Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August 12, 2012

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

Feb 9, 2012

POSIX Authorization Data
draft-ietf-krb-wg-pad-01

Abstract

This document proposes a Kerberos Authorization Data element containing user and group directory information similar to that provided by RFC 2307, typically used by POSIX and POSIX-like systems in the course of login type activities.

Table of Contents

1. Introduction 3
2. Requirements Language 3
3. Use-Case: Cross-Realm Directory Services 3
4. A POSIX Authorization Structure for Kerberos V5 4
4.1. Attributes 4
4.2. PAD-Realm 5
4.3. PAD-DNS-Domain 5
4.4. PAD-Short-Domain 5
4.5. PAD-UDID 5
4.6. PAD-Posix-Username 6
4.7. PAD-Posix-UID 6
4.8. PAD-Posix-GID 6
4.9. PAD-Posix-Gecos 6
4.10. PAD-Posix-Homedir 6
4.11. PAD-Posix-Shell 6
4.12. PAD-Fullname 7
4.13. PAD-AlternateNames 7
4.14. PAD-Groups 7
4.15. PAD Mapped Attributes 8
4.16. RFC2307 references for Directory Services backed KDCs . . 8
4.16.1. PAD-Posix-Username as 'uid' 8
4.16.2. PAD-Posix-UID as 'uidNumber' 8
4.16.3. PAD-Posix-GID as 'gidNumber' 9
4.16.4. PAD-Posix-Gecos as 'gecos' 9
4.16.5. PAD-Posix-Homedir as 'homeDirectory' 9
4.16.6. PAD-Posix-Shell as 'loginShell' 9
5. Validation 9
6. Encoding 9
6.1. PAD Format 9
7. Data Structures and Extensions 11
7.1. AD-ID-ANCHOR 11
7.2. AD-PAD-DATA 12
7.3. GSS-API Authenticator Extension 12
8. Assigned numbers 13
9. Timeouts Considerations 13
10. IANA Considerations 13
11. Security Considerations 14
12. Acknowledgements 14
13. References 14
13.1. Normative References 14
13.2. Informative References 15
Appendix A. Additional Stuff 15
Authors' Addresses 15

1. Introduction

There is an increasing need today for Kerberos to support the delivery and processing of authorization information pertaining to the principals seeking access to the servers. Kerberos today is used extensively for authentication to directory services within the Enterprise. In many cases, a directory service is implemented as a distributed database system organized across multiple realms. As such, when a client in one realm seeks access to a directory service component located within a different realm, information regarding both the identity of the client and the permissions associated with that client must be communicated across the realms. Currently there does not exist a common and standardized structure in Kerberos (V5) for conveying access control or authorization information.

This draft proposes an authorization data element for Kerberos that contains information that is useful for dynamically updating the user and group directory information in a POSIX system, usually in the course of a login type activity.

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. Use-Case: Cross-Realm Directory Services

In this section we discuss one of the primary use-case scenarios for the POSIX Authorization Data (PAD) structure within Kerberos V5. In this use-case a client principal is seeking to access a service in a different realm. Since the remote service does not have authorization information regarding the client, it needs to obtain it either from querying the directory service in its own realm or the directory service located in the client's realm. It is here that a common PAD structure becomes necessary and invaluable in order to achieve a high-degree of interoperability between directory services in distinct realms.

In this use-case a client principal C1 in realm R1 is seeking access to services (or servers) located in a different realm R2. In accessing local service S1 in realm R1 the client must first be authenticated by KDC1 in that realm. A directory service (e.g. LDAP) called D1 is used in realm R1 to perform authorization of the client, after the client has been authenticated by KDC1.

When the client prinicipal later seeks to access services or resources S2 in realm R2, following the usual Kerberos flow the client must first obtain a cross-realm TGT from KDC1 (in realm R1) and then present it to KDC2 (in realm R2) in order to obtain a service-ticket for S2. However, one immediate issue is the fact that service S2 does not have authorization information regarding the permissions or privileges of client C1 in realm R1. The service S2 could query its own directory service D2 to obtain authorization information pertaining to client C1. In the absence of such information in D2, the service S2 could then perform a cross-realm query to the directory services D1 operating in realm R1.

However, this cross-realm query from S2 to D1 is not only inefficient, but it also implies knowledge of multiple heterogenous systems by all actors. Two different realms may rely on completely different infrastructures for user information storage, ranging from different LDAP implementations with different schema conventions to NIS, SQL databases, flat files, and so on. Every service in the realm R2 would have to know what information system is in use in R1, how to reach it, how to read and eventually how to map data from it. Moreover security related aspects on the authentication of S2 by the directory D1, the authorization of S2 to make such a query, the protection of responses from D1 to S2, and so on, would have to be addressed.

