

Internet Engineering Task Force (IETF)
Request for Comments: 7861
Updates: 5403
Category: Standards Track
ISSN: 2070-1721

W. Adamson
NetApp
N. Williams
Cryptonector
November 2016

Remote Procedure Call (RPC) Security Version 3

Abstract

This document specifies version 3 of the Remote Procedure Call (RPC) security protocol (RPCSEC_GSS). This protocol provides support for multi-principal authentication of client hosts and user principals to a server (constructed by generic composition), security label assertions for multi-level security and type enforcement, structured privilege assertions, and channel bindings. This document updates RFC 5403.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7861>.

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.

Table of Contents

1. Introduction and Motivation	2
1.1. Requirements Language	3
1.2. Added Functionality	4
1.3. XDR Code Extraction	5
2. The RPCSEC_GSSv3 Protocol	6
2.1. Compatibility with RPCSEC_GSSv2	6
2.2. Version Negotiation	6
2.3. New Reply Verifier	7
2.4. XDR Code Preliminaries	8
2.5. RPCSEC_GSS_BIND_CHANNEL Operation	10
2.6. New auth_stat Values	10
2.7. New Control Procedures	10
2.7.1. New Control Procedure - RPCSEC_GSS_CREATE	12
2.7.2. New Control Procedure - RPCSEC_GSS_LIST	20
2.8. Extensibility	21
3. Operational Recommendation for Deployment	21
4. Security Considerations	21
5. IANA Considerations	22
5.1. New RPC Authentication Status Numbers	22
5.2. Structured Privilege Name Definitions	23
5.2.1. Initial Registry	24
5.2.2. Updating Registrations	24
6. References	25
6.1. Normative References	25
6.2. Informative References	26
Acknowledgments	26
Authors' Addresses	26

1. Introduction and Motivation

The original Remote Procedure Call (RPC) security protocol (RPCSEC_GSS) [RFC2203] provided for authentication of RPC clients and servers to each other using the Generic Security Service Application Programming Interface (GSS-API) [RFC2743]. The second version of RPCSEC_GSS [RFC5403] added support for channel bindings [RFC5056].

Existing GSS-API mechanisms are insufficient for communicating certain authorization and authentication information to a server. The GSS-API and its mechanisms certainly could be extended to address this shortcoming. However, it is addressed here at the application layer, i.e., in RPCSEC_GSS.

A major motivation for version 3 of RPCSEC_GSS (RPCSEC_GSSv3) is to add support for multi-level (labeled) security and server-side copy for NFSv4.

Multi-Level Security (MLS) is a traditional model where subjects (processes) are given a security level (Unclassified, Secret, Top Secret, etc.) and objects (files) are given security labels that mandate the access of the subject to the object (see Section 9.1 of [RFC7862]).

Labeled NFS (see Section 9 of [RFC7862]) uses an MLS policy with Mandatory Access Control (MAC) systems as defined in [RFC4949]. Labeled NFS stores MAC file object labels on the NFS server and enables client Guest Mode MAC as described in Section 9.5.3 of [RFC7862]. RPCSEC_GSSv3 label assertions assert client MAC process subject labels to enable Full Mode MAC when combined with Labeled NFS as described in Section 9.5.1 of [RFC7862].

A traditional inter-server file copy entails the user gaining access to a file on the source, reading it, and writing it to a file on the destination. In secure NFSv4 inter-server server-side copy (see Section 4 of [RFC7862]), the user first secures access to both source and destination files and then uses NFSv4.2-defined RPCSEC_GSSv3 structured privileges to authorize the destination to copy the file from the source on behalf of the user.

Multi-principal authentication can be used to address shared cache poisoning attacks (see Section 9 of [AFS-RXGK]) on the client cache by a user. As described in Section 7 of [AFS-RXGK], multi-user machines with a single cache manager can fetch and cache data on a user's behalf and re-display it for another user from the cache without refetching the data from the server. The initial data acquisition is authenticated by the first user's credentials, and if only that user's credentials are used, it may be possible for a malicious user or users to "poison" the cache for other users by introducing bogus data into the cache.

Another use of the multi-principal assertion is the secure conveyance of privilege information for processes running with more (or even with less) privilege than the user normally would be accorded.

1.1. 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].

1.2. Added Functionality

RPCSEC_GSS version 3 (RPCSEC_GSSv3) is the same as RPCSEC_GSSv2 [RFC5403], except that the following assertions of authority have been added:

- o Security labels for Full Mode security type enforcement, and other labeled security models (see Section 9.5.1 of [RFC7862]).
- o Application-specific structured privileges. These allow an RPC application client to pass structured information to the corresponding application code in a server to control the use of the privilege and/or the conditions in which the privilege may be exercised. For an example, see server-side copy as described in [RFC7862].
- o Multi-principal authentication of the client host and user to the server, done by binding two RPCSEC_GSS handles.
- o Simplified channel binding.

Assertions of labels and privileges are evaluated by the server, which may then map the asserted values to other values, all according to server-side policy. See [RFC7862].

An option for enumerating server-supported Label Format Specifiers (LFSs) is provided. See Section 9.1 of [RFC7862].

Note that there is no RPCSEC_GSS_CREATE payload that is REQUIRED to implement. RPCSEC_GSSv3 implementations are feature driven. Besides implementing the RPCSEC_GSS_CREATE operation and payloads for the desired features, all RPCSEC_GSSv3 implementations MUST implement:

- o The new RPCSEC_GSS version number (Section 2.2).
- o The new reply verifier (Section 2.3).
- o The new auth_stat values (Section 2.6).

