TAPS
Internet-Draft
Intended status: Informational
Expires: September 14, 2017

S. Gjessing
M. Welzl
University of Oslo

March 13, 2017

A Minimal Set of Transport Services for TAPS Systems
draft-gjessing-taps-minset-04

Abstract

This draft recommends a minimal set of IETF Transport Services offered by end systems supporting TAPS, and gives guidance on choosing among the available mechanisms and protocols. It is based on the set of transport features given in the TAPS document draft-ietf-taps-transports-usage-03.

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 http://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 September 14, 2017.

Table of Contents

1. Introduction

An application has an intended usage and demands for transport services, and the task of any system that implements TAPS is to offer these services to its applications, i.e. the applications running on top of TAPS, without binding them to a particular transport protocol. Currently, the set of transport services that most applications use is based on TCP and UDP; this limits the ability for the network stack to make use of features of other protocols. For example, if a protocol supports out-of-order message delivery but applications always assume that the network provides an ordered bytestream, then the network stack can never utilize out-of-order message delivery: doing so would break a fundamental assumption of the application.

By exposing the transport services of multiple transport protocols, a TAPS system can make it possible to use these services without having to statically bind an application to a specific transport protocol. The first step towards the design of such a system was taken by [RFC8095], which surveys a large number of transports, and [TAPS2], which identifies the specific transport features that are exposed to applications by the protocols TCP, MPTCP, UDP(-Lite) and SCTP as well as the LEDBAT congestion control mechanism. The present draft is based on these documents and follows the same terminology (also listed below).

The number of transport features of current IETF transports is large, and exposing all of them has a number of disadvantages: generally, the more functionality is exposed, the less freedom a TAPS system has to automate usage of the various functions of its available set of transport protocols. Some functions only exist in one particular protocol, and if an application would use them, this would statically tie the application to this protocol, counteracting the purpose of a TAPS system. Also, if the number of exposed features is exceedingly large, a TAPS system might become very hard to use for an application programmer. Taking [TAPS2] as a basis, this document therefore develops a minimal set of transport features, removing the ones that

could be harmful to the purpose of a TAPS system but keeping the ones that must be retained for applications to benefit from useful transport functionality.

Applications use a wide variety of APIs today. The point of this document is to identify transport features that must be reflected in *all* network APIs in order for the underlying functionality to become usable everywhere. For example, it does not help an application that talks to a middleware if only the Berkeley Sockets API is extended to offer "unordered message delivery". Instead, the middleware would have to expose the "unordered message delivery" transport feature to its applications (alternatively, there may be

interesting ways for certain types of middleware to try to use some of the transport features that we describe here without exposing them to applications, based on knowledge about the applications -- but this is not the general case). In most situations, in the interest of being as flexible and efficient as possible, the best choice will be for a middleware or library to expose all of the transport features that are recommended as a "minimal set" here. As an example considering only TCP and UDP, a middleware or library that only exposes TCP's reliable bytestream cannot make use of UDP (unless it implements extra functionality on top of UDP) -- doing so could break a fundamental assumption that applications make about the data they send and receive.

This document approaches the construction of a minimal set of transport features in the following way:

1. Categorization: the superset of transport features from [TAPS2] is presented, and transport features are categorized for later reduction.
2. Reduction: a shorter list of transport features is derived from the categorization in the first step. This removes all transport features that do not require application-specific knowledge or cannot be implemented with TCP.
3. Discussion: the resulting list shows a number of peculiarities that are discussed, to provide a basis for constructing the minimal set.
4. Construction: Based on the reduced set and the discussion of the transport features therein, a minimal set is constructed.

2. Terminology

The following terms are used throughout this document, and in subsequent documents produced by TAPS that describe the composition and decomposition of transport services.

Transport Feature: a specific end-to-end feature that the transport layer provides to an application. Examples include confidentiality, reliable delivery, ordered delivery, message-versus-stream orientation, etc.

Transport Service: a set of Transport Features, without an association to any given framing protocol, which provides a complete service to an application.

Transport Protocol: an implementation that provides one or more different transport services using a specific framing and header format on the wire.

Transport Service Instance: an arrangement of transport protocols with a selected set of features and configuration parameters that implements a single transport service, e.g., a protocol stack (RTP over UDP).

Application: an entity that uses the transport layer for end-to-end delivery data across the network (this may also be an upper layer protocol or tunnel encapsulation).

Application-specific knowledge: knowledge that only applications have.

Endpoint: an entity that communicates with one or more other endpoints using a transport protocol.

Connection: shared state of two or more endpoints that persists across messages that are transmitted between these endpoints.

Socket: the combination of a destination IP address and a destination port number.

3. Step 1: Categorization -- The Superset of Transport Features

Following [TAPS2], we divide the transport features into two main groups as follows:

1. CONNECTION related transport features

- ESTABLISHMENT
- AVAILABILITY
- MAINTENANCE
- TERMINATION

2. DATA Transfer Related transport features
   - Sending Data
   - Receiving Data
   - Errors

Because QoS is out of scope of TAPS, this document assumes a "best effort" service model [RFC5290], [RFC7305]. Applications using a TAPS system can therefore not make any assumptions about e.g. the time it will take to send a message. We also assume that TAPS applications have no specific requirements that need knowledge about the network, e.g. regarding the choice of network interface or the end-to-end path. Even with these assumptions, there are certain requirements that are strictly kept by transport protocols today, and these must also be kept by a TAPS system. Some of these requirements relate to transport features that we call "Functional".

Functional transport features provide functionality that cannot be used without the application knowing about them, or else they violate assumptions that might cause the application to fail. For example, unordered message delivery is a functional transport feature: it cannot be used without the application knowing about it because the application's assumption could be that messages arrive in order.

Failure includes any change of the application behavior that is not performance oriented, e.g. security.

"Change DSCP" and "Disable Nagle algorithm" are examples of transport features that we call "Optimizing": if a TAPS system autonomously decides to enable or disable them, an application will not fail, but a TAPS system may be able to communicate more efficiently if the application is in control of this optimizing transport feature. These transport features require application-specific knowledge (e.g., about delay/bandwidth requirements or the length of future data blocks that are to be transmitted).

The transport features of IETF transport protocols that do not require application-specific knowledge and could therefore be transparently utilized by a TAPS system are called "Automatable".

Finally, some transport features are aggregated and/or slightly changed in the TAPS API. These transport features are marked as "ADDED". The corresponding transport features are automatable, and they are listed immediately below the "ADDED" transport feature.

In this description, transport services are presented following the nomenclature "CATEGORY.[SUBCATEGORY].SERVICENAME.PROTOCOL", equivalent to "pass 2" in [TAPS2]. The PROTOCOL name "UDP(-Lite)" is used when transport features are equivalent for UDP and UDP-Lite; the PROTOCOL name "TCP" refers to both TCP and MPTCP. We also sketch how some of the TAPS transport services can be implemented. For all transport features that are categorized as "functional" or "optimizing", and for which no matching TCP primitive exists in "pass 2" of [TAPS2], a brief discussion on how to fall back to TCP is included.

