DTNRG
Internet-Draft
Intended status: Experimental
Expires: May 23, 2009

H. Kruse
S. Ostermann
Ohio University
November 19, 2008

UDP Convergence Layers for the DTN Bundle and LTP Protocols
draft-irtf-dtnrg-udp-clayer-00

Status of this Memo

Abstract

This document specifies the use of the User Datagram Protocol (UDP) as a method for transporting DTN protocol data over the Internet. The specification covers both the use of UDP as a convergence layer for the Bundle Protocol as well as the use of UDP to carry LTP segments.

Table of Contents

1. Introduction

Delay/Disruption Tolerant Network (DTN) communication protocols include the Bundle Protocol described in RFC 5050 [RFC5050], which provides reliable transmission of application data blocks (bundles) with optional intermediate custody transfer, and the Licklider Transport Protocol (LTP), RFCs 5325 [RFC5325], 5326 [RFC5326], and 5327 [RFC5327] which can be used to transmit bundles reliably and efficiently over a point to point link. It is often desirable to test these protocols over Internet Protocol links. draft-irtf-dtnrg-tcp-clayer [I-D.irtf-dtnrg-tcp-clayer] defines a method for transporting bundles over TCP. This draft specifies the convergence layer for transmitting either bundles or LTP blocks over UDP.

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].

2. General UDP Considerations

2.1. UDP Checksums are Required

Both the core bundle protocol specification and core LTP specification assume that they are transmitting over an erasure channel, i.e. a channel that either delivers packets correctly or not at all. The UDP CL transmitter therefore MUST NOT disable UDP checksums, and the UDP CL receiver MUST NOT disable checking of received UDP checksums.

Even when UDP checksums are enabled a small probablity of UDP packet corruption remains. In some environments it may be acceptable for LTP or the bundle protocol to occasionally receive corrcupted input.

In general, however, a UDP CL implementation SHOULD use optional security extensions available in the bundle protocol or LTP to protect against message corruption (see the protocol specific sections below).

2.2. Congestion Control

UDP operates on a packet by packet, best effort delivery basis. It provides no congestion control. When the bundle protocol or LTP are operated over UDP, the lack of congestion control can interfere with other traffic in the network, and will be particularly harmful to traffic that does obey congestion control. If the UDP CL is used to

send more than a very small number of packets at a time, it either SHOULD NOT be used outside an isolated network, or it MUST implement congestion control procedures as outlined in RFC 5405.

2.3. How and Where to Deal with Fragmentation

The bundling protocol allows bundles with sizes limited only by node resource constraints. In IPv4, the maximum size of a UDP datagram is nearly 64KB. In IPv6, when using jumbograms [RFC2675], UDP datagrams can be up to 4GB in size [RFC2147]. It is well understood that sending large IP datagrams that must be fragmented by the network has enormous efficiency penalties [Kent88]. The primary efficiency penalty is increased loss probability. When a large datagram is broken into a number of fragments, the original datagram can only be recreated if all the fragments arrive at the ultimate destination for reassembly. When transmitted over a network with a packet loss probability of 2%, for example, a single, unfragmented datagram will arrive with probability 98%; a large datagram fragmented into 10 fragments will have all of its fragments arrive with probability 98%**10, giving a datagram arrival probability of only 81.7%. The higher-level protocol using UDP for delivery can retransmit lost UDP datagrams, but cannot retransmit lost IP datagram fragments. Therefore, retransmitting large, lost datagrams because of a small number of missing fragments can require many more packets than retransmitting a number of smaller, unfragmented datagrams because only the missing pieces need to be retransmitted. The other efficiency penalty paid by fragmentation that would be significant for DTN is the resources (time, complexity, and memory) required for IP reassembly.

When an LTP CL is using UDP for datagram delivery, it SHOULD NOT create segments that will result in UDP datagrams that will need to be fragmented, as discussed above. When using UDP directly as a CL, the software SHOULD NOT directly encapsulate large bundles into large UDP datagrams that would need to be fragmented, as discussed above. In the latter case, the bundle protocol specification provides a bundle fragmentation concept [RFC5050] that allows a large bundle to be divided into bundle fragments, each of which SHOULD be created of sufficiently small size that it can then be encapsulated into a UDP datagram that will not need to be fragmented.

Without information from elsewhere in the networking stack about path MTU, the protocol can assume a minimum path MTU that would allow 512 bytes of UDP data [RFC0791] over IPv4 or (1280-(UDP and IP header sizes)) bytes [RFC1883] over IPv6.