RPCSEC_GSSv3 targets implementing a desired feature MUST also implement the RPCSEC_GSS_LIST operation, and the RPCSEC_GSS_CREATE operation replies for unsupported features as follows:

- o For label assertions, the target indicates no support by returning the new RPCSEC_GSS_LABEL_PROBLEM auth_stat value (see Section 2.7.1.3).
- o For structured privilege assertions, the target indicates no support by returning the new RPCSEC_GSS_UNKNOWN_MESSAGE auth_stat value (see Section 2.7.1.4).
- o For multi-principal authentication (Section 2.7.1.1), the target indicates no support by not including an rgss3_gss_mp_auth value in the rgss3_create_res.
- o For channel bindings (Section 2.7.1.2), the target indicates no support by not including an rgss3_chan_binding value in the rgss3_create_res.

1.3. XDR Code Extraction

This document contains the External Data Representation (XDR) [RFC4506] definitions for the RPCSEC_GSSv3 protocol. The XDR description is provided in this document in a way that makes it simple for the reader to extract it into a form that is ready to compile. The reader can feed this document in the following shell script to produce the machine-readable XDR description of RPCSEC_GSSv3:

```
<CODE BEGINS>
```

```
#!/bin/sh
grep "^ *///" | sed 's?^ */// ??' | sed 's?^ *///$??'
```

```
<CODE ENDS>
```

That is, if the above script is stored in a file called "extract.sh" and this document is in a file called "spec.txt", then the reader can do:

```
<CODE BEGINS>
```

```
sh extract.sh < spec.txt > rpcsec_gss_v3.x
```

```
<CODE ENDS>
```

The effect of the script is to remove leading white space from each line, plus a sentinel sequence of "///".

2. The RPCSEC_GSSv3 Protocol

RPCSEC_GSS version 3 (RPCSEC_GSSv3) is very similar to RPCSEC_GSS version 2 (RPCSEC_GSSv2) [RFC5403]. The difference is that the new support for assertions and channel bindings is implemented via a different mechanism.

The entire RPCSEC_GSSv3 protocol is not presented here. Only the differences between RPCSEC_GSSv3 and RPCSEC_GSSv2 are shown.

RPCSEC_GSSv3 is implemented as follows:

- o A client uses an existing RPCSEC_GSSv3 context handle established in the usual manner (see Section 5.2 of [RFC2203]) to protect RPCSEC_GSSv3 exchanges; this will be termed the "parent" handle.
- o The server issues a "child" RPCSEC_GSSv3 handle in the RPCSEC_GSS_CREATE response, which uses the underlying GSS-API security context of the parent handle in all subsequent exchanges that use the child handle.
- o An RPCSEC_GSSv3 child handle MUST NOT be used as the parent handle in an RPCSEC_GSS3_CREATE control message.

2.1. Compatibility with RPCSEC_GSSv2

The functionality of RPCSEC_GSSv2 [RFC5403] is fully supported by RPCSEC_GSSv3, with the exception of the RPCSEC_GSS_BIND_CHANNEL operation, which is not supported when RPCSEC_GSSv3 is in use (see Section 2.5).

2.2. Version Negotiation

An initiator that supports version 3 of RPCSEC_GSS simply issues an RPCSEC_GSS request with the `rgc_version` field set to `RPCSEC_GSS_VERS_3`. If the target does not recognize `RPCSEC_GSS_VERS_3`, the target will return an RPC error per Section 5.1 of [RFC2203].

The initiator MUST NOT attempt to use an RPCSEC_GSS handle returned by version 3 of a target with version 1 or version 2 of the same target. The initiator MUST NOT attempt to use an RPCSEC_GSS handle returned by version 1 or version 2 of a target with version 3 of the same target.

2.3. New Reply Verifier

A new reply verifier is needed for RPCSEC_GSSv3 because of a situation that arises from the use of the same GSS context by child and parent handles. Because the RPCSEC_GSSv3 child handle uses the same GSS context as the parent handle, a child and parent RPCSEC_GSSv3 handle could have the same RPCSEC_GSS sequence numbers. Since the reply verifier of previous versions of RPCSEC_GSS computes a Message Integrity Code (MIC) on just the sequence number, this provides opportunities for man-in-the-middle attacks.

This issue is addressed in RPCSEC_GSS version 3 by computing the verifier using exactly the same input as the information used to compute the request verifier, except that the mtype is changed from CALL to REPLY. The new reply verifier computes a MIC over the following RPC reply header data:

```
unsigned int xid;
msg_type mtype; /* set to REPLY */
unsigned int rpcvers;
unsigned int prog;
unsigned int vers;
unsigned int proc;
opaque_auth cred; /* binds the RPCSEC_GSS handle */
```

2.4. XDR Code Preliminaries

The following code fragment replaces the corresponding preliminary code shown in Figure 1 of [RFC5403]. The values in the code fragment in Section 2.6 are additions to the `auth_stat` enumeration. Subsequent code fragments are additions to the code for version 2 that support the new procedures defined in version 3.

