
Kerberos Concepts

Release 1.11

MIT

CONTENTS

| | | |
|----------|--|-----------|
| 1 | keytab | 1 |
| 1.1 | Default keytab | 1 |
| 1.2 | Default client keytab | 1 |
| 2 | replay cache | 3 |
| 2.1 | Background information | 3 |
| 2.2 | Default rcache type | 3 |
| 2.3 | Performance issues | 4 |
| 3 | stash file | 5 |
| 4 | Supported date and time formats | 7 |
| 4.1 | Time duration | 7 |
| 4.2 | getdate time | 7 |
| 4.3 | Absolute time | 8 |
| | Index | 11 |

KEYTAB

A keytab (short for “key table”) stores long-term keys for one or more principals. Keytabs are normally represented by files in a standard format, although in rare cases they can be represented in other ways. Keytabs are used most often to allow server applications to accept authentications from clients, but can also be used to obtain initial credentials for client applications.

Keytabs are named using the format *type : value*. Usually *type* is `FILE` and *value* is the absolute pathname of the file. Other possible values for *type* are `SRVTAB`, which indicates a file in the deprecated Kerberos 4 `srvtab` format, and `MEMORY`, which indicates a temporary keytab stored in the memory of the current process.

A keytab contains one or more entries, where each entry consists of a timestamp (indicating when the entry was written to the keytab), a principal name, a key version number, an encryption type, and the encryption key itself.

A keytab can be displayed using the *klist(1)* command with the `-k` option. Keytabs can be created or appended to by extracting keys from the KDC database using the *kadmin(1)* *ktadd* command. Keytabs can be manipulated using the *ktutil(1)* and *k5srvutil(1)* commands.

1.1 Default keytab

The default keytab is used by server applications if the application does not request a specific keytab. The name of the default keytab is determined by the following, in decreasing order of preference:

1. The **KRB5_KTNAME** environment variable.
2. The **default_keytab_name** profile variable in *libdefaults*.
3. The hardcoded default, *DEFKTNAME*.

1.2 Default client keytab

The default client keytab is used, if it is present and readable, to automatically obtain initial credentials for GSSAPI client applications. The principal name of the first entry in the client keytab is used by default when obtaining initial credentials. The name of the default client keytab is determined by the following, in decreasing order of preference:

1. The **KRB5_CLIENT_KTNAME** environment variable.
2. The **default_client_keytab_name** profile variable in *libdefaults*.
3. The hardcoded default, *DEFCKTNAME*.

REPLAY CACHE

A replay cache (or “rcache”) keeps track of all authenticators recently presented to a service. If a duplicate authentication request is detected in the replay cache, an error message is sent to the application program.

The replay cache interface, like the credential cache and *keytab* interfaces, uses *type:value* strings to indicate the type of replay cache and any associated cache naming data to use.

2.1 Background information

Some Kerberos or GSSAPI services use a simple authentication mechanism where a message is sent containing an authenticator, which establishes the encryption key that the client will use for talking to the service. But nothing about that prevents an eavesdropper from recording the messages sent by the client, establishing a new connection, and re-sending or “replaying” the same messages; the replayed authenticator will establish the same encryption key for the new session, and the following messages will be decrypted and processed. The attacker may not know what the messages say, and can’t generate new messages under the same encryption key, but in some instances it may be harmful to the user (or helpful to the attacker) to cause the server to see the same messages again a second time. For example, if the legitimate client sends “delete first message in mailbox”, a replay from an attacker may delete another, different “first” message. (Protocol design to guard against such problems has been discussed in [RFC 4120](#).)

Even if one protocol uses further protection to verify that the client side of the connection actually knows the encryption keys (and thus is presumably a legitimate user), if another service uses the same service principal name, it may be possible to record an authenticator used with the first protocol and “replay” it against the second.

The replay cache mitigates these attacks somewhat, by keeping track of authenticators that have been seen until their five-minute window expires. Different authenticators generated by multiple connections from the same legitimate client will generally have different timestamps, and thus will not be considered the same.

This mechanism isn’t perfect. If a message is sent to one application server but a man-in-the-middle attacker can prevent it from actually arriving at that server, the attacker could then use the authenticator (once!) against a different service on the same host. This could be a problem if the message from the client included something more than authentication in the first message that could be useful to the attacker (which is uncommon; in most protocols the server has to indicate a successful authentication before the client sends additional messages), or if the simple act of presenting the authenticator triggers some interesting action in the service being attacked.

2.2 Default rcache type

There is currently only one implemented kind of replay cache, called **dfi**. It stores replay data in one file, occasionally rewriting it to purge old, expired entries.

The default type can be overridden by the **KRB5RCACHETYPE** environment variable.