We designate some transport features as "automatable" on the basis of a broader decision that affects multiple transport features:

- Most transport features that are related to multi-streaming were designated as "automatable". This was done because the decision on whether to use multi-streaming or not does not depend on application-specific knowledge. This means that a connection that is exhibited to an application could be implemented by using a single stream of an SCTP association instead of mapping it to a complete SCTP association or TCP connection. This could be achieved by using more than one stream when an SCTP association is first established (CONNECT.SCTP parameter "outbound stream count"), maintaining an internal stream number, and using this stream number when sending data (SEND.SCTP parameter "stream number"). Closing or aborting a connection could then simply free the stream number for future use. This is discussed further in Section 5.2.

- All transport features that are related to using multiple paths or the choice of the network interface were designated as "automatable". Choosing a path or an interface does not depend on application-specific knowledge. For example, "Listen" could always listen on all available interfaces and "Connect" could use the default interface for the destination IP address.

3.1. CONNECTION Related Transport Features

ESTABLISHMENT:

- Connect
  Protocols: TCP, SCTP, UDP(-Lite)
  Functional because the notion of a connection is often reflected in applications as an expectation to be able to communicate after a "Connect" succeeded, with a communication sequence relating to this transport feature that is defined by the application protocol.
  Implementation: via CONNECT.TCP, CONNECT.SCTP or CONNECT.UDP(-Lite).

- Specify which IP Options must always be used
  Protocols: TCP
  Automatable because IP Options relate to knowledge about the network, not the application.

- Request multiple streams
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Limit the number of inbound streams
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Specify number of attempts and/or timeout for the first establishment message
  Protocols: TCP, SCTP
  Functional because this is closely related to potentially assumed reliable data delivery for data that is sent before or during connection establishment.
  Implementation: Using a parameter of CONNECT.TCP and CONNECT.SCTP.

- Obtain multiple sockets
  Protocols: SCTP
  Automatable because the usage of multiple paths to communicate to

the same end host relates to knowledge about the network, not the

application.

- Disable MPTCP
  Protocols: MPTCP
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.
  Implementation: via a boolean parameter in CONNECT.MPTCP.
  Fall-back to TCP: Do nothing.

- Specify which chunk types must always be authenticated
  Protocols: SCTP
  Functional because this has a direct influence on security.
  Implementation: via a parameter in CONNECT.SCTP.
  Fall-back to TCP: TBD: this relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered by [TAPS2].

- Indicate (and/or obtain upon completion) an Adaptation Layer via an adaptation code point
  Protocols: SCTP
  Functional because it allows to send extra data for the sake of identifying an adaptation layer, which by itself is application-specific.
  Implementation: via a parameter in CONNECT.SCTP.
  Fall-back to TCP: not possible.

- Request to negotiate interleaving of user messages
  Protocols: SCTP
  Automatable because it requires using multiple streams, but requesting multiple streams in the CONNECTION.ESTABLISHMENT category is automatable.
  Implementation: via a parameter in CONNECT.SCTP.

- Hand over a message to transfer (possibly multiple times) before connection establishment
  Protocols: TCP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via a parameter in CONNECT.TCP.

- Hand over a message to transfer during connection establishment
  Protocols: SCTP
  Functional because this can only work if the message is limited in size, making it closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via a parameter in CONNECT.SCTP.

- Enable UDP encapsulation with a specified remote UDP port number
  Protocols: SCTP
  Automatable because UDP encapsulation relates to knowledge about the network, not the application.

AVAILABILITY:

- Listen
  Protocols: TCP, SCTP, UDP(-Lite)
  Functional because the notion of accepting connection requests is often reflected in applications as an expectation to be able to communicate after a "Listen" succeeded, with a communication sequence relating to this transport feature that is defined by the application protocol.
  ADDED. This differs from the 3 automatable transport features below in that it leaves the choice of interfaces for listening open.
  Implementation: by listening on all interfaces via LISTEN.TCP (not providing a local IP address) or LISTEN.SCTP (providing SCTP port number / address pairs for all local IP addresses).

- Listen, 1 specified local interface
  Protocols: TCP, SCTP, UDP(-Lite)
  Automatable because decisions about local interfaces relate to knowledge about the network and the Operating System, not the application.

- Listen, N specified local interfaces
  Protocols: SCTP, UDP(-Lite)
  Automatable because decisions about local interfaces relate to knowledge about the network and the Operating System, not the application.

- Listen, all local interfaces
  Protocols: TCP, SCTP, UDP(-Lite)
  Automatable because decisions about local interfaces relate to knowledge about the network and the Operating System, not the application.

- Specify which IP Options must always be used
  Protocols: TCP
  Automatable because IP Options relate to knowledge about the network, not the application.

- Disable MPTCP
  Protocols: MPTCP

Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the

application.

- Specify which chunk types must always be authenticated
  Protocols: SCTP
  Functional because this has a direct influence on security.
  Implementation: via a parameter in CONNECT.SCTP.
  Fall-back to TCP: TBD: this relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered by [TAPS2].

- Obtain requested number of streams
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Limit the number of inbound streams
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Indicate (and/or obtain upon completion) an Adaptation Layer via an adaptation code point
  Protocols: SCTP
  Functional because it allows to send extra data for the sake of identifying an adaptation layer, which by itself is application-specific.
  Implementation: via a parameter in LISTEN.SCTP.
  Fall-back to TCP: not possible.

- Request to negotiate interleaving of user messages
  Protocols: SCTP
  Automatable because it requires using multiple streams, but requesting multiple streams in the CONNECTION.ESTABLISHMENT category is automatable.
  Implementation: via a parameter in LISTEN.SCTP.

MAINTENANCE:

- Change timeout for aborting connection (using retransmit limit or time value)
  Protocols: TCP, SCTP
  Functional because this is closely related to potentially assumed reliable data delivery.
  Implementation: via CHANGE-TIMEOUT.TCP or CHANGE-TIMEOUT.SCTP.

- Suggest timeout to the peer
  Protocols: TCP
  Functional because this is closely related to potentially assumed reliable data delivery.
  Implementation: via CHANGE-TIMEOUT.TCP.

- Disable Nagle algorithm
  Protocols: TCP, SCTP
  Optimizing because this decision depends on knowledge about the size of future data blocks and the delay between them.
  Implementation: via DISABLE-NAGLE.TCP and DISABLE-NAGLE.SCTP.

- Request an immediate heartbeat, returning success/failure
  Protocols: SCTP
  Automatable because this informs about network-specific knowledge.

- Notification of Excessive Retransmissions (early warning below abortion threshold)
  Protocols: TCP
  Optimizing because it is an early warning to the application, informing it of an impending functional event.
  Implementation: via ERROR.TCP.

- Add path
  Protocols: MPTCP, SCTP
  MPTCP Parameters: source-IP; source-Port; destination-IP; destination-Port
  SCTP Parameters: local IP address
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.

- Remove path
  Protocols: MPTCP, SCTP
  MPTCP Parameters: source-IP; source-Port; destination-IP; destination-Port
  SCTP Parameters: local IP address
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.

- Set primary path
  Protocols: SCTP
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.

- Suggest primary path to the peer
  Protocols: SCTP
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.

