

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: September 12, 2016

R. Van Rein  
ARPA2.net  
March 11, 2016

Declaring Kerberos Realm Names in DNS (\_kerberos TXT)  
draft-vanrein-dnstxt-krb1-07

## 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 is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on September 12, 2016.

## Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect 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.

Internet-Draft

\_kerberos TXT

March 2016

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Defining _kerberos TXT Resource Records . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Querying Kerberos Realm Names . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Publishing Kerberos Realm Names . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Efficiency Considerations . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Privacy Considerations . . . . .	<a href="#">8</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">8</a>
<a href="#">8.</a>	IANA Considerations . . . . .	<a href="#">9</a>
<a href="#">9.</a>	References . . . . .	<a href="#">9</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">9</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">10</a>
<a href="#">Appendix A.</a>	Acknowledgements . . . . .	<a href="#">10</a>
	Author's Address . . . . .	<a href="#">10</a>

[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. Multiple TXT records may be

defined as alternatives in contexts that welcome this; they suggest iterating over the values until a ticket can be procured.

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, 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 Kerberos realm name MUST NOT have an additional trailing dot unless it is actually part of the Kerberos realm name. The realm name is never relative to the DNS name at which the TXT record was found. Finally, there is no required relationship (such as partial overlap) between the realm name and the DNS name at which the TXT record was found.

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>.

Realm names in Kerberos are often domain-styled [[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.

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.

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 too many middle boxes suppress unknown DNS resource record types; we distinguish the particular use specified here by always prefixing a fixed `_kerberos` label to the DNS name that is being investigated. An example declaration of realm name `EXAMPLE.ORG` for a server named `imap.example.org` would be:

```
_kerberos.imap.example.org. IN TXT "EXAMPLE.ORG"
```

The RDATA for this TXT record has already been shown above, in the generic RDATA section notation.

As another example, an IMAP server may be specified through SRV records [[RFC2782](#)], in which case the domain holding those records needs to hold the TXT record, rather than the pointed-to server name, for example:

```
_imap._tcp.example.org. IN SRV 10 10 143 imap.example.org.  
_kerberos.example.org.  IN TXT "EXAMPLE.ORG"
```

In operational contexts where both the SRV record resolution as the direct host name may be applied, it is also possible to define both kinds of TXT records; it depends on the application and its configuration in which single location it will look for the Kerberos realm name.

Multiple TXT records may be supplied for the same `_kerberos`-prefixed DNS name; these represent alternatives that can be tried by the DNS client. Since DNS delivers the resource records as a set without any particular order, the DNS client can choose freely in what order to process these records.

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.

It is worth noting that the authority of a `_kerberos` TXT record is determined by the domain name holding it, not the realm contained in it. Since there is no defined relationship between these two pieces of information, proper security may require reverse lookups in DNS and/or Kerberos.

### 3. 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.

When applications directly know their server host name, perhaps because it is mentioned in a URL or in a ticket as a service principal name, they will lookup the TXT record in DNS at the server host name, prefixed with a `_kerberos` label.

When applications locate their servers through a domain name, for example via MX [[RFC1035](#)] or SRV records [[RFC2782](#)], a choice between the use of the domain name or the appointed server host names is to be made. In these situations, the TXT query MUST use only the domain name, without adding prefixes such as `_Service._Proto` for SRV records, but with an additional label `_kerberos` prefixed. The appointed servers for a domain service MUST NOT be used to locate TXT records, since these may reside under operational control foreign to the service domain, which could introduce a security problem; in

addition, the possible shared use of a host in the SRV records of many domains would mean that they would have to specify alternatives, where it is not always certain that clients will be able to process more than one such alternative. The domain name holding the SRV record is closer to the client's context and provides a better-targeted location for the `_kerberos` TXT record.

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 absence 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.

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.

#### 4. Publishing Kerberos Realm Names

To facilitate the dynamic lookup procedures described in this specification, `_kerberos` TXT records should generally be published at the domain names and server host names that a Kerberos application may want to approach. Their location is prefixed with an extra `_kerberos` label to distinguish the specific use of the TXT record defined herein.

As an exception to this general rule, server host names may be exempted if they are always looked up through a domain name. Note however, that most query mechanisms that lookup domain names through special resource record types have fallbacks to plain server host name lookups. When such direct server host names are supported in a deployment of those applications, usually meaning that these server host names may occur in Kerberos service principal names, and when the realm name cannot be detected through a preferred mechanism for

such server host names, then an additional `_kerberos` TXT record located at the server name is also going to be helpful.

There may be situations where a domain represents a Kerberos realm, but uses domain-style references such as MX and SRV records to point to foreign server host names. Such foreign server host names are often unsuitable to accompany with a `_kerberos` TXT record, so it is advisable to either not refer to the server host names directly in service tickets, or to see to it that a realm name will be supplied by the KDC.

These instructions are somewhat flexible, as a result of the pre-existing static mechanisms for finding a realm name for a given service principal name. Assuming that the more flexible approach of `_kerberos` TXT records is preferred, a useful functionality test would be to access the targeted applications from Kerberos clients without any static configuration.

Zones that intend to provide applications with Kerberos realm names through `_kerberos` TXT records MUST be protected by DNSSEC to make them usable to those applications.

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.

## [5.](#) Efficiency Considerations

The `_kerberos` TXT records are intended to define realm names for servers whose DNS information is not statically mapped in a Kerberos setup. This may simplify operational control of such static mappings. It may also introduce more dynamicity, which is expected to be useful for such things as realm crossover.

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.

The use of `_kerberos` TXT records helps the Kerberos client and/or its KDC to focus on the realm for which a service ought to have a name defined. This generally provides a deterministic path, reducing the effort of guessing.

The definition of `_kerberos` TXT records under a domain rather than under a server host name (for cases where a choice exists, such as for applications that use SRV or MX records) should limit the choice of realm names, because the domain is not farther away from the client's domain/realm than the server host name.

Constraining the number of paths to a target realm to one is especially useful for situations with realm crossover when this involves public-key cryptography, because such algorithms are much slower than the symmetric algorithms that are normally used for Kerberos.

## [6.](#) Privacy Considerations

The information contained in a `_kerberos` TXT record for a domain is barely more than the domain's DNS name already holds; the `_kerberos` TXT record spells out the case-sensitive realm name, and implies the use of Kerberos. The use of Kerberos is usually also implied by the presence of certain SRV records.

The `_kerberos` TXT records also marks the names of Kerberised servers. Internal use may prohibit such disclosure, but in those use cases it is often possible to rely on existing mechanisms for guessing a realm name, including simply using the realm name under which a client logged on with Kerberos.

## [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, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), DOI 10.17487/RFC2782, February 2000, <<http://www.rfc-editor.org/info/rfc2782>>.
- [RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", [RFC 3597](#), DOI 10.17487/RFC3597, September 2003, <<http://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, <<http://www.rfc-editor.org/info/rfc4033>>.

Internet-Draft

\_kerberos TXT

March 2016

- [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>>.
- [RFC4343] Eastlake 3rd, D., "Domain Name System (DNS) Case Insensitivity Clarification", [RFC 4343](#), DOI 10.17487/RFC4343, January 2006, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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 added useful input.

Author's Address

Rick van Rein  
ARPA2.net  
Haarlebrink 5  
Enschede, Overijssel 7544 WP  
The Netherlands

Email: rick@openfortress.nl

Van Rein

Expires September 12, 2016

[Page 10]