

Network Working Group
INTERNET-DRAFT
Category: Proposed Standard
<[draft-ietf-radext-tcp-transport-00.txt](#)>
Expires: June 11, 2009
11 December 2008

A. DeKok
FreeRADIUS

RADIUS Over TCP
draft-ietf-radext-tcp-transport-00

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on June 11, 2009.

Copyright Notice

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Copyright (c) 2008 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.

Abstract

The Remote Authentication Dial In User Server (RADIUS) Protocol has traditionally used the User Datagram Protocol (UDP) as it's underlying transport layer. This document defines RADIUS over the Transport Control Protocol (TCP).

Table of Contents

1.	Introduction	3
1.1.	Benefits of Reliable Transport	3
1.2.	Drawbacks of Reliable Transport	4
1.3.	Terminology	4
1.4.	Requirements Language	5
2.	Changes to RADIUS	5
2.1.	Packet Format	5
2.2.	TCP Ports	5
2.3.	Management Information Base (MIB)	6
2.4.	Interaction with RadSec	6
2.4.1.	Applicability	6
2.5.	RADIUS Proxies	7
2.6.	TCP Specific Issues	7
2.6.1.	Duplicates and Retransmissions	8
2.6.2.	Shared Secrets	9
2.6.3.	Malformed Packets and Unknown Clients	9
2.6.4.	Limits of the ID Field	10
2.6.5.	EAP Sessions	10
2.6.6.	TCP Applications are not UDP Applications	11
3.	Diameter Considerations	11
4.	IANA Considerations	11
5.	Security Considerations	11
6.	References	12
6.1.	Normative References	12
6.2.	Informative References	12

1. Introduction

The RADIUS Protocol has been defined in [\[RFC2865\]](#) as using the User Datagram Protocol (UDP) for the underlying transport layer. While there are a number of benefits to using UDP as outlined in [\[RFC2865\] Section 2.4](#), there are also some limitations:

- * Unreliable transport. As a result, systems using RADIUS have to implement application-layer timers and re-transmissions, as described in [\[RFC5080\] Section 2.2.1](#).
- * Packet fragmentation. [\[RFC2865\] Section 3](#) permits RADIUS packets to up to 4096 octets in length. These packets are larger than the default Internet MTU (576), resulting in fragmentation of the packets at the IP layer. Transport of fragmented UDP packets appearsto be a poorly tested code path on network devices. Some devices appear to be incapable of transporting fragmented UDP packets, making it difficult to deploy RADIUS in a network where those devices are deployed.
- * Connectionless transport. Neither clients no servers can reliably detect when the other is down. This information has be deduced from the absence of a reply to a request.

As RADIUS is widely deployed, and has been widely deployed for well over a decade, these issues are relatively minor. However, new systems may be interested in choosing a different set of trade-offs than those outlined in [\[RFC2865\] Section 2.4](#). For those systems, we define RADIUS over TCP.

1.1. Benefits of Reliable Transport

There are a number of benefits to using a reliable transport. For example, when RADIUS is used to carry EAP conversations [\[RFC3579\]](#), the EAP exchanges may involve 10 round trips at the RADIUS application layer. If we assume a 0.1% probability of packet loss in each direction, then approximately 2% ($1 - 0.999^{20}$) of the authentication attempts will have a lost packet. If we assume a 0.01% packet loss, then 0.2% of authentication attempts will result in a lost packet.

These lost packets require the supplicant and/or the NAS to re-transmit packets at the application layer. The difficulty with this approach is that retransmission implementations have historically been poor. Some implementations retransmit packets, others do not. Some implementations are incapable of detecting EAP retransmissions, and will instead treat the retransmitted packet as an error.

These retransmissions have a high likelihood of causing the entire

authentication session to fail. For systems with millions to tens of millions of users, such a high authentication failure rate (0.2% to 2%) may be unacceptable.