<CODE BEGINS>

```
/// /*
///  * Copyright (c) 2016 IETF Trust and the persons
///  * identified as the authors. All rights reserved.
///  *
///  * The authors of the code are identified in RFC 2203,
///  * RFC 5403, and RFC 7861.
///  *
///  * Redistribution and use in source and binary forms,
///  * with or without modification, are permitted
///  * provided that the following conditions are met:
///  *
///  * o Redistributions of source code must retain the above
///  *   copyright notice, this list of conditions and the
///  *   following disclaimer.
///  *
///  * o Redistributions in binary form must reproduce the
///  *   above copyright notice, this list of
///  *   conditions and the following disclaimer in
///  *   the documentation and/or other materials
///  *   provided with the distribution.
///  *
///  * o Neither the name of Internet Society, IETF or IETF
///  *   Trust, nor the names of specific contributors, may be
///  *   used to endorse or promote products derived from this
///  *   software without specific prior written permission.
///  *
///  * THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS
///  * AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED
///  * WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
///  * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
///  * FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO
///  * EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE
///  * LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL,
///  * EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
///  * NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR
///  * SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS
///  * INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
///  * LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,
```



```

///  *   OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING
///  *   IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
///  *   ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
///  */
///
///  /*
///  * This code was derived from RFC 2203, RFC 5403,
///  * and RFC 7861. Please reproduce this note if possible.
///  */
///
///  enum rpc_gss_service_t {
///      /* Note: The enumerated value for 0 is reserved. */
///      rpc_gss_svc_none          = 1,
///      rpc_gss_svc_integrity     = 2,
///      rpc_gss_svc_privacy       = 3,
///      rpc_gss_svc_channel_prot  = 4
///  };
///
///  enum rpc_gss_proc_t {
///      RPCSEC_GSS_DATA           = 0,
///      RPCSEC_GSS_INIT           = 1,
///      RPCSEC_GSS_CONTINUE_INIT  = 2,
///      RPCSEC_GSS_DESTROY        = 3,
///      RPCSEC_GSS_BIND_CHANNEL   = 4, /* Not used */
///      RPCSEC_GSS_CREATE         = 5, /* New */
///      RPCSEC_GSS_LIST           = 6  /* New */
///  };
///
///  struct rpc_gss_cred_vers_1_t {
///      rpc_gss_proc_t    gss_proc; /* Control procedure */
///      unsigned int      seq_num;  /* Sequence number */
///      rpc_gss_service_t service;  /* Service used */
///      opaque            handle<>; /* Context handle */
///  };
///
///  const RPCSEC_GSS_VERS_1 = 1;
///  const RPCSEC_GSS_VERS_2 = 2;
///  const RPCSEC_GSS_VERS_3 = 3; /* New */
///
///  union rpc_gss_cred_t switch (unsigned int rgc_version) {
///      case RPCSEC_GSS_VERS_1:
///      case RPCSEC_GSS_VERS_2:
///      case RPCSEC_GSS_VERS_3: /* New */
///          rpc_gss_cred_vers_1_t rgc_cred_v1;
///      };
///

```

<CODE ENDS>

As seen above, the `RPCSEC_GSSv3` credential has the same format as the `RPCSEC_GSSv1` [RFC2203] and `RPCSEC_GSSv2` [RFC5403] credential. Setting the `rgc_version` field to 3 indicates that the initiator and target support the new `RPCSEC_GSSv3` control procedures.

2.5. `RPCSEC_GSS_BIND_CHANNEL` Operation

`RPCSEC_GSSv3` provides a channel-binding assertion that replaces the `RPCSEC_GSSv2` `RPCSEC_GSS_BIND_CHANNEL` operation.

The `RPCSEC_GSS_BIND_CHANNEL` operation is not supported on `RPCSEC_GSS` version 3 handles. If a server receives an `RPCSEC_GSS_BIND_CHANNEL` operation on an `RPCSEC_GSSv3` handle, it MUST return a reply status of `MSG_ACCEPTED` with an `accept_stat` of `PROC_UNAVAIL` [RFC5531].

2.6. New `auth_stat` Values

`RPCSEC_GSSv3` requires the addition of several values to the `auth_stat` enumerated type definition. The use of these new `auth_stat` values is explained throughout this document.

```
enum auth_stat {  
    ...  
    /*  
     * RPCSEC_GSSv3 errors  
     */  
    RPCSEC_GSS_INNER_CREDPROBLEM = 15,  
    RPCSEC_GSS_LABEL_PROBLEM      = 16,  
    RPCSEC_GSS_PRIVILEGE_PROBLEM = 17,  
    RPCSEC_GSS_UNKNOWN_MESSAGE   = 18  
};
```

2.7. New Control Procedures

There are two new `RPCSEC_GSSv3` control procedures: `RPCSEC_GSS_CREATE` and `RPCSEC_GSS_LIST`.

The `RPCSEC_GSS_CREATE` procedure binds any combination of assertions -- multi-principal authentication, labels, structured privileges, or channel bindings -- to a new `RPCSEC_GSSv3` context returned in the `rgss3_create_res` `rcr_handle` field.

The `RPCSEC_GSS_LIST` procedure queries the target for supported assertions.

RPCSEC_GSS version 3 control messages are similar to the RPCSEC_GSS version 1 and version 2 RPCSEC_GSS_DESTROY control message (see Section 5.4 of [RFC2203]) in that the sequence number in the request must be valid and the header checksum in the verifier must be valid. As in RPCSEC_GSS version 1 and version 2, the RPCSEC_GSS version 3 control messages may contain call data following the verifier in the body of the NULLPROC procedure. In other words, they look a lot like an RPCSEC_GSS data message with the header procedure set to NULLPROC.

