Workgroup: Network Working Group
Internet-Draft: draft-vanrein-dnstxt-krb1-11
Published: 8 November 2022
Intended Status: Informational
Expires: 12 May 2023
Authors: R. Van Rein
ARPA2.net

Declaring Kerberos Realm Names in DNS (_kerberos TXT)

Abstract

This specification defines a method to determine Kerberos realm names for services that are known by their DNS name. Currently, such information can only be found in static mappings or through educated guesses. DNS can make this process more flexible, provided that DNSSEC is used to assure authenticity of resource records.

Status of This Memo

This Internet-Draft will expire on 12 May 2023.

Copyright Notice

Table of Contents

1. Introduction

When a Kerberos client contacts a service, it needs to acquire a service ticket, and for that it needs to contact the KDC for a realm under which the service is run. To map a service name into a realm name and then into a KDC, clients tend to use static mappings or educated guesses; the client's KDC may or may not be involved in this process. Through DNS, the static mappings could be replaced by dynamic lookups, and migrate from local client configuration into the hands of the party administrating a server's presence in DNS. This brings improved flexibility and centralisation, which is operationally desirable.

Two mappings are needed for a client to contact a service. One is a mapping from the FQDN of a service to its realm name; the other is a mapping from the realm name to the Kerberos-specific services such as the KDC. The latter mapping is published in SRV records [RFC4120] and such traffic is usually protected by Kerberos itself. The first mapping however, has hitherto not been standardised and is ill-advised over unsecured DNS because the published information is then neither validated by DNS nor does it lead to a protocol that could provide end-to-end validation for it.

With the recent uprise of DNSSEC, it is now possible to make a reliable judgement on the authenticity of data in DNS, which enables the standardisation of the first mapping in the form of resource records under DNSSEC.

This specification defines a method to publish and process Kerberos realm names in TXT resource records. These records hold a case-sensitive string with the realm name. This has been informally described and practiced, but generally considered insecure; adding DNSSEC means that much of this existing practice can now be trusted.

It is suggested to use the name "_kerberos TXT" to informally refer to the style of using DNS that is introduced in this specification.

2. Defining _kerberos TXT Resource Records

This specification uses the TXT resource record type in DNS to represent a Kerberos realm name. The corresponding RDATA format is as follows:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
/                   REALMNAME                   /
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

The REALMNAME is represented as a <character-string> [RFC1035] which starts with a single-byte length, followed by as many bytes of realm name as the length byte's value. The RDATA field therefore has a length of 1 up to 256 bytes, to hold a realm name of 0 up to 255 bytes. For instance, a realm EXAMPLE.ORG would be represented with the following RDATA, written in the notation for unknown resource record types [RFC3597]:

```
\# 12 ( 0b 45 58 41 4d 50 4c 45 2e 4f 52 47 )
```

The REALMNAME represents a Kerberos realm name [Section 6.1 of [RFC4120]], not a DNS name; invalid names SHOULD be ignored. The empty string is considered an invalid REALMNAME, and it should be noted that a REALMNAME may exceed the size constraints of a DNS name.

The TXT record can hold one or more <character-string> values in an ordered sequence, and implementations of this specification MUST NOT reject TXT records with multiple <character-string>s. This specification only describes the meaning of the first <character-string> as a REALMNAME, and leaves the interpretation of further <character-string>s to future specifications. Until these specifications are adopted, master zone files SHOULD NOT introduce these extra <character-string>s. If such future specifications intend to specify Kerberos aspects that do not include a realm name, then they can mention an invalid realm name such as an empty <character-string>.

Though any style of realm name may be published as _kerberos TXT, it is common for realm names in Kerberos to follow the domain style [Section 6.1 of [RFC4120]], in which case they look like DNS names but are case sensitive; unlike the DNS names used as lookup keys in the DNS hierarchy, the REALMNAME format follows the <character-string> format in being case-sensitive. Even for domain-style realm names, there is no required relationship (such as partial overlap) between the realm name and the DNS name at which a TXT record is found.