Using TCP as an underlying reliable transport means that the RADIUS implementations can remove all of the application-layer retransmissions, and instead rely on the Operating System (OS) kernel's well-tested TCP transport.

1.2. Drawbacks of Reliable Transport

No protocol is perfect for all uses. RADIUS over TCP has some drawbacks, as noted in [\[RFC2865\] Section 2.4.](#) [\[RFC3539\] Section 2](#) discusses further issues with using TCP as a transport for Authentication, Authorization, and/or Accounting (AAA) protocols such as RADIUS.

The impact of these issues is discussed in more detail, below.

1.3. Terminology

This document uses the following terms:

Network Access Server (NAS)

A device that provides an access service for a user to a network.

RADIUS server

A RADIUS authentication, authorization, and/or accounting (AAA) server is an entity that provides one or more AAA services to a NAS.

RADIUS proxy

A RADIUS proxy acts as a RADIUS server to the NAS, and a RADIUS client to the RADIUS server.

RADIUS request packet

A packet originated by a RADIUS client to a RADIUS server. e.g. Access-Request, Accounting-Request, CoA-Request, or Disconnect-Request.

RADIUS response packet

A packet sent by a RADIUS server to a RADIUS client, in response to a RADIUS request packet. e.g. Access-Accept, Access-Reject, Access-Challenge, Accounting-Response, CoA-ACK, etc.

1.4. 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 [[RFC2119](#)].

2. Changes to RADIUS

Adding TCP as a RADIUS transport has a number of impacts on the protocol, on applications using the protocol, and on networks that deploy the protocol. This section outlines those impacts, and defines behaviors.

2.1. Packet Format

The RADIUS packet format is unchanged from [[RFC2865](#)], [[RFC2866](#)], and [[RFC5176](#)]. Specifically, all of the following portions of RADIUS MUST be unchanged when using RADIUS over TCP:

- * Packet format
- * Permitted codes
- * Request Authenticator calculation
- * Response Authenticator calculation
- * Minimum packet length
- * Maximum packet length
- * Attribute format
- * Vendor-Specific Attribute (VSA) format
- * Permitted data types
- * Calculations of dynamic attributes such as CHAP-Challenge, or Message-Authenticator.

The changes to RADIUS implementations required to implement this specification are largely limited to the code that sends and receives packets on the network.

2.2. TCP Ports

IANA has already assigned TCP ports for RADIUS transport, as outlined below:

- * radius 1812/udp
- * radius-acct 1813/tcp
- * radius-dynauth 3799/tcp

These ports are unused by existing RADIUS applications. Implementations SHOULD use the assigned values as the default ports for RADIUS over TCP.

The early deployment of RADIUS was done using UDP port number 1645, which conflicts with the "datametrics" service. Implementations using RADIUS over TCP MUST NOT use TCP ports 1645 or 1646 as the default ports for this specification.

2.3. Management Information Base (MIB)

The MIB definitions in [[RFC4668](#)], [[RFC4669](#)], [[RFC4670](#)], [[RFC4671](#)], [[RFC4672](#)], and [[RFC4673](#)] each contain only one reference to UDP. These references are in the DESCRIPTION field of the MIB definition, and are in the form of "The UDP port" or "the UDP destination port".

Implementations of RADIUS over TCP MAY re-use these MIBs to perform statistics counting for RADIUS over TCP connections. However, implementors are warned that there is no way for these MIBs to distinguish between packets sent over UDP or over TCP transport. Similarly, there is no requirement in RADIUS that the RADIUS services offered over UDP on a particular IP address and port are identical to the RADIUS services offered over TCP on a particular IP address and the same (numerical) port.

2.4. Interaction with RadSec

IANA has already assigned TCP ports for RadSec transport, as outlined below:

* radsec	2083/tcp
----------	----------

This value SHOULD be used as the default port for RADIUS over TLS (i.e. RadSec). The "radius" port (1812/tcp) SHOULD NOT be used for RadSec.

