Workgroup: Internet Engineering Task Force
Internet-Draft: draft-ietf-xml2rfc-template-06
Published: 25 June 2020
Intended Status: Informational
Expires: 27 December 2020
Authors: D.J. Massameno, Ed.
Yale University

RADIUS Extensions for Server Load Balancing

Abstract

This document describes a method for a Network Access Server (NAS) to dynamically discover all available RADIUS servers. It defines a new message type within the STATUS-SERVER message, which is requested by the NAS and provided by the RADIUS server. The NAS is then able to load balance its RADIUS messages across multiple RADIUS servers based on priority and weight supplied by the initial server.

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 https://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 27 December 2020.

Table of Contents

1. Introduction
  1.1. Requirements Language
  1.2. Terminology
2. Overall Message Exchange Summary
  2.1. Attributes Needed for Status-Server-LB
  2.2. Table of Attributes
  2.3. Required Status-Server-LB Attributes
    2.3.1. NAS-IP Address and/or NAS-Identifier
    2.3.2. Message-Authenticator
3. LB-Request Attribute
4. LB-Response Attribute
  4.1. LB-Response Attribute Format
5. The SVR-Record TLV
  5.1. SVR-Record-IPv4
  5.2. SVR-Record-IPv6
  5.3. Table of Sub-Attributes
6. Sub-Attributes Needed for SVR-Record TLV
  6.1. LB-TTL
  6.2. LB-Priority
  6.3. LB-Weight
  6.4. LB-IPv4
  6.5. LB-IPv6
7. Load Balancing Rules
  7.1. Session-Based Load Balancing
  7.2. Load Balancing Weight
8. Security Considerations
  8.1. Clear Text Transmission
  8.2. Reconnaissance
  8.3. MD5
9. NAS identifying Initial RADIUS Servers
10. IANA Considerations
11. References
Author's Address

1. Introduction

Modern networks require Authentication, Authorization and Accounting (AAA) services for a wide range of deployment scenarios. Many of these scenarios are mission critical and require fault tolerance and increased up-time. Most network equipment can be configured to access multiple back-end RAIDUS servers. When one server fails the equipment switches to the other RADIUS server.

The configuration of multiple RADIUS servers within the Network Access Server (RADIUS Client) currently has a number of limitations within contemporary implementations. There may be a limitation in the number of RADIUS servers that can be easily configured and

maintained. Also, until a failure is detected, the Network Access Server will likely only use one RADIUS server, even if multiple are configured.

To relieve these limitations, some installations choose to use a load balancer between the Network Access Server and the RADIUS server. This has the advantage of supporting an arbitrarily large number of RADIUS servers. The load balancer can be configured to distribute the load evenly based on a defined algorithm. Proportional load distribution may be a desirable property when trying to scale out to multiple back-end RADIUS servers for the purposes of increasing capacity.

The RADIUS extensions in this document achieve the load balancing property without using a separate load balancing device. With the removal of the external load balancer the operational complexity of the entire system will decrease. Also, as opposed to a third-party device, the NAS and RADIUS servers are the best devices to determine the operational status of the necessary components, thereby assisting in fault detection and avoidance.

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. Terminology

This document frequently uses the following terms:

session
: Each service provided by the NAS to a user attempting to connect (a dial-in user in the original RADIUS specifications) constitutes a session. The beginning of the session is defined as the point where service is first provided and the end of the session defined as the point where service is ended. A user may have multiple sessions in parallel or series if supported by the NAS.

calling-station
: This is the user or device that wishes access to the network. It connects to the NAS and, in most cases, presents identifying information.

NAS
: This is a Network Access Server. This is the device that receives the incoming connection from a calling-station that wishes access to the network. Some vendors call this the Network Access Device

(NAD). The NAS then communicates with the RADIUS Server on behalf of the calling-station.

RADIUS Server

This is the machine that implements the server side of the RADIUS protocol.

AAA services

The RADIUS protocol serves the functions of Authentication (identity), Authorization (what the user is allowed to do and how their connection should be configured), and Accounting (a record of the actions taken for the connection).

PDU

The Protocol Data Unit is the organization of data in a formal specification that is serialized and transmitted between entities on a network.

