conditions> element containing an <AudienceRestriction> identifying the IdP as a desired audience for the assertion(s) to be issued. In the event that the specific IdP to be used is unknown, the constant "urn:oasis:names:tc:SAML:2.0:conditions:delegation" may be used as a stand-in, per Section 2.3.2 of [SAMLECP20].

Upon receipt of an assertion satisfying this property, and containing a <SubjectConfirmation> element that the acceptor can satisfy, the security context will have its deleg_state flag (GSS_C_DELEG_FLAG) set to TRUE.

The IdP, if it issues a delegated assertion to the acceptor, MUST include in the SOAP response to the initiator a <samlec:Delegated> SOAP header block, indicating that delegation was enabled. It has no content, other than mandatory SOAP attributes (an example follows):

```
<samlec:Delegated xmlns:samlec="urn:ietf:params:xml:ns:samlec"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/"
  S:mustUnderstand="1"
  S:actor="http://schemas.xmlsoap.org/soap/actor/next" />
```

Upon receipt of such a header block, the initiator MUST fail the establishment of the security context if it did not request delegation in its initial client response to the acceptor. It SHOULD signal this failure to the acceptor with a SOAP fault message in its final client response.

As noted previously, the exact means of client authentication to the IdP is formally out of scope of this mechanism. This extends to the use of a delegation assertion as a means of authentication by an acceptor acting as an initiator. In practice, some profile of

[WSS-SAML] is used to attach the assertion and a confirmation proof to the SOAP message from the client to the IdP.

5.2. GSS-API Channel Binding

GSS-API channel binding [RFC5554] is a protected facility for exchanging a cryptographic identifier for an enclosing channel between the initiator and acceptor. The initiator sends channel binding data and the acceptor confirms that channel binding data has been checked.

The acceptor SHOULD accept any channel binding provided by the initiator if null channel bindings are passed into `gss_accept_sec_context`. Protocols such as HTTP Negotiate [RFC4559] depend on this behavior of some Kerberos implementations.

The exchange and verification of channel binding information is described by [SAMLECP20].

5.3. Session Key Derivation

Some GSS-API features (discussed in the following sections) require a session key be established as a result of security context establishment. In the common case of a "bearer" assertion in SAML, a mechanism is defined to communicate a key to both parties via the IdP. In other cases such as assertions based on "holder of key" confirmation bound to a client-controlled key, there may be additional methods defined in the future, and extension points are provided for this purpose.

Information defining or describing the session key, or a process for deriving one, is communicated between the initiator and acceptor using a `<samlec:SessionKey>` element, defined by the XML schema in Appendix A. This element is a SOAP header block. The content of the element further depends on the specific use in the mechanism. The Algorithm XML attribute identifies a mechanism for key derivation. It is omitted to identify the use of an IdP-generated key (see following section) or will contain a URI value identifying a derivation mechanism defined outside this specification. Each header block's `mustUnderstand` and `actor` attributes MUST be set to "1" and "`http://schemas.xmlsoap.org/soap/actor/next`" respectively.

In the acceptor's first response message containing its SAML request, one or more `<samlec:SessionKey>` SOAP header blocks MUST be included. The element MUST contain one or more `<EncType>` elements containing the number of a supported encryption type defined in accordance with [RFC3961]. Encryption types should be provided in order of preference by the acceptor.

In the final client response message, a single <samlec:SessionKey> SOAP header block MUST be included. A single <EncType> element MUST be included to identify the chosen encryption type used by the initiator.

All parties MUST support the "aes128-cts-hmac-sha1-96" encryption type, number 17, defined by [RFC3962].

Further details depend on the mechanism used, one of which is described in the following section.

5.3.1. Generated by Identity Provider

The IdP, if issuing a bearer assertion for use with this mechanism, SHOULD provide a generated key for use by the initiator and acceptor. This key is used as pseudorandom input to the "random-to-key" function for a specific encryption type defined in accordance with [RFC3961]. The key is base64-encoded and placed inside a <samlec:GeneratedKey> element. The IdP does not participate in the selection of the encryption type and simply generates enough pseudorandom bits to supply key material to the other parties.

The resulting <samlec:GeneratedKey> element is placed within the <saml:Advice> element of the assertion issued. The identity provider MUST encrypt the assertion (implying that it MUST have the means to do so, typically knowledge of a key associated with the RP). If multiple assertions are issued (allowed, but not typical), the element need only be included in one of the assertions issued for use by the relying party.

A copy of the element is also added as a SOAP header block in the response from the IdP to the client (and then removed when constructing the response to the acceptor).

If this mechanism is used by the initiator, then the <samlec:SessionKey> SOAP header block attached to the final client response message will identify this via the omission of the Algorithm attribute and will identify the chosen encryption type using the <samlec:EncType> element:

```
<samlec:SessionKey xmlns:samlec="urn:ietf:params:xml:ns:samlec"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/"
  S:mustUnderstand="1"
  S:actor="http://schemas.xmlsoap.org/soap/actor/next">
  <samlec:EncType>17</samlec:EncType>
</samlec:SessionKey>
```

Both the initiator and acceptor MUST execute the chosen encryption type's random-to-key function over the pseudorandom value provided by the <samlec:GeneratedKey> element. The result of that function is used as the protocol and session key. Support for subkeys from the initiator or acceptor is not specified.

5.3.2. Alternate Key Derivation Mechanisms

In the event that a client is proving possession of a secret or private key, a formal key agreement algorithm might be supported. This specification does not define such a mechanism, but the <samlec:SessionKey> element is extensible to allow for future work in this space by means of the Algorithm attribute and an optional <ds:KeyInfo> child element to carry extensible content related to key establishment.

However a key is derived, the <samlec:EncType> element will identify the chosen encryption type, and both the initiator and acceptor MUST execute the encryption type's random-to-key function over the result of the key agreement or derivation process. The result of that function is used as the protocol key.

5.4. Per-Message Tokens

The per-message tokens SHALL be the same as those for the Kerberos V5 GSS-API mechanism [RFC4121] (see Section 4.2 and sub-sections).

The replay_det_state (GSS_C_REPLAY_FLAG), sequence_state (GSS_C_SEQUENCE_FLAG), conf_avail (GSS_C_CONF_FLAG) and integ_avail (GSS_C_INTEG_FLAG) security context flags are always set to TRUE.

The "protocol key" SHALL be a key established in a manner described in the previous section. "Specific keys" are then derived as usual as described in Section 2 of [RFC4121], [RFC3961], and [RFC3962].

The terms "protocol key" and "specific key" are Kerberos V5 terms [RFC3961].

SAML20EC is PROT_READY as soon as the SAML response message has been seen.

5.5. Pseudo-Random Function (PRF)

The GSS-API has been extended with a Pseudo-Random Function (PRF) interface in [RFC4401]. The purpose is to enable applications to derive a cryptographic key from an established GSS-API security context. This section defines a GSS_Pseudo_random that is applicable for the SAML20EC GSS-API mechanism.

The `GSS_Pseudo_random()` [RFC4401] SHALL be the same as for the Kerberos V5 GSS-API mechanism [RFC7802]. There is no acceptor-asserted sub-session key, thus `GSS_C_PRF_KEY_FULL` and `GSS_C_PRF_KEY_PARTIAL` are equivalent. The protocol key to be used for the `GSS_Pseudo_random()` SHALL be the same as the key defined in the previous section.

5.6. GSS-API Principal Name Types for SAML EC

Services that act as SAML relying parties are typically identified by means of a URI called an "entityID". Clients that are named in the <Subject> element of a SAML assertion are typically identified by means of a <NameID> element, which is an extensible XML structure containing, at minimum, an element value that names the subject and a Format attribute.

In practice, a GSS-API client and server are unlikely to know in advance the name of the initiator as it will be expressed by the SAML IdP upon completion of authentication. It is also generally incorrect to assume that a particular acceptor name will directly map into a particular RP entityID, because there is often a layer of naming indirection between particular services on hosts and the identity of a relying party in SAML terms.

To avoid complexity, and avoid unnecessary use of XML within the naming layer, the SAML EC mechanism relies on the common/expected name types used for acceptors and initiators, `GSS_C_NT_HOSTBASED_SERVICE` and `GSS_C_NT_USER_NAME`. The mechanism provides for validation of the host-based service name in conjunction with the SAML exchange. It does not attempt to solve the problem of mapping between an initiator "username", the user's identity while authenticating to the IdP, and the information supplied by the IdP to the acceptor. These relationships must be managed through local policy at the initiator and acceptor.

SAML-based information associated with the initiator SHOULD be expressed to the acceptor using GSS-API naming extensions [RFC6680], in a similar manner to [RFC7056].

5.6.1. User Naming Considerations

The `GSS_C_NT_USER_NAME` form represents the name of an individual user. Clients often rely on this value to determine the appropriate credentials to use in authenticating to the IdP, and supply it to the server for use by the acceptor.

Upon successful completion of this mechanism, the server MUST construct the authenticated initiator name based on the <saml:NameID>

element in the assertion it successfully validated. The name is constructed as a UTF-8 string in the following form:

```
name = element-value "!" Format "!" NameQualifier  
      "!" SPNameQualifier "!" SPProvidedID
```

The "element-value" token refers to the content of the <saml:NameID> element. The other tokens refer to the identically named XML attributes defined for use with the element. If an attribute is not present, which is common, it is omitted (i.e., replaced with the empty string). The Format value is never omitted; if not present, the SAML-equivalent value of "urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified" is used.

Not all SAML assertions contain a <saml:NameID> element. In the event that no such element is present, including the exceptional cases of a <saml:BaseID> element or a <saml:EncryptedID> element that cannot be decrypted, the GSS_C_NT_ANONYMOUS name type MUST be used for the initiator name.

As noted in the previous section, it is expected that most applications able to rely on SAML authentication would make use of naming extensions to obtain additional information about the user based on the assertion. This is particularly true in the anonymous case, or in cases in which the SAML name is pseudonymous or transient in nature. The ability to express the SAML name in GSS_C_NT_USER_NAME form is intended for compatibility with applications that cannot make use of additional information.

5.6.2. Service Naming Considerations

The GSS_C_NT_HOSTBASED_SERVICE name form represents a service running on a host; it is textually represented as "service@host". This name form is required by most SASL profiles and is used by many existing applications that use the Kerberos GSS-API mechanism. As described in the SASL mechanism's Section 4.7, such a name is used directly by this mechanism as the effective AssertionConsumerService "location" associated with the service and applied in IdP verification of the request against the claimed SAML entityID.

6. Example

Suppose the user has an identity at the SAML IdP saml.example.org and a Jabber Identifier (jid) "somenode@example.com", and wishes to authenticate his XMPP connection to xmpp.example.com (and example.com and example.org have established a SAML-capable trust relationship). The authentication on the wire would then look something like the following:

Step 1: Client initiates stream to server:

```
<stream:stream xmlns='jabber:client'
xmlns:stream='http://etherx.jabber.org/streams'
to='example.com' version='1.0'>
```

Step 2: Server responds with a stream tag sent to client:

```
<stream:stream
xmlns='jabber:client' xmlns:stream='http://etherx.jabber.org/streams'
id='some_id' from='example.com' version='1.0'>
```

Step 3: Server informs client of available authentication mechanisms:

```
<stream:features>
  <mechanisms xmlns='urn:ietf:params:xml:ns:xmpp-sasl'>
    <mechanism>DIGEST-MD5</mechanism>
    <mechanism>PLAIN</mechanism>
    <mechanism>SAML20EC</mechanism>
  </mechanisms>
</stream:features>
```

Step 4: Client selects an authentication mechanism and sends the initial client response (it is base64 encoded as specified by the XMPP SASL protocol profile):

```
<auth xmlns='urn:ietf:params:xml:ns:xmpp-sasl' mechanism='SAML20EC'>
biwsLCw=
</auth>
```

The initial response is "n,,," which signals that channel binding is not used, there is no authorization identity, and the client does not support key-based confirmation, or want mutual authentication or delegation.