2.4.1. Applicability

As noted in [[RFC3539](#)] [Section 2.1](#), for systems originating low numbers of RADIUS request packets, inter-packet spacing is often larger than the RTT. In those situations, RADIUS over TCP SHOULD NOT be used.

In general, RADIUS clients generating small amounts of RADIUS traffic SHOULD NOT use TCP. This suggestion will usually apply to most NASes, and to most clients that originate CoA-Request and Disconnect-Request packets.

RADIUS over TCP is most applicable to RADIUS proxies that exchange a large volume of packets with RADIUS clients and servers (10's to 1000's of packets per second). In those situations, RADIUS over TCP is a good fit, and may result in increased network stability and

performance.

2.5. RADIUS Proxies

As RADIUS is a "hop by hop" protocol, a RADIUS proxy effectively shields the client from any information about downstream servers. While the client may be able to deduce the operational state of the local server (i.e. proxy), it cannot make any determination about the operational state of the downstream servers.

If a request is proxied through intermediate proxies, it is not possible to detect which of the later hops is responsible for the absence of a reply. An intermediate proxy also cannot signal that the outage lies in a later hop because RADIUS does not have the ability to carry such signalling information. This issue is further exacerbated by some proxy implementations that do not reply to a client if they do not receive a reply to a proxied request.

When UDP was used as a transport protocol, the absence of a reply can cause a client to deduce (incorrectly) that the proxy is unavailable. The client could then fail over to another server, or conclude that no "live" servers are available. This situation is made even worse when requests are sent through a proxy to multiple destinations. Failures in one destination may result in service outages for other destinations, if the client erroneously believes that the proxy is unresponsive.

For RADIUS over TCP, the continued existence of the TCP connection SHOULD be used to deduce that the service on the other end of the connection is still responsive. Further, the application layer watchdog defined in [\[RFC3539\] Section 3.4](#) enables clients to determine that the server is "live", even though it may not have responded recently to other, non-watchdog requests.

RADIUS clients using RADIUS over TCP MUST NOT decide that a connection is down until the application layer watchdog algorithm has marked it DOWN ([\[RFC3539\] Appendix A](#)). RADIUS clients using RADIUS over TCP MUST NOT decide that a RADIUS server is unresponsive until all TCP connections to it have been marked DOWN.

2.6. TCP Specific Issues

The guidelines defined in [\[RFC3539\]](#) for implementing an AAA protocol operating over a reliable transport MUST be followed by implementors of this specification.

The Application Layer Watchdog defined in [\[RFC3539\] Section 3.4](#) MUST be used. The Status-Server packet [\[STATUS\]](#) MUST be used as the

application layer watchdog message. Implementations MUST reserve one RADIUS ID per connection for the application layer watchdog message. This restriction is described further below.

Implementations MUST NOT confuse UDP and TCP transport. That is, RADIUS clients and servers MUST be treated as unique based on a key of (IP address, port, transport protocol). Implementations MUST be configurable to have different shared secrets for UDP and TCP to the same destination IP address and numerical port.

This requirement does not forbid the traditional practice of using primary and secondary servers in a fail-over relationship. Instead, it requires that two services sharing an IP address and numerical port, but differing in transport protocol, MUST be treated as independent services for the purpose of fail-over, load-balancing, etc.

Whenever the underlying operating system permits the use of TCP keepalive socket options, their use is RECOMMENDED.

2.6.1. Duplicates and Retransmissions

As TCP is a reliable transport, implementors of this specification MUST NOT retransmit RADIUS packets over the same TCP connection. Similarly, if there is no response to a RADIUS packet over one TCP connection, implementations MUST NOT retransmit that packet over a different TCP connection to the same destination IP address and port.

However, if the TCP connection is broken or closed, the above requirement can be relaxed somewhat. RADIUS request packets that have not yet received a response MAY be transmitted by a RADIUS client over a new TCP connection. As this procedure involves using a new source port, the ID of the packet MAY change. If the ID changes, any security attributes such as Message-Authenticator MUST be recalculated.