2. Overall Message Exchange Summary

The RADIUS protocol [RFC2865] defines a PDU for transporting messages within UDP. The operation of the Code, Identifier, Length, and Authenticator fields are specified in RFC2865. The operation of the Status-Server message is specified in [RFC5997].

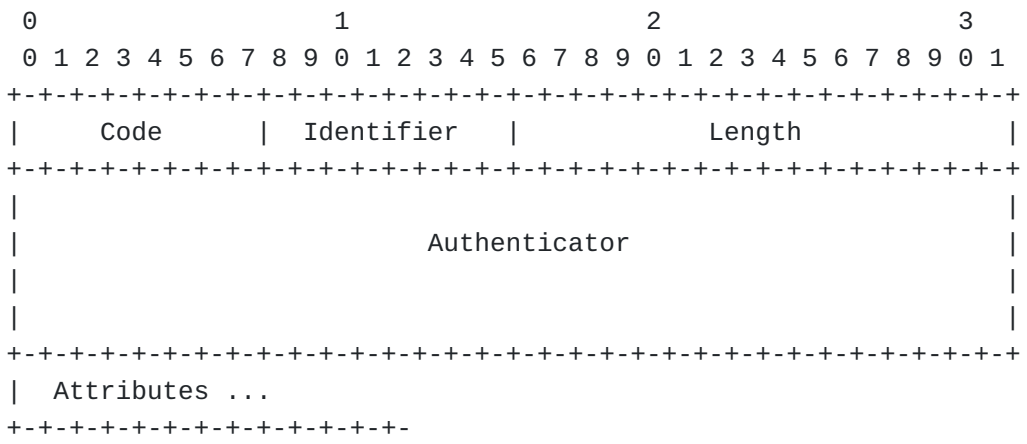

Figure 1

Code

12 for Status-Server. This is the code used in RFC5997.

Identifier

The Identifier field MUST be changed whenever the content of the Attributes field changes, and whenever a valid reply has been

received for a previous request. For retransmissions, the Identifier MUST remain unchanged.

Authenticator

The Request Authenticator value MUST be changed each time a new Identifier is used. The Authenticator does not authenticate the identity of the NAS or the RADIUS server. The Message-Authenticator (Attribute 80) [RFC3579] MUST be used to authenticate both sides of the message exchange.

Length

The Length field is two octets. It indicates the length of the packet including the Code, Identifier, Length, Authenticator and Attribute fields. Octets outside the range of the Length field MUST be treated as padding and ignored on reception. If the packet is shorter than the Length field indicates, it MUST be silently discarded. The minimum length is 20 and maximum length is 4096.

Attributes

The Attribute field is variable in length, and contains the list of Attributes that are required for the type of service, as well as any desired optional Attributes.

2.1. Attributes Needed for Status-Server-LB

The conversation between the NAS and the RADIUS server for the purposes of the load-balance function involves a sending an LB-Request attribute to the server. The server then responds with an LB-Response attribute. Both MUST contain a RADIUS attribute Message-Authenticator [RFC3579].

2.2. Table of Attributes

The following table provides a guide to which attributes may be found in which kinds of packets, and in what quantity.

LB-Request	LB-Response	#	Attribute
0-1	0	4	NAS-IP-Address [Note 1]
0-1	0	32	NAS-Identifier [Note 1]
1	1	80	Message-Authenticator
1	1	191	LB-Request / LB-Response

Figure 2

[Note 1] A Status-Server message MUST contain either a NAS-IP-Address or a NAS-Identifier (or both).

0 This attribute MUST NOT be present in packet. 0-1 Zero or one instance of this attribute MAY be present in packet. 1 Exactly one instance of this attribute MUST be present in packet.

2.3. Required Status-Server-LB Attributes

The required attributes for a valid LB-Request message are outlined here.

2.3.1. NAS-IP Address and/or NAS-Identifier

The NAS-IP-Address or the NAS-Identifier or both attributes are required in an LB-Request message. These are specified in [RFC2865] as Attributes 4 and 32 respectively. This will identify the NAS to the RADIUS server.