- Configure Path Switchover
  Protocols: SCTP
  Automatable because the usage of multiple paths to communicate to the same end host relates to knowledge about the network, not the application.

- Obtain status (query or notification)
  Protocols: SCTP, MPTCP
  SCTP parameters: association connection state; destination transport address list; destination transport address reachability states; current local and peer receiver window size; current local congestion window sizes; number of unacknowledged DATA chunks; number of DATA chunks pending receipt; primary path; most recent SRTT on primary path; RTO on primary path; SRTT and RTO on other destination addresses; MTU per path; interleaving supported yes/no
  MPTCP parameters: subflow-list (identified by source-IP; source-Port; destination-IP; destination-Port)
  Automatable because these parameters relate to knowledge about the network, not the application.

- Specify DSCP field
  Protocols: TCP, SCTP, UDP(-Lite)
  Optimizing because choosing a suitable DSCP value requires application-specific knowledge.
  Implementation: via SET_DSCP.TCP / SET_DSCP.SCTP / SET_DSCP.UDP(-Lite)

- Notification of ICMP error message arrival
  Protocols: TCP, UDP(-Lite)
  Optimizing because these messages can inform about success or failure of functional transport features (e.g., host unreachable relates to "Connect")
  Implementation: via ERROR.TCP or ERROR.UDP(-Lite).

- Obtain information about interleaving support
  Protocols: SCTP
  Automatable because it requires using multiple streams, but requesting multiple streams in the CONNECTION.ESTABLISHMENT category is automatable.
  Implementation: via a parameter in GETINTERL.SCTP.

- Change authentication parameters
  Protocols: SCTP
  Functional because this has a direct influence on security.
  Implementation: via SETAUTH.SCTP.
  Fall-back to TCP: TBD: this relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered by [TAPS2].

- Obtain authentication information
  Protocols: SCTP
  Functional because authentication decisions may have been made by the peer, and this has an influence on the necessary application-level measures to provide a certain level of security.
  Implementation: via GETAUTH.SCTP.
  Fall-back to TCP: TBD: this relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered by [TAPS2].

- Reset Stream
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Notification of Stream Reset
  Protocols: STCP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Reset Association
  Protocols: SCTP
  Functional because it affects "Obtain a message delivery number", which is functional.
  Implementation: via RESETASSOC.SCTP.
  Fall-back to TCP: not possible.

- Notification of Association Reset
  Protocols: STCP
  Functional because it affects "Obtain a message delivery number", which is functional.
  Implementation: via RESETASSOC-EVENT.SCTP.
  Fall-back to TCP: not possible.

- Add Streams
  Protocols: SCTP
  Automatable because using multi-streaming does not require application-specific knowledge.

  Implementation: see Section 5.2.

- Notification of Added Stream
  Protocols: STCP
  Automatable because using multi-streaming does not require application-specific knowledge.
  Implementation: see Section 5.2.

- Choose a scheduler to operate between streams of an association
  Protocols: SCTP
  Optimizing because the scheduling decision requires application-

specific knowledge. However, if a TAPS system would not use this, or wrongly configure it on its own, this would only affect the performance of data transfers; the outcome would still be correct within the "best effort" service model.
Implementation: using SETSTREAMSCHEDULER.SCTP.
Fall-back to TCP: do nothing.

- Configure priority or weight for a scheduler
  Protocols: SCTP
  Optimizing because the priority or weight requires application-specific knowledge. However, if a TAPS system would not use this, or wrongly configure it on its own, this would only affect the performance of data transfers; the outcome would still be correct within the "best effort" service model.
  Implementation: using CONFIGURESTREAMSCHEDULER.SCTP.
  Fall-back to TCP: do nothing.

- Configure send buffer size
  Protocols: SCTP
  Automatable because this decision relates to knowledge about the network and the Operating System, not the application (see also the discussion in Section 5.4).

- Configure receive buffer (and rwnd) size
  Protocols: SCTP
  Automatable because this decision relates to knowledge about the network and the Operating System, not the application.

- Configure message fragmentation
  Protocols: SCTP
  Automatable because fragmentation relates to knowledge about the network and the Operating System, not the application.
  Implementation: by always enabling it with CONFIG_FRAGMENTATION.SCTP and auto-setting the fragmentation size based on network or Operating System conditions.

- Configure PMTUD
  Protocols: SCTP
  Automatable because Path MTU Discovery relates to knowledge about the network, not the application.

- Configure delayed SACK timer
  Protocols: SCTP
  Automatable because the receiver-side decision to delay sending SACKs relates to knowledge about the network, not the application (it can be relevant for a sending application to request not to delay the SACK of a message, but this is a different transport feature).

- Set Cookie life value
  Protocols: SCTP
  Functional because it relates to security (possibly weakened by keeping a cookie very long) versus the time between connection establishment attempts. Knowledge about both issues can be application-specific.
  Fall-back to TCP: the closest TCP functionality is the cookie in TCP Fast Open; for this, [RFC7413] states that the server "can expire the cookie at any time to enhance security" and section 4.1.2 describes an example implementation where updating the key on the server side causes the cookie to expire; however, this is different from this transport feature because SCTP's cookie life value is set on the client side, not the server side. The TCP client has no control of this value. Thus, the recommended fall-back implementation is to do nothing.

- Set maximum burst
  Protocols: SCTP
  Automatable because it relates to knowledge about the network, not the application.

- Configure size where messages are broken up for partial delivery
  Protocols: SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Fall-back to TCP: do nothing. Since TCP does not deliver messages, partial or not, this will have no effect on TCP.

- Disable checksum when sending
  Protocols: UDP
  Functional because application-specific knowledge is necessary to decide whether it can be acceptable to lose data integrity.
  Implementation: via CHECKSUM.UDP.
  Fall-back to TCP: do nothing.

- Disable checksum requirement when receiving
  Protocols: UDP
  Functional because application-specific knowledge is necessary to decide whether it can be acceptable to lose data integrity.
  Implementation: via CHECKSUM_REQUIRED.UDP.
  Fall-back to TCP: do nothing.

- Specify checksum coverage used by the sender
  Protocols: UDP-Lite
  Functional because application-specific knowledge is necessary to decide for which parts of the data it can be acceptable to lose data integrity.
  Implementation: via SET_CHECKSUM_COVERAGE.UDP-Lite.
  Fall-back to TCP: do nothing.

- Specify minimum checksum coverage required by receiver
  Protocols: UDP-Lite
  Functional because application-specific knowledge is necessary to decide for which parts of the data it can be acceptable to lose data integrity.
  Implementation: via SET_MIN_CHECKSUM_COVERAGE.UDP-Lite.
  Fall-back to TCP: do nothing.

- Specify DF field
  Protocols: UDP(-Lite)
  Optimizing because the DF field can be used to carry out Path MTU Discovery, which can lead an application to choose message sizes that can be transmitted more efficiently.
  Implementation: via MAINTENANCE.SET_DF.UDP(-Lite) and SEND_FAILURE.UDP(-Lite).
  Fall-back to TCP: do nothing. With TCP the sender is not in control of transport message sizes, making this functionality irrelevant.

