Transport-Services Internet Draft Intended status: Informational Expires: June 5, 2014 M.J. Montpetit MIT I. Zhovnirovsky QFactor December 5, 2013

# Use Cases and Requirements draft-montpetit-transport-use-cases-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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/lid-abstracts.txt

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

This Internet-Draft will expire on June 5, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) 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.

Montpetit

Expires June 5, 2014

[Page 1]

Internet-Draft <u>draft-montpetit-transport-use-case</u>

December 2013

Abstract

This document describes some of the use cases that will allow choosing different transport services. It also presents their derived requirements and conditions of utilization, when available, to an application attached to an Internet subnetwork. The use cases assume that a transport "choice" is based on the requirements of the application. The fallback is considered to be TCP/UDP.

Table of Contents

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

Montpetit

Expires June 5, 2014

[Page 2]

#### Internet-Draft <u>draft-montpetit-transport-use-case</u>

December 2013

#### **1**. Introduction

Different transport services are now being developed in the Internet Community. A number of these are used to address performance requirements that standard transport such as the TCP and UDP protocol meet only partially especially over wireless. 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 responds to a need to define which transport service is best suited for an application and which API(s) would allow the use of those services. The APIs are considered to be native to the platform. The applications will request a connection by specifying a set of requirements. The will then API 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 middleware to work better without requiring a change to existing API and 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 enhanced transport services will be used in many of the cases.

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 <u>RFC-2119[1]</u>. 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 <u>RFC-2119</u> 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.

```
<u>3.1</u>. Web Surfing/Browsing
```

<Will add text>

Montpetit

Expires June 5, 2014

[Page 3]

# <u>3.1.1</u>. Faster Browsing via multi-streaming for HTTP prior to 2.0 <u>3.1.1.1</u>. 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 alleviated with QUIC [9], a UDP-based transport protocol that operates under SPDY [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.

# 3.1.2. 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 SCTP's multi-streaming. In [10], 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 fewer overheads. This requires hiding the use of SCTP (or any other multi-streamingcapable transport, e.g. Minion [7]) from applications, and it needs a signal 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.

Montpetit

Expires June 5, 2014

[Page 4]

# 3.1.3. Other HTTP acceleration mechanisms

<To be provided>

#### <u>3.2</u>. 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 and increase user experience.

# <u>3.2.1</u>. Adaptive Video Streaming

<To be provided>

### <u>3.2.2</u>. 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 bitrate 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 couple of seconds ahead of the play-out time. If the network's capacity would allow filling it much further ahead, this could have multiple benefits:

- if congestion happens later, the play-out would stay uninterrupted for a longer period or the risk of interruptions would be reduced further
- 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 prefetching 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 scavengerlike 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.

Montpetit

Expires June 5, 2014

[Page 5]

Internet-Draft <u>draft-montpetit-transport-use-case</u>

December 2013

#### <u>3.2.2.2</u>. 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 and must therefore also be possible on the fly.

# <u>3.3</u>. Real-time Communications

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

#### <u>3.3.2</u>. IP based real-time video

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

<reviewing rmt documentation and more inputs to come>

#### 3.3.2.1. Derived Requirements

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

Montpetit

Expires June 5, 2014

[Page 6]

#### <u>3.3.3</u>. Gaming

Online games require very low delay and because of that use small packets (and in this share requirements with 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>

# <u>3.3.3.2</u>. Derived Requirements

<to come>

#### <u>3.4</u>. Photo/video Sharing <to come>

#### **<u>3.5</u>**. 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] >

# <u>3.6</u>. 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.

Montpetit

Expires June 5, 2014

[Page 7]

### 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.

#### <u>3.7.2</u>. 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.

#### <u>3.8.1</u>. Description

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

#### <u>**3.8.2</u>**. Derived Requirements</u>

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>

<u>4</u>. Security Considerations None relating to the document. Specific use case security requirements to be detailed later.

Montpetit

Expires June 5, 2014

[Page 8]

# 5. IANA Considerations

**This document makes no request of IANA.** Note to RFC Editor: this section may be removed on publication as an RFC.

# <u>6</u>. Conclusion

<Will add a conclusion>

#### 7. References

#### 7.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, March 1997.

# <u>7.2</u>. Informative References

- [2] Mike Belshe, Robert Peon, "SPDY Protocol", <u>draft-mbelshe-httpbis-spdy-00</u>, February 2012.
- [3] A. Bittau, Dan Boneh, M. Hamburg, H. Handley, D. Mazieres, Q. Slack, "Cryptographic protection of TCP Streams", <u>draft-bittau-tcp-crypt-03</u>, September 3, 2012.
- [4] DTCP Slide deck: http://www.ietf.org/proceedings/80/slides/iccrg-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: <u>http://irtf.org/icnrg</u>.
- [7] Jana Iyengar, Stuart Cheschire, Josh Graessley, "Minion -Service Model and Conceptual API", <u>draft-iyengar-minion-</u> <u>concept-02</u>, October 21 2013 (work in progress).
- [8] Michael Sharf, Alan Ford, "Multipath TCP (MPTCP) Application Interface Considerations", <u>RFC 6897</u>, March 2013.

Montpetit Expires June 5, 2014

[Page 9]

- [9] QUIC Slide Deck at IETF88, <u>http://www.ietf.org/proceedings/88/slides/slides-88-tsvarea-10.pdf</u>
- [10] Michael Welzl, Florian Niederbacher, Stein Gjessing: "Beneficial Transparent Deployment of SCTP: the Missing Pieces", IEEE GlobeCom 2011, 5-9 December 2011, Houston, Texas.

<others to come>

# 8. Acknowledgments

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

Copyright (c) 2013 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@gfactor.com

Montpetit

Expires June 5, 2014

[Page 10]