The placement of the replay cache file is determined by the following:

1. The **KRB5RCACHEDIR** environment variable;
2. If KRB5RCACHEDIR is unspecified, on UNIX, the library will fall back to the environment variable **TMPDIR**, and then to a temporary directory determined at configuration time such as */tmp* or */var/tmp*; on Windows, it will check the environment variables *TEMP* and *TMP*, and fall back to the directory *C:*.

2.3 Performance issues

Several known minor performance issues that may occur when replay cache is enabled on the Kerberos system include: delays due to writing the authenticator data to disk slowing down response time for very heavily loaded servers, and delays during the rewrite that may be unacceptable to high-performance services.

For use cases where replays are adequately defended against for all protocols using a given service principal name, or where performance or other considerations outweigh the risk of replays, the special replay cache type “none” can be specified:

```
KRB5RCACHETYPE=none
```

It doesn't record any information about authenticators, and reports that any authenticator seen is not a replay.

STASH FILE

The stash file is a local copy of the master key that resides in encrypted form on the KDC's local disk. The stash file is used to authenticate the KDC to itself automatically before starting the *kadmind(8)* and *krb5kdc(8)* daemons (e.g., as part of the machine's boot sequence). The stash file, like the keytab file (see *keytab_file*) is a potential point-of-entry for a break-in, and if compromised, would allow unrestricted access to the Kerberos database. If you choose to install a stash file, it should be readable only by root, and should exist only on the KDC's local disk. The file should not be part of any backup of the machine, unless access to the backup data is secured as tightly as access to the master password itself.

Note: If you choose not to install a stash file, the KDC will prompt you for the master key each time it starts up. This means that the KDC will not be able to start automatically, such as after a system reboot.

SUPPORTED DATE AND TIME FORMATS

4.1 Time duration

This format is used to express a time duration in the Kerberos configuration files and user commands. The allowed formats are:

| Format | Example | Value |
|------------------------------------|--------------------|--------------------|
| <code>h:m[:s]</code> | <code>36:00</code> | 36 hours |
| <code>NdNhNmNs</code> | <code>8h30s</code> | 8 hours 30 seconds |
| <code>N (number of seconds)</code> | <code>3600</code> | 1 hour |

Here *N* denotes a number, *d* - days, *h* - hours, *m* - minutes, *s* - seconds.

Note: The time interval should not exceed 2147483647 seconds.

Examples:

Request a ticket valid for one hour, five hours, 30 minutes and 10 days respectively:

```
kinit -l 3600
kinit -l 5:00
kinit -l 30m
kinit -l "10d 0h 0m 0s"
```

4.2 getdate time

Some of the `kadmin` and `kdb5_util` commands take a date-time in a human-readable format. Some of the acceptable date-time strings are:

| | Format | Example |
|--|--------------|----------|
| Date mm/dd/yyyy dd-mm-yyyy yyyy-mm-dd | mm/dd/yy | 07/27/12 |
| | Jul 27, 2012 | |
| | 2012-07-27 | |
| Absolute time hh:mm[:ss] | HH:mm[:ss]pp | 08:30 PM |
| | 20:30 | |
| Relative time | N tt | 30 sec |
| Time zone z | Z | EST |
| | -0400 | |

(See *Abbreviations used in this document.*)

Examples:

Create a principal that expires on the date indicated:

```
addprinc test1 -expire "3/27/12 10:00:07 EST"
addprinc test2 -expire "January 23, 2015 10:05pm"
addprinc test3 -expire "22:00 GMT"
```

Add a principal that will expire in 30 minutes:

```
addprinc test4 -expire "30 minutes"
```

4.3 Absolute time

This rarely used date-time format can be noted in one of the following ways:

| Format | Example | Value |
|------------------------|----------------------|--------------------------|
| yyyymmddhhmmss | 20141231235900 | |
| yyyy.mm.dd.hh.mm.ss | 2014.12.31.23.59.00 | |
| yymmddhhmmss | 141231235900 | One minute before 2015 |
| yy.mm.dd.hh.mm.ss | 14.12.31.23.59.00 | |
| dd-month-yyyy:hh:mm:ss | 31-Dec-2014:23:59:00 | |
| hh:mm:ss | 20:00:00 | 8 o'clock in the evening |
| hhmmss | 200000 | |

(See *Abbreviations used in this document.*)

Example

Set the default expiration date to July 27, 2012 at 20:30

```
default_principal_expiration = 20120727203000
```

4.3.1 Abbreviations used in this document

month : locale's month name or its abbreviation;

dd : day of month (01-31);

HH : hours (00-12);

hh : hours (00-23);

mm : in time - minutes (00-59); in date - month (01-12);

N : number;

pp : AM or PM;

ss : seconds (00-60);

tt : time units (hours, minutes, min, seconds, sec);

yyyy : year;

yy : last two digits of the year;
Z : alphabetic time zone abbreviation;
z : numeric time zone;

Note:

- If the date specification contains spaces, you may need to enclose it in double quotes;
 - All keywords are case-insensitive.
-

INDEX

R

RFC

RFC 4120#section-10, 3