Step 5: Server sends a challenge to client in the form of a SOAP envelope containing its SAML <AuthnRequest>:


```
<challenge xmlns='urn:iETF:params:xml:ns:xmpp-sasl'>
PFM6RW52ZWxvcGUKICAgIHhtbG5zOnNhbWw9InVybjpvYXNpczpuYW1lczp0YzptQU1MOjIuMDphc3NlcnRpb24iCiAgICB4bWxuczpzYW1scD0idXJuOm9hc2lwZm5hbnVzZnRjOjNBtUw6Mi4wOnByb3RvY29sIgogICAgeGlSbnM6Uz0iaHR0cDovL3NjaGVtYXMueGlsc29hcC5vcmcvc29hcC9lb3lnZlg9wZS8iPgogIDxtOkhlYWRLcj4KICAgIDxwYW9zOlJlclXVlc3QgeGlSbnM6GFvcz0idXJuOmxpYmVydHk6cGFvczoymDAzLTAAIGogICAgICBTZXNZYWdlSUQ9ImMzYTdRM0GI5YzjkIIbTOml1c3RVbmRlcnNOYW5kPSIxIgogICAgICBTOMfjdG9yPSJodHRWOi8vc2NoZWlhcy54bWxzczFwLm9yZy9zb2FwL2FjdG9yL25leHQiCiAgICAgIAgIHJlc3BvbnlQ29uc3VtZXJVUkw9InhtCHBAeGlwcC5leGFtcGxlLnNvbSIKICAgICAgc2VydmljZT0idXJuOm9hc2lwZm5hbnVzZnRjOjNBtUw6Mi4wOnByb2ZpbGVzOlNTTzplyY3AiLz4KICAgIDxlyA6Um5xdWVzc2AoagICAgICB4bWxuczplyY3A9InVybjpvYXNpczpuYW1lczp0YzptQU1MOjIuMDpwcm9maWxlc3puTU086ZWNWigogICAgICBTOMfjdG9yPSJodHRWOi8vc2NoZWlhcy54bWxzczFwLm9yZy9zb2FwL2FjdG9yL25leHQiCiAgICAgIAgIFM6bXVzdFVuZGVyc3RhbmQ9IjeiIFByb3ZpZGVyTmFtZT0iSmFiYmVyIGF0IGV4YW1wbGUuY29tIj4KICAgICAgPHNhbw6SXnzdwvyPmh0dBHzOi8veGlwcC5leGFtcGxlLnNvbTwvc2FtbDpjC3NlZXI+CiAgICAgICA8L2VjcDpSZXF1ZXN0PgogICAgPHNhbwXlyzptZXNZaw9us2V5IHhtbG5zOnNhbWxlyz0idXJuOmllldGY6cGFyYW1zonhtbDpuc3pzYW1sZWMiCiAgICAgIAgIHhtbG5zOlM9Imh0dHA6Ly9zy2h1bWFwZLNhtbHNvYXAub3JnL3NvYXAavZW52ZWxvcGUvIggogICAgICBTOMl1c3RVbmRlcnNOYW5kPSIxIgogICAgICBTOMfjdG9yPSJodHRWOi8vc2NoZWlhcy54bWxzczFwLm9yZy9zb2FwL2FjdG9yL25leHQiPgogICAgICA8c2FtbGVjOKvuY1R5cGU+MTc8L3NhbWxlyzpfbmNUexBlPgogICAgICA8c2FtbGVjOKvuY1R5cGU+MTg8L3NhbWxlyzpfbmNUexBlPgogICAgPHNhbWxlyzptZXNZaw9us2V5PgogIDdvUZcpIZWFkfXZI+CiAgPFM6Uz0keT4KICAgIDxxY1slcdPdBDXRoblJlcXVlc3QKICAgICAgISuq9ImMzYTdRM0GI5YzjkIIbWZxJzaW9uPSIylClJaIElzc3Vls2V5ZGFudD0imJAwnNy0xmi0xmFYXMTozOTokNFoiCiAgICAgIEFzc2VydGlvbkNvb3N1bWVyu2VydmljZVVSTDD0ieGlwcEB4bXBwLmV4YW1wbGUuY29tIj4KICAgICAgPHNhbw6SXnzdwvyIHhtbG5zOnNhbWw9InVybjpvYXNpczpuYW1lczp0YzptQU1MOjIuMDphc3NlcnRpb24iPgogICAgICAgahR0cHM6Ly94bXBwLmV4YW1wbGUuY29tCiAgICAgIDwvc2FtbDpjC3NlZXI+CiAgICAgIDxxY1slcdPoYwl1SURqb2xpY3kgQWxsbc3dcmVhdGU9InRydwUiCiAgICAgICAgRm9ybWF0PSJlcm46b2FzaXM6bmFtZXNM6dGM6U0FNTDoYLja6bmFtZWlkLWZvcmlhdDpwZXJzaXN0ZW50Ii8+CiAgICAgIDxxY1slcdPdSZXF1ZXN0ZWRBdXRobKNvb3N1bWVydHk6cGFyaXNbbyoIjZXhhY3QiPgogICAgICAgPHNhbw6QXV0ag5Db250ZXh0Q2xhc3NSZWY+CicAgICAgICB1cm46b2FzaXM6bmFtZXNM6dGM6U0FNTDoYLja6YWM6Y2xhc3Nlc3puYXNzd29yZFByb3RlY3RlZFryYW5zcG9ydAoagICAgICAgPC9zyW1sOkFlldGhuQ29udGV4dEV4dENsYXNzUmVmPgogICAgICA8L3NhbWxwOlJlcXVlc3RlZEFlldGhuQ29udGV4dD4gcicAgICAgICA8L3NhbWxwOkFlldGhuUmVxdWVzdD4KICA8L1M6QM9ket4KPC9TOkVu dmVsb3BlPgo=
</challenge>
```

The Base64 [RFC4648] decoded envelope:

```
<S:Envelope
  xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
  xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/">
  <S:Header>
    <paos:Request xmlns:paos="urn:liberty:paos:2003-08"
      messageID="c3a4f8b9c2d" S:mustUnderstand="1"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next"
      responseConsumerURL="xmpp@xmpp.example.com"
      service="urn:oasis:names:tc:SAML:2.0:profiles:SSO:ecp"/>
    <ecp:Request
      xmlns:ecp="urn:oasis:names:tc:SAML:2.0:profiles:SSO:ecp"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next"
      S:mustUnderstand="1" ProviderName="Jabber at example.com">
      <saml:Issuer>https://xmpp.example.com</saml:Issuer>
    </ecp:Request>
    <samlec:SessionKey xmlns:samlec="urn:ietf:params:xml:ns:samlec"
      xmlns:S="http://schemas.xmlsoap.org/soap/envelope/"
      S:mustUnderstand="1"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next">
      <samlec:EncType>17</samlec:EncType>
      <samlec:EncType>18</samlec:EncType>
    </samlec:SessionKey>
  </S:Header>
  <S:Body>
    <samlp:AuthnRequest
      ID="c3a4f8b9c2d" Version="2.0" IssueInstant="2020-12-10T11:39:34Z"
      AssertionConsumerServiceURL="xmpp@xmpp.example.com">
      <saml:Issuer xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">
        https://xmpp.example.com
      </saml:Issuer>
      <samlp:NameIDPolicy AllowCreate="true"
        Format="urn:oasis:names:tc:SAML:2.0:nameid-format:persistent"/>
      <samlp:RequestedAuthnContext Comparison="exact">
        <saml:AuthnContextClassRef>
          urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport
        </saml:AuthnContextClassRef>
      </samlp:RequestedAuthnContext>
    </samlp:AuthnRequest>
  </S:Body>
</S:Envelope>
```

Step 5 (alt): Server returns error to client:

```
<failure xmlns='urn:ietf:params:xml:ns:xmpp-sasl'>
  <incorrect-encoding/>
</failure>
</stream:stream>
```

Step 6: Client relays the request to IdP in a SOAP message transmitted over HTTP (over TLS). The HTTP portion is not shown, so the use of Basic Authentication is assumed. The body of the SOAP envelope is exactly the same as received in the previous step.

```
<S:Envelope
  xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
  xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/">
  <S:Body>
    <samlp:AuthnRequest>
      <!-- same as above -->
    </samlp:AuthnRequest>
  </S:Body>
</S:Envelope>
```

Step 7: IdP responds to client with a SOAP response containing a SAML <Response> containing a short-lived SSO assertion (shown as an encrypted variant in the example). A generated key is included in the assertion and in a header for the client.

```
<S:Envelope
  xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
  xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/">
  <S:Header>
    <ecp:Response S:mustUnderstand="1"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next"
      AssertionConsumerServiceURL="xmpp@xmpp.example.com"/>
    <samlec:GeneratedKey xmlns:samlec="urn:ietf:params:xml:ns:samlec">
      3w1wSBKUosRLsU69xGK7dg==
    </samlec:GeneratedKey>
  </S:Header>
  <S:Body>
    <samlp:Response ID="d43h94r389309r" Version="2.0"
      IssueInstant="2020-12-10T11:42:34Z" InResponseTo="c3a4f8b9c2d"
      Destination="xmpp@xmpp.example.com">
      <saml:Issuer>https://saml.example.org</saml:Issuer>
      <samlp:Status>
        <samlp:StatusCode
          Value="urn:oasis:names:tc:SAML:2.0:status:Success"/>
        </samlp:Status>
      <saml:EncryptedAssertion>
        <!-- contents elided, copy of samlec:GeneratedKey in Advice -->
      </saml:EncryptedAssertion>
    </samlp:Response>
  </S:Body>
</S:Envelope>
```

Step 8: Client sends SOAP envelope containing the SAML <Response> as a response to the SASL server's challenge:

```
<response xmlns='urn:ietf:params:xmml:sasasl'>  
PFM6RW52ZWxvcGUKICAgIHhtbG5zOnNhbWw9InVybjpvYXNpczpYWl1cD0YzP  
QU1MOjIuMDphc3NlcnpRpb24iCiAgICB4bWxuczpzYWlsdD0idXJuOm9hc2lwZm5h  
bWVvZnRjOlNBtUw6Mi4wOnByb3RvY29sIgogICAgeGlslbnM6Uz0iaHR0cDovL3Nj  
agVtYXMueGlsc29hcC5vcmcvc29hcC9lbzlbnG9wZS8iPgogIDxtOkhlYWRlclj4K  
ICAgIDxwYW9zOlJlc3BvbnlIHhtbG5zOnBhb3M9InVybjpsawJlcnR5OnBhb3M6  
MjAwMy0wOCIKICAgICAgUzphY29hYyOiaHR0cDovL3NjaGVtYXNMueGlsc29hcC5v  
cmcvzc29hcC9hY3Rvcni9uZXh0IGogICAgICBTOMllc3RVbmRlcnN0YWE6PSXIiBY  
cwZUb011c3NhZ2VJRDI0NmMZYTRMOGI5YzJkIi8+CiAgICAgICA8c2FtbGVjOlNlc3Np  
b25LZXkgGeGlslbnM6c2FtbGVjPSJlcm44aWV0ZjpwYXJhbXB6eGlsc29hcC5vcmcvzc29h  
cc9lbzlbnG9wZS8iCiAgICAgIFM6bXVzdFVuZGVyc3RhbmQ9IjEjeiCiAgICAgIFM6  
YWNOb3I9Imh0dHA6Ly9zY2hlbnWFzfLnhtbHNvYXAub3JnL3NvYXAvaWYNOb3IvbmV4  
dCI+CiAgICAgIDxzYWlsZWMM6RW5jVHlwZT5hZXMXmJgty3RzLWhtYWMtc2hhMS05  
Njwvc2FtbGVjOkVuY1R5cGU+CIAgICA8c2FtbGVjOlNlc3Npb25LZXk+CIAgPC9T  
OkhlYWRlclj4KICAgUzPcb2R5PgogICAgPHNhbWxwOlJlc3BvbnlIELEPSJKNDNO  
OTryMZgm5MaA5ciIqVmVyc2lvbiOIMi4wIGogICAgICAgICAgICelzc3VlSW5zdGFudD0i  
MJAwNy0xMj0iXMFQxMTOM0mjgoNFoiIELuUmVzcG9uc2Vubz0iYzNhNGY4YjljMmQi  
CIAgICAgICAgRGVzdGluYXRpb249InhtCHBAeGlwc5leGFtcGxlLmNvbSI+CIAg  
ICAgIDxzYWlsOkklzc3Vlcj5odHRwczovL3NhbWwuZXhxbXBzZS5vcmc8L3NhbWw6  
SXNZdWVyPgogICAgICA8c2FtbHA6U3RhdHVzPgogICAgICAgIDxzYWlsdDpdTdGF0  
dXNDb2RlCiAgICAgICAgICAgIFZhbnVlPSJlcm44b2FzaXM6bmFtZXMM6dGM6U0FN  
TDoyLjA6c3RhbdHVzOlNlY2Nlc3MiLz4KICAgICAgPC9zYWlsdDpdTdGF0dXM+CIAg  
ICAgIDxzYWlsOkVuY3J5cHRlZEZFzc2VydGlvdj4KICAgICAgICA8IS0tIGNvbnlRl  
bnRzIGVsawRLZCWgy29weSBvziBzYWlsZWMM6R2VuZXJhdGVKS2V5IGluIEFKdmllj  
ZSATLT4KICAgICAgPC9zYWlsOkVuY3J5cHRlZEZFzc2VydGlvdj4KICAgIDwvc2Ft  
bHA6UmVzcG9uc2U+CIAgPC9TOkJvZHk+CjwwUzpfbnZlbg9wZTAK  
</response>
```

The Base64 [RFC4648] decoded envelope:

```

<S:Envelope
  xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
  xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/">
  <S:Header>
    <paos:Response xmlns:paos="urn:liberty:paos:2003-08"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next"
      S:mustUnderstand="1" refToMessageID="6c3a4f8b9c2d"/>
    <samlec:SessionKey xmlns:samlec="urn:ietf:params:xml:ns:samlec"
      xmlns:S="http://schemas.xmlsoap.org/soap/envelope/"
      S:mustUnderstand="1"
      S:actor="http://schemas.xmlsoap.org/soap/actor/next">
      <samlec:EncType>17</samlec:EncType>
    <samlec:SessionKey>
  </S:Header>
  <S:Body>
    <samlp:Response ID="d43h94r389309r" Version="2.0"
      IssueInstant="2020-12-10T11:42:34Z" InResponseTo="c3a4f8b9c2d"
      Destination="xmpp@xmpp.example.com">
      <saml:Issuer>https://saml.example.org</saml:Issuer>
      <samlp:Status>
        <samlp:StatusCode
          Value="urn:oasis:names:tc:SAML:2.0:status:Success"/>
      </samlp:Status>
      <saml:EncryptedAssertion>
        <!-- contents elided, copy of samlec:GeneratedKey in Advice -->
      </saml:EncryptedAssertion>
    </samlp:Response>
  </S:Body>
</S:Envelope>