This use-case illustrates the need for a common PAD structure to address this cross-realm authorization problem. In particular, the PAD structure for the cross-realm access to remote services needs to be contained or carried within cross-realm TGTs and service-tickets. Such a PAD structure needs to carry enough authorization information such that a decision can be made by service S2 in realm R2 regarding the access request originating from the client principal C1 within realm R1.

4. A POSIX Authorization Structure for Kerberos V5

4.1. Attributes

The following attributes are defined in this document:

- PAD-Realm

- PAD-DNS-Domain

- PAD-Short-Domain

- o PAD-UDID
- o PAD-Posix-Username
- o PAD-Posix-UID
- o PAD-Posix-GID
- o PAD-Posix-Gecos
- o PAD-Posix-Homedir
- o PAD-Posix-Shell
- o PAD-Fullname
- o PAD-AlternateNames
- o PAD-Groups

These are each defined and discussed further below.

4.2. PAD-Realm

The full Realm Name of the Realm the authorization information applies to.

4.3. PAD-DNS-Domain

The DNS Domain name associated to the Realm.

4.4. PAD-Short-Domain

A short domain name that uniquely identifies, within the set of trusted realms, the domain the principal belongs to. The short Domain name is useful for representation purposes in the OS. A KDC MUST verify this name is unique and correctly represnt a remote realm within its own realm and is allowed to change or remove this field during validation. This may be done to resolve name conflicts in large trust relationships.

4.5. PAD-UDID

A UDID is a Unique Domain Identifier. Ideally it universally identifies the domain as the one the following local identifiers belongs to. This is used to differentiate between local identifiers belonging to different domains/realms.

The UDID size can be dependent on the specific Domain type and imlementation. However it SHOULD be not less than 96 bits long so that chances of conflicts are relatively low. A 96 bit long identifier allows to construct a 128bit account identifier by concatenating the UDID to the local account Identifier (32bit quantity in POSIX).

For the purpose of this document the UDID is a compeltely opaque number and implementations SHOULD not try to perform any enforcement on the format of this number on receiving it.

4.6. PAD-Posix-Username

This is the user name that correspond to the kerberos principal, this is the name that SHOULD be used by the OS to represent the user. The OS may decide to prefix or suffix this name with the PAD-Domain or PAD-Realm or PAD-Short-Domain names to avoid name conflicts with local accounts.

4.7. PAD-Posix-UID

This is the UID Number associated to the user. This number is local to the domain identified by PAD-UDID.

4.8. PAD-Posix-GID

This is the Primary GID Number associated to the user. This number is local to the domain identified by PAD-UDID.

4.9. PAD-Posix-Gecos

The Gecos field for the User associated to the Principal if available. Can be omitted. If not available PAD-Fullname can be used instead.

4.10. PAD-Posix-Homedir

The home directory path relative to the local system, if available. If not available local defined defaults apply.

4.11. PAD-Posix-Shell

The default shell for the user, defined as the path of the binary relative to the local filesystem, if available. If not available local defined defaults apply.

4.12. PAD-Fullname

The full name of the user if available.

4.13. PAD-AlternateNames

Alternate names can be used by application to identify a user by means that differ from the user principal. Names are in string form and utf8 encoded [UTF-8]. In order to allow applications to recognize the name type without guesswork, alternate names are prefixed with a string followed by the colon ':' character and the name, without any space or other separation character. The following Alternate names are currently recognized: EMAIL, OS, OPENID, OAUTH It is allowed to include multiple alternate names of the same type. The order in which they are provided represent the priority within the same name type, if applications need to choose between names.

(TODO: need discussion on whether these needs labeled prefixes or explicit attributes for each alternate representation etc...)

4.14. PAD-Groups

This is a structured attribute and defines the groups the principal is member of.

The first value in the structure represents the domain UDID and is optional. If missing the domain UDID is assumed to be the one defined in the PAD-UDID attribute.

Then an array of values that define the groups as follows. Each group value includes 3 subvalues:

- (1) Name: This is the name of the group.

- (2) Type: Optional, type of group

- (3) ID: group ID.

If the type is missing it is assumed that the group is of type "Posix Group" and the follwing ID is required and represents the gid number. The type is represented through a simplified OID like type where only 2 levels are defined. 0.0 Is reserved for posix groups, and the 0 prefix is reserved to official RFX use. Additional Prefixes can be assigned to organizations that request it for their purposes. Assignment TBD.

Multiple PAD-Groups attributes can be present at the same time. A trusting KDC can augment the original user's set of groups by adding

a new PAD-Groups structure that contains groups local to the trusting domain. In this case the domain UDID is required. The domain UDID is used for gid number conflict resolution when the PAC is transmitted between services of different realms.