If a TCP connection is broken or closed, any cached RADIUS response packets ([\[RFC5080\] Section 2.2.2](#)) associated with that connection MUST be discarded. A RADIUS server SHOULD stop processing any "live" requests associated with that TCP connection. No response to these requests cannot be sent over the TCP connection, so any further processing is pointless. A RADIUS proxy that has a client close it's TCP connection SHOULD silently discard any responses it receives to a proxied requests that is associated with the original client request.

Despite the above requirement, RADIUS servers SHOULD still perform duplicate detection on received packets, as described in [\[RFC5080\] Section 2.2.2](#). This effort can prevent duplicate processing of

packets from non-conformant clients.

As noted above, RADIUS packets SHOULD NOT be re-transmitted to the same destination IP and numerical port, but over a different transport layer. There is no guarantee in RADIUS that the two ports are in any way related. This requirement does not forbid the practice of putting multiple servers into a fail-over or load-balance pool.

Much of the discussion in this section can be summarized by the following requirement. RADIUS requests MAY be re-transmitted verbatim only if the following 5-tuple (Client IP address, Client port, Transport Protocol, Server IP address, Server port) remains the same. If any field of that 5-tuple changes, the packet MUST NOT be considered to be a re-transmission. Instead, the packet MUST be considered to be a new request, and be treated accordingly. (e.g. header calculations, packet signatures, associated timers and counters, etc.)

The above requirement is necessary, but not sufficient in all cases. Other specifications give additional situations where the packet is to be considered as a new request. Those recommendations MUST be followed.

2.6.2. Shared Secrets

The use of shared secrets in calculating the Response Authenticator, and other attributes such as User-Password or Message-Authenticator [[RFC3579](#)] MUST be unchanged from previous specifications.

Clients and servers MUST be able to store and manage shared secrets based on the key described above, of (IP address, port, transport).

2.6.3. Malformed Packets and Unknown Clients

The original specifications say that an implement should "silently discard" a packet in a number of circumstances. This action has no further consequences for UDP transport, as the "next" packet is completely independent of the previous one.

When TCP is used as a transport, decoding the "next" packet on a connection depends on the proper decoding of the previous packet. As a result, the behavior with respect to discarded packets has to change.

Implementations of this specification SHOULD treat the "silently discard" texts referenced above as "silently discard and close the connection." Specifically, the TCP connection MUST be closed if any

of the following circumstances are seen:

- * Packet from an unknown client (using the key as defined above)
- * Packet with an invalid code field
- * Packet that is less than the minimim RADIUS packet length
- * Packet that is more than the minimim RADIUS packet length
- * A packet that is otherwise malformed, e.g. Attribute Length of 0 or 1
- * Packet where the Request Authenticator fails validation (if applicable)
- * Packet where the Response Authenticator fails validation
- * Packet where the Message-Authenticator fails validation
- * Response packets that do not match any outstanding request

These requirements minimize the possibilty for a misbehaving client or server to wreak havoc on the network.

2.6.4. Limits of the ID Field

The RADIUS ID field is one octet in size. As a result, any one TCP connection can have only 256 "in flight" RADIUS packets at a time. If more than 256 simultaneous "in flight" packets are required, additional TCP connections will need to be opened. This limitation is also noted in [\[RFC3539\] Section 2.4](#).

An additional limit is the requirement to send a Status-Server packet over the same TCP connection as is used for normal requests. As noted in [\[STATUS\]](#), the response to a Status-Server packet is either an Access-Accept, or an Accounting-Response. If all IDs were allocated to normal requests, then there would be no free Id to use for the Status-Server packet, and it could not be sent over the connection.

Implementations SHOULD reserve ID zero on each TCP connection for Status-Server packets. This value was picked arbitrarily, as there is no reason to choose any one value over another for this use.