The client MUST use one of the following security services to protect the RPCSEC_GSS_CREATE or RPCSEC_GSS_LIST control message:

- o rpc_gss_svc_integrity
- o rpc_gss_svc_privacy

Specifically, the client MUST NOT use rpc_gss_svc_none.

RPCSEC_GSS_LIST can also use rpc_gss_svc_channel_prot (see RPCSEC_GSSv2 [RFC5403]) if the request is sent using an RPCSEC_GSSv3 child handle with channel bindings enabled as described in Section 2.7.1.2.

2.7.1. New Control Procedure - RPCSEC_GSS_CREATE

<CODE BEGINS>

```

    /// struct rgss3_create_args {
    ///     rgss3_gss_mp_auth    *rca_mp_auth;
    ///     rgss3_chan_binding    *rca_chan_bind_mic;
    ///     rgss3_assertion_u     rca_assertions<>;
    /// };
    ///
    /// struct rgss3_create_res {
    ///     opaque                rcr_handle<>;
    ///     rgss3_gss_mp_auth    *rcr_mp_auth;
    ///     rgss3_chan_binding    *rcr_chan_bind_mic;
    ///     rgss3_assertion_u     rcr_assertions<>;
    /// };
    ///
    /// enum rgss3_assertion_type {
    ///     LABEL = 0,
    ///     PRIVS = 1
    /// };
    ///
    /// union rgss3_assertion_u
    ///     switch (rgss3_assertion_type atype) {
    /// case LABEL:
    ///     rgss3_label    rau_label;
    /// case PRIVS:
    ///     rgss3_privs    rau_privs;
    /// default:
    ///     opaque          rau_ext<>;
    /// };
    ///

```

<CODE ENDS>

The call data for an RPCSEC_GSS_CREATE request consists of an `rgss3_create_args`, which binds one or more items of several kinds to the returned `rcr_handle` RPCSEC_GSSv3 context handle (the child handle):

- o Multi-principal authentication: another RPCSEC_GSS context handle
- o A channel binding
- o Authorization assertions: labels and/or privileges

The reply to this message consists of either an error or an `rgss3_create_res` structure. As noted in Sections 2.7.1.3 and 2.7.1.4, successful `rgss3_assertions` are enumerated in `rcr_assertions` and are REQUIRED to be enumerated in the same order as they appeared in the `rca_assertions` argument.

Upon a successful `RPCSEC_GSS_CREATE`, both the client and the server need to associate the resultant child `rcr_handle` context handle with the parent context handle in their GSS context caches so as to be able to reference the parent context given the child context handle.

`RPCSEC_GSSv3` child handles MUST be destroyed upon the destruction of the associated parent handle.

Server implementation and policy MAY result in labels, privileges, and identities being mapped to concepts and values that are local to the server. Server policies should take into account the identity of the client and/or user as authenticated via the GSS-API.

2.7.1.1. Multi-Principal Authentication

<CODE BEGINS>

```
///
/// struct rgss3_gss_mp_auth {
///     opaque          rgmp_handle<>; /* Inner handle */
///     opaque          rgmp_rpcheader_mic<>;
/// };
```

<CODE ENDS>

`RPCSEC_GSSv3` clients MAY assert a multi-principal authentication of the RPC client host principal and a user principal. This feature is needed, for example, when an RPC client host wishes to use authority assertions that the server may only grant if a user and an RPC client host are authenticated together to the server. Thus, a server may refuse to grant requested authority to a user acting alone (e.g., via an unprivileged user-space program) or to an RPC client host acting alone (e.g., when an RPC client host is acting on behalf of a user) but may grant requested authority to an RPC client host acting on behalf of a user if the server identifies the user and trusts the RPC client host.

It is assumed that an unprivileged user-space program would not have access to RPC client host credentials needed to establish a GSS-API security context authenticating the RPC client host to the server; therefore, an unprivileged user-space program could not create an `RPCSEC_GSSv3` `RPCSEC_GSS_CREATE` message that successfully binds an RPC client host and a user security context.

In addition to the parent handle (Section 2), the multi-principal authentication call data has an `RPCSEC_GSS` version 3 handle referenced via the `rgmp_handle` field termed the "inner" handle. Clients using `RPCSEC_GSSv3` multi-principal authentication MUST use an `RPCSEC_GSSv3` context handle that corresponds to a GSS-API security context that authenticates the RPC client host for the parent handle. The inner context handle of the multi-principal authentication assertion MUST use an `RPCSEC_GSSv3` context handle that corresponds to a GSS-API security context that authenticates the user. The reverse (parent handle authenticates user, inner context handle authenticates an RPC client host) MUST NOT be used. Other multi-principal parent and inner context handle uses might eventually make sense, but they would need to be introduced in a new revision of the `RPCSEC_GSS` protocol.

The child context handle returned by a successful multi-principal assertion binds the inner `RPCSEC_GSSv3` context handle to the parent `RPCSEC_GSS` context handle and MUST be treated by servers as authenticating the GSS-API initiator principal authenticated by the inner context handle's GSS-API security context. This principal may be mapped to a server-side notion of user or principal.