2.4. Keep Alive Option

It may be desirable for a UDP CL to send "keep-alive" packets during extended idle periods. This may be needed to refresh a contact table entry at the destination, or to maintain an address mapping in a NAT or a dynamic access rule in a firewall. Therefore, a UDP CL MAY send a UDP packet containing exactly 4 octets of zero bits. A UDP CL receiving such a packet MUST discard this packet; the receiving CL may then perform local maintenance of its state tables, these maintenance functions are not covered in this draft. Note that "real" CL packets will always contain more than 4 octets of information (either the bundle or the LTP header); keep-alives will therefore never be mistaken for actual data packets.

3. Bundle Protocol over UDP Convergence Layer

In general, the use of the bundle protocol over a UDP CL is discouraged. Bundles can be of (almost) arbitrary length, and the bundle protocol does not include an effective retransmission mechanism. Whenever possible the bundle protocol SHOULD be operated over the TCP Convergence Layer or over LTP.

If a UDP CL is used for transmission of bundles, every UDP packet MUST contain exactly one bundle or four zero octets as a keep-alive. The UDP CL SHOULD use avalailable operating system services to obtain the largest supported UDP packet size, and MAY use the default UDP packet size limit if path-specific information is not available. For bundles that are too large for the supported UDP packet size, the bundle protocol fragmentation process SHOULD be used transmit the large bundle.

The UDP CL for bundle protocol use SHOULD use the IANA assigned port 4556/UDP; the use of other port numbers is implementation specific.

4. LTP over UDP Convergence Layer

LTP is designed as a point to point protocol within DTN, and it provides intrinsic acknowledgement and retransmission facilities. Transmission of LTP over a UDP CL is therefore the most appropriate choice. When a UDP CL is used to transmit LTP data, every UDP packet MUST contain exactly one LTP segment or four zero octets as a keep-alive. The UDP CL SHOULD use avalailable operating system services to obtain the largest supported UDP packet size, and MAY use the default UDP packet size limit if path-specific information is not available. LTP MUST perform segmentation in such a way as to insure that every LTP segments fits into a UDP packet.

The UDP CL for LTP SHOULD use the IANA assigned port 1113/UDP; the use of other port numbers is implementation specific.

5. Acknowledgements

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

This memo describes the use of UDP to transport DTN applications

data. Hosts may be in a position of having to accept and process UDP packet from unknown sources; the DTN Endpoint ID can be discovered only after the bundle has been retrieved from the UDP packet. Hosts SHOULD use authentication methods available in the DTN specifications to prevent malicious hosts from inserting unknown data into the application.

Hosts need to listen for and process UDP data on the known LTP or bundle protocol ports. A denial of service scenario exists where a malicious host send UDP packets at a high rate, forcing the receiving hosts to use its resources to process and attempt to authenticate these data. Whenever possible, hosts SHOULD use IP address filtering to limit the origin of packets to known hosts.

8. References

8.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC1883] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, December 1995.

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

[RFC2147] Borman, D., "TCP and UDP over IPv6 Jumbograms", RFC 2147, May 1997.

[RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", RFC 2675, August 1999.

[RFC5050] Scott, K. and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.

[RFC5325] Burleigh, S., Ramadas, M., and S. Farrell, "Licklider Transmission Protocol - Motivation", RFC 5325, September 2008.

[RFC5326] Ramadas, M., Burleigh, S., and S. Farrell, "Licklider

Transmission Protocol - Specification", RFC 5326, September 2008.

[RFC5327] Farrell, S., Ramadas, M., and S. Burleigh, "Licklider Transmission Protocol - Security Extensions", RFC 5327, September 2008.

8.2. Informative References

[I-D.irtf-dtnrg-tcp-clayer] Demmer, M. and J. Ott, "Delay Tolerant Networking TCP Convergence Layer Protocol", draft-irtf-dtnrg-tcp-clayer-02 (work in progress), November 2008.

[Kent88] Kent, C. and J. Mogul, "Fragmentation considered harmful.", 1988, <http://doi.acm.org/10.1145/55482.55524>.

Authors' Addresses

Hans Kruse
Ohio University
292 Lindley Hall
Athens, OH 45701
United States

Phone: +1 740 593 4891
Email: kruse@ohiou.edu

Shawn Ostermann
Ohio University
Stoecker Engineering Center
Athens, OH 45701
United States

Phone: +1 740 593 1566
Email: ostermann@eecs.ohiou.edu

Full Copyright Statement

Intellectual Property