In fact, the <character-string> format is a binary format, and DNS notation \DDD [Section 5.1 of [RFC1035]] exists to put arbitrary bytes in the string notation. This binary format leaves the door ajar for future internationalisation of Kerberos realm names. Realm names are defined with the KerberosString type [Section 5.2.1 of [RFC4120]] which is an ASN.1 GeneralString, but its specification currently advises to constrain the use of this string type to an IA5String (basically using only the first 128 codes of the ASCII table) to avoid interoperability problems. After the <character-string>'s length byte, the REALMNAME holds the value of the GeneralString, but not its preceding ASN.1 tag and length.

It is worth noting that the ESC "%" "G" prefix [TODO:xref target="ISO2022"/] can be used to introduce an UTF8String in a GeneralString, and that implementations exist that insert UTF8String values in KerberosString fields without even that escape. All this precedes formal standardisation of internationalisation, but it suggests that the RDATA definition for TXT can be supportive of future internationalisation of realm names, even if the current advised use is limited to the value of an IA5String.

It is possible to create a TXT record for any _kerberos-prefixed DNS name, but this specification only provides query procedures for host names and domain names. The use with a domain name has the additional use of denoting the precise spelling for a realm name under its DNS-mapped name. DNS-mapped names currently would not modify more than the case of a DNS name, and even that is only done as the result of DNS compression [RFC4343]; but in a future with internationalised realm names there might be more to reconstruct, in which case this facility is likely to be helpful.

The format for the resource data in master zone files is standard for DNS [RFC1035]. The TXT record is a general record and was not especially designed for this purpose. The reason to use it nonetheless is that it is an existing practice; the particular use specified here is distinguished from comments in TXT records by always prefixing a _kerberos label to a DNS name. An example declaration of realm name EXAMPLE.ORG for a server named imap.example.org would be:

```
imap.example.org.            IN AAAA  2001:db8::143
_kerberos.imap.example.org.  IN TXT   "EXAMPLE.ORG"
```

The RDATA for this TXT record is shown above, in the generic RDATA section notation.

3. Publishing Kerberos Realm Names

Zones that intend to provide applications with Kerberos realm names through _kerberos TXT records SHOULD protect them with DNSSEC.

Operators SHOULD NOT define more than one valid realm name for a given domain or host name.

Note that _kerberos TXT records with wildcard names will not work. All host names and most domain names define at least one resource record (of any type) with the name that the wildcard should cover. These defined names cause the wildcards to be suppressed [RFC4592] from DNS responses, even when querying a non-existent TXT record.

4. Querying Kerberos Realm Names

This section defines a procedure for determining the Kerberos realm names for a server with a given host name or domain name, as well as for a DNS-mapped realm name. This specification does not impose any restriction on the additional use of other-than-DNS methods for for obtaining a realm name.

When applications know their server host name, perhaps because it is mentioned in a URL or in a ticket as a service principal name, or when applications know a domain name for which they intend to learn the realm name, they resolve the TXT record in DNS for the name, prefixed with a _kerberos label.

Since DNS in general cannot be considered secure, the client MUST validate DNSSEC and it MUST dismiss any DNS responses that are Insecure, Bogus or Indeterminate [Section 5 of [RFC4033]]. Only the remaining Secure responses are to be taken into account. This specification does not require that the DNS client validates the responses by itself, but a deployment of _kerberos TXT records SHOULD NOT accept DNS responses from a trusted validating DNS resolver over untrusted communication channels.

In addition to the above, the absense of a _kerberos DNS record may be meaningful for security decisions. If such cases, the only denial of existence of the _kerberos TXT records MUST be authenticated denial.

Only the first lt;character-string> of a _kerberos TXT record is considered; any further ones are silently ignored under this specification. In addition, invalid realm names such as they empty string are silently ignored.

To give one possible implementation, a Kerberos client or its KDC may send DNS queries with the Authentic Data (AD) bit set to enable DNSSEC [Section 5.7 of [RFC6840]], and thereby request that the

Authenticated Data bit is set in the response to indicate [RFC3655] the Secure state for answer and authority sections of the response. When the DNS traffic to and from the validating resolver is protected, for instance because the validating resolver is reached over a loopback interface, then the Kerberos client or its KDC has implemented the requirements for Secure use of the answer and authority sections in DNS responses.