```

Step 9: Server informs client of successful authentication:

```
<success xmlns='urn:ietf:params:xml:ns:xmpp-sasl' />
```

Step 9 (alt): Server informs client of failed authentication:

```

<failure xmlns='urn:ietf:params:xml:ns:xmpp-sasl'>
  <temporary-auth-failure/>
</failure>
</stream:stream>

```

Step 10: Client initiates a new stream to server:

```
<stream:stream xmlns='jabber:client'  
xmlns:stream='http://etherx.jabber.org/streams'  
to='example.com' version='1.0'>
```

Step 11: Server responds by sending a stream header to client along with any additional features (or an empty features element):

```
<stream:stream xmlns='jabber:client'  
xmlns:stream='http://etherx.jabber.org/streams'  
id='c2s_345' from='example.com' version='1.0'>  
<stream:features>  
  <bind xmlns='urn:ietf:params:xml:ns:xmpp-bind' />  
  <session xmlns='urn:ietf:params:xml:ns:xmpp-session' />  
</stream:features>
```

Step 12: Client binds a resource:

```
<iq type='set' id='bind_1'>  
  <bind xmlns='urn:ietf:params:xml:ns:xmpp-bind'>  
    <resource>someresource</resource>  
  </bind>  
</iq>
```

Step 13: Server informs client of successful resource binding:

```
<iq type='result' id='bind_1'>  
  <bind xmlns='urn:ietf:params:xml:ns:xmpp-bind'>  
    <jid>somenode@example.com/someresource</jid>  
  </bind>  
</iq>
```

Please note: line breaks were added to the base64 for clarity.

7. Security Considerations

This section will address only security considerations associated with the use of SAML with SASL applications. For considerations relating to SAML in general, the reader is referred to the SAML specification and to other literature. Similarly, for general SASL Security Considerations, the reader is referred to that specification.

Version 2.0 of the Enhanced Client or Proxy Profile [SAMLECP20] adds optional support for channel binding and use of "Holder of Key" subject confirmation. The former is strongly recommended for use with this mechanism to detect "Man in the Middle" attacks between the client and the RP without relying on the commercial TLS infrastructure that does not provide the level of assurance desired by sensitive SAML applications. The latter may be impractical in many cases, but is a valuable way of strengthening client authentication, protecting against phishing, and improving the overall mechanism.

7.1. Risks Left Unaddressed

The adaptation of a web-based profile that is largely designed around security-oblivious clients and a bearer model for security token validation results in a number of basic security exposures that should be weighed against the compatibility and client simplification benefits of this mechanism.

When channel binding is not used, protection against "Man in the Middle" attacks is left solely to lower layer protocols such as TLS, and the development of user interfaces able to implement that has not been effectively demonstrated. Failure to detect a MITM can result in phishing of the user's credentials if the attacker is between the client and IdP, or the theft and misuse of a short-lived credential (the SAML assertion) if the attacker is able to impersonate a RP. SAML allows for source address checking as a minor mitigation to the latter threat, but this is often impractical. IdPs can mitigate to some extent the exposure of personal information to RP attackers by encrypting assertions with authenticated keys.

7.2. User Privacy

The IdP is aware of each RP that a user logs into. There is nothing in the protocol to hide this information from the IdP. It is not a requirement to track the activity, but there is nothing technically that prohibits the collection of this information. Servers should be aware that SAML IdPs will track - to some extent - user access to their services. This exposure extends to the use of session keys generated by the IdP to secure messages between the parties, but note that when bearer assertions are involved, the IdP can freely impersonate the user to any relying party in any case.

It is also out of scope of the mechanism to determine under what conditions an IdP will release particular information to a relying party, and it is generally unclear in what fashion user consent could be established in real time for the release of particular

information. The SOAP exchange with the IdP does not preclude such interaction, but neither does it define that interoperably.

7.3. Collusion between RPs

Depending on the information supplied by the IdP, it may be possible for RPs to correlate data that they have collected. By using the same identifier to log into every RP, collusion between RPs is possible. SAML supports the notion of pairwise, or targeted/directed, identity. This allows the IdP to manage opaque, pairwise identifiers for each user that are specific to each RP. However, correlation is often possible based on other attributes supplied, and is generally a topic that is beyond the scope of this mechanism. It is sufficient to say that this mechanism does not introduce new correlation opportunities over and above the use of SAML in web-based use cases.

8. IANA Considerations

8.1. GSS-API and SASL Mechanism Registration

The IANA is requested to assign a new entry for this GSS mechanism in the sub-registry for SMI Security for Mechanism Codes, whose prefix is iso.org.dod.internet.security.mechanisms (1.3.6.1.5.5) and to reference this specification in the registry.

The IANA is requested to register the following SASL profile:

SASL mechanism profiles: SAML20EC and SAML20EC-PLUS

Security Considerations: See this document

Published Specification: See this document

For further information: Contact the authors of this document.

Owner/Change controller: the IETF

Note: None

8.2. XML Namespace Name for SAML-EC

A URN sub-namespace for XML constructs introduced by this mechanism is defined as follows:

URI: urn:ietf:params:xml:ns:samlec

Specification: See Appendix A of this document.

Description: This is the XML namespace name for XML constructs introduced by the SAML Enhanced Client SASL and GSS-API Mechanisms.

Registrant Contact: the IESG

9. References

9.1. Normative References

- [OASIS.saml-bindings-2.0-os]
Cantor, S., Hirsch, F., Kemp, J., Philpott, R., and E. Maler, "Bindings for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-bindings-2.0-os, March 2005.
- [OASIS.saml-core-2.0-os]
Cantor, S., Kemp, J., Philpott, R., and E. Maler, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-core-2.0-os, March 2005.
- [OASIS.saml-profiles-2.0-os]
Hughes, J., Cantor, S., Hodges, J., Hirsch, F., Mishra, P., Philpott, R., and E. Maler, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard OASIS.saml-profiles-2.0-os, March 2005.
- [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>>.
- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, DOI 10.17487/RFC2743, January 2000, <<https://www.rfc-editor.org/info/rfc2743>>.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, DOI 10.17487/RFC3961, February 2005, <<https://www.rfc-editor.org/info/rfc3961>>.
- [RFC3962] Raeburn, K., "Advanced Encryption Standard (AES) Encryption for Kerberos 5", RFC 3962, DOI 10.17487/RFC3962, February 2005, <<https://www.rfc-editor.org/info/rfc3962>>.

- [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>>.
- [RFC4121] Zhu, L., Jaganathan, K., and S. Hartman, "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2", RFC 4121, DOI 10.17487/RFC4121, July 2005, <<https://www.rfc-editor.org/info/rfc4121>>.
- [RFC4401] Williams, N., "A Pseudo-Random Function (PRF) API Extension for the Generic Security Service Application Program Interface (GSS-API)", RFC 4401, DOI 10.17487/RFC4401, February 2006, <<https://www.rfc-editor.org/info/rfc4401>>.
- [RFC4422] Melnikov, A., Ed. and K. Zeilenga, Ed., "Simple Authentication and Security Layer (SASL)", RFC 4422, DOI 10.17487/RFC4422, June 2006, <<https://www.rfc-editor.org/info/rfc4422>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC5554] Williams, N., "Clarifications and Extensions to the Generic Security Service Application Program Interface (GSS-API) for the Use of Channel Bindings", RFC 5554, DOI 10.17487/RFC5554, May 2009, <<https://www.rfc-editor.org/info/rfc5554>>.
- [RFC5801] Josefsson, S. and N. Williams, "Using Generic Security Service Application Program Interface (GSS-API) Mechanisms in Simple Authentication and Security Layer (SASL): The GS2 Mechanism Family", RFC 5801, DOI 10.17487/RFC5801, July 2010, <<https://www.rfc-editor.org/info/rfc5801>>.
- [RFC6680] Williams, N., Johansson, L., Hartman, S., and S. Josefsson, "Generic Security Service Application Programming Interface (GSS-API) Naming Extensions", RFC 6680, DOI 10.17487/RFC6680, August 2012, <<https://www.rfc-editor.org/info/rfc6680>>.
- [RFC7056] Hartman, S. and J. Howlett, "Name Attributes for the GSS-API Extensible Authentication Protocol (EAP) Mechanism", RFC 7056, DOI 10.17487/RFC7056, December 2013, <<https://www.rfc-editor.org/info/rfc7056>>.

- [RFC7617] Reschke, J., "The 'Basic' HTTP Authentication Scheme", RFC 7617, DOI 10.17487/RFC7617, September 2015, <<https://www.rfc-editor.org/info/rfc7617>>.
- [RFC7802] Emery, S. and N. Williams, "A Pseudo-Random Function (PRF) for the Kerberos V Generic Security Service Application Program Interface (GSS-API) Mechanism", RFC 7802, DOI 10.17487/RFC7802, March 2016, <<https://www.rfc-editor.org/info/rfc7802>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [SAMLECP20] Cantor, S., "SAML V2.0 Enhanced Client or Proxy Profile Version 2.0", OASIS Committee Specification OASIS.sstc-saml-ecp-v2.0-cs01, August 2013.
- [W3C.soap11] Box, D., Ehnebuske, D., Kakivaya, G., Layman, A., Mendelsohn, N., Nielsen, H., Thatte, S., and D. Winer, "Simple Object Access Protocol (SOAP) 1.1", W3C Note soap11, May 2000, <<http://www.w3.org/TR/SOAP/>>.

9.2. Informative References

- [OASIS.saml-metadata-2.0-os] Cantor, S., Moreh, J., Philpott, R., and E. Maler, "Metadata for the Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-metadata-2.0-os, March 2005.
- [RFC1939] Myers, J. and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, DOI 10.17487/RFC1939, May 1996, <<https://www.rfc-editor.org/info/rfc1939>>.
- [RFC3501] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, DOI 10.17487/RFC3501, March 2003, <<https://www.rfc-editor.org/info/rfc3501>>.

- [RFC4559] Jaganathan, K., Zhu, L., and J. Brezak, "SPNEGO-based Kerberos and NTLM HTTP Authentication in Microsoft Windows", RFC 4559, DOI 10.17487/RFC4559, June 2006, <<https://www.rfc-editor.org/info/rfc4559>>.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, DOI 10.17487/RFC6120, March 2011, <<https://www.rfc-editor.org/info/rfc6120>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.
- [W3C.REC-xmlschema-1]
Thompson, H., Beech, D., Maloney, M., and N. Mendelsohn, "XML Schema Part 1: Structures", W3C REC-xmlschema-1, May 2001, <<http://www.w3.org/TR/xmlschema-1/>>.
- [WSS-SAML]
Monzillo, R., "Web Services Security SAML Token Profile Version 1.1.1", OASIS Standard OASIS.wss-SAMLSecurityTokenProfile, May 2012.

Appendix A. XML Schema

The following schema formally defines the "urn:ietf:params:xml:ns:saml" namespace used in this document, in conformance with [W3C.REC-xmlschema-1] While XML validation is optional, the schema that follows is the normative definition of the constructs it defines. Where the schema differs from any prose in this specification, the schema takes precedence.

```
<schema
  targetNamespace="urn:ietf:params:xml:ns:samlec"
  xmlns="http://www.w3.org/2001/XMLSchema"
  xmlns:ds="http://www.w3.org/2000/09/xmldsig#"
  xmlns:S="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:samlec="urn:ietf:params:xml:ns:samlec"
  elementFormDefault="unqualified"
  attributeFormDefault="unqualified"
  blockDefault="substitution"
  version="1.0">

  <import namespace="http://www.w3.org/2000/09/xmldsig#" />
  <import namespace="http://schemas.xmlsoap.org/soap/envelope/" />

  <element name="SessionKey" type="samlec:SessionKeyType" />
  <complexType name="SessionKeyType">
    <sequence>
      <element ref="samlec:EncType" maxOccurs="unbounded" />
      <element ref="ds:KeyInfo" minOccurs="0" />
    </sequence>
    <attribute ref="S:mustUnderstand" use="required" />
    <attribute ref="S:actor" use="required" />
    <attribute name="Algorithm" />
  </complexType>

  <element name="EncType" type="integer" />

  <element name="GeneratedKey" type="samlec:GeneratedKeyType" />
  <complexType name="GeneratedKeyType">
    <simpleContent>
      <extension base="base64Binary">
        <attribute ref="S:mustUnderstand" />
        <attribute ref="S:actor" />
      </extension>
    </simpleContent>
  </complexType>

  <element name="Delegated" type="samlec:DelegatedType" />
  <complexType name="DelegatedType">
    <sequence />
    <attribute ref="S:mustUnderstand" use="required" />
    <attribute ref="S:actor" use="required" />
  </complexType>

</schema>
```

Appendix B. Acknowledgments

The authors would also like to thank Klaas Wierenga, Sam Hartman, Nico Williams, Jim Basney, Venkat Yekkirala, and Ben Kaduk for their contributions.