- Specify TTL/Hop count field
  Protocols: UDP(-Lite)
  Automatable because a TAPS system can use a large enough system default to avoid communication failures. Allowing an application to configure it differently can produce notifications of ICMP error message arrivals that yield information which only relates to knowledge about the network, not the application.

- Obtain TTL/Hop count field
  Protocols: UDP(-Lite)
  Automatable because the TTL/Hop count field relates to knowledge about the network, not the application.

- Specify ECN field
  Protocols: UDP(-Lite)
  Automatable because the ECN field relates to knowledge about the network, not the application.

- Obtain ECN field
  Protocols: UDP(-Lite)
  Automatable because the ECN field relates to knowledge about the network, not the application.

- Specify IP Options
  Protocols: UDP(-Lite)
  Automatable because IP Options relate to knowledge about the network, not the application.

- Obtain IP Options
  Protocols: UDP(-Lite)
  Automatable because IP Options relate to knowledge about the network, not the application.

- Enable and configure a "Low Extra Delay Background Transfer"
  Protocols: A protocol implementing the LEDBAT congestion control mechanism
  Optimizing because whether this service is appropriate or not depends on application-specific knowledge. However, wrongly using this will only affect the speed of data transfers (albeit including other transfers that may compete with the TAPS transfer in the network), so it is still correct within the "best effort" service model.
  Implementation: via CONFIGURE.LEDBAT and/or SET_DSCP.TCP / SET_DSCP.SCTP / SET_DSCP.UDP(-Lite) [LBE-draft].
  Fall-back to TCP: do nothing.

TERMINATION:

- Close after reliably delivering all remaining data, causing an event informing the application on the other side
  Protocols: TCP, SCTP
  Functional because the notion of a connection is often reflected in applications as an expectation to have all outstanding data

delivered and no longer be able to communicate after a "Close" succeeded, with a communication sequence relating to this transport feature that is defined by the application protocol.
Implementation: via CLOSE.TCP and CLOSE.SCTP.

- Abort without delivering remaining data, causing an event informing the application on the other side
  Protocols: TCP, SCTP

Functional because the notion of a connection is often reflected in applications as an expectation to potentially not have all outstanding data delivered and no longer be able to communicate after an "Abort" succeeded, with a communication sequence relating to this transport feature that is defined by the application protocol.
Implementation: via ABORT.TCP and ABORT.SCTP.

- Timeout event when data could not be delivered for too long
  Protocols: TCP, SCTP
  Functional because this notifies that potentially assumed reliable data delivery is no longer provided.
  Implementation: via TIMEOUT.TCP and TIMEOUT.SCTP.

3.2. DATA Transfer Related Transport Features

3.2.1. Sending Data

- Reliably transfer data, with congestion control
  Protocols: TCP, SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via SEND.TCP and SEND.SCTP.

- Reliably transfer a message, with congestion control
  Protocols: SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via SEND.SCTP and SEND.TCP. With SEND.TCP, messages will not be identifiable by the receiver. Inform the application of the result.

- Unreliably transfer a message
  Protocols: SCTP, UDP(-Lite)
  Optimizing because only applications know about the time criticality of their communication, and reliably transfering a message is never incorrect for the receiver of a potentially unreliable data transfer, it is just slower.
  ADDED. This differs from the 2 automatable transport features below in that it leaves the choice of congestion control open.
  Implementation: via SEND.SCTP or SEND.UDP or SEND.TCP. With SEND.TCP, messages will not be identifiable by the receiver.
  Inform the application of the result.

- Unreliably transfer a message, with congestion control
  Protocols: SCTP
  Automatable because congestion control relates to knowledge about

  the network, not the application.

- Unreliably transfer a message, without congestion control
  Protocols: UDP(-Lite)
  Automatable because congestion control relates to knowledge about the network, not the application.

- Configurable Message Reliability
  Protocols: SCTP
  Optimizing because only applications know about the time criticality of their communication, and reliably transfering a message is never incorrect for the receiver of a potentially unreliable data transfer, it is just slower.
  Implementation: via SEND.SCTP.
  Fall-back to TCP: By using SEND.TCP and ignoring this configuration: based on the assumption of the best-effort service model, unnecessarily delivering data does not violate application expectations. Moreover, it is not possible to associate the requested reliability to a "message" in TCP anyway.

- Choice of stream
  Protocols: SCTP
  Automatable because it requires using multiple streams, but requesting multiple streams in the CONNECTION.ESTABLISHMENT category is automatable. Implementation: see Section 5.2.

- Choice of path (destination address)
  Protocols: SCTP
  Automatable because it requires using multiple sockets, but obtaining multiple sockets in the CONNECTION.ESTABLISHMENT category is automatable.

- Choice between unordered (potentially faster) or ordered delivery of messages
  Protocols: SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via SEND.SCTP.
  Fall-back to TCP: By using SEND.TCP and always sending data ordered: based on the assumption of the best-effort service model, ordered delivery may just be slower and does not violate application expectations. Moreover, it is not possible to associate the requested delivery order to a "message" in TCP anyway.

- Request not to bundle messages
  Protocols: SCTP
  Optimizing because this decision depends on knowledge about the

  size of future data blocks and the delay between them.
  Implementation: via SEND.SCTP.
  Fall-back to TCP: By using SEND.TCP and DISABLE-NAGLE.TCP to disable the Nagle algorithm when the request is made and enable it again when the request is no longer made.

- Specifying a "payload protocol-id" (handed over as such by the receiver)
  Protocols: SCTP
  Functional because it allows to send extra application data with every message, for the sake of identification of data, which by itself is application-specific.
  Implementation: SEND.SCTP.
  Fall-back to TCP: not possible.

- Specifying a key id to be used to authenticate a message
  Protocols: SCTP
  Functional because this has a direct influence on security.
  Implementation: via a parameter in SEND.SCTP.

Fall-back to TCP: TBD: this relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered by [TAPS2].

- Request not to delay the acknowledgement (SACK) of a message
  Protocols: SCTP
  Optimizing because only an application knows for which message it wants to quickly be informed about success / failure of its delivery.
  Fall-back to TCP: do nothing.

3.2.2. Receiving Data

- Receive data (with no message delineation)
  Protocols: TCP
  Functional because a TAPS system must be able to send and receive data.
  Implementation: via RECEIVE.TCP

- Receive a message
  Protocols: SCTP, UDP(-Lite)
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via RECEIVE.SCTP and RECEIVE.UDP(-Lite).
  Fall-back to TCP: not possible.

- Choice of stream to receive from
  Protocols: SCTP
  Automatable because it requires using multiple streams, but requesting multiple streams in the CONNECTION.ESTABLISHMENT category is automatable.
  Implementation: see Section 5.2.

- Information about partial message arrival
  Protocols: SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via RECEIVE.SCTP.

Fall-back to TCP: do nothing: this information is not available with TCP.

- Obtain a message delivery number
  Protocols: SCTP
  Functional because this number can let applications detect and, if desired, correct reordering. Whether messages are in the correct order or not is closely tied to properties of the data that an application sends or expects to receive.
  Implementation: via RECEIVE.SCTP.
  Fall-back to TCP: not possible.