When no Secure DNS responses are received when the DNS query times out, then the TXT query MUST be terminated without extracting realm names from DNS. This termination MAY be done immediately upon receiving Secure denial for the requested TXT record. TXT query termination need not be fatal; non-DNS procedures may exist to find a realm name, including the current practice of static mappings and educated guessing.

5. Efficiency Considerations

The lookup of _kerberos TXT records can be done by the Ticket Granting Service of a KDC, which can respond with a Server Referral [Section 8 of [RFC6806]] to Kerberos clients that enable canonicalization. This can be used for clients that are not setup to query DNS as specified above, and that will assume that a service is running under the client's realm. The caching of DNS records, their validation and possibly realm-crossover caching at the KDC can all benefit the response time for future lookups by other Kerberos clients.

6. Privacy Considerations

This specification barely publishes new information in DNS, with the exception of markation of Kerberised services. When this is considered unattractive from a privacy viewpoint, it may be better to rely on the existing static tables for spreading this information in a more controlled manner.

7. Security Considerations

There is no restriction for _kerberos TXT records to mention realm names that map back to DNS names in a disjoint part of the DNS hierarchy. The records could therefore specify realm names for a service even if the service is not recognised by the realm. The KDC for the appointed realm would be very clear about that when trying to procure a service ticket, so there is no anticipated security issue with such misguided use of _kerberos TXT records.

The general point is that the use of DNSSEC makes Kerberos accept authentic information from the party that publishes the _kerberos TXT record, and that party could specify improper realm names or drop realm names that are vital to the client. This is not expected to be a security risk either; the party publishing the _kerberos TXT record is the same party that publishes the service's records, namely its DNS operator. By publishing the service's record in DNS, this operator already has potential control over service denial and other man-in-the-middle attacks, so the _kerberos TXT record does not add any new powers of abuse.

When an external attacker would be permitted to spoof a _kerberos TXT record in a victim's DNS, then it could be possible for that attacker to convince the client that the attacker is the authentic provider for the service. Additional spoofing of host name references could then complete the attack. This has been mitigated by strictly requiring Secure validation results from a DNSSEC-aware resolver for all _kerberos TXT records.

8. IANA Considerations

None.

9. References

9.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <https://www.rfc-editor.org/info/rfc3597>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The Kerberos Network Authentication Service (V5)", RFC 4120, DOI 10.17487/RFC4120, July 2005, <https://www.rfc-editor.org/info/rfc4120>.

[RFC4343] Eastlake 3rd, D., "Domain Name System (DNS) Case Insensitivity Clarification", RFC 4343, DOI 10.17487/RFC4343, January 2006, <https://www.rfc-editor.org/info/rfc4343>.

[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, <https://www.rfc-editor.org/info/rfc6806>.

9.2. Informative References

[RFC3655] Wellington, B. and O. Gudmundsson, "Redefinition of DNS Authenticated Data (AD) bit", RFC 3655, DOI 10.17487/RFC3655, November 2003, <https://www.rfc-editor.org/info/rfc3655>.

[RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", RFC 4592, DOI 10.17487/RFC4592, July 2006, <https://www.rfc-editor.org/info/rfc4592>.

[RFC6840] Weiler, S., Ed. and D. Blacka, Ed., "Clarifications and Implementation Notes for DNS Security (DNSSEC)", RFC 6840, DOI 10.17487/RFC6840, February 2013, <https://www.rfc-editor.org/info/rfc6840>.

Appendix A. Acknowledgements

Thanks are due to the Kitten Workgroup for discussions during the creation of this document. Especially Greg Hudson, Nico Williams and Viktor Dukhovni have provided useful input.

This work was conducted under a grant from the programme "[veilig] door innovatie" from the government of the Netherlands. It has also been liberally supported by the NLnet Foundation.

Author's Address

Rick van Rein
ARPA2.net
Haarlebrink 5
Enschede

Email: rick@openfortress.nl