Appendix C. Changes

This section to be removed prior to publication.

- o 20, address nits and easy fixes from Ben Kaduk's AD review
- o 19, update obsoleted references
- o 15,16,17,18 avoid expiration
- o 14, address some minor comments
- o 13, clarify SAML metadata usage, adding a recommended Binding value alongside the backward-compatibility usage of PAOS
- o 12, clarifying comments based on WG feedback, with a normative change to use enctype numbers instead of names
- o 11, update EAP Naming reference to RFC
- o 10, update SAML ECP reference to final CS
- o 09, align delegation signaling to updated ECP draft
- o 08, more corrections, added a delegation signaling header
- o 07, corrections, revised section on delegation
- o 06, simplified session key schema, moved responsibility for random-to-key to the endpoints, and defined advertisement of session key algorithm and encypes by acceptor
- o 05, revised session key material, added requirement for random-to-key, revised XML schema to capture enctype name, updated GSS naming reference
- o 04, stripped down the session key material to simplify it, and define an IdP-brokered keying approach, moved session key XML constructs from OASIS draft into this one
- o 03, added TLS key export as a session key option, revised GSS naming material based on list discussion

- o 02, major revision of GSS-API material and updated references
- o 01, SSH language added, noted non-assumption of HTTP error handling, added guidance on life of security context.
- o 00, Initial Revision, first WG-adopted draft. Removed support for unsolicited SAML responses.

Authors' Addresses

Scott Cantor
Shibboleth Consortium
1050 Carmack Rd
Columbus, Ohio 43210
United States

Phone: +1 614 247 6147
Email: cantor.2@osu.edu

Margaret Cullen
Painless Security
4 High St, Suite 134
North Andover, Massachusetts 01845
United States

Phone: +1 781 405 7464
Email: mrcullen42@painless-security.com

Simon Josefsson
SJD AB
Hagagatan 24
Stockholm 113 47
SE

Email: simon@josefsson.org
URI: <http://josefsson.org/>

Internet Engineering Task Force
Internet-Draft
Updates: 4120 (if approved)
Intended status: Standards Track
Expires: April 6, 2015

S. Sorce, Ed.
Red Hat
T. Yu, Ed.
T. Hardjono, Ed.
MIT Kerberos Consortium
October 3, 2014

Kerberos Authorization Data Container Authenticated by Multiple MACs
draft-ietf-krb-wg-cammac-11

Abstract

Abstract: This document specifies a Kerberos Authorization Data container that supersedes AD-KDC-ISSUED. It allows for multiple Message Authentication Codes (MACs) or signatures to authenticate the contained Authorization Data elements. This document updates RFC 4120.

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 April 6, 2015.

Copyright Notice

Copyright (c) 2014 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
2. Requirements Language	2
3. Motivations	2
4. Encoding	4
5. Usage	6
6. Assigned numbers	6
7. IANA Considerations	6
8. Security Considerations	6
9. Acknowledgements	7
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Authors' Addresses	8

1. Introduction

This document specifies a new Authorization Data container for Kerberos, called AD-CAMMAC (Container Authenticated by Multiple MACs), that supersedes AD-KDC-ISSUED. This new container allows both the receiving application service and the Key Distribution Center (KDC) itself to verify the authenticity of the contained authorization data. The AD-CAMMAC container can also include additional verifiers that "trusted services" can use to verify the contained authorization data.

2. 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].

3. Motivations

The Kerberos protocol allows clients to submit arbitrary authorization data for a KDC to insert into a Kerberos ticket. These client-requested authorization data allow the client to express authorization restrictions that the application service will interpret. With few exceptions, the KDC can safely copy these client-requested authorization data to the issued ticket without necessarily inspecting, interpreting, or filtering their contents.

The AD-KDC-ISSUED authorization data container specified in RFC 4120 [RFC4120] is a means for KDCs to include positive or permissive

(rather than restrictive) authorization data in service tickets in a way that the service named in a ticket can verify that the KDC has issued the contained authorization data. This capability takes advantage of a shared symmetric key between the KDC and the service to assure the service that the KDC did not merely copy client-requested authorization data to the ticket without inspecting them.

The AD-KDC-ISSUED container works well for situations where the flow of authorization data is from the KDC to the service. However, protocol extensions such as Constrained Delegation (S4U2Proxy [MS-SFU]) require that a service present to the KDC a service ticket that the KDC previously issued, as evidence that the service is authorized to impersonate the client principal named in that ticket. In the S4U2Proxy extension, the KDC uses the evidence ticket as the basis for issuing a derivative ticket that the service can then use to impersonate the client. The authorization data contained within the evidence ticket constitute a flow of authorization data from the application service to the KDC. The properties of the AD-KDC-ISSUED container are insufficient for this use case because the service knows the symmetric key for the checksum in the AD-KDC-ISSUED container. Therefore, the KDC has no way to detect whether the service has tampered with the contents of the AD-KDC-ISSUED container within the evidence ticket.

The new AD-CAMMAC authorization data container specified in this document improves upon AD-KDC-ISSUED by including additional verifier elements. The svc-verifier element of the CAMMAC has the same functional and security properties as the ad-checksum element of AD-KDC-ISSUED; the svc-verifier allows the service to verify the integrity of the AD-CAMMAC contents as it already could with the AD-KDC-ISSUED container. The kdc-verifier and other-verifiers elements are new to AD-CAMMAC and provide its enhanced capabilities.

The kdc-verifier element of the AD-CAMMAC container allows a KDC to verify the integrity of authorization data that it previously inserted into a ticket, by using a key that only the KDC knows. The KDC thus avoids recomputing all of the authorization data for the issued ticket; this operation might not always be possible when that data includes ephemeral information such as the strength or type of authentication method used to obtain the original ticket.

The verifiers in the other-verifiers element of the AD-CAMMAC container are not required, but can be useful when a lesser-privileged service receives a ticket from a client and needs to extract the CAMMAC to demonstrate to a higher-privileged "trusted service" on the same host that it is legitimately acting on behalf of that client. The trusted service can use a verifier in the other-

verifiers element to validate the contents of the CAMMAC without further communication with the KDC.

4. Encoding

The Kerberos protocol is defined in [RFC4120] using Abstract Syntax Notation One (ASN.1) [X.680] and using the ASN.1 Distinguished Encoding Rules (DER) [X.690]. For consistency, this specification also uses ASN.1 for specifying the layout of AD-CAMMAC. The ad-data of the AD-CAMMAC authorization data element is the ASN.1 DER encoding of the AD-CAMMAC ASN.1 type specified below.

```
KerberosV5CAMMAC DEFINITIONS EXPLICIT TAGS ::= BEGIN
```

```
AD-CAMMAC ::= SEQUENCE {
    elements          [0] AuthorizationData,
    kdc-verifier      [1] Verifier-MAC OPTIONAL,
    svc-verifier      [2] Verifier-MAC OPTIONAL,
    other-verifiers   [3] SEQUENCE (SIZE (1..MAX))
                        OF Verifier OPTIONAL
}
```

```
Verifier ::= CHOICE {
    mac      Verifier-MAC,
    ...
}
```

```
Verifier-MAC ::= SEQUENCE {
    identifier [0] PrincipalName OPTIONAL,
    kvno       [1] UInt32 OPTIONAL,
    enctype    [2] Int32 OPTIONAL,
    mac        [3] Checksum
}
```

```
END
```

elements:

A sequence of authorization data elements issued by the KDC. These elements are the authorization data that the verifier fields authenticate.

Verifier:

A CHOICE type that currently contains only one alternative: Verifier-MAC. Future extensions might add support for public-key signatures.

Verifier-MAC:

Contains an RFC 3961 [RFC3961] Checksum (MAC) computed over the ASN.1 DER encoding of the AuthorizationData value in the elements field of the AD-CAMMAC. The identifier, kvno, and enctype fields help the recipient locate the key required for verifying the MAC. For the kdc-verifier and the svc-verifier, the identifier, kvno and enctype fields are often obvious from context and MAY be omitted. For the kdc-verifier, the MAC is computed differently than for the svc-verifier and the other-verifiers, as described later. The key usage for computing the MAC (Checksum) is 64.

kdc-verifier:

A Verifier-MAC where the key is a long-term key of the local Ticket-Granting Service (TGS). The checksum type is the required checksum type for the enctype of the TGS key. In contrast to the other Verifier-MAC elements, the KDC computes the MAC in the kdc-verifier over the ASN.1 DER encoding of the EncTicketPart of the surrounding ticket, but where the AuthorizationData value in the EncTicketPart contains the AuthorizationData value contained in the CAMMAC instead of the AuthorizationData value that would otherwise be present in the ticket. This altered Verifier-MAC computation binds the kdc-verifier to the other contents of the ticket, assuring the KDC that a malicious service has not substituted a mismatched CAMMAC received from another ticket.

svc-verifier:

A Verifier-MAC where the key is the same long-term service key that the KDC uses to encrypt the surrounding ticket. The checksum type is the required checksum type for the enctype of the service key used to encrypt the ticket. This field MUST be present if the service principal of the ticket is not the local TGS, including when the ticket is a cross-realm TGT.

other-verifiers:

A sequence of additional verifiers. In each additional Verifier-MAC, the key is a long-term key of the principal name specified in the identifier field. The PrincipalName MUST be present and be a valid principal in the realm. KDCs MAY add one or more "trusted service" verifiers. Unless otherwise administratively configured, the KDC SHOULD determine the "trusted service" principal name by replacing the service identifier component of the sname of the surrounding ticket with "host". The checksum is computed using a long-term key of the identified principal, and the checksum type is the required checksum type for the enctype of that long-term key. The kvno and enctype SHOULD be specified to disambiguate which of the long-term keys of the trusted service is used.

5. Usage

Application servers and KDCs MAY ignore the AD-CAMMAC container and the authorization data elements it contains. For compatibility with older Kerberos implementations, a KDC issuing an AD-CAMMAC SHOULD enclose it in an AD-IF-RELEVANT container unless the KDC knows that the application server is likely to recognize it.

6. Assigned numbers

The ad-type number for AD-CAMMAC is 96.

The key usage number for the Verifier-MAC checksum is 64.

7. IANA Considerations

[RFC Editor: please remove this section prior to publication.]

There are no IANA considerations in this document. Any numbers assigned in this document are not in IANA-controlled number spaces.

8. Security Considerations

Although authorization data are generally conveyed within the encrypted part of a ticket and are thereby protected by the existing encryption scheme used for the surrounding ticket, some authorization data requires the additional protection provided by the CAMMAC.

Some protocol extensions such as S4U2Proxy allow the KDC to issue a new ticket based on an evidence ticket provided by the service. If the evidence ticket contains authorization data that needs to be preserved in the new ticket, then the KDC MUST revalidate it.

Extracting a CAMMAC from a ticket for use as a credential removes it from the context of the ticket. In the general case, this could turn it into a bearer token, with all of the associated security implications. Also, the CAMMAC does not itself necessarily contain sufficient information to identify the client principal. Therefore, application protocols that rely on extracted CAMMACs might need to duplicate a substantial portion of the ticket contents and include that duplicated information in the authorization data contained within the CAMMAC. The extent of this duplication would depend on the security properties required by the application protocol.

The method for computing the kdc-verifier does not bind it to any authorization data within the ticket but outside of the CAMMAC. At least one (non-standard) authorization data type, AD-SIGNEDPATH,

attempts to bind to other authorization data in a ticket, and it is very difficult for two such authorization data types to coexist.

To minimize ticket size when embedding CAMMACs in Kerberos tickets, a KDC MAY omit the kdc-verifier from the CAMMAC when it is not needed. In this situation, the KDC cannot always determine whether the CAMMAC contents are intact. The KDC MUST NOT create a new CAMMAC from an existing one unless the existing CAMMAC has a valid kdc-verifier, with two exceptions.

Only KDCs for the local realm have knowledge of the local TGS key, so the outer encryption of a local TGT is sufficient to protect the CAMMAC of a local TGT from tampering, assuming that all of the KDCs in the local realm consistently filter out CAMMAC authorization data submitted by clients. The KDC MAY create a new CAMMAC from an existing CAMMAC lacking a kdc-verifier if that CAMMAC is contained within a local TGT and all of the local realm KDCs are configured to filter out CAMMAC authorization data submitted by clients.

An application service might not use the S4U2Proxy extension, or the realm policy might disallow the use of S4U2Proxy by that service. In such situations where there is no flow of authorization data from the service to the KDC, the application service could modify the CAMMAC contents, but such modifications would have no effect on other services. Because of the lack of security impact, the KDC MAY create a new CAMMAC from an existing CAMMAC lacking a kdc-verifier if it is inserting the new CAMMAC into a service ticket for the same service principal as the ticket that contained the existing CAMMAC, but MUST NOT place a kdc-verifier in the new CAMMAC.