Multi-principal binding is done by including an assertion of type `rgss3_gss_mp_auth` in the `RPCSEC_GSS_CREATE` `rgss3_create_args` call data. The inner context handle is placed in the `rgmp_handle` field. A MIC of the RPC header, up to and including the credential, is computed using the GSS-API security context associated with the inner context handle and is placed in the `rgmp_rpcheader_mic` field. Note that the `rgmp_rpcheader_mic` only identifies the client host GSS context by its context handle (the parent context handle) in the RPC header.

An `RPCSEC_GSS_CREATE` control procedure with a multi-principal authentication payload MUST use the `rpc_gss_svc_privacy` security service for protection. This prevents an attacker from intercepting the `RPCSEC_GSS_CREATE` control procedure, reassigning the (parent) context handle, and stealing the user's identity.

The target verifies the multi-principal authentication by first confirming that the parent context used is an RPC client host context; the target then verifies the `rgmp_rpcheader_mic` using the GSS-API security context associated with the `rgmp_handle` field.

On successful verification, the `rgss3_gss_mp_auth` field in the `rgss3_create_res` reply MUST be filled in with the inner RPCSEC_GSSv3 context handle as the `rgmp_handle` and a MIC computed over the RPC reply header (see Section 2.3) using the GSS-API security context associated with the inner handle.

On failure, the `rgss3_gss_mp_auth` field is not sent (`rgss3_gss_mp_auth` is an optional field). A `MSG_DENIED` reply to the `RPCSEC_GSS_CREATE` call is formulated as usual.

As described in Section 5.3.3.3 of [RFC2203], the server maintains a list of contexts for the clients that are currently in session with it. When a client request comes in, there may not be a context corresponding to its handle. When this occurs on an `RPCSEC_GSS3_CREATE` request processing of the parent handle, the server rejects the request with a reply status of `MSG_DENIED` with the `reject_stat` of `AUTH_ERROR` and with an `auth_stat` value of `RPCSEC_GSS_CREDPROBLEM`.

A new value, `RPCSEC_GSS_INNER_CREDPROBLEM`, has been added to the `auth_stat` type. With a multi-principal authorization request, the server must also have a context corresponding to the inner context handle. When the server does not have a context handle corresponding to the inner context handle of a multi-principal authorization request, the server sends a reply status of `MSG_DENIED` with the `reject_stat` of `AUTH_ERROR` and with an `auth_stat` value of `RPCSEC_GSS_INNER_CREDPROBLEM`.

When processing the multi-principal authentication request, if the `GSS_VerifyMIC()` call on the `rgmp_rpcheader_mic` fails to return `GSS_S_COMPLETE`, the server sends a reply status of `MSG_DENIED` with the `reject_stat` of `AUTH_ERROR` and with an `auth_stat` value of `RPCSEC_GSS_INNER_CREDPROBLEM`.

2.7.1.2. Channel Binding

<CODE BEGINS>

```
///  
///  typedef opaque rgss3_chan_binding<>;  
///
```

<CODE ENDS>

RPCSEC_GSSv3 provides a different way to do channel binding than RPCSEC_GSSv2 [RFC5403]. Specifically:

- a. RPCSEC_GSSv3 builds on RPCSEC_GSSv1 by reusing existing, established context handles rather than providing a different RPC security flavor for establishing context handles.
- b. Channel-bindings data is not hashed because there is now general agreement that it is the secure channel's responsibility to produce channel-bindings data of manageable size.

(a) is useful in keeping RPCSEC_GSSv3 simple in general, not just for channel binding. (b) is useful in keeping RPCSEC_GSSv3 simple specifically for channel binding.

Channel binding is accomplished as follows. The client prefixes the channel-bindings data octet string with the channel type as described in [RFC5056]; then, the client calls GSS_GetMIC() to get a MIC of the resulting octet string, using the parent RPCSEC_GSSv3 context handle's GSS-API security context. The MIC is then placed in the rca_chan_bind_mic field of RPCSEC_GSS_CREATE arguments (rgss3_create_args).

If the rca_chan_bind_mic field of the arguments of an RPCSEC_GSS_CREATE control message is set, then the server MUST verify the client's channel-binding MIC if the server supports this feature. If channel-binding verification succeeds, then the server MUST generate a new MIC of the same channel bindings and place it in the rcr_chan_bind_mic field of the RPCSEC_GSS_CREATE rgss3_create_res results. If channel-binding verification fails or the server doesn't support channel binding, then the server MUST indicate this in its reply by not including an rgss3_chan_binding value in rgss3_create_res (rgss3_chan_binding is an optional field).

The client MUST verify the result's rcr_chan_bind_mic value by calling GSS_VerifyMIC() with the given MIC and the channel-bindings data (including the channel-type prefix). If client-side channel-binding verification fails, then the client MUST call

RPCSEC_GSS_DESTROY. If the client requested channel binding but the server did not include an `rcr_chan_binding_mic` field in the results, then the client MAY continue to use the resulting context handle as though channel binding had never been requested. If the client considers channel binding critical, it MUST call `RPCSEC_GSS_DESTROY`.

As per `RPCSEC_GSSv2` [RFC5403]:

Once a successful [channel-binding] procedure has been performed on an `RPCSEC_GSSv3` context handle, the initiator's implementation may map application requests for `rpc_gss_svc_none` and `rpc_gss_svc_integrity` to `rpc_gss_svc_channel_prot` credentials. And if the secure channel has privacy enabled, requests for `rpc_gss_svc_privacy` can also be mapped to `rpc_gss_svc_channel_prot`.