2.3.2. Message-Authenticator

In a normal RADIUS access-request message the Request Authenticator field is hashed with the identity material from the calling-station and the RADIUS shared secret. In an LB-Request message there is no calling-station, so this mechanism cannot be used.

RFC3579 specifies Attribute 80 that computes an MD5 hash across the entire RADIUS PDU combined with the shared secret. This mechanism must be used in all LB-Request and LB-Response PDUs.

3. LB-Request Attribute

The conversation between the NAS and the server to implement the Status-Server-LB protocol MUST include a RADIUS message with the LB-Request attribute. This message informs the server that the NAS would like to discover all RADIUS servers that are available to handle RADIUS authentication requests. The NAS anticipates a Status-Server-LB response in the form of an LB-Response PDU.

The Attributes field in the RADIUS message shall be arranged as follows.

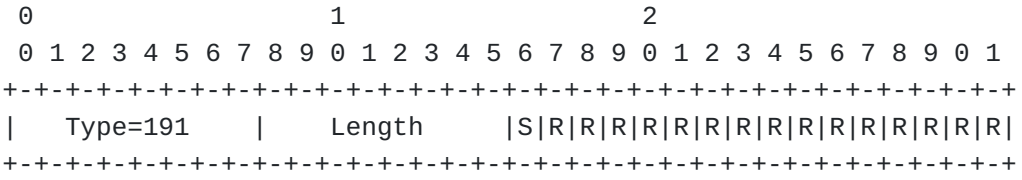

Figure 3

Type

191 for Status-Server Load Balancing.

Length
= 4

The length field calculates the length of the attribute, which includes the type, length and capabilities field.

Server-Status-LB Bit (S-bit)
This bit is to indicate the client can process Status-Server-LB as described in this document. It MUST be set to indicate compliance with this standard.

R-Bit
These bits are reserved for future capabilities of the protocol. These MUST be set to zero on transmission and ignored on receipt.

A NAS that sends an LB-Request attribute but does not receive an LB-Response attribute MUST continue normally as if it had not sent the LB-Request attribute.

4. LB-Response Attribute

When the server receives a Status-Server packet from the NAS and it contains an LB-Request attribute it SHOULD respond with a Status-Server message that contains an LB-Response attribute. In scenarios where the administrator does not want to convey load-balancer information to the NAS the RADIUS server MAY choose to not respond.

If the initial Status-Server message included attributes other than the LB-Request attribute the server MAY choose to respond but simply omit the LB-Response attribute.

4.1. LB-Response Attribute Format

The Attributes field in the RADIUS message shall be arranged as follows.

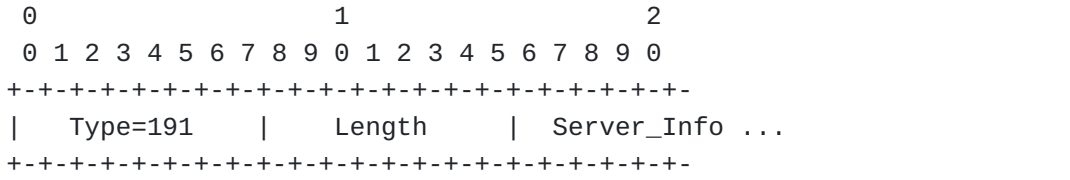

Figure 4

Type
191 for Status-Server Load Balancing.

Length
>= 3

The length field calculates the length in octets of the Type field, Length field, and the concatenation of all the server_info PDUs.

Server_Info

The String field is one or more server_info PDUs. Each PDU defines the status of a single server and its defining characteristics.

5. The SVR-Record TLV

The SVR-Record (Server Record) TLV is a family of Type-Length-Value attributes that holds multiple sub-attributes as described in Section 5. Each SVR-Record type supports a particular address family. SVR-Record-IPv4 and SVR-Record-IPv6 are defined in this document. Other address families may be supported by future standards.

In order to support more than six SVR-Records in one RADIUS packet these attributes are allocated in the Long-Extended-Type Attribute defined in [RFC6929].

5.1. SVR-Record-IPv4

Description

This attribute indicates a SVR-Record that contains information about an IPv4 RADIUS server. This attribute conforms to the TLV-Data type described in [RFC8044].