The kdc-verifier in CAMMAC does not bind the service principal name to the CAMMAC contents, because the service principal name is not part of the EncTicketPart. An entity that has access to the keys of two different service principals can decrypt a ticket for one service and encrypt it in the key of the other service, altering the svc-verifier to match. Both the kdc-verifier and the svc-verifier would still validate, but the KDC never issued this fabricated ticket. The impact of this manipulation is minor if the CAMMAC contents only communicate attributes related to the client. If an application requires an authenticated binding between the service principal name and the CAMMAC or ticket contents, the KDC MUST include in the CAMMAC some authorization data element that names the service principal.

9. Acknowledgements

Shawn Emery, Sam Hartman, Greg Hudson, Ben Kaduk, Zhanna Tsitkov, and Kai Zheng provided helpful technical and editorial feedback on earlier versions of this document.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, February 2005.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.
- [X.680] ISO, , "Information technology -- Abstract Syntax Notation One (ASN.1): Specification of basic notation -- ITU-T Recommendation X.680 (ISO/IEC International Standard 8824-1:2008)", 2008.
- [X.690] ISO, , "Information technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER) -- ITU-T Recommendation X.690 (ISO/IEC International Standard 8825-1:2008)", 1997.

10.2. Informative References

- [MS-SFU] Microsoft, "[MS-SFU]: Kerberos Protocol Extensions: Service for User and Constrained Delegation Protocol", January 2013, <<http://msdn.microsoft.com/en-us/library/cc246071.aspx>>.

Authors' Addresses

Simo Sorce (editor)
Red Hat

Email: ssorce@redhat.com

Tom Yu (editor)
MIT Kerberos Consortium

Email: tlyu@mit.edu

Thomas Hardjono (editor)
MIT Kerberos Consortium

Email: hardjono@mit.edu

KERBEROS WORKING GROUP
Internet-Draft
Intended status: Standards Track
Expires: July 18, 2013

Johansson
SUNET
January 14, 2013

An information model for Kerberos version 5
draft-ietf-krb-wg-kdc-model-16

Abstract

This document describes an information model for Kerberos version 5 from the point of view of an administrative service. There is no standard for administrating a kerberos 5 KDC. This document describes the services exposed by an administrative interface to a KDC.

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 18, 2013.

Copyright Notice

Copyright (c) 2013 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
2. Requirements notation	4
3. Information model demarcation	6
4. Information model specification	7
4.1. Principal	7
4.1.1. Principal: Attributes	7
4.1.2. Principal: Associations	8
4.2. KeySet	9
4.2.1. KeySet: Attributes	9
4.2.2. KeySet: Associations	9
4.3. Key	9
4.3.1. Key: Attributes	10
4.3.2. Key: Associations	10
4.3.3. Key: Remarks	11
4.4. Policy	11
4.4.1. Policy: Attributes	11
4.4.2. Mandatory-to-implement Policy	12
5. Implementation Scenarios	13
5.1. LDAP backend to KDC	13
5.2. LDAP frontend to KDC	13
5.3. SOAP	13
5.4. Netconf	13
6. Security Considerations	14
7. IANA Considerations	15
8. Acknowledgments	16
9. References	17
9.1. Normative References	17
9.2. Informative References	17
Author's Address	18

1. Introduction

The Kerberos version 5 authentication service described in [RFC4120] describes how a Key Distribution Center (KDC) provides authentication to clients. The standard does not stipulate how a KDC is managed and several "kadmin" servers have evolved. This document describes the services required to administer a KDC and the underlying information model assumed by a kadmin-type service.

The information model is written in terms of "attributes" and "services" or "interfaces" but the use of these particular words must not be taken to imply any particular modeling paradigm. Neither an object oriented model nor an LDAP [RFC4510] schema is intended. The author has attempted to describe in natural language the intended semantics and syntax of the components of the model. An LDAP schema (for instance) based on this model will be more precise in the expression of the syntax while preserving the semantics of this model.

Implementations of this document MAY decide to change the names used (e.g. principalName). If so an implementation MUST provide a name to name mapping to this document. In particular schema languages may have different conventions for caseing, eg camelCase vs use of '_' and '-' to separate 'words' in a name. Implementations MUST call out such conventions explicitly.

Implementations of this document MUST be able to support default values for attributes as well as the ability to specify syntax for attribute values.

2. Requirements notation

This document uses the standard key words ("MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL") that are defined in [RFC2119] but with modifications to those definitions as described below. The reason for this (which was discussed extensively in the kerberos WG) is as follows:

This document describes an information model for kerberos 5 but does not directly describe any mapping onto a particular schema- or modelling language. Hence an implementation of this model consists of a mapping to such a language - e.g. an LDAP or SQL schema. The standard normative key word therefore require precise definition:

The terms MUST or REQUIRED means that schema implementing this model must have a way to represent a feature (i.e that it is mandatory to implement in schema) but that unless otherwise specified the feature may represent an optional element in the chosen schema definition language.

However MUST also means that a KDC or administrative interface implementing this information model MUST provide the feature and associated behavior consistent with schema.

For instance, `principalLastFailedAuthentication` (cf below) represents the last time an authentication failed for a principal. In an LDAP schema (for instance) this may be represented as an optional attribute even though all KDCs implementing this specification must support this attribute.

The terms MAY or OPTIONAL means that the feature is optional to implement by a KDC or administrative interface implementing this information model. It also means that the feature is optional to implement in schema.

Implementors of schema should be aware that unless there is a way to represent critical but optional elements in the schema definition language confusion may arise when optional elements are used but not understood by all implementations in a particular deployment.

The expression "MUST NOT be OPTIONAL" means that a feature is mandatory to implement ("MUST" cf above) and that additionally it must not be marked optional in the schema language. In particular this means that the feature is both mandatory to implement and must be present in all representations of the object to which it applies.

The term SHOULD or RECOMMENDED means that the consequences of not

implementing the feature as if it was described with the "MUST" keyword must be carefully weighed before choosing a different course. In particular this implies that interoperability concerns may arise from not following the recommended practice in schema that implements this model.

The context will determine if the "SHOULD" key word applies to schema, or to underlying behaviour of the KDC or both. For instance, `principalIsDisabled` (cf below) SHOULD default to FALSE implies both a recommendation for the behaviour of KDCs aswell as a recommendation for the representation of that behaviour in schema.

3. Information model demarcation

The information model specified in the next chapter describes objects, properties of those objects and relations between those objects. These elements comprise an abstract view of the data represented in a KDC. It is important to understand that the information model is not a schema. In particular the way objects are compared for equality beyond that which is implied by the specification of a syntax is not part of this specification. Nor is ordering specified between elements of a particular syntax.

Further work on Kerberos will undoubtedly prompt updates to this information model to reflect changes in the functions performed by the KDC. Such extensions to the information model should always use a normative reference to the relevant RFCs detailing the change in KDC function.

This model describes a number of elements related to password policy management. Not all of the elements in this model are unique to Kerberos; an LDAP implementation of this model should incorporate existing LDAP schema where functional overlap exists, rather than defining additional Kerberos-specific elements.

4. Information model specification

4.1. Principal

The fundamental entity stored in a KDC is the principal. The Principal is associated to keys and generalizes the "user" concept. The Principal **MUST** be implemented in full and **MUST NOT** be **OPTIONAL** in an implementation

4.1.1. Principal: Attributes

4.1.1.1. principalName

The principalName **MUST** uniquely identify the Principal within the administrative context of the KDC. The principalName **MUST** be equivalent to the string representation of the Principal name (section 2.1.1 of [RFC1964]) including, if applicable for the name type, the realm.

The attribute **MAY** be multi-valued if the implementation supports aliases and/or enterprise names. In that case exactly one of the principalName values **MAY** be designated the canonical principalName and if the implementation supports encetypes which require salt then exactly one of the values of principalName **MAY** be designated as the canonical salting principalName.

Implementations (i.e. schema) that support enterprise names and/or aliases **SHOULD** provide for efficient lookup of Principal objects based on alias/enterprise name.

4.1.1.2. principalNotUsedBefore

The Principal **MUST NOT** be used before this date. The syntax of the attribute **MUST** be Internet Date/Time Format from [RFC3339]. The attribute **MUST** be single-valued.

4.1.1.3. principalNotUsedAfter

The Principal **MUST NOT** be used after this date. The syntax of the attribute **MUST** be Internet Date/Time Format from [RFC3339]. The attribute **MUST** be single-valued.

4.1.1.4. principalIsDisabled

A boolean attribute used to disable a Principal. The attribute **SHOULD** default to boolean **FALSE**.

4.1.1.5. `principalLastCredentialChangeTime`

This single-valued attribute contains the time of the last successful change of credential (e.g. password or private key) associated with this Principal. The syntax of the attribute MUST be Internet Date/Time Format from [RFC3339].

4.1.1.6. `principalCreateTime`

This single-valued attribute contains the time and date when this Principal was created. The syntax of the attribute MUST be Internet Date/Time Format from [RFC3339].

4.1.1.7. `principalModifyTime`

This single-valued attribute contains the time and date when this Principal was last modified excluding credentials change. The syntax of the attribute MUST be Internet Date/Time Format from [RFC3339].

4.1.1.8. `principalMaximumTicketLifetime`

This single-valued attribute contains the time in seconds representing the maximum lifetime for tickets issued for this Principal.

4.1.1.9. `principalMaximumRenewableTicketLifetime`

This single-valued attribute contains the delta time in seconds representing the maximum amount of time a ticket may be renewed for this Principal.

4.1.1.10. `principalAllowedEncatype`

This OPTIONAL multi-valued attribute lists the encatypes allowed for this principal. If empty or absent any encatype supported by the implementation is allowed for this Principal.

This attribute is intended as a policy attribute and restricts all uses of encatypes including server, client, and session keys. Data models MAY choose to use policy objects in order to represent more complex decision mechanisms.

4.1.2. Principal: Associations

Each Principal MAY be associated with 0 or more KeySet and MAY be associated with 0 or more Policies. The KeySet is represented as an object in this model since it has attributes associated with it (the key version number). In typical situations the Principal is

associated with exactly 1 KeySet but implementations MUST NOT assume this case, i.e. an implementation of this standard MUST be able to handle the general case of multiple KeySet associated with each principal. Multiple KeySets may for instance be useful when performing a key rollover for a principal.

4.2. KeySet

In Kerberos principals are associated with zero or more symmetric secret keys, and each key has a key version number (kvno) and enctype. In this model we group keys by kvno into KeySet objects. A Principal can have zero or more KeySet objects associated with it, each of which MUST have one or more keys. Each KeySet is associated with exactly one principal. Schemas derived from this model MAY lack a direct analogue of KeySet as described in this document.

It is expected that most Kerberos implementations will use a special-purpose interface for setting and changing Principal passwords and keys.

If a server supports an enctype for a Principal that enctype must be present in at least one key for the Principal in question. For any given enctype a KeySet MUST NOT contain more than one Key with that enctype.

The security of Kerberos 5 depends absolutely on the confidentiality and integrity of the Key objects stored in the KDC. Implementations of this standard MUST facilitate, to the extent possible, an administrator's ability to place more restrictive access controls on KeySets than on other Principal data, and to arrange for more secure backup for KeySets.

4.2.1. KeySet: Attributes

4.2.1.1. kvno

Also known as the key version number. This is a single-valued attribute containing a non-negative integer. This number is incremented by one each time a key in the KeySet is changed.

4.2.2. KeySet: Associations

To each KeySet MUST be associated a set of 1 or more Keys.

4.3. Key

Implementations of this model MUST NOT REQUIRE keys to be represented.

4.3.1. Key: Attributes

4.3.1.1. keyEncryptionType

The enctype SHOULD be represented as an enumeration of the encetypes supported by the KDC using the string name ("encryption type") of the enctype from the IANA registry of Kerberos Encryption Type Numbers. One example is 'aes128-cts-hmac-sha1-96'.

4.3.1.2. keyValue

The binary representation of the key data. This MUST be a single-valued octet string.

4.3.1.3. keySaltValue

The binary representation of the key salt. This MUST be a single-valued octet string.

4.3.1.4. keyStringToKeyParameter

This MUST be a single-valued octet string representing an opaque parameter associated with the enctype. This parameter is specified in the "string-to-key" method in section 3 of [RFC3961].

4.3.1.5. keyNotUsedBefore

This key MUST NOT be used before this date. The syntax of the attribute MUST be semantically equivalent with the standard ISO date format. This MUST be a single-valued attribute.

4.3.1.6. keyNotUsedAfter

This key MUST NOT be used after this date. The syntax of the attribute MUST be semantically equivalent with the standard ISO date format. This MUST be a single-valued attribute.

4.3.1.7. keyIsDisabled

This is a boolean attribute which SHOULD be set to false by default. If this attribute is true the key MUST NOT be used. This is used to temporarily disable a key.

4.3.2. Key: Associations

None

4.3.3. Key: Remarks

The security of the keys is an absolute requirement for the operation of Kerberos 5. If keys are implemented adequate protection from unauthorized modification and disclosure **MUST** be available and **REQUIRED** by the implementation.