It is tempting to extend RADIUS to permit more than 256 outstanding packets on one connection. However, doing so will likely require fundamental changes to the RADIUS protocol, and as such, are outside of the scope of this specification.

2.6.5. EAP Sessions

When RADIUS clients send EAP requests using RADIUS over TCP, they SHOULD choose the same TCP connection for all packets related to one EAP conversation. A simple method that may often work is hashing the contents of the Calling-Station-Id attribute, which normally contains

the MAC address. The output of that hash can be used to select a particular TCP connection.

It may be difficult to implement this suggestion in practice, as busy servers may allocate all RADIUS IDs in one TCP connection in the time between two subsequent EAP packets. It is difficult to suggest simple and reasonable methods to address this issue.

2.6.6. TCP Applications are not UDP Applications

Implementors should be aware that programming a robust TCP application can be a very different process than programming a robust UDP application. We RECOMMEND that implementors of this specification familiarize themselves with TCP application programming concepts. We RECOMMEND also that existing TCP applications be examined with an eye to robustness, performance, scalability, etc.

Clients and servers SHOULD implement configurable connection limits. Allowing an unlimited number of connections may result in resource exhaustion.

Further discussion of implementation issues is outside of the scope of this document.

3. Diameter Considerations

This document defines TCP as a transport layer for RADIUS. It defines no new RADIUS attributes or codes. The only interaction with Diameter is in a RADIUS to Diameter, or in a Diameter to RADIUS gateway. The RADIUS side of such a gateway MAY implement RADIUS over TCP, but this change has no effect on Diameter.

4. IANA Considerations

This document requires no action by IANA.

5. Security Considerations

As the RADIUS packet format, signing, and client verification are unchanged from prior specifications, all of the security issues outlined in previous specifications for RADIUS over UDP are also applicable here.

As noted above, clients and servers SHOULD support configurable connection limits. Allowing an unlimited number of connections may result in resource exhaustion.

There are no (at this time) other known security issues for RADIUS

over TCP transport.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", [RFC 2865](#), June 2000.
- [RFC3539] Aboba, B. et al., "Authentication, Authorization and Accounting (AAA) Transport Profile", [RFC 3539](#), June 2003.

6.2. Informative References

- [RFC2866] Rigney, C., "RADIUS Accounting", [RFC 2866](#), June 2000.
- [RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", [RFC 3579](#), September 2003.
- [RFC4668] Nelson, D, "RADIUS Authentication Client MIB for IPv6", [RFC 4668](#), August 2006.
- [RFC4669] Nelson, D, "RADIUS Authentication Server MIB for IPv6", [RFC 4669](#), August 2006.
- [RFC4670] Nelson, D, "RADIUS Accounting Client MIB for IPv6", [RFC 4670](#), August 2006.
- [RFC4671] Nelson, D, "RADIUS Accounting Server MIB for IPv6", [RFC 4671](#), August 2006.
- [RFC4672] Nelson, D, "RADIUS Dynamic Authorization Client MIB", [RFC 4672](#), August 2006.
- [RFC4673] Nelson, D, "RADIUS Dynamic Authorization Server MIB", [RFC 4673](#), August 2006.
- [RFC5080] Nelson, D. and DeKok, A, "Common Remote Authentication Dial In User Service (RADIUS) Implementation Issues and Suggested Fixes", [RFC 5080](#), December 2007.
- [RFC5176] Chiba, M. et al., "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", [RFC 5176](#),

January 2008.

[STATUS] DeKok, A., "Use of Status-Server Packets in the Remote Authentication Dial In User Service (RADIUS) Protocol", [draft-ietf-radext-status-server-02.txt](#), November 2008.

Acknowledgments

None at this time.

Authors' Addresses

Alan DeKok
The FreeRADIUS Server Project
<http://freeradius.org/>

Email: aland@freeradius.org

Open issues

Open issues relating to this document are tracked on the following web site:

<http://www.drizzle.com/~aboba/RADEXT/>