A summary of the SVR-Record_ipv4 Attribute format is shown below. The fields are transmitted from left to right.

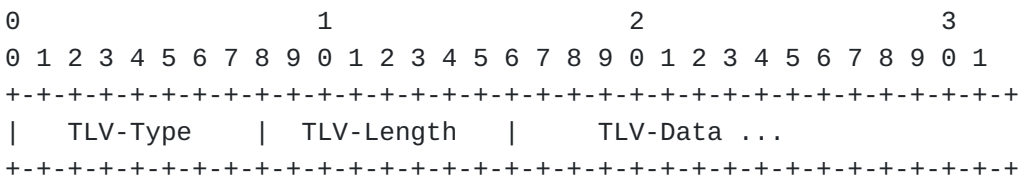

Figure 5

Type

245.TBD1 for SVR-Record-IPv4.

Length

The TLV-Length field is one octet and indicates the length of this TLV, including the TLV-Type, TLV-Length, and TLV-Value fields. It MUST have a value between 3 and 255. If a client or

server receives a TLV with an invalid TLV-Length, then the attribute that encapsulates that TLV MUST be considered to be an invalid attribute and is handled as per [RFC6929], Section 2.8.

TLVs having a TLV-Length of two (2) MUST NOT be sent; omit the entire TLV instead.

TLV-Data

The TLV-Data for the SVR-Record-IPv6 attribute indicates the usage of one RADIUS server that has an IPv6 address. It must include sub-attributes for LB-TTL, LB-priority, LB-Weight and LB-IPv4.

5.2. SVR-Record-IPv6

Description

This attribute indicates an SVR-Record that contains information about an IPv6 RADIUS server. This attribute conforms to the TLV-Data type described in RFC8044.

A summary of the SVR-Record-IPv6 Attribute format is shown below. The fields are transmitted from left to right.

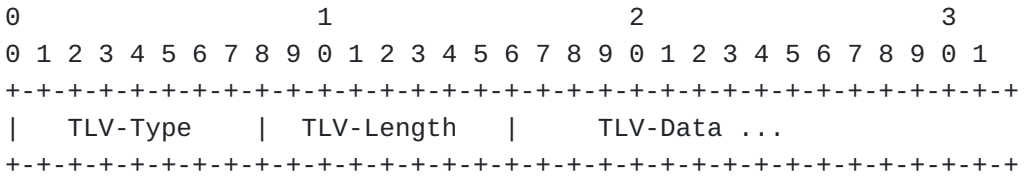

Figure 6

Type

245.TBD2 for SVR-Record-IPv6.

Length

The TLV-Length field is one octet and indicates the length of this TLV, including the TLV-Type, TLV-Length, and TLV-Value fields. It MUST have a value between 3 and 255. If a client or server receives a TLV with an invalid TLV-Length, then the attribute that encapsulates that TLV MUST be considered to be an invalid attribute and is handled as per [RFC6929], Section 2.8.

TLVs having a TLV-Length of two (2) MUST NOT be sent; omit the entire TLV instead.

TLV-Data

The TLV-Data for the SVR-Record-IPv6 attribute indicates the usage of one RADIUS server that has an IPv6 address. It must include sub-attributes for LB-TTL, LB-priority, LB-Weight and LB-IPv6.

5.3. Table of Sub-Attributes

The following table provides a guide to which sub-attributes may be found in which kinds of packets and in what quantity.

SVR-Record-IPv4	SVR-Record-IPv6	#	Sub-Attribute
1	1	1	LB-TTL
1	1	2	LB-priority
1	1	3	LB-Weight
1	0	4	LB-IPv4
0	1	5	LB-IPv6

Figure 7

The following table defines the meaning of the above table entries.

0	This attribute MUST NOT be present in TLV.
0+	Zero or more instances of this attribute MAY be present in packet.
0-1	Zero or one instance of this attribute MAY be present in packet.
1	Exactly one instance of this attribute MUST be present in packet.

Figure 8

6. Sub-Attributes Needed for SVR-Record TLV

The server_info field contains multiple SVR-Records. Each SVR-Record will contain multiple sub-fields that are documented in this section.

6.1. LB-TTL

Description

This attribute indicates how long the NAS should consider the SVR-Record valid.