3.2.3. Errors

This section describes sending failures that are associated with a specific call to in the "Sending Data" category (Section 3.2.1).

- Notification of send failures
  Protocols: SCTP, UDP(-Lite)
  Functional because this notifies that potentially assumed reliable data delivery is no longer provided.
  ADDED. This differs from the 2 automatable transport features below in that it does not distinugish between unsent and unacknowledged messages.
  Implementation: via SENDFAILURE-EVENT.SCTP and SEND_FAILURE.UDP(-Lite).
  Fall-back to TCP: do nothing: this notification is not available and will therefore not occur with TCP.

- Notification of an unsent (part of a) message
  Protocols: SCTP, UDP(-Lite)
  Automatable because the distinction between unsent and unacknowledged is network-specific.

- Notification of an unacknowledged (part of a) message
  Protocols: SCTP
  Automatable because the distinction between unsent and unacknowledged is network-specific.

- Notification that the stack has no more user data to send
  Protocols: SCTP
  Optimizing because reacting to this notification requires the application to be involved, and ensuring that the stack does not run dry of data (for too long) can improve performance.
  Fall-back to TCP: do nothing. See also the discussion in Section 5.4.

- Notification to a receiver that a partial message delivery has been aborted
  Protocols: SCTP
  Functional because this is closely tied to properties of the data that an application sends or expects to receive.
  Fall-back to TCP: do nothing. This notification is not available and will therefore not occur with TCP.

4. Step 2: Reduction -- The Reduced Set of Transport Features

By hiding automatable transport features from the application, a TAPS system can gain opportunities to automate the usage of network-related functionality. This can facilitate using the TAPS system for the application programmer and it allows for optimizations that may not be possible for an application. For instance, system-wide configurations regarding the usage of multiple interfaces can better be exploited if the choice of the interface is not entirely up to the application. Therefore, since they are not strictly necessary to expose in a TAPS system, we do not include automatable transport features in the reduced set of transport features. This leaves us with only the transport features that are either optimizing or functional.

A TAPS system should be able to fall back to TCP or UDP if alternative transport protocols are found not to work. Here we only consider falling back to TCP. For some transport features, it was identified that no fall-back to TCP is possible. This eliminates the possibility to use TCP whenever an application makes use of one of these transport features. Thus, we only keep the functional and optimizing transport features for which a fall-back to TCP is possible in our reduced set. "Reset Association" and "Notification of Association Reset" are only functional because of their relationship to "Obtain a message delivery number", which is

functional. Because "Obtain a message delivery number" does not have a fall-back to TCP, none of these three transport features are included in the reduced set.

4.1. CONNECTION Related Transport Features

ESTABLISHMENT:

- Connect
- Specify number of attempts and/or timeout for the first establishment message
- Specify which chunk types must always be authenticated
- Hand over a message to transfer (possibly multiple times) before connection establishment
- Hand over a message to transfer during connection establishment

AVAILABILITY:

- Listen
- Specify which chunk types must always be authenticated

MAINTENANCE:

- Change timeout for aborting connection (using retransmit limit or time value)
- Suggest timeout to the peer
- Disable Nagle algorithm
- Notification of Excessive Retransmissions (early warning below abortion threshold)
- Specify DSCP field
- Notification of ICMP error message arrival
- Change authentication parameters
- Obtain authentication information
- Choose a scheduler to operate between streams of an association
- Configure priority or weight for a scheduler
- Set Cookie life value
- Configure size where messages are broken up for partial delivery
- Disable checksum when sending
- Disable checksum requirement when receiving
- Specify checksum coverage used by the sender
- Specify minimum checksum coverage required by receiver
- Specify DF field
- Enable and configure a "Low Extra Delay Background Transfer"

TERMINATION:

- Close after reliably delivering all remaining data, causing an event informing the application on the other side
- Abort without delivering remaining data, causing an event informing the application on the other side

- Timeout event when data could not be delivered for too long

4.2. DATA Transfer Related Transport Features

4.2.1. Sending Data

- Reliably transfer data, with congestion control
- Reliably transfer a message, with congestion control
- Unreliably transfer a message
- Configurable Message Reliability
- Choice between unordered (potentially faster) or ordered delivery of messages
- Request not to bundle messages
- Specifying a key id to be used to authenticate a message
- Request not to delay the acknowledgement (SACK) of a message

4.2.2. Receiving Data

- Receive data (with no message delineation)
- Information about partial message arrival

4.2.3. Errors

This section describes sending failures that are associated with a specific call to in the "Sending Data" category (Section 3.2.1).

- Notification of send failures
- Notification that the stack has no more user data to send
- Notification to a receiver that a partial message delivery has been aborted

5. Step 3: Discussion

The reduced set in the previous section exhibits a number of peculiarities, which we will discuss in the following.

5.1. Sending Messages, Receiving Bytes

There are several transport features related to sending, but only a single transport feature related to receiving: "Receive data (with no message delineation)" (and, strangely, "information about partial

message arrival"). Notably, the transport feature "Receive a message" is also the only non-automatable transport feature of UDP(-Lite) that had to be removed because no fall-back to TCP is possible. It is also represents the only way that UDP(-Lite) applications can receive data today.

For the transport to operate on messages, it only needs be informed about them as they are handed over by a sending application; on the receiver side, receiving a message only differs from receiving a bytestream in that the application is told where messages begin and end in the former case but not in the latter. The receiving application can still operate on these messages as long as it does not rely on the transport layer to inform it about message boundaries.

For example, if an application requests to transfer fixed-size messages of 100 bytes with partial reliability, this needs the receiving application to be prepared to accept data in chunks of 100 bytes. If, then, some of these 100 byte messages are missing (e.g., if SCTP with Configurable Reliability is used), this is the expected application behavior. With TCP, no messages would be missing, but this is also correct for the application, and possible retransmission delay is acceptable within the best effort service model. Still, the receiving application would separate the byte stream into 100-byte chunks.

Note that this usage of messages does not require all messages to be equal in size. Many application protocols use some form of Type-Length-Value (TLV) encoding, e.g. by defining a header including length fields; another alternative is the use of byte stuffing methods such as COBS [COBS]. If an application needs message numbers, e.g. to restore the correct sequence of messages, these must also be encoded by the application itself, as the sequence number related transport features of SCTP are no longer provided (in the interest of enabling a fall-back to TCP).

For the implementation of a TAPS system, this has the following consequences:

- Because the receiver-side transport leaves it up to the application to delineate messages, messages must always remain intact as they are handed over by the transport receiver. Data

can be handed over at any time as they arrive, but the byte stream must never "skip ahead" to the beginning of the next message.

- With SCTP, a "partial flag" informs a receiving application that a message is incomplete. Then, the next receive calls will only deliver remaining parts of the same message (i.e., no messages or partial messages will arrive on other streams until the message is complete) (see Section 8.1.20 in [RFC6458]). This can facilitate the implementation of the receiver buffer in the receiving application, but then such an application does not support message interleaving (which is required by stream schedulers). However, receiving a byte stream from multiple SCTP streams requires a per-stream receiver buffer anyway, so this potential benefit is lost and the "partial flag" (the transport feature "Information about

partial message arrival") becomes unnecessary for a TAPS system. With it, the transport features "Configure size where messages are broken up for partial delivery" and "Notification to a receiver that a partial message delivery has been aborted" become unnecessary too.

