

Transport-Services
Internet Draft
Intended status: Informational
Expires: August 3, 2014

M.J. Montpetit
MIT
I. Zhovnirovsky
QFactor
February 3, 2014

Use Cases and Requirements
draft-montpetit-transport-use-cases-01

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), 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 5, 2014.

Copyright Notice

Copyright (c) 2014 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

This document describes some of the services and application use cases that may require specific transport services over an Internet subnetwork. It also presents, when available, some derived requirements and conditions of utilization that will allow to choose the most appropriate transport service. The use cases assume that choice of transport service is fully based on the requirements. In all cases however, the fallback is considered to be TCP or UDP.

Table of Contents

1. Introduction	3
2. Conventions used in this document.	3
3. Application Use Cases and Derived Requirements	3
3.1. Web Surfing/Browsing	3
3.2. Streaming Video.	5
3.3. Real-time Communications	6
3.4. Photo/video Sharing.	7
3.5. Data centers/storage	7
3.6. Internet of Things (IoT)	7
3.7. Public Safety.	7
3.8. Information Centric Networks/Content Centric Networks.	8
4. Security Considerations.	8
5. IANA Considerations.	9
6. Conclusions.	9
7. References	9
8. Acknowledgments.	10

1. Introduction

Different transport services are now being developed in the Internet Community to address the changing nature of services or applications and of the an Internet subnetworks themselves. For example, a number of these transports services are used to address performance requirement that standards such as the TCP and UDP protocol meet only partially either because of impairments in wireless transmission or the need for added security. This document defines the use cases most likely to necessitate the use of one of those transport service and to list their derived requirements. It will help define which transport service is best suited for an application and which API(s) would enable the use of those In general, the applications will request a connection by specifying a set of requirements. The the API will connect it to the best-fit "transport", falling back to TCP/UDP if nothing is compliant. Hence the use cases will help define how to update applications or middleware to perform in a better way without requiring any changes to existing APIs. They also envision some futuristic situations where applications would use a newer enhanced and common API, to benefit from novel network services. In this document the actual underlying link layer is not specified but it is assumed that, for example, wireless transport services could be used in many instances.

2. Conventions used in this document

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\[1\]](#). In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC-2119](#) significance.

3. Application Use Cases and Derived Requirements

In this section we define the use cases for a number of applications that require one (1) or more transport service. The application should be able to use an API based on derived requirements such as bandwidth or delay or security, beyond the usual fallback of TCP and UDP.

3.1. Web Surfing/Browsing

Web surfing and browsing has been identified as applications that will most likely profit from transport services will improve the capacity occupancy, reduce slow start due to dropped packets or multiple connections and other inefficiencies due to for example Head-of-Line blocking.

3.1.1. Faster Browsing via multi-streaming for HTTP prior to 2.0

3.1.1.1. Description

Using a new TCP connection for every HTTP request is known to be inefficient; it causes delay for connection setup and requires a TCP connection to spend most of its time in slow start, operating at a point that may be far from the capacity of the path between the sender and receiver. HTTP/1.1 allows persistently using the same TCP connection for multiple HTTP requests, but with a fixed sequence that is defined by the browser, thereby creating Head-Of-Line (HOL) blocking delay if e.g. the first of several such requests requires accessing a database. SPDY and HTTP/2.0 seek to improve upon the performance of HTTP/1.1 by asynchronously multiplexing multiple application-level (web) flows onto a single TCP connection. However, HOL blocking delay in the order of a Round-Trip Time (RTT) can still occur whenever a packet is dropped, as TCP only allows delivering data in the exact order in which they arrive. This problem is addressed by QUIC [9], a UDP-based transport protocol that operates under the SPDY protocol for browsing [2].

The SPDY/QUIC solution is not without issues. UDP may sometimes be blocked; some middle-boxes rate-limit UDP more than TCP [5]. Running a transport protocol in user space may have issues in very-low-RTT environments due to less precise Operating System timing. Not all clients and servers support both protocols. It seems that being able to provide similar benefits without requiring changing the HTTP implementation could be beneficial for many web browsers and servers.