A summary of the User-Name Attribute format is shown below. The fields are transmitted from left to right.

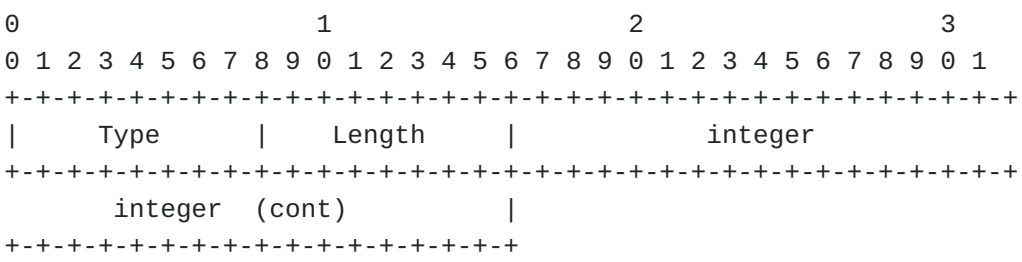

Figure 9

Type

1 for LB-TTL

Length

6

integer

This field indicates the number of seconds this SVR-Record should be in-use and considered valid. If the TTL is 0, the entry SHOULD be removed from the cache immediately. If the value is 0xffffffff, the recipient can decide locally how long to store the mapping. It conforms to the integer data type specified in RFC8044.

6.2. LB-Priority

Description

This attribute indicates the priority of the SVR-Record.

A summary of the LB-priority Attribute format is shown below. The fields are transmitted from left to right.

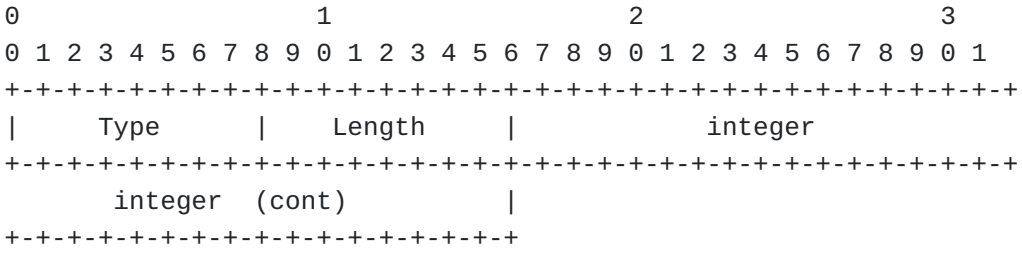

Figure 10

Type
2 for LB-Priority

Length
6

integer
This field indicates the priority of this target host. A NAS MUST attempt to use the target RADIUS server with the lowest-numbered priority it can reach; target servers with the same priority SHOULD be used in an order defined by the LB-Weight field. An SVR-Record with a lower-numbered LB-priority should always used be before an SVR-Record of a higher-numbered LB-priority, regardless of LB-Weight. It conforms to the integer data type specified in RFC8044.

6.3. LB-Weight

Description
This attribute indicates the weighting of the SVR-Record relative to other SVR-Record of the same priority.

A summary of the LB-Weight Attribute format is shown below. The fields are transmitted from left to right.

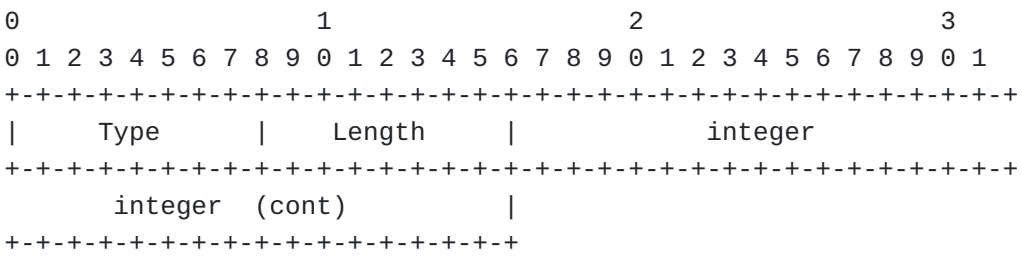

Figure 11

Type
3 for LB-Weight

Length
6

integer
The weight field specifies a relative weight for entries with the same priority. Larger weights SHOULD be given a proportionately higher probability of being used for AAA services. SVR-Record with a lower-numbered LB-priority should always be used before SVR-Record of a higher-numbered LB-priority, regardless of LB-

Weight. It conforms to the integer data type specified in RFC8044.