- From the above, a TAPS system should always support message interleaving because it enables the use of stream schedulers and comes at no additional implementation cost on the receiver side. Stream schedulers operate on the sender side. Hence, because a TAPS sender-side application may talk to an SCTP receiver that does not support interleaving, it cannot assume that stream schedulers will always work as expected.

5.2. Stream Schedulers Without Streams

We have already stated that multi-streaming does not require application-specific knowledge. Potential benefits or disadvantages of, e.g., using two streams over an SCTP association versus using two separate SCTP associations or TCP connections are related to knowledge about the network and the particular transport protocol in use, not the application. However, the transport features "Choose a scheduler to operate between streams of an association" and "Configure priority or weight for a scheduler" operate on streams. Here, streams identify communication channels between which a scheduler operates, and they can be assigned a priority. Moreover, the transport features in the MAINTENANCE category all operate on assocations in case of SCTP, i.e. they apply to all streams in that assocation.

With only these semantics necessary to represent, the interface to a TAPS system becomes easier if we rename connections into "TAPS flows" (the TAPS equivalent of a connection which may be a transport connection or association, but could also become a stream of an existing SCTP association, for example) and allow assigning a "Group Number" to a TAPS flow. Then, all MAINTENANCE transport features can be said to operate on flow groups, not connections, and a scheduler also operates on the flows within a group.

For the implementation of a TAPS system, this has the following consequences:

- Streams may be identified in different ways across different protocols. The only multi-streaming protocol considered in this document, SCTP, uses a stream id. The transport association below still uses a Transport Address (which includes one port number) for each communicating endpoint. To implement a TAPS system without exposed streams, an application must be given an identifier for each TAPS flow (akin to a socket), and depending on whether streams are used or not, there will be a 1:1 mapping between this identifier and local ports or not.
- In SCTP, a fixed number of streams exists from the beginning of an association; streams are not "established", there is no handshake or any other form of signaling to create them: they can just be used. They are also not "gracefully shut down" -- at best, an "SSN Reset Request Parameter" in a "RE-CONFIG" chunk [RFC6525] can be used to inform the peer that of a "Stream Reset", as a rough equivalent of an "Abort". This has an impact on the semantics connection establishment and teardown (see Section 6.1).
- To support stream schedulers, a receiver-side TAPS system should always support message interleaving because it comes at no additional implementation cost (because of the receiver-side stream reception discussed in Section 5.1). Note, however, that Stream schedulers operate on the sender side. Hence, because a TAPS sender-side application may talk to a native TCP-based receiver-side application, it cannot assume that stream schedulers will always work as expected.

5.3. Early Data Transmission

There are two transport features related to transferring a message

early: "Hand over a message to transfer (possibly multiple times) before connection establishment", which relates to TCP Fast Open [RFC7413], and "Hand over a message to transfer during connection establishment", which relates to SCTP's ability to transfer data together with the COOKIE-Echo chunk. Also without TCP Fast Open, TCP can transfer data during the handshake, together with the SYN packet -- however, the receiver of this data may not hand it over to the application until the handshake has completed. This functionality is commonly available in TCP and supported in several implementations, but the TCP specification does not specify how to provide it to applications.

The amount of data that can successfully be transmitted before or during the handshake depends on various factors: the transport protocol, the use of header options, the choice of IPv4 and IPv6 and the Path MTU. A TAPS system should therefore allow a sending application to query the maximum amount of data it can possibly transmit before or during connection establishment, respectively.

5.4. Sender Running Dry

The transport feature "Notification that the stack has no more user data to send" relates to SCTP's "SENDER DRY" notification. Such notifications can, in principle, be used to avoid having an unnecessarily large send buffer, yet ensure that the transport sender always has data available when it has an opportunity to transmit it. This has been found to be very beneficial for some applications

[WWDC2015]. However, "SENDER DRY" truly means that the buffer has emptied -- i.e., when it notifies the sender, it is already too late, the transport protocol already missed an opportunity to send data. Some modern TCP implementations now include the unspecified "TCP_NOTSENT_LOWAT" socket option proposed in [WWDC2015], which limits the amount of unsent data that TCP can keep in the socket buffer; this allows to specify at which buffer filling level the socket becomes writable, rather than waiting for the buffer to run empty.

SCTP has means to configure the sender-side buffer too: the automatable Transport Feature "Configure send buffer size" provides this functionality, but only for the complete buffer, which includes both unsent and unacknowledged data. SCTP does not allow to control

these two sizes separately. A TAPS system should allow for uniform access to "TCP_NOTSENT_LOWAT" as well as the "SENDER DRY" notification.

5.5. Capacity Profile

The transport features:
- Disable Nagle algorithm
- Enable and configure a "Low Extra Delay Background Transfer"
- Specify DSCP field

all relate to a QoS-like application need such as "low latency" or "scavenger". In the interest of flexibility of a TAPS system, they could therefore be offered in a uniform, more abstract way, where a TAPS system could e.g. decide by itself how to use combinations of LEDBAT-like congestion control and certain DSCP values, and an application would only specify a general "capacity profile" (a description of how it wants to use the available capacity). A need for "lowest possible latency at the expense of overhead" could then translate into automatically disabling the Nagle algorithm.

In some cases, the Nagle algorithm is best controlled directly by the application because it is not only related to a general profile but also to knowledge about the size of future messages. For fine-grain control over Nagle-like functionality, the "Request not to bundle messages" is available.

5.6. Security

Both TCP and SCTP offer authentication. SCTP allows to configure which of SCTP's chunk types must always be authenticated -- if this is exposed as such, it creates an undesirable dependency on the transport protocol. Generally, to an application it is relevant whether the transport protocol authenticates its own control data, the user data, or both, and a TAPS system should therefore allow to

configure and query these three cases.

TBD -- more to come in the next version. This relates to the TCP Authentication Option in Section 7.1 of [RFC5925], which is not currently covered.

Set Cookie life value -- TBD in the next version: SCTP is client-

side, TCP is server-side.

5.7. Packet Size

UDP(-Lite) has a transport feature called "Specify DF field". This yields an error message in case of sending a message that exceeds the Path MTU, which is necessary for a UDP-based application to be able to implement Path MTU Discovery (a function that UDP-based applications must do by themselves). This is the only transport feature related to packet sizes. UDP applications typically make use of IP-layer functionality to obtain the size of the link MTU; it would therefore seem that offering such functionality to TAPS applications could be useful, albeit in a transport protocol independent way.

This also relates to the fact that the choice of path is automatable: if a TAPS system can switch a path at any time, unknown to an application, yet the application intends to do Path MTU Discovery, this could yield very inefficient behavior. Thus, a TAPS system should probably avoid automatically switching paths, and inform the application about any unavoidable path changes, when applications request to disallow fragmentation with the "Specify DF field" feature.