3.1.1.2. Derived requirements

It is possible to let applications that open the same set of (seemingly) TCP connections between the same pair of hosts benefit from for example the multi streaming features of the Stream Control Transport Protocol(SCTP) [10]. In [11], this was prototypically done with middle-boxes that terminated the TCP connections and mapped them onto streams of a single SCTP association. This did not require any change to the applications themselves. If a similar function was provided by the Operating System, directly underneath the socket interface, the same benefits could be attained with less overhead. This requires hiding the use of SCTP (or any other multi-streaming- capable transport, e.g. Minion [7]) from applications, and it may need signaling inside the protocol in use to tell the host on the receiving end that these streams must be mapped back onto what an application believes to be TCP sockets.

3.1.3. Other HTTP acceleration mechanisms

<To be provided>

3.2. Streaming Video

Streaming video is one of the dominant traffic types on the Internet. Hence the choice of an appropriate transport service for video may become essential in the future to minimize congestion especially in edge subnetworks and improve user experience.

3.2.1. Adaptive Video Streaming

<To be provided>

3.2.2. Scavenger-Type Pre-Fetching Of Streaming Media

3.2.2.1. Description

Streaming media applications use a play-out buffer on the receiver (consumer) side to avoid playback interruptions when the network bandwidth temporarily drops below the bit-rate of the stream. As visualized with the playback bar in most popular streaming systems such as e.g. YouTube, this buffer is typically filled a number of seconds ahead of the play-out time. If the network's capacity could allow filling it much further ahead, this could have multiple benefits for the viewer:

- if congestion happens later, the play-out would stay uninterrupted for a longer period or the risk of interruptions would be reduced
- if a user decides to skip ahead within the length of the already buffered data, this request could be honored instantly as the data are already available in the receiving host

This is, however, not typically done. One reason to avoid pre-fetching much more than users need most of the time is that this may be seen as a waste of network capacity that could otherwise be left up to other applications. If, however, this was done in a scavenger-like mode, using a congestion control that backs off as soon as other flows compete for network capacity, this problem would disappear, and it would probably be beneficial to pre-fetch as much as possible at any time.

3.2.2.2. Derived Requirements

A transport protocol that is used for streaming media should, upon request from the application, switch to a low-priority congestion control mechanism such as LEDBAT on the fly. Switching back from low-priority to "normal" congestion control must happen when e.g. the user moves the playback bar beyond the end of the playback buffer; this must therefore also be possible on the fly.

3.3. Real-time Communications

Real time communications can be divided into three (3) categories, namely voice, video and gaming that may have very different transport requirements.

3.3.1. Voice over IP

3.3.1.1. Description

Voice over IP uses the real time transport protocol (RTP) over UDP transport. The transport requirements of VoIP have been defined and solutions implemented including when there are interactions with intermediate boxes including firewalls and NAT boxes.

3.3.1.2. Derived Requirements

In order for a transport service to provide VoIP services it requires minimal delay, in-order delivery and adequate error mitigation. And while UDP is the preferred protocol newer transports (such as Minion or others) could provide better performance.

3.3.2. IP based real-time video

3.3.2.1. Description

TV over the Internet Protocol (IPTV) requires multicast streaming as well as minimal caching, except for so called "trick plays" (rewind, pause and fast forward) over linear transmission. The now concluded reliable multicast working group (rmt) has defined a number of mechanisms to provide IPTV services using acknowledgements and forward error correction.

3.3.2.2. Derived Requirements

In order for a transport service to provide IPTV services it requires minimal delay, in-order delivery and very good error mitigation.

3.3.3. Gaming

Online games require very low delay and they use very small packets to ensure this (small packet are also features of VoIP). It is also beneficial to think that games could use tunneling and other multiplexing mechanisms to ensure minimal delay and synchronization between players.

3.3.3.1 Description

<to come>

3.3.3.2. Derived Requirements

<to come>

3.4. Photo/video Sharing

<to come>

3.5. Data centers/storage

There is a lot of work being performed in TCP for data centers (DCs) and as such DCs require certain characteristics from the transport layer.

3.5.1. Description

<reviewing DCTCP-slides][4] for inputs>

3.5.2 Derived Requirements

<to come>

3.6. Internet of Things (IoT)

While the Internet of Things has yes to define specific requirements for transport services a number of protocols to collect sensor data have been proposed, such a IBM Machine to Machine communication protocol call Message Queuing Telemetry Transport (MQTT) which covers part of the IoT spectrum of application.