Any `RPCSEC_GSSv3` child context handle that has been bound to a secure channel in this way SHOULD be used only with the `rpc_gss_svc_channel_prot` and SHOULD NOT be used with `rpc_gss_svc_none` or `rpc_gss_svc_integrity` -- if the secure channel does not provide privacy protection, then the client MAY use `rpc_gss_svc_privacy` where privacy protection is needed or desired.

2.7.1.3. Label Assertions

<CODE BEGINS>

```

    /// struct rgss3_label {
    ///     rgss3_lfs      rl_lfs;
    ///     opaque         rl_label<>;
    /// };
    ///
    /// struct rgss3_lfs {
    ///     unsigned int rlf_lfs_id;
    ///     unsigned int rlf_pi_id;
    /// };
    ///

```

<CODE ENDS>

The client discovers, via the `RPCSEC_GSS_LIST` control message, which LFSs the server supports. Full Mode MAC is enabled when an `RPCSEC_GSS` version 3 process subject label assertion is combined with a file object label provided by the NFSv4.2 `sec_label` attribute.

Label encoding is specified to mirror the NFSv4.2 `sec_label` attribute described in Section 12.2.4 of [RFC7862]. The LFS is an identifier used by the client to establish the syntactic format of the security

label and the semantic meaning of its components. The Policy Identifier (PI) is an optional part of the definition of an LFS that allows clients and the server to identify specific security policies. The opaque label field (rgss3_label) is dependent on the MAC model to interpret and enforce.

If a label itself requires privacy protection (i.e., requires that the user can assert that the label is a secret), then the client MUST use the rpc_gss_svc_privacy protection service for the RPCSEC_GSS_CREATE request.

RPCSEC_GSSv3 clients MAY assert a set of subject security labels in some LFS by binding a label assertion to the RPCSEC_GSSv3 child context handle. This is done by including an assertion of type rgss3_label in the RPCSEC_GSS_CREATE rgss3_create_args rca_assertions call data. The label assertion payload is the set of subject labels asserted by the calling NFS client process. The resultant child context is used for NFS requests asserting the client process subject labels. The NFS server process that handles such requests then asserts the (client) process subject label(s) as it attempts to access a file that has associated Labeled NFS object labels.

Servers that support labeling in the requested LFS MAY map the requested subject label to a different subject label as a result of server-side policy evaluation.

The labels that are accepted by the target and bound to the RPCSEC_GSSv3 context MUST be enumerated in the rcr_assertions field of the rgss3_create_res RPCSEC_GSS_CREATE reply.

Servers that do not support labeling or that do not support the requested LFS reject the label assertion with a reply status of MSG_DENIED, a reject_status of AUTH_ERROR, and an auth_stat of RPCSEC_GSS_LABEL_PROBLEM.

2.7.1.4. Structured Privilege Assertions

<CODE BEGINS>

```
///  
/// typedef opaque utf8string<>; /* UTF-8 encoding */  
/// typedef utf8string utf8str_cs; /* Case-sensitive UTF-8 */  
///  
/// struct rgss3_privs {  
///     utf8str_cs      rp_name<>;  
///     opaque          rp_privilege<>;  
/// };
```

<CODE ENDS>

A structured privilege is a capability defined by a specific RPC application. To support the assertion of this privilege, by a client using the application, in a server that also supports the application, the application may define a private data structure that is understood by clients and servers implementing the RPC application.

RPCSEC_GSSv3 clients MAY assert a structured privilege by binding the privilege to the RPCSEC_GSSv3 context handle. This is done by including an assertion of type `rgss3_privs` in the `RPCSEC_GSS_CREATE` `rgss3_create_args` `rca_assertions` call data.

The privilege is identified by the description string that is used by `RPCSEC_GSSv3` to identify the privilege and communicate the private data between the relevant RPC application-specific code without needing to be aware of the details of the structure used. Thus, as far as `RPCSEC_GSSv3` is concerned, the defined structure is passed between client and server as opaque data encoded in the `rpc_gss3_privs` `rp_privilege` field.

Encoding, server verification, and any server policies for structured privileges are described by the RPC application definition. The `rp_name` field of `rpc_gss3_privs` carries the description string used to identify and list the privilege. The `utf8str_cs` definition is from [RFC7530].

A successful structured privilege assertion MUST be enumerated in the `rcr_assertions` field of the `rgss3_create_res` `RPCSEC_GSS_CREATE` reply.

If a server receives a structured privilege assertion that it does not recognize, the assertion is rejected with a reply status of `MSG_DENIED`, a `reject_status` of `AUTH_ERROR`, and an `auth_stat` of `RPCSEC_GSS_UNKNOWN_MESSAGE`.

It is assumed that a client asserting more than one structured privilege to be bound to a context handle would not require all the privilege assertions to succeed.

The server MUST NOT reject `RPCSEC_GSS_CREATE` requests containing supported structured privilege assertions, even if some of those assertions are rejected (e.g., for local policy reasons).

If a server receives an `RPCSEC_GSS_CREATE` request containing one or more unsupported structured privilege assertions, the request MUST be rejected with a reply status of `MSG_DENIED`, a `reject_status` of `AUTH_ERROR`, and an `auth_stat` of `RPCSEC_GSS_PRIVILEGE_PROBLEM`.