PAD-Groups are optional attributes and the KDC, upon PAC revalidation, may decide to remove the original attributes that do not belong to the KDC security domain in order to save space or to censor information to avoid disclosing data to services.

4.15. PAD Mapped Attributes

In POSIX, users and groups ID are not universally unique, and different Realms (even different machines within an authorization realm actually) may have overlapping and conflicting IDs. If this is the case, a trusting KDC may decide to re-map IDs coming from a foreign Realm to help services with uid/gid mapping and avoid ID conflicts that can lead to serious security issues. The original IDs are generally preserved.

If multiple PAD buffers are received and one of them contains a PAD-UDID that is recognized by the application to be the local security domain identifier, then only the mapped attributes in this buffer SHOULD be used for authorization purposes.

4.16. RFC2307 references for Directory Services backed KDCs

A few attributes contain the keyword 'Posix' in their name. These attributes are usually represented by RFC2307 in Directory Services. If the primary store for these attributes is a Directory the following equivalence with RFC2307 defined attributes can be used.

4.16.1. PAD-Posix-Username as 'uid'

The PAD-Posix-Username is the User ID, and its syntax is equivalent to the attribute named 'uid' in RFC 2307. This attribute is defined in RFC 4519 (2.39). The attribute is defined as multivalued in RFC 4519 but in this context only a single value is allowed. To define aliases refer to the attribute PAD-AlternateNames.

4.16.2. PAD-Posix-UID as 'uidNumber'

The PAD-Posix-UID is the User's Unique Identifier Number, and its syntax is equivalent to the attribute named 'uidNumber' in RFC 2307.

4.16.3. PAD-Posix-GID as 'gidNumber'

The PAD-Posix-GID is the User's Primary Group Identifier Number, and its syntax is equivalent to the attribute named 'gidNumber' in RFC 2307.

4.16.4. PAD-Posix-Gecos as 'gecos'

The PAD-Posix-Gecos is the User's Common Name, although, traditionally, this field has been used to convey additional information beyond the user's full name. Its syntax is equivalent to the attribute named 'gecos' in RFC 2307.

4.16.5. PAD-Posix-Homedir as 'homeDirectory'

The PAD-Posix-Homedir is the User's LOCAL home directory. Its syntax is equivalent to the attribute named 'homeDirectory' in RFC 2307.

4.16.6. PAD-Posix-Shell as 'loginShell'

The PAD-Posix-Shell is the User's preferred login shell. Its syntax is equivalent to the attribute named 'loginShell' in RFC 2307.

5. Validation

The PAD information is used by a client to perform authorization, therefore this information is highly sensitive and must be validated to insure no tampering has occured.

Therefore AD-PAD elements MUST always be transmitted contained within an AD-CAMMAC element

6. Encoding

The Kerberos protocol is defined in [RFC4120] using Abstract Syntax Notation One (ASN.1) [X680]. As such, this specification also uses the ASN.1 syntax for specifying both the abstract layout of the PAD attributes, as well as their encodings.

6.1. PAD Format

The information carried in the PAD needs to be augmented by some identifying information in order to tie the PAD data to a specific identity within the Kerberos Realm.

In order to allow additional authorization data to be tied together

and at the same time always verifiable we propose that the PAD is delivered as an AD element within a AD-CAMMAC.

An AuthorizationData element of type AD-ID-ANCHOR is used to bind the PAD to the ticket and the authorization data within the PAD to the specific principal. This element MUST always be present and SHOULD be validated. If this element is not available the PAD data MUST be discarded and considered untrustworthy.

The AD-ID-ANCHOR includes the full principal name, the realm, the expiration time and an optional session ID.

The ad-type for AD-PAD-ANCHOR is (TBD).

The AD-PAD-DATA include the attributes described in paragraph 4.

The ad-type for AD-PAD-DATA is (TBD).

The final structure used to deliver the PAD Data looks loosely like the following diagram.

```
==AD-CAMMAC=================================
|-------------------------------------------|
| KDC Signature (Checksum)                  |
|-------------------------------------------|
| Service Signature (Checksum)              |
|-------------------------------------------|
| Trusted Service Signature (Optional)      |
|-------------------------------------------|
| Asymmetric Key KDC Signature (Optional)   |
|-------------------------------------------|
| /-AuthorizationData SEQUENCE:----------\ |
| |                                      | |
| | 0:  --AD-ID-ANCHOR----               | |
| |    | Realm            |              | |
| |    | PrincipalName    |              | |
| |    | expiration time  |              | |
| |    | session ID       |              | |
| |     ------------------               | |
| |                                      | |
| | 1:  --AD-PAD-DATA------              | |
| |    | PAD Attributes ...|             | |
| |    | ..                |             | |
| |     -------------------              | |
| | ....    ....                         | |
| | X:  --AD-XXXXXXX-------              | |
| |    | ..                |             | |
| |     -------------------              | |
| \--------------------------------------/ |
============================================
```