<More to come>

3.7. Public Safety

In the past few years there has been a push for public safety systems to use Internet based protocols and private and public networks to ensure availability and resiliency.

3.7.1. Description

Public safety services are similar to the ones already defined above (browsing, caching, real time services, streaming etc.). However, the nature of these services requires strong security, high reliability and 'always on' nature.

3.7.2. Derived Requirements

In order for transport services to be used in public safety strong encryption like the one proposed by tcpcrypt [3] may be useful. In addition low delay is essential. Reliability the and always-on requirements can also be enhanced by mechanisms like FEC. Some requirements are shared with multipath TCP [8].

3.8. Information Centric Networks/Content Centric Networks

Information Centric (ICN) and Content Centric (CCN) Networking are emerging interest area in the Internet and define new mechanisms to connect to resources. The IRTF has chartered an ICN working group [6] that will investigate the issues like scalability and naming conventions to ensure deployability.

3.8.1. Description

An ICN does not require a different transport but will signal transport requirements in a different manner.

3.8.2. Derived Requirements

In the ICN/CCN work the access to a specific resource via the transport mechanism will be in the form of a publish/subscribe. It is for the transport to derive what this means in terms of performance.

<To be completed>

4. Security Considerations

None relating to the document. Specific use case security requirements to be detailed later.

5. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

6. Conclusion

This document presents a number of use cases for applications that could benefit from transport services other than the usual TCP and UDP protocols. While there is no underlying architecture to the document it is assumed that the transport services would be chosen based on specific requirements and using standard APIs. It is also assumed that APIs would not require major changes with potential changes applied to applications or middle-boxes.

7. References

7.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

7.2. Informative References

- [2] Mike Belshe, Robert Peon, "SPDY Protocol", [draft-mbelshe-httpbis-spdy-00](#), February 2012.
- [3] A. Bittau, Dan Boneh, M. Hamburg, H. Handley, D. Mazieres, Q. Slack, "Cryptographic protection of TCP Streams", [draft-bittau-tcp-crypt-03](#), September 3, 2012.
- [4] DTCP Slide deck:
<http://www.ietf.org/proceedings/80/slides/iccr-3.pdf>
- [5] Seppo Hatonen, Aki Nyrhinen, Lars Eggert, Stephen Strowes, Pasi Sarolahti and Markku Kojo. "An Experimental Study of Home Gateway Characteristics." Proc. ACM SIGCOMM Internet Measurement Conference (IMC), Melbourne, Australia, November 1-3, 2010.
- [6] Information Centric Network IRTF charter:
<http://irtf.org/icnrg>.

- [7] Jana Iyengar, Stuart Cheschire, Josh Graessley, "Minion - Service Model and Conceptual API", [draft-iyengar-minion-concept-02](#), October 21 2013 (work in progress).
- [8] Michael Scharf, Alan Ford, "Multipath TCP (MPTCP) Application Interface Considerations", [RFC 6897](#), March 2013.
- [9] QUIC Slide Deck at IETF88, <http://www.ietf.org/proceedings/88/slides/slides-88-tsvarea-10.pdf>
- [10] R. Stewart, ed. "Stream Control Transport Protocol", [RFC 4960](#), September 2007.
- [11] Michael Welzl, Florian Niederbacher, Stein Gjessing: "Beneficial Transparent Deployment of SCTP: the Missing Pieces", IEEE GlobeCom 2011, 5-9 February 2011, Houston, Texas.

<others to come>

8. Acknowledgments

The author would like to thank Michael Welzl, Jose Saldana, Lars Eggert and Martin Sustrik for use case suggestions and technical discussions prior to and during the preparation of this draft.

Copyright (c) 2014 IETF Trust and the persons identified as authors.

All rights reserved.

Authors' Addresses

Marie-Jose Montpetit

MIT

Email: mariejo@mit.edu

Igor Zhovnirovsky

QFactor Communications

Email: iez@qfactor.com

Note:

Changes between version 00 and version 01

Minor changes: Rewrote the Abstract, fixed the introduction, corrected spelling and some sentences, added a conclusion and a reference to SCTP.