4.4. Policy

Implementations **SHOULD** implement policy but **MAY** allow them to be **OPTIONAL**. The Policy should be thought of as a 'typed hole'. i.e. an opaque binary value paired with an identifier of type of data contained in the binary value. Both attributes (type and value) must be present.

4.4.1. Policy: Attributes

4.4.1.1. policyIdentifier

The policyIdentifier **MUST** be globally unique. Possible types of identifiers include:

An Object Identifier (OID) [RFC4517]

A URI [RFC3986]

A UUID [RFC4122]

Implementations of this specification are expected to assign globally unique identifiers to the list of standard policy below in accordance with best-practice for identifier-management for the schema-language used.

4.4.1.2. policyIsCritical

This boolean attribute indicates that the KDC **MUST** be able to correctly interpret and apply this policy for the Principal to be used.

4.4.1.3. policyContent

This is an optional single opaque binary value used to store a representation of the policy. In general a policy cannot be fully expressed using attribute-value pairs. The policyContent is **OPTIONAL** in the sense that an implementation **MAY** use it to store an opaque value for those policy-types which are not directly representable in that implementation.

4.4.1.4. policyUse

This is an optional single enumerated string value used to describe the use of the policy. Implementations SHOULD provide this attribute and MUST (if the attribute is implemented) describe the enumerated set of possible values. The intent is that this attribute be useful in providing an initial context-based filtering.

4.4.2. Mandatory-to-implement Policy

All implementations that represent Policy objects MUST be able to represent the policies listed in this section. Implementations are not required to use the same underlying data-representation for the policyContent binary value but SHOULD use the same OIDs as the policyIdentifier. In general the expression of policy may require a Turing-complete language. This specification does not attempt to model policy expression language.

4.4.2.1. Password Quality Policy

Password quality policy controls the requirements placed by the KDC on new passwords.

4.4.2.2. Password Management Policy

Password management policy controls how passwords are changed.

4.4.2.3. Keying Policy

A keying policy specifies the association of enctypes with new principals, e.g. when a Principal is created one of the applicable keying policies is used to determine the set of keys to associate with the principal.

4.4.2.4. Ticket Flag Policy

A ticket flag policy specifies the ticket flags allowed for tickets issued for a principal.

5. Implementation Scenarios

There are several ways to implement an administrative service for Kerberos 5 based on this information model. In this section we list a few of them.

5.1. LDAP backend to KDC

Given an LDAP schema implementation of this information model it would be possible to build an administrative service by back-ending the KDC to a directory server where principals and keys are stored. Using the security mechanisms available on the directory server keys are protected from access by anyone apart from the KDC. Administration of the principals, policy, and other non-key data is done through the directory server while the keys are modified using the set/change password protocol [I-D.ietf-krb-wg-kerberos-set-passwd].

5.2. LDAP frontend to KDC

An alternative way to provide a directory interface to the KDC is to implement an LDAP-frontend to the KDC which exposes all non-key objects as entries and attributes. As in the example above all keys are modified using the set/change password protocol [I-D.ietf-krb-wg-kerberos-set-passwd]. In this scenario the implementation would typically not use a traditional LDAP implementation but treat LDAP as an access protocol to data in the native KDC database.

5.3. SOAP

Given an XML schema implementation of this information model it would be possible to build a SOAP interface to the KDC. This demonstrates the value of creating an abstract information model which is mappable to multiple schema representations.

5.4. Netconf

Given a YAML implementation of this information model it would be possible to create a Netconf-based interface to the KDC, enabling management of the KDC from standard network management applications.

6. Security Considerations

This document describes an abstract information model for Kerberos 5. The Kerberos 5 protocol depends on the security of the keys stored in the KDC. The model described here assumes that keys **MUST NOT** be transported in the clear over the network and furthermore that keys are treated as write-only attributes that **SHALL** only be modified (using the administrative interface) by the change-password protocol [I-D.ietf-krb-wg-kerberos-set-passwd].

Exposing the object model of a KDC typically implies that objects can be modified and/or deleted. In a KDC not all principals are created equal, so that for instance deleting `krbtgt/EXAMPLE.COM@EXAMPLE.COM` effectively disables the `EXAMPLE.COM` realm. Hence access control is paramount to the security of any implementation. This document does not mandate access control. This only implies that access control is beyond the scope of the standard information model, i.e. that access control may not be accessible via any protocol based on this model. If access control objects are exposed via an extension to this model the presence of access control may in itself provide points of attack by giving away information about principals with elevated rights etc.

7. IANA Considerations

This document has no IANA actions.

8. Acknowledgments

The author wishes to extend his thanks to Love Hoernquist-Aestrand and Sam Hartman for their important contributions to this document.

9. References

9.1. Normative References

- [RFC1964] Linn, J., "The Kerberos Version 5 GSS-API Mechanism", RFC 1964, June 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3339] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.
- [RFC3961] Raeburn, K., "Encryption and Checksum Specifications for Kerberos 5", RFC 3961, February 2005.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC4517] Legg, S., "Lightweight Directory Access Protocol (LDAP): Syntaxes and Matching Rules", RFC 4517, June 2006.

9.2. Informative References

- [I-D.ietf-krb-wg-kerberos-set-passwd]
Williams, N., "Kerberos Set/Change Key/Password Protocol Version 2", draft-ietf-krb-wg-kerberos-set-passwd-08 (work in progress), November 2008.
- [RFC4510] Zeilenga, K., "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map", RFC 4510, June 2006.

Author's Address

Leif Johansson
Swedish University Network
Thulegatan 11
Stockholm

Email: leifj@sunset.se
URI: <http://www.sunet.se>

Network Working Group
Internet-Draft
Updates: 1964, 2743, 2744
(if approved)
Intended status: Standards Track
Expires: April 30, 2015

N. Williams
Cryptonector
October 27, 2014

Generic Naming Attributes for the Generic Security Services Application
Programming Interface (GSS-API)
draft-williams-kitten-generic-naming-attributes-02

Abstract

This document specifies several useful generic naming attributes for use with the Generic Security Services Application Programming Interface (GSS-API) Naming Extensions specified in RFC6680.

These attributes allow applications to extract discrete components of a GSS-API "mechanism name" (MN) object: issuer (e.g., realm name, domain name, certification authority name), service and host names (for host-based service names), user names, and others.

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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 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 and Motivation	3
1.1.	Naming Constraints	3
1.2.	Conventions used in this document	4
2.	Generic Attributes	5
2.1.	Concrete Attributes	5
2.1.1.	Issuer Name	5
2.1.2.	Trust Validation Path	6
2.1.3.	User Name	6
2.1.4.	Service Name	6
2.1.5.	Host Name	6
2.1.6.	Domain Name	7
2.2.	Prefix Attributes	7
2.2.1.	GSS_C_ATTR_GENERIC_UNCONSTRAINED	7
2.2.2.	GSS_C_ATTR_GENERIC_UNCONSTRAINED_OK	8
2.2.3.	GSS_C_ATTR_GENERIC_FAST	8
3.	Local Name Attributes	9
3.1.	GSS_C_ATTR_LOCAL_LOGIN_USER	9
4.	Suggested Mechanism-Specific Name Attributes (INFORMATIONAL)	10
4.1.	Suggested Kerberos-Specific Name Attributes	10
4.1.1.	Kerberos Transit Path Constraint Semantics	10
4.2.	Suggested PKU2U-Specific Name Attributes	11
5.	Generic Issuer Name Type	12
5.1.	Kerberos Realm Name Type	12
5.2.	PKIX Issuer Name Type	12
6.	Security Considerations	13
7.	IANA Considerations	14
8.	References	15
8.1.	Normative References	15
8.2.	Informative References	15
	Author's Address	16

1. Introduction and Motivation

The Generic Security Services Application Programming Interface (GSS-API) [RFC2743] allows applications -and application protocol specifications- to use various security mechanisms in a generic way. There are some shortcomings of this API that preclude a fully-generic treatment of security mechanisms. This document builds on the naming extensions to the GSS-API [RFC6680] to correct some of those shortcomings.

In RFC6680 we introduced an interface by which to access "attributes" of names, but we did not specify any attributes. This document specifies some such attributes. Some of the new attributes are specifically intended to make it possible to use the GSS-API in a mechanism-generic way in common use cases where it is otherwise not possible to do so.

For example, some applications need to be able to observe the discrete elements of a peer principal's host-based service name, but they generally could only do so by parsing mechanism-specific display syntaxes or exported name token formats. Such applications are inherently not generic: they can only function correctly when used with security mechanism whose principal naming conventions/formats the applications understand.

More generally, we use the the extended naming interface to introduce an attribute model of principal naming.

1.1. Naming Constraints

This document also introduces a notion of naming constraints, not unlike PKIX's [RFC5280]. Naming constraints apply to "issuers" of principal names and/or their attributes. For example, to Kerberos [RFC4120] realms, to PKIX certification authorities, to identity providers (IdPs), and so on. The goal is allow specification of policies which constrain the set of principal names that a given issuer can issue credentials for.

For example, the Kerberos realm FOO.EXAMPLE would generally not be expected to issue credentials to host-based principals in domains other than "foo.example".

For each concrete attribute specified below there are several ways to inquire a NAME's value for that attribute:

1. with naming constraint checking, providing no output if naming constraints are violated;

2. with naming constraint checking, providing an output indicator of naming constraint violations;
3. without naming constraint checking;
4. any of the above with "fast" (no slow I/O involved) naming constraint checking.

(1) is the default behavior. The others are obtained by adding an appropriate prefix to the attribute name.

Existing security mechanisms may not have any formal notion of naming constraints, but it is common to have some naming constraint conventions nonetheless. For example, Kerberos realm naming conventions are that realm names should mirror Domain Name System (DNS) [RFC1035] domain names, and that hostnames embedded in Kerberos principal names should a) be fully-qualified, b) within the domain corresponding to the DNS domain name derived from the realm's name. Or a Kerberos implementation might lookup a host's realm and check that it matches the principal's realm. Naming constraints should be formalized for all GSS-API security mechanisms.

1.2. Conventions used in this document

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. Generic Attributes

We add a number of generic name attributes, to be used via the GSS-API extended naming facility [RFC6680]. Some of these attributes can be used as prefixes of other attributes, that is, they can be used to modify the semantics of other attributes (see section 6 of RFC6680).

We also provide C bindings for these attributes, namely, the same symbolic names that we provide for the generic attributes.

Note: in all cases the display form of each attribute SHALL consist of text using the character set, codeset, and encoding from the caller's locale.

2.1. Concrete Attributes

These attributes generally have a single value each. Only one of these attributes can also be used a prefix: the issuer name attribute.

2.1.1. Issuer Name

We add an attribute by which to obtain a name of an issuer of a mechanism name (MN) or of an attribute of an MN. The API name for this attribute is `GSS_C_ATTR_GENERIC_ISSUENAME`, and its actual attribute name is "urn:ietf:id:ietf-kitten-name-attrs-00-issuename".

The display form of issuer names is mechanism-specific.

The non-display form of issuer names SHALL be the exported name token form of the issuer's name. Not all mechanisms will support issuer names as MNs, therefore implementations MAY output a null non-display value.

For example, for the Kerberos mechanism [RFC4121] an issuer name would generally (but not always!) be a Kerberos realm name, probably displayed as just the realm name. (But note that there is not yet a Kerberos realm name as MN specification. We will specify one separately.)

This attribute can be used as prefix of other attributes. When used as a prefix, this attribute indicates that the application wishes to know the name of the issuer of the prefixed attribute of the given MN.

2.1.2. Trust Validation Path

We add an attribute by which to obtain the trust validation path for a given authenticated MN. The API for this attribute is `GSS_C_ATTR_GENERIC_TRUST_PATH`, and its actual attribute name is `urn:ietf:id:ietf-kitten-name-attrs-01-trust-path`.

This attribute has zero or more ordered values. The interpretation of the trust validation path will vary somewhat by mechanism. For PKIX-based mechanisms this is the list of issuers in the trust validation path for the given MN's cert. For Kerberos this is the list of realms traversed from the MN to the local name of a security context. The MN's immediate issuer is not included. In the case of Kerberos, the issuer of the local MN is also not included. For Kerberos the trust validation path is the realm transit path of the Ticket used to establish a security context, but may also include PKIX trust validation paths (e.g., if PKINIT is used).

The display and non-display forms of trust validation path values is as for issuer names; see Section 2.1.1.

2.1.3. User Name

We add an attribute by which to obtain the component of an MN naming a user. The API name for this attribute is `GSS_C_ATTR_GENERIC_USERNAME`, and its actual attribute name is `"urn:ietf:id:ietf-kitten-name-attrs-00-username"`.

The display form of user names is mechanism-specific.

The non-display form of user names is mechanism-specific.

2.1.4. Service Name

We add an attribute by which to obtain the component of an MN naming a service as part of a host- or domain-based service name. The API name for this attribute is `GSS_C_ATTR_GENERIC_SERVICENAME`, and its actual attribute name is `"urn:ietf:id:ietf-kitten-name-attrs-00-servicename"`.

The non-display form of the service name SHALL be the UTF-8 encoding of the service name.