Section 4.9.1.1 of [RFC7862] ("Inter-Server Copy via ONC RPC with `RPCSEC_GSSv3`") shows an example of structured privilege definition and use.

2.7.2. New Control Procedure - `RPCSEC_GSS_LIST`

<CODE BEGINS>

```

    /// enum rgss3_list_item {
    ///     LABEL = 0,
    ///     PRIVS = 1
    /// };
    ///
    /// struct rgss3_list_args {
    ///     rgss3_list_item      rla_list_what<>;
    /// };
    ///
    /// union rgss3_list_item_u
    ///     switch (rgss3_list_item itype) {
    /// case LABEL:
    ///     rgss3_label          rli_labels<>;
    /// case PRIVS:
    ///     rgss3_privs          rli_privs<>;
    /// default:
    ///     opaque                rli_ext<>;
    /// };
    ///
    /// typedef rgss3_list_item_u rgss3_list_res<>;
    ///

```

<CODE ENDS>

The call data for an `RPCSEC_GSS_LIST` request consists of a list of integers (`rla_list_what`) indicating what assertions are to be listed, and the reply consists of an error or the requested list.

The result of requesting a list of `rgss3_list_item` LABEL objects is a list of LFSs supported by the server. The client can then use the LFS list to assert labels via the `RPCSEC_GSS_CREATE` label assertions. See Section 2.7.1.3.

2.8. Extensibility

Assertion types may be added in the future by adding arms to the `"rgss3_assertion_u"` union (Section 2.7.1) and the `"rgss3_list_item_u"` union (Section 2.7.2). Examples of other potential assertion types include:

- o Client-side assertions of identity:
 - * Primary client/user identity.
 - * Supplementary group memberships of the client/user, including support for specifying deltas to the membership list as seen on the server.

3. Operational Recommendation for Deployment

`RPCSEC_GSSv3` is a superset of `RPCSEC_GSSv2` [RFC5403], which in turn is a superset of `RPCSEC_GSSv1` [RFC2203], and so can be used in all situations where `RPCSEC_GSSv2` is used, or where `RPCSEC_GSSv1` is used and channel-bindings functionality is not needed. `RPCSEC_GSSv3` should be used when the new functionality is needed.

4. Security Considerations

This entire document deals with security issues.

The `RPCSEC_GSSv3` protocol allows for client-side assertions of data that is relevant to server-side authorization decisions. These assertions must be evaluated by the server in the context of whether the client and/or user are authenticated, whether multi-principal authentication was used, whether the client is trusted, what ranges of assertions are allowed for the client and the user (separately or together), and any relevant server-side policy.

The security semantics of assertions carried by `RPCSEC_GSSv3` are application protocol-specific.

Note that `RPCSEC_GSSv3` is not a complete solution for labeling: it conveys the labels of actors but not the labels of objects. RPC application protocols may require extending in order to carry object label information.

There may be interactions with NFSv4's callback security scheme and NFSv4.1's [RFC5661] GSS SSV (Secret State Verifier) mechanisms. Specifically, the NFSv4 callback scheme requires that the server initiate GSS-API security contexts, which does not work well in practice; in the context of client-side processes running as the same user but with different privileges and security labels, the NFSv4 callback security scheme seems particularly unlikely to work well. NFSv4.1 has the server use an existing, client-initiated RPCSEC_GSS context handle to protect server-initiated callback RPCs. The NFSv4.1 callback security scheme lacks all the problems of the NFSv4 scheme; however, it is important that the server pick an appropriate RPCSEC_GSS context handle to protect any callbacks. Specifically, it is important that the server use RPCSEC_GSS context handles that authenticate the client to protect any callbacks related to server state initiated by RPCs protected by RPCSEC_GSSv3 contexts.

As described in Section 2.10.10 of [RFC5661], the client is permitted to associate multiple RPCSEC_GSS handles with a single SSV GSS context. RPCSEC_GSSv3 handles will work well with SSV in that the man-in-the-middle attacks described in Section 2.10.10 of [RFC5661] are solved by the new reply verifier (Section 2.3). Using an RPCSEC_GSSv3 handle backed by a GSS-SSV mechanism context as a parent handle in an RPCSEC_GSS_CREATE call, while permitted, is complicated by the lifetime rules of SSV contexts and their associated RPCSEC_GSS handles.

5. IANA Considerations

This section uses terms that are defined in [RFC5226].

5.1. New RPC Authentication Status Numbers

The following new RPC Authentication Status Numbers have been added to the IANA registry:

- o RPCSEC_GSS_INNER_CREDPROBLEM (15) "No credentials for multi-principal assertion inner context user". See Section 2.7.1.1.
- o RPCSEC_GSS_LABEL_PROBLEM (16) "Problem with label assertion". See Section 2.7.1.3.
- o RPCSEC_GSS_PRIVILEGE_PROBLEM (17) "Problem with structured privilege assertion". See Section 2.7.1.4.
- o RPCSEC_GSS_UNKNOWN_MESSAGE (18) "Unknown structured privilege assertion". See Section 2.7.1.4.

5.2. Structured Privilege Name Definitions

IANA has created a registry called the "RPCSEC_GSS Structured Privilege Names Registry".