6.4. LB-IPv4

Description

This attribute indicates the IPv4 address of the SVR-Record.

A summary of the LB-IPv4 Attribute format is shown below. The fields are transmitted from left to right.

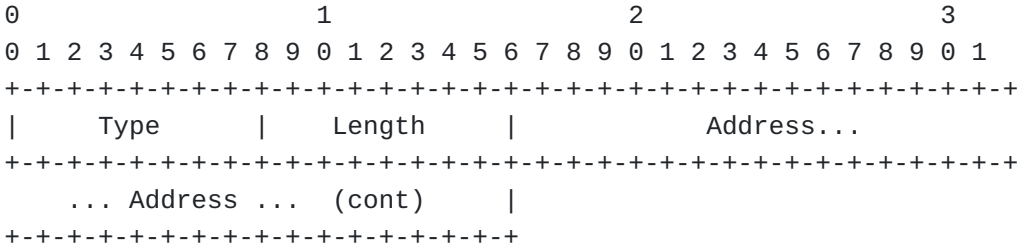

Figure 12

Type

4 for LB-IPv4

Address

6

Address

This is the IPv4 address of the SVR-Record. While taking into consideration the LB-priority and LB-Weight attributes the NAS SHOULD attempt to use this address as a RADIUS server. All considerations for client and server authentication mechanisms MUST still be observed. It conforms to the ipv4addr data type specified in RFC8044.

6.5. LB-IPv6

Description

This attribute indicates the IPv6 address of the SVR-Record.

A summary of the LB-IPv6 Attribute format is shown below. The fields are transmitted from left to right.

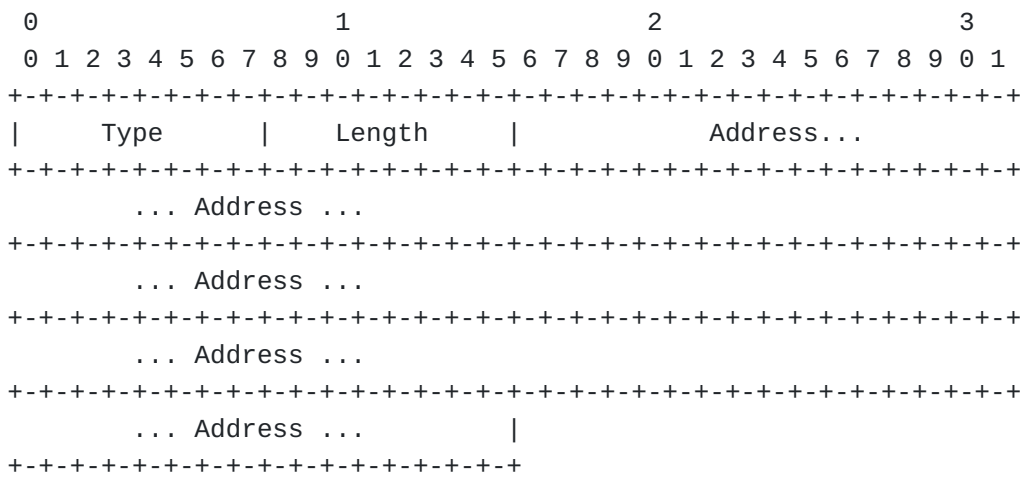

Figure 13

Type
5 for LB-IPv4

Length
18

Address
This is the IPv6 address of the SVR-Record. While taking into consideration the LB-priority and LB-Weight attributes the NAS SHOULD attempt to use this address as a RADIUS server. All considerations for client and server authentication mechanisms MUST still be observed. It conforms to the ipv6addr data type specified in RFC8044.

7. Load Balancing Rules

When there are multiple SVR-Records available with the same LB-priority, a non-zero weight, and excluding those SVR-Records with an inferior LB-priority, the NAS MUST distribute the AAA messages across those servers.