2.1.5. Host Name

We add an attribute by which to obtain the component of an MN naming a host as part of a host- or domain-based service name. The API name for this attribute is `GSS_C_ATTR_GENERIC_HOSTNAME`, and its actual

attribute name is "urn:ietf:id:ietf-kitten-name-attrs-00-hostname".

The display form of a host name MAY be stylized and SHOULD NOT be A-labels. [RFC5890].

The non-display form of host names SHOULD be a character string as described in [RFC1123], and SHOULD NOT be U-labels [RFC5890].

2.1.6. Domain Name

We add an attribute by which to obtain the component of an MN naming a domain as part of a domain-based service name. The API name for this attribute is GSS_C_ATTR_GENERIC_DOMAINNAME, and its actual attribute name is "urn:ietf:id:ietf-kitten-name-attrs-00-domainname".

The display form of a domain name MAY be stylized and SHOULD NOT be A-labels. [RFC5890].

The non-display form of domain names SHOULD be a character string as described in [RFC1123], and SHOULD NOT be U-labels [RFC5890].

2.2. Prefix Attributes

GSS_Get_name_attribute() using attributes described in the preceding section SHALL fail if there are any name constraints that can be applied to the issuers of those names and, in applying those constraints, it is discovered that the issuer was not permitted to issue credentials for the MN.

For example, a Kerberos realm named "FOO.EXAMPLE" might not be expected to issue credentials (tickets, keys) to host-based service names for hosts not ending in ".foo.example" or which are not "foo.example".

Several generic attribute prefixes are described below for overriding this behavior.

2.2.1. GSS_C_ATTR_GENERIC_UNCONSTRAINED

This attribute prefix, named GSS_C_ATTR_GENERIC_UNCONSTRAINED in the API, and with an actual name of "urn:ietf:id:ietf-kitten-name-attrs-00-gen-unconstrained", indicates that the application wants the value of the prefixed attribute without any name constraint checking.

2.2.2. GSS_C_ATTR_GENERIC_UNCONSTRAINED_OK

This attribute prefix, named GSS_C_ATTR_GENERIC_UNCONSTRAINED_OK in the API, and with an actual name of "urn:ietf:id:ietf-kitten-name-attrs-00-gen-unconstrained-ok", indicates that the application wants the value of the prefixed attribute regardless of any applicable naming constraints, but to indicate the name constraint status via the 'authenticated' output parameter of the GSS_Get_name_attribute() interface.

2.2.3. GSS_C_ATTR_GENERIC_FAST

This attribute prefix, named GSS_C_ATTR_GENERIC_FAST in the API, and with an actual name of "urn:ietf:id:ietf-kitten-name-attrs-00-gen-fast", indicates that the application requires that the mechanism not perform any slow operations (e.g., connecting to a directory for the purposes of name constraint validation) in obtaining the prefixed attribute of the given MN.

3. Local Name Attributes

Normally an Internet specification would not be expected to specify any local name attributes of GSS names. However, there is one common and very useful local name attribute, which we specify below. Implementations are free to use different names for this attribute or exclude it altogether -- it is a local name attribute, after all.

3.1. GSS_C_ATTR_LOCAL_LOGIN_USER

This attribute, with suggested API symbolic name `GSS_C_ATTR_LOCAL_LOGIN_USER`, and suggested actual name "local-login-user", requests a local user name corresponding to the given MN, if any.

Obtaining the local user name corresponding to an MN may require complex name mapping or lookup operations that are completely implementation-defined.

4. Suggested Mechanism-Specific Name Attributes (INFORMATIONAL)

[[anchor1: This section should really be split out into separate Internet-Drafts. It is here only because the author lacks the time at the moment of writing to create such separate I-Ds.]]

[[anchor2: Actually, we should probably make this section normative. It's easier than publishing a larger number of RFCs...]]

4.1. Suggested Kerberos-Specific Name Attributes

- o realm (corresponding to issuer name)
- o component 0 (first component of a principal name)
- o component 1 (second component of a principal name)
- o ..
- o component 9 (tenth component of a principal name; ten is enough)
- o components (ordered set of all components of a principal name)
- o specific authorization data elements
- o PKINIT client certificate
- o session key enctype
- o encypes involved in transit path (this would only be available to initiators)

4.1.1. Kerberos Transit Path Constraint Semantics

For initiator MNs obtained by acceptors from established security contexts, the trust path SHALL be the uncompressed domain- and X.500-style realm names from the initiator's Ticket's 'transited' field, plus the issuer names from the AD-INITIAL-VERIFIED-CAS authorization-data element (if it's in an AD-KDC-ISSUED or similar) if PKINIT [RFC4556] was used.

For acceptor MNs obtained by initiators from established security contexts, the trust path SHALL be the realms traversed -including realms issuing referrals- to obtain a service ticket for the target acceptor.

For MNs for the local end of a security context, the trust path SHALL be empty. This means that GSS_Get_name_attribute() will return empty

value sets; for the C bindings the `gss_get_name_attribute()` function will return zero in the 'more' output parameter and empty values.

For initiator MNs as seen by acceptors, if the initiator's Ticket has the TRANSIT-POLICY-CHECKED flag set, and if local transit path policy is missing, then the GSS_C_ATTR_GENERIC_TRUST_PATH attribute will be considered authenticated -- the trust path will be considered to meet constraints.

Otherwise, if the acceptor has local transit path policy then the GSS_C_ATTR_GENERIC_TRUST_PATH attribute will be considered authenticated -- the trust path will be considered to meet constraints.

In all other cases the GSS_C_ATTR_GENERIC_TRUST_PATH attribute will be considered not authenticated.

4.2. Suggested PKU2U-Specific Name Attributes

[[anchor3: Add reference to PKU2U.]]

- o issuer CA name
- o certificate trust validation path to a trust anchor
- o certificate
- o certificate subject public key
- o certificate subject public key algorithm
- o certificate subject name
- o certificate subject alternate names
- o specific certificate extensions
- o certificate algorithm names
- o session key enctype

5. Generic Issuer Name Type

We add a GSS name-type for use in representing issuer names, designated symbolically as GSS_C_NT_ISSUER. Its query syntax is unspecified and mechanism-specific.

At least initially the common use of this name-type will be for representation of issuer names using the GSS_C_ATTR_GENERIC_ISSUERNAME GSS name attribute (see Section 2.1.1).

5.1. Kerberos Realm Name Type

No name-type is needed in the Kerberos protocol for realm names. Because all three forms of Kerberos realm-names (DOMAIN, X.500, and OTHER) and unambiguously distinguishable from each other, we also do not add a Kerberos-specific GSS name-type.

The query and display syntax of GSS_C_NT_ISSUER names for Kerberos is just a realm name prefixed with an '@'. We prefix the realm name with '@' to take advantage of an otherwise useless ambiguity in the query and display form of Kerberos mechanism principal names [RFC1964], namely that zero-component, and one-zero-length component principal names display identically, therefore those are useless name forms in Kerberos (they would be useless anyways); we appropriate this otherwise useless name form as the query and display syntax of Kerberos realm names. For example, "@FOO.EXAMPLE".

In the unlikely event that a name of GSS_C_NT_ISSUER type is used as a GSS initiator or acceptor principal, the actual Kerberos principal name should be an appropriate TGS principal name. More specific information for such use-cases will be provided by any future application protocol specifications that use them.

5.2. PKIX Issuer Name Type

[[anchor4: A name type for PKIX issuers is needed, even when dealing with Kerberos, since X.500-style realm names may be involved, as well as real PKIX CA names from PKINIT/PKCROSS. We'll need a mechanism OID for a generic PKIX mechanism (even if it isn't specified!) for the exported name tokens! One that Kerberos, PKU2U and other PKIX mechanisms can share.]]

6. Security Considerations

[Add text regarding name constraint checking and explaining the default-to-safe design of the generic name attributes defined in section 2.]

7. IANA Considerations

[Add text regarding the registration and assignment of the name attributes described in the preceding sections. In particular we should want these attributes' names to not reflect an Internet-Draft name, but an RFC number.]

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.
- [RFC1964] Linn, J., "The Kerberos Version 5 GSS-API Mechanism", RFC 1964, June 1996.
- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, January 2000.
- [RFC4556] Zhu, L. and B. Tung, "Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)", RFC 4556, June 2006.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010.
- [RFC6680] Williams, N., Johansson, L., Hartman, S., and S. Josefsson, "Generic Security Service Application Programming Interface (GSS-API) Naming Extensions", RFC 6680, August 2012.

8.2. Informative References

- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.
- [RFC4121] Zhu, L., Jaganathan, K., and S. Hartman, "The Kerberos Version 5 Generic Security Service Application Program Interface (GSS-API) Mechanism: Version 2", RFC 4121, July 2005.
- [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, May 2008.

Author's Address

Nicolas Williams
Cryptonector, LLC

Email: nico@cryptonector.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2015

N. Williams
Cryptonector
October 27, 2014

Public Key-Based Kerberos Cross Realm Path Traversal Protocol Using
Kerberized Certification Authorities (kx509) and PKINIT
draft-williams-kitten-krb5-pkcross-05

Abstract

This document specifies a protocol for obtaining cross-realm Kerberos tickets using existing, related protocols: kerberized certification authorities (kx509) and public key cryptography initial authentication in Kerberos (PKINIT). The resulting protocol has a number of desirable properties, primarily that it allows Kerberos to scale to large numbers of realms.

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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 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	3
1.1.	Conventions used in this document	3
2.	The PKCROSS Protocol	4
2.1.	Client-Driven PKCROSS	4
2.2.	TGS-Driven PKCROSS	4
2.2.1.	Issuing cross-realm TGTs issued for PKCROSS-keyed cross-realm TGS principals	5
2.2.2.	Handling impatient clients	5
2.3.	Stapled DANE	6
2.4.	Validation	6
2.5.	Transit Path	6
2.5.1.	Transit path representation	7
2.6.	Exchange of Long-Term Cross-Realm Symmetric Keys	7
3.	Security Properties	9
3.1.	Automatic Cross-Realm Keying	9
3.2.	Scalability	9
3.2.1.	Simplified trust routing	9
3.2.2.	Simplified trust path validation	9
3.3.	Privacy Protection relative to home realm	10
4.	Application Programming Interface Considerations	11
4.1.	GSS-API Considerations	11
5.	Security Considerations	12
5.1.	Loss of Cross-Realm Principal Trust Establishment Information	12
5.2.	On the Need for a Common Transit Path Policy Language	12
5.3.	On the Need for Trust Routing	13
6.	IANA Considerations	14
7.	TODO	15
8.	Acknowledgements	16
9.	References	17
9.1.	Normative References	17
9.2.	Informative References	17
	Author's Address	19

1. Introduction

Kerberos [RFC4120] supports meshes of many realms. The individual relationships between realms must be manually keyed, usually with keys derived from passwords. A full mesh wouldn't scale, therefore the protocol calls for hierarchical trust universes. In practice non-hierarchical but also non-fully-meshed relationships are used, and these generally require distribution of trust routing information to clients, services, and KDCs. With referrals it is possible to reduce the need for client-side trust routing information, but KDCs still need it, as do services (unless they accept KDC trust path policy and the KDC applies it via the TRANSITED-POLICY-CHECKED ticket flag).

These manually-exchanged keys are very difficult to rollover safely, and when they are changed the result is often outages -- controlled outages where foreseen, but outages nonetheless.

Manual cross-realm keying does not scale, and has very poor security properties. We seek to remediate this using public key cryptography, building on existing Kerberos specifications.

Distribution of trust routing (traditionally known as "capaths") and trust path validation (also "capaths") information is difficult; there is no standard protocol for it. Maintenance of it is a thoroughly manual process.

Many years ago there was a proposal for exchanging cross-realm keys using a public key infrastructure (PKI) [RFC5280]; that proposal went by the name "PKCROSS". We appropriate that long-dead proposal's name, but the protocol specified here is very different from the original proposal.

PKCROSS can make Kerberos scale to large numbers of realms, will remove the need for manual keying of cross-realm TGS principals, will further reduce the need for maintenance and distribution of trust routing information, and will tend to reduce the complexity of trust path validation.

1.1. Conventions used in this document

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 PKCROSS Protocol

We provide two variants of the PKCROSS protocol: one that is client-driven, and another that is driven by a Ticket Granting Service (TGS) on behalf of its clients. The latter is based on the former, with the TGS acting as a client. We begin with the client-driven case. DNS-Based Authentication of Named Entities (DANE) [RFC6698] can and should be used for realm CA certificate validation.

2.1. Client-Driven PKCROSS

A Kerberos client in with a ticket-granting ticket (TGT) for any one source realm (usually but not necessarily the client's own realm) wishing to acquire a TGT for a destination realm may use this protocol instead of the traditional cross-realm ticket-granting service (TGS) exchanges as follows:

1. Generate private key to a public key cryptosystem;
2. Request a certificate from the kx509 [RFC6717] service run by the source realm;
3. Request a TGT from the destination realm using PKINIT [RFC4556] and the client certificate obtained in step #2.

If the destination realm issues the requested Ticket then it SHOULD include the client's certificate in an AD-CLIENT-CERTIFICATE authorization-data element, and it MUST do so if it does not validate the client's certificate to an acceptable trust anchor. The AD-CLIENT-CERTIFICATE authorization-data MUST be in a KDC-signed authorization-data container [XXX add reference to CAMMAC].