6. Step 4: Construction -- the Minimal Set of Transport Features

Based on the categorization, reduction and discussion in the previous sections, this section presents the minimal set of transport features that is offered by end systems supporting TAPS. They are described in an abstract fashion, i.e. they can be implemented in various different ways. For example, information that is provided to an application can either be offered via a primitive that is polled, or via an asynchronous notification.

Future versions of this document will probably describe the transport features in this section in greater detail; for now, we only specify how they differ from the transport features they are based upon. We carry out an additional simplification: CONNECTION.ESTABLISHMENT "Specify number of attempts and/or timeout for the first establishment message" and CONNECTION.MAINTENANCE "Change timeout for

aborting connection (using retransmit limit or time value)" are essentially the same, just applied upon connection establishment or during the lifetime of a connection. The same is the case for CONNECTION.ESTABLISHMENT "Specify which chunk types must always be authenticated" and CONNECTION.MAINTENANCE "Change authentication parameters". We therefore state that connections (called TAPS flows) must be instantiated before connecting them, and allow configurations to be carried out before connecting (in cases where this is not allowed by the transport protocol, a TAPS system will have to internall delay this configuration until the flow has been connected).

6.1. Flow Creation, Connection and Termination

A TAPS flow must be "created" before it is connected, to allow for initial configurations to be carried out. All configuration parameters in Section 6.2 and Section 6.3 can be used initially, although some of them may only take effect when the flow has been connected. Configuring a flow early helps a TAPS system make the right decisions. In particular, the "group number" can influence the the TAPS system to implement a TAPS flow as a stream of a multi-streaming protocol's existing association or not.

A created flow can be queried for the maximum amount of data that an application can possibly expect to have transmitted before or during connection establishment. An application can also give the flow a message for transmission before or during connection establishment, and specify which case is preferred (before / during). In case of transmission before establishment, the receiving application must be prepared to potentially receive multiple copies of the message.

To be compatible with multiple transports, including streams of a multi-streaming protocol (used as if they were transports themselves), the semantics of opening and closing need to be the most restrictive subset of all of them. For example, TCP's support of half-closed connections can be seen as a feature on top of the more restrictive "ABORT"; this feature cannot be supported because not all protocols used by a TAPS system (including streams of an association) support half-closed connections.

After creation, a flow can be actively connected to the other side using "Connect", or passively listen for incoming connection requests with "Listen". Note that "Connect" may or may not trigger a notification on the listening side. It is possible that the first notification on the listening side is the arrival of the first data that the active side sends (a receiver-side TAPS system could handle this by continuing a blocking "Listen" call, immediately followed by issuing "Receive", for example). This also means that the active

opening side is assumed to be the first side sending data.

A flow can be actively closed, i.e. terminated after reliably delivering all remaining data, or aborted, i.e. terminated without delivering remaining data. A timeout can be configured to abort a flow when data could not be delivered for too long. Because half-closed connections are not supported, when a TAPS host receives a notification that the peer is closing or aborting the flow, the other side may not be able to read outstanding data. This means that unacknowledged data residing in the TAPS system's send buffer may have to be dropped from that buffer upon arrival of a notification to close or abort the flow from the peer. In case of SCTP streams, "Stream Reset" (a "SSN Reset Request Parameter" in a "RE-CONFIG" chunk [RFC6525]) can be used to notify a peer of an intention to close a flow.

6.2. Flow Group Configuration

A flow group can be configured with a number of transport features, and there are some notifications to applications about a flow group. Here we list transport features and notifications that are taken from Section 4 unchanged, with the exception that some of them can also be applied initially (before a flow is connected).

Timeout, error notifications:
- Change timeout for aborting connection (using retransmit limit or time value)
- Suggest timeout to the peer
- Notification of Excessive Retransmissions (early warning below abortion threshold)
- Notification of ICMP error message arrival

Checksums:
- Disable checksum when sending
- Disable checksum requirement when receiving
- Specify checksum coverage used by the sender
- Specify minimum checksum coverage required by receiver

Others:
- Choose a scheduler to operate between flows of a group

The following transport features are new or changed, based on the discussion in Section 5:

- Capacity profile
  This describes how an application wants to use its available capacity. Choices can be "lowest possible latency at the expense of overhead", "scavenger", and some more values that help determine the DSCP value for a flow (e.g. similar to table 1 in

  [I-D.ietf-tsvwg-rtcweb-qos]). (details TBD)

- Authentication
  TBD in the next version: Different from SCTP's original transport features, this will only allow to configure authenticating the whole transport, all control information, or user data (not to distinguish between various SCTP chunks, to avoid this protocol dependency). It will also have to be made in line with TCP Authentication [RFC5925]. For SCTP, this functionality will be based on the transport features "Change authentication parameters", "Obtain authentication information" and the initially available "Specify which chunk types must always be authenticated". Note that SCTP also allows per-message configuration via "Specifying a key id to be used to authenticate a message", which may affect Section 6.4.

- Set Cookie life value
  TBD in the next version (not yet sure how to handle the client vs. server semantics of SCTP and TCP, respectively)

6.3. Flow Configuration

A flow can be assigned a priority or weight for a scheduler.

6.4. Data Transfer

6.4.1. The Sender

This section discusses how to send data after flow establishment. Section 6.1 discusses the possiblity to hand over a message to send before or during establishment.

For compatibility with TCP receiver semantics, we define an "Application-Framed Bytestream". This is a bytestream where the sending application optionally informs the transport about frame

boundaries and required properties per frame (configurable order and reliability, or embedding a request not to delay the acknowledgement of a frame). Whenever the sending application specifies per-frame properties that relax the notion of reliable in-order delivery of bytes, it must assume that the receiving application is 1) able to determine frame boundaries, provided that frames are always kept intact, and 2) able to accept these relaxed per-frame properties. Any signaling of such information to the peer is up to an application-layer protocol and considered out of scope of this document.

Here we list per-frame properties that a sender can optionally

configure if it hands over a delimited frame for sending with congestion control, taken from Section 4:

- Configurable Message Reliability
- Choice between unordered (potentially faster) or ordered delivery of messages
- Request not to bundle messages
- Request not to delay the acknowledgement (SACK) of a message

Additionally, an application can hand over delimited frames for unreliable transmission without congestion control (note that such applications should perform congestion control in accordance with [RFC2914]). Then, none of the per-frame properties listed above have any effect, but it is possible to use the transport feature "Specify DF field" to allow/disallow fragmentation.

AUTHOR'S NOTE: do folks agree with this design? It ties fragmentation to UDP only, because we called SCTP's "Configure message fragmentation" transport feature "automatable". It is indeed questionable whether applications need control over fragmentation when they work with SCTP -- doing so creates a complication for app writers that may not be necessary, especially when messages can be interleaved.

Following Section 5.7, there are two new transport features and a notification:

- Query maximum unfragmented frame size
  This is optional for a TAPS system to offer, and if it is offered, it informs the sender about the maximum expected size of a data frame that it can send without fragmentation. This can aid

applications implementing Path MTU Discovery.

- Query maximum transport frame size
  Irrespective of fragmentation, there is a size limit for the messages that can be handed over to SCTP or UDP(-Lite); because a TAPS system is independent of the transport, it must allow a TAPS application to query this value -- the maximum size of a frame in an Application-Framed-Bytestream.