7.1. Session-Based Load Balancing

RADIUS servers may cache user data after retrieving that data from a back-end database. If a NAS queries a RADIUS server for a particular user the server cache will be populated. If the NAS then uses the same RADIUS server for subsequent queries for the same user it will represent a cache hit. Sending the query to a different RADIUS server may represent a cache miss. A cache miss may be an expensive operation in terms of time and other server resources. Under these Status-Server-LB rules the NAS MUST send all RADIUS messages relative to a particular session to the same RADIUS server to maximize the probability of a cache-hit.

A NAS may terminate multiple sessions from multiple calling-stations. It SHOULD use whatever means is available from the conversation with the calling-station to uniquely identify sessions. Examples include, but are not limited to, the RADIUS attributes User Name and Calling-Station-Id.

7.2. Load Balancing Weight

The weight field specifies a relative weight for entries with the same priority. Larger weights SHOULD be given a proportionately higher probability of being used for AAA services. The NAS should attempt to maintain the proper distribution of sessions based on LB-Weight, but MUST retain the properties of Session-Based Load Balancing described in Section 7.1

8. Security Considerations

RADIUS Server Load Balancing has many of the same security implications as the base RADIUS protocol.

8.1. Clear Text Transmission

The RADIUS protocol is unencrypted clear-text on the wire. The Message-Authenticator attribute is required and protects the RADIUS message from tampering, but it does not encrypt the data. The Server Load Balance extensions to RADIUS do not communicate any user identity information or user-authentication materials. Being able to view the LB-Request or LB Response PDU in clear-text does not compromise any of this data.

8.2. Reconnaissance

If an attacker could impersonate a NAS then they would be able to use the LB-Request to gain a list of available RADIUS servers. This is additional information the attacker may not have access to otherwise. Each RADIUS message is authenticated with the Message-Authenticator attribute, so the attacker would need access to the shared secret to correctly submit a valid LB-Request.

8.3. MD5

The RADIUS Response Authenticator and the Message-Authenticator attribute both rely on the integrity of the MD5 algorithm. If an attacker is able to reverse-engineer the shared secret by using a weakness in the MD5 algorithm, then the Message-Authenticator attribute will no longer provide message integrity.

This document recommends the development of better mechanisms for authenticating messages within the RADIUS protocol using more modern encryption standards.

9. NAS identifying Initial RADIUS Servers

For a NAS to use the RADIUS Server Load Balancing service it must be able to contact an initial RADIUS server. Options include static configuration of an initial seed RADIUS server. Other implementations may use a DNS SRV record of the form _radius._udp.name. The format and use of the SRV record is described in [RFC2782].

Implementation-specific mechanisms may be employed but are generally outside the scope of this document.

10. IANA Considerations

This implementation used Attribute 191 for the LB-Request and LB-Response PDU. An official registration is requested from IANA.

The Vendor Specified Attribute 26 may be used to encapsulate the LB-Request and LB-Response PDU where no vendor interoperability is required.

11. References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", BCP 14, RFC 2865, DOI 10.17487/RFC2865, June 2000, <https://www.rfc-editor.org/info/rfc2865>.

[RFC3579] Aboba, B., Calhoun, P., and W. Simpson, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", BCP 14, RFC 3579, DOI 10.17487/RFC3579, September 2003, <https://www.rfc-editor.org/info/rfc3579>.

[RFC5997] Dekok, A., "Use of Status-Server Packets in the Remote Authentication Dial In User Service (RADIUS) Protocol", BCP 14, RFC RFC5997, DOI 10.17487/RFC5997, August 2010, <https://www.rfc-editor.org/info/rfc5997>.

[RFC6929] Dekok, A., "Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions", BCP 14, RFC 6929, DOI

10.17487/RFC6929, April 2013, <https://www.rfc-editor.org/info/rfc6929>.

[RFC8044] Dekok, A., "Data Types in RADIUS", BCP 14, RFC 8044, DOI 10.17487/RFC8044, January 2017, <https://www.rfc-editor.org/info/rfc8044>.

Author's Address

Daniel Massameno (editor)
Yale University
150 Munson Street
New Haven, CT 06492
United States of America

Email: dan.massameno@yale.edu