[[anchor1: QUESTION: Should the PKINIT request in step #3 be a TGS-REQ with PKINIT pre-auth data?]]

[[anchor2: QUESTION: Should the PKINIT request in step #3 be required to be used within a FAST tunnel?]]

2.2. TGS-Driven PKCROSS

A TGS can bootstrap ephemeral cross-realm trust principals on behalf of its clients. This allows the cost of PKCROSS to be amortized over many clients, and it allows participation by clients that do not support client-driven PKCROSS (or whose PKCROSS requests are rejected by the target).

In this mode the TGS uses the client-driven PKCROSS protocol, modified as follows:

- o the TGS's client certificate MUST have an id-pkinit-san Subject Alternative Name (SAN) identifying the source TGS as `krbtgt/SOURCE@SOURCE`
- o the TGS's client certificate MUST have an Extended Key Usage (EKU) of `id-pkcross-issuer` (TBD)

The resulting TGT -which we shall term an "issuer TGT" (ITGT)- and its session key can then be used by the source TGS to create cross-realm TGTs for the source-to-target trust principal ("`krbtgt/TARGET@SOURCE`").

This ITGT will be used to mint tickets as described below.

2.2.1. Issuing cross-realm TGTs issued for PKCROSS-keyed cross-realm TGS principals

Cross-realm TGTs issued by a source TGS using an ITGT will not be quite like normal Kerberos Tickets: their encrypted part contains an AP-REQ using the ITGT acquired by the source TGS, and this AP-REQ is "encrypted" with the null enctype. The AP-REQ's Authenticator MUST contain an authorization-data element that carries a) the name of the client principal, b) the session key that the client should be using with the cross-realm TGTs issued.

```
AD-PKCROSS-TGT-INFO ::= SEQUENCE {  
    cname [0] Principal,      -- the client's realm is the  
                                -- crealm from the ITGT's EncTicketPart  
    key   [1] EncryptionKey  
}
```

Figure 1: AD-PKCROSS-TGT-INFO

2.2.2. Handling impatient clients

Because the process of acquiring an ITGT might be slow, a TGS doing so on behalf of a client could use a mechanism for instructing the client to be patient. Existing clients would not handle a new error code by waiting, therefore there is not much that can be done to keep an impatient client from retrying at another KDC.

The existing `KDC_ERR_SVC_UNAVAILABLE` error code cannot be used as often this causes the client to immediately retry the request at another KDC. A new error code for indicating estimated time to completion of request would be handy, but out of scope for this document.

Note that there is a denial of service (DoS) attack by clients on

willing source KDCs: the clients can ask the KDCs to acquire cross-realm ITGTs for many target realms. Ideally the quality of service for the Kerberos authentication service (AS) with PKINIT (and/or other slow pre-authentication mechanisms) should be separate from that of the Kerberos TGS co-located with it, and the PKCROSS-capable TGS as well, so as to be able to throttle low-priority requests when under load.

2.3. Stapled DANE

[[anchor3: TBD. We should use Google's serialization of DNS RRsets needed for DANE validation. We will need a label for the TLSA RRs for kx509 issuers.]]

2.4. Validation

KDCs processing PKINIT requests crossing realms MUST apply either or both of:

- o PKIX certificate validation
- o DANE certificate validation

KDCs MUST reject PKINIT requests from clients of foreign realms whose certificates cannot be validated, unless the client request the anonymous principal name in the target's realm.

2.5. Transit Path

The combined Kerberos/PKIX/DNSSEC transit path MUST be represented in any tickets issued using PKCROSS (see below). As usual, each realm's KDCs in the mix can set the transit policy checked flag if a client's transit path is acceptable per the realm's KDCs' local policy.

Two validation mechanisms are available: all PKIX [RFC5280] validation methods, and DANE [RFC6698]. DANE validation records SHOULD be stapled onto the client certificates by the issuing kx509 CA; alternatively, clients can staple <http://src.chromium.org/viewvc/chrome/trunk/src/net/base/dnssec_chain_verifier.cc?pathrev=167227> onto their PKINIT requests using an authorization-data element, AD-PKINIT-CLIENT-DANE.

Additionally, when PKIX certificate validation is used, the trust path should be encoded in an AD-INITIAL-VERIFIED-CAS authorization data element, per-PKINIT.

2.5.1. Transit path representation

The notional transit path for a ticket issued by a target realm's KDCs includes:

- o the source realm (never expressed in the 'transited' field of Kerberos Tickets)
- o all realms in the ITGT's transited field (in the TGS-driven PKCROSS case)
- o all issuers in the validation path for the kx509-issued certificate, which are
 - * all issuers in the certificate's PKIX validation path when PKIX validation is used
 - * all DNS zone domainnames transited from the source realm's domainname to the root zone
- o the target realm (also never expressed in the 'transited' field)

When using DANE for validation of the issuer's certificate the target SHOULD represent the transit path as hierarchical from the source realm's domain to the root domain, then direct from there to the target's realm.

The notional transit path for a given client principal MUST be encoded as usual, using the Kerberos X.500 and domain-style representations of PKIX issuer names and DNS domainnames as faithfully to the original as possible.

[[anchor4: QUESTION: Do we need a 100% faithful representation of the transit path?]]

2.6. Exchange of Long-Term Cross-Realm Symmetric Keys

A KDC can acquire a TGT using PKCROSS whose session key then becomes the long-lived, persistent symmetric key for a cross-realm principal from the source realm to the target realm ("krbtgt/TARGET@SOURCE").

To do this the KDC MUST set the USE-SESSION-KEY-AS-REALM-KEY KDCOptions flag (TBD) in its request for an ITGT from the target realm. As usual, the target realm's KDC MUST validate the client principal's certificate. The target realm's KDC MUST NOT return a TGS-REP until the new principal is committed to its principal database, and MUST set the endtime of the ITGT to the time at which the source realm may begin using the new symmetrically-keyed

principal.

The source realm's KDC MUST commit the new principal to its principal database and MUST NOT begin using the new principal's long-term keys until the new principal is available to all KDCs for the source realm and the endtime of the ITGT passes.

Target KDCs SHOULD require manual pre-approval of such new cross-realm principals. In small, isolated environments a KDC MAY be configured to pre-approve all such new principals.

By default, source KDCs SHOULD NOT automatically request long-term keying of cross-realm principals.

3. Security Properties

The proposed PKCROSS protocol has several useful properties described below.

3.1. Automatic Cross-Realm Keying

No more manual keying of cross-realm principals via exchanging passwords in-person on a telephone call (or similar).

3.2. Scalability

Kerberos with commonplace symmetrically-keyed hierarchical cross-realm trusts can scale to a large universe of realms, but only if there are top-level realms that are willing to pair-wise trust and "child" realms. Such top-level realms do not exist in practice, leading to an $O(N^2)$ scaling problem for most two-label realms.

Leveraging a PKI, such as a PKIX PKI [RFC5280] or a DNSSEC PKI [RFC4033] removes the need for either top-level realms (which are not likely to ever be operated as commercial or even non-profit entities) or $O(N^2)$ pair-wise cross-realm symmetric keying.

The cost of this is having to add PKI trust paths to Kerberos trust paths (though the resulting trust path length need not be much different than before).

3.2.1. Simplified trust routing

For clients, relying on referrals (and TGS-driven PKCROSS) and/or client-driven PKCROSS will greatly reduce the need for client-side trust routing information.

Even KDCs won't need trust routing information.

3.2.2. Simplified trust path validation

For services that accept hierarchical trust paths, PKCROSS will greatly reduce the complexity of trust path / transit validation. Such services that also trust DANE/DNSSEC will need no trust path validation information for any clients using PKCROSS to reach the service. In many cases a very simple policy expressed in terms of whitelists of top-level domains (TLDs) or near top-level domains traversed, trust anchor sets, and trust of all zero-length transit paths, will suffice.

3.3. Privacy Protection relative to home realm

This protocol protects the privacy of client principals vis-a-vis their home realms, when the clients use the client-driven PKCROSS protocol.

This feature is generally and naturally available in PKI, and as this protocol is based on a kerberized certification authority, this protocol inherits this privacy feature from PKI.

The realms visited by the client may, of course, inform the client's home realm, but in the event that they don't, the client does gain this small measure of privacy. Of course, the privacy-conscious client SHOULD attach an OCSP Response [RFC6960] to its PKINIT request, per [RFC4557].

4. Application Programming Interface Considerations

Improved scalability for Kerberos realm traversal implies larger Kerberos universes, and the larger a universe of trust the more important it is to have useful and expressive local policy for evaluating the trustworthiness of any given transit path. Because in most applications local policy should be a component external to the application, there is mostly no impact on APIs here. However, an implementation may wish to provide applications with interfaces for specifying policies, either named or by value.

4.1. GSS-API Considerations

The naming attributes [RFC6680] defined in [I-D.williams-kitten-generic-naming-attributes] provide access to information about transit paths.

Note that information about how PKCROSS was used to establish symmetrically-keyed cross-realm principals is lost and will not appear in the transit path in tickets issued by KDCs reached via such cross-realm principals.

5. Security Considerations

[[anchor5: All the security considerations of Kerberos and PKI apply. Security considerations are discussed throughout this document.]]

Scaling up the universe of realms reachable via any trust path necessarily dilutes trust overall, but not for specific paths. On the other hand, by shortening transit path lengths trust can be improved, though some short transit paths will have been symmetrically keyed using this PKCROSS protocol and therefore will be longer than they appear to be. These are subjective notions of trust, of course.

5.1. Loss of Cross-Realm Principal Trust Establishment Information

Once a cross-realm principal is symmetrically keyed the transit path used to automatically key that principal will no longer appear in subsequent cross-realm tickets issued by the target.

The Kerberos transit path encodes only realmnames (including X.500-style names, thus PKIX certificate subject and issuer names), and lacks any public key information that might be useful for pinning. However, the certificate validation path for each realm in a transit path SHOULD be included in the transit path.

5.2. On the Need for a Common Transit Path Policy Language

There are no standard ways to express authorization policies for trust transit paths for either Kerberos nor PKI. A standard language for this would be extremely useful. Such a language should allow for the expression of policies for both, clients and services. Such a language should allow for the expression of complex realm/domain/other naming, and should allow for HSTS-style pinning [add references -Nico]. Such a language should allow for multiple paths where desired, and should allow for more than path rejection: it should also allow for reducing the entitlements assigned to a peer/realm for authorization purposes.

The need for a standard transit path policy expression language is not new, and such a language is broadly and generally needed. Therefore such a language is outside this document's scope.

PKCROSS can greatly simplify the process of validating a ticket's trust path, first by shortening the number of realms involved to two (in the typical case), maybe three, second by making the actual trust path (PKI or DANE) hierarchical, thus hopefully leaving much less policy to express in a transit path policy language: whitelists of domains and sub-domains, perhaps. But a common language would still

be desirable.

5.3. On the Need for Trust Routing

A common language for trust routing is not necessary in a purely hierarchical world, as in DANE. But since it's likely that there will be some non-hierarchical, non-zero-length transit paths in many deployments for a long time to come, a common language for trust routing would be desirable as well. Routing protocols normally used for network addresses could be used for discovery and distribution of trust routing information as well. But note that there are subtle differences between trust routing and trust path validation, even though in traditional Kerberos deployments the same information is used for both, with the trust validation policy effectively being that the client must have taken the shortest, highest-priority path specified in "capaths" configuration.

6. IANA Considerations

[[anchor6: Allocate the new KDCOptions flag (USE-SESSION-KEY-AS-REALM-KEY) and authorization-data element (AD-CLIENT-CERTIFICATE), as well as the new EKU id-pkcross-issuer.]]

7. TODO

- o Provide a normative reference for DANE stapling.

8. Acknowledgements

Although the author arrived at this "kx509 + PKINIT == PKCROSS" idea independently, it is not an original idea. Henry Hotz and Jeffrey Altman each conceived the same idea years earlier. It is a relatively obvious idea when taking into account efforts to bridge disparate security mechanisms and credentials infrastructures.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.
- [RFC4556] Zhu, L. and B. Tung, "Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)", RFC 4556, June 2006.
- [RFC4557] Zhu, L., Jaganathan, K., and N. Williams, "Online Certificate Status Protocol (OCSP) Support for Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)", RFC 4557, June 2006.
- [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, May 2008.
- [RFC6680] Williams, N., Johansson, L., Hartman, S., and S. Josefsson, "Generic Security Service Application Programming Interface (GSS-API) Naming Extensions", RFC 6680, August 2012.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, August 2012.
- [RFC6717] Hotz, H. and R. Allbery, "kx509 Kerberized Certificate Issuance Protocol in Use in 2012", RFC 6717, August 2012.
- [I-D.williams-kitten-generic-naming-attributes]
Williams, N., "Generic Naming Attributes for the Generic Security Services Application Programming Interface (GSS-API)", draft-williams-kitten-generic-naming-attributes-01 (work in progress), August 2013.

9.2. Informative References

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

- [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, June 2013.

Author's Address

Nicolas Williams
Cryptonector, LLC

Email: nico@cryptonector.com