- Notify the application of a path change
  If an application has disallowed fragmentation via the "Specify DF field" transport feature, this notification may optionally tell it that a path has changed (with a means to identify the path, so that the application can e.g. tell two flipping paths apart from completely diverse path changes). This informs the application that it may have to repeat Path MTU Discovery, and it can have relevance for application-level congestion control. For MPTCP and SCTP, a TAPS system can implement this functionality using the "Obtain status (query or notification)" transport feature.

There are two more sender-side notifications. These are unreliable, i.e. a TAPS system cannot be assumed to implement them, but they may occur:

- Notification of send failures
  A TAPS system may inform a sender application of a failure to send a specific frame. This was taken over unchanged from Section 4.

- Notification of draining below a low water mark
  A TAPS system can notify a sender application when the TAPS system's filling level of the buffer of unsent data is below a configurable threshold in bytes. Even for TAPS systems that do implement this notification, supporting thresholds other than 0 is optional.

"Notification of draining below a low water mark" is a generic notification that tries to enable uniform access to "TCP_NOTSENT_LOWAT" as well as the "SENDER DRY" notification (as discussed in Section 5.4 -- SCTP's "SENDER DRY" is a special case where the threshold is 0). Note that this threshold and its notification should operate across the buffers of the whole TAPS system, i.e. also any potential buffers that the TAPS system itself may use on top of the transport's send buffer.

6.4.2. The Receiver

A receiving application obtains an Application-Framed Bytestream. Similar to TCP's receiver semantics, it is just stream of bytes. If frame boundaries were specified by the sender, a TAPS system will still not inform the receiving application about them, but frames themselves will always stay intact (partial frames are not supported - see Section 5.1). Different from TCP's semantics, there is no guarantee that all bytes in the bytestream are received, and that all of them are in the same sequence in which they were handed over by the sender. If an application is aware of frame delimiters in the bytestream, and if the sender-side application has informed the TAPS system about these boundaries and about potentially relaxed requirements regarding the sequence of frames or per-frame reliability, frames within the receiver-side bytestream may be out-of-order or missing.

7. Conclusion

By decoupling applications from transport protocols, a TAPS system provides a different abstraction level than the Berkeley sockets interface. As with high- vs. low-level programming languages, a higher abstraction level allows more freedom for automation below the interface, yet it takes some control away from the application

programmer. This is the design trade-off that a TAPS system developer is facing, and this document provides guidance on the design of this abstraction level. Some transport features are currently rarely offered by APIs, yet they must be offered or they can never be used ("functional" transport features). Other transport features are offered by the APIs of the protocols covered here, but not exposing them in a TAPS API would allow for more freedom to automate protocol usage in a TAPS system.

The minimal set presented in this document is an effort to find a middle ground that can be recommended for TAPS systems to implement, on the basis of the transport features discussed in [TAPS2]. This middle ground eliminates a large number of transport features on the basis that they do not require application-specific knowledge, but rather rely on knowledge about the network or the Operating System.

This leaves us with an unanswered question about how exactly a TAPS system should automate using all these transport features.

The answers are different for every case. In some cases, it may be best to not entirely automate the decision making, but leave it up to a system-wide policy. For example, when multiple paths are available, a system policy could guide the decision on whether to connect via a WiFi or a cellular interface. Such high-level guidance could also be provided by application developers, e.g. via a primitive that lets applications specify such preferences. As long as this kind of information from applications is treated as advisory, it will not lead to a permanent protocol binding and does therefore not limit the flexibility of a TAPS system. Decisions to add such primitives are therefore left open to TAPS system designers.

8. Acknowledgements

The authors would like to thank the participants of the TAPS Working Group and the NEAT research project for valuable input to this document. We especially thank Michael Tuexen for help with TAPS flow connection establishment/teardown and Gorry Fairhurst for his suggestions regarding fragmentation and packet sizes. This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 644334 (NEAT). The views expressed are solely those of the author(s).

9. IANA Considerations

XX RFC ED - PLEASE REMOVE THIS SECTION XXX

This memo includes no request to IANA.

10. Security Considerations

Authentication, confidentiality protection, and integrity protection are identified as transport features by [RFC8095]. As currently deployed in the Internet, these features are generally provided by a protocol or layer on top of the transport protocol; no current full-featured standards-track transport protocol provides all of these transport features on its own. Therefore, these transport features

are not considered in this document, with the exception of native authentication capabilities of TCP and SCTP for which the security considerations in [RFC5925] and [RFC4895] apply.

11. References

11.1. Normative References

[RFC8095] Fairhurst, G., Ed., Trammell, B., Ed., and M. Kuehlewind, Ed., "Services Provided by IETF Transport Protocols and Congestion Control Mechanisms", RFC 8095, DOI 10.17487/RFC8095, March 2017, <http://www.rfc-editor.org/info/rfc8095>.

[TAPS2] Welzl, M., Tuexen, M., and N. Khademi, "On the Usage of Transport Features Provided by IETF Transport Protocols", draft-ietf-taps-transports-usage-03 (work in progress), March 2017.

11.2. Informative References

[COBS] Cheshire, S. and M. Baker, "Consistent Overhead Byte Stuffing", September 1997, <http://stuartcheshire.org/papers/COBSforToN.pdf>.

[I-D.ietf-tsvwg-rtcweb-qos] Jones, P., Dhesikan, S., Jennings, C., and D. Druta, "DSCP Packet Markings for WebRTC QoS", draft-ietf-tsvwg-rtcweb-qos-18 (work in progress), August 2016.

[LBE-draft] Bless, R., "A Lower Effort Per-Hop Behavior (LE PHB)", draft-tsvwg-le-phb-00 (work in progress), October 2016.

[RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <http://www.rfc-editor.org/info/rfc2914>.

[RFC4895] Tuexen, M., Stewart, R., Lei, P., and E. Rescorla,

"Authenticated Chunks for the Stream Control Transmission Protocol (SCTP)", RFC 4895, DOI 10.17487/RFC4895, August 2007, <http://www.rfc-editor.org/info/rfc4895>.

[RFC5290] Floyd, S. and M. Allman, "Comments on the Usefulness of Simple Best-Effort Traffic", RFC 5290, DOI 10.17487/RFC5290, July 2008, <http://www.rfc-editor.org/info/rfc5290>.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <http://www.rfc-editor.org/info/rfc5925>.

[RFC6458] Stewart, R., Tuexen, M., Poon, K., Lei, P., and V. Yasevich, "Sockets API Extensions for the Stream Control Transmission Protocol (SCTP)", RFC 6458, DOI 10.17487/RFC6458, December 2011, <http://www.rfc-editor.org/info/rfc6458>.

[RFC6525] Stewart, R., Tuexen, M., and P. Lei, "Stream Control Transmission Protocol (SCTP) Stream Reconfiguration", RFC 6525, DOI 10.17487/RFC6525, February 2012, <http://www.rfc-editor.org/info/rfc6525>.

[RFC7305] Lear, E., Ed., "Report from the IAB Workshop on Internet Technology Adoption and Transition (ITAT)", RFC 7305, DOI 10.17487/