Structured privilege assertions (Section 2.7.1.4) are defined by a specific RPC application. The namespace identifiers for these assertions (the `rp_name`) are defined as string names. The RPCSEC_GSSv3 protocol does not define the specific assignment of the namespace for these structured privilege assertion names. The IANA registry promotes interoperability where common interests exist. While RPC application developers are allowed to define and use structured privileges as needed, they are encouraged to register structured privilege assertion names with IANA.

The registry is to be maintained using the Standards Action policy as defined in Section 4.1 of [RFC5226].

Under the RPCSEC_GSS version 3 specification, the name of a structured privilege can in theory be up to $2^{32} - 1$ bytes in length, but in practice RPC application clients and servers will be unable to handle a string that long. IANA should reject any assignment request with a structured privilege name that exceeds 128 UTF-8 characters. To give the IESG the flexibility to set up bases of assignment of Experimental Use, the prefix "EXPE" is Reserved. The structured privilege with a zero-length name is Reserved.

The prefix "PRIV" is allocated for Private Use. A site that wants to make use of unregistered named attributes without risk of conflicting with an assignment in IANA's registry should use the prefix "PRIV" in all of its structured privilege assertion names.

Because some RPC application clients and servers have case-insensitive semantics, the fifteen additional lower-case and mixed-case permutations of each of "EXPE" and "PRIV" are Reserved (e.g., "expe", "expE", and "exPe" are Reserved). Similarly, IANA must not allow two assignments that would conflict if both structured privilege names were converted to a common case.

The registry of structured privilege names is a list of assignments, each containing three fields for each assignment.

1. A US-ASCII string name that is the actual name of the structured privilege. This name must be unique. This string name can be 1 to 128 UTF-8 characters long.
2. A reference to the specification of the RPC-application-defined structured privilege. The reference can consume up to 256 bytes (or more if IANA permits).
3. The point of contact of the registrant. The point of contact can consume up to 256 bytes (or more if IANA permits).

5.2.1. Initial Registry

The initial registry consists of the three structured privileges defined in [RFC7862].

1. NAME: `copy_to_auth`, REFERENCE: RFC 7862, CONTACT: William A.(Andy) Adamson, andros@netapp.com
2. NAME: `copy_from_auth`, REFERENCE: RFC 7862, CONTACT: William A.(Andy) Adamson, andros@netapp.com
3. NAME: `copy_confirm_auth`, REFERENCE: RFC 7862, CONTACT: William A.(Andy) Adamson, andros@netapp.com

5.2.2. Updating Registrations

The registrant is always permitted to update the point of contact field. To make any other change will require Expert Review or IESG Approval.

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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2203] Eisler, M., Chiu, A., and L. Ling, "RPCSEC_GSS Protocol Specification", RFC 2203, DOI 10.17487/RFC2203, September 1997, <<http://www.rfc-editor.org/info/rfc2203>>.
- [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>>.
- [RFC4506] Eisler, M., Ed., "XDR: External Data Representation Standard", STD 67, RFC 4506, DOI 10.17487/RFC4506, May 2006, <<http://www.rfc-editor.org/info/rfc4506>>.
- [RFC5056] Williams, N., "On the Use of Channel Bindings to Secure Channels", RFC 5056, DOI 10.17487/RFC5056, November 2007, <<http://www.rfc-editor.org/info/rfc5056>>.
- [RFC5403] Eisler, M., "RPCSEC_GSS Version 2", RFC 5403, DOI 10.17487/RFC5403, February 2009, <<http://www.rfc-editor.org/info/rfc5403>>.
- [RFC5661] Shepler, S., Ed., Eisler, M., Ed., and D. Noveck, Ed., "Network File System (NFS) Version 4 Minor Version 1 Protocol", RFC 5661, DOI 10.17487/RFC5661, January 2010, <<http://www.rfc-editor.org/info/rfc5661>>.
- [RFC7530] Haynes, T., Ed., and D. Noveck, Ed., "Network File System (NFS) Version 4 Protocol", RFC 7530, DOI 10.17487/RFC7530, March 2015, <<http://www.rfc-editor.org/info/rfc7530>>.
- [RFC7862] Haynes, T., "Network File System (NFS) Version 4 Minor Version 2 Protocol", RFC 7862, DOI 10.17487/RFC7862, November 2016, <<http://www.rfc-editor.org/info/rfc7862>>.

6.2. Informative References

- [AFS-RXGK] Wilkinson, S. and B. Kaduk, "Integrating rxgk with AFS", Work in Progress, draft-wilkinson-afs3-rxgk-afs-08, May 2015.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.
- [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>>.
- [RFC5531] Thurlow, R., "RPC: Remote Procedure Call Protocol Specification Version 2", RFC 5531, DOI 10.17487/RFC5531, May 2009, <<http://www.rfc-editor.org/info/rfc5531>>.

Acknowledgments

Andy Adamson would like to thank NetApp, Inc. for its funding of his time on this project.

We thank Lars Eggert, Mike Eisler, Ben Kaduk, Bruce Fields, Tom Haynes, and Dave Noveck for their most helpful reviews.

Authors' Addresses

William A. (Andy) Adamson
NetApp
3629 Wagner Ridge Ct.
Ann Arbor, MI 48103
United States of America

Phone: +1 734 665 1204
Email: andros@netapp.com

Nico Williams
cryptonector.com
13115 Tamayo Dr.
Austin, TX 78729
United States of America

Email: nico@cryptonector.com