Figure 1: PAD Format

7. Data Structures and Extensions

7.1. AD-ID-ANCHOR

The AD-ID-ANCHOR is intended to provide a means to bind data, carried in a AD-CAMMAC element, to a specific Identity (Principal), and optionally to a specific Ticket by using the session ID element.

```
AD-ID-ANCHOR            ::=SEQUENCE {
    p-realm             [0] Realm,
    p-name              [1] PrincipalName,
    expiration          [2] KerberosTime,
    session-id          [3] TBD
}
```

p-realm, p-name
The realm and name of the principal the authorization data elements apply to.

expiration
The Expitration Date of the Authorization Data. Normally this is the same as the original TGT expiration date.

session-id
A random number that uniquely ties any following ticket this PAD Data is associated to with the original TGT Released to the user

7.2. AD-PAD-DATA

The AD-PAD-DATA data is intended to provide a means for a Kerberos principal credentials to carry authorization data that the receiving service can use to perform authorization decisions.

```
AD-PAD-ANCHOR           ::=SEQUENCE {
     TBD
}
```

7.3. GSS-API Authenticator Extension

The Authenticator Checksum as defined in RFC 4121 limit the size of delegated credentials in the KRB_CRED message to a size of 64KiB.

In order to be able to transfer larger messgaes an extension is defined. This extension is used in stead of the Dlght/Deleg fields, and the Dlght and Deleg fileds MUST not be included when this extensions is appended to the authenticator.

The extension SHALL have the following format which is drafted according to [draft-ietf-krb-wg-gss-cb-hash-agility]:

Octet	Name	Description
0..3	ExtN	A 16bit value identifying the extension. Represented in big-endian order; Contains the hex value 0xXXXXXXXX.
4..7	Length	The length of the Extended Delegation field. Represented in big-endian order;
8..N	Data	A KRB_CRED message (N = Length + 8)

A new flag GSS_C_EXT_DELEG_FLAG with Value X is also defined. This flag is used instead of GSS_C_DELEG_FLAG when the delegated credentials are larger then 64KiB and cannot fit in the starndard Deleg field.

Implementors SHOULD use this Extensions and this flag only if the KRB_CRED message is larger than 64KiB and use the standard Deleg field otherwise.

8. Assigned numbers

TBD

9. Timeouts Considerations

Current implementations depend on very strict timeouts on obtaining AS Replies. In popular implementations the client will timeout if it doesn't receive a reply within 1 second. Adding authorization data may involve lookups to external (to the KDC) data sources. Implementors should consider whether the current timeout is still reasonbale in light of the additional processing KDCs may be required to do.

10. IANA Considerations

TBD.

11. Security Considerations

Although it is anticipated that the PAD structure itself will be carried within a ticket and thereby protected using the existing encryption methods on that ticket, there are a number of issues that have bearings on the security of the entire Kerberos realm as a whole. Some of these issues are as follows:

- o UID and GID Collisions: There is always the possibilty of collison of numbers repressing a UID and a GID. This problem can be remedied to a large degree by realms using an appropriate range selection policy and algorithms.

- o When collisions are detected the KDC or, alternatively, the receiving Service MUST be able to remap IDs so that they do not conflict with locally defined IDs

- o Transit-domain issues: The PAC must be signed by the KDC that is attaching it to a ticket with 2 different signatures. The service signature so that the service can verify its KDC validated the contents. The KDC signature, so that the OS can ask the KDC to confirm the PAD has not been modified by a less trusted service. An optional asymmetric key signature is also allowed if Keys are available in order to avoid additional roundtrips. For cross-realm tickets the "service" signature is made with the cross-realm key. When a KDC receives a PAD it is allowed to modify it in any way. It can filter out information or add information (like group memberships defined locally). A KDC may also decide to change information in different ways depending on what service it is targeted to.

12. Acknowledgements

TBD.

13. References

13.1. Normative References

[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.

[RFC4120] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, "The

Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.

13.2. Informative References

[MIT-Athena] Steiner, J., Neuman, B., and J. Schiller, "Kerberos: An Authentication Service for Open Network Systems. In Proceedings of the Winter 1988 Usenix Conference. February.", 1988.

[MS-PAC] Microsoft, "Microsoft MS-PAC: Privilege Attribute Certificate Data Structure (v20100711)", July 2010.

[POSIX] The Open Group, "Portable Operating System Interface (POSIX.1-2008)", 2008.

[RFC1510] Kohl, J. and B. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2307] Howard, L., "An Approach for Using LDAP as a Network Information Service", RFC 2307, March 1998.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[X.690] ISO, "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 8825-1:1998)", 1997.

Appendix A. Additional Stuff

This becomes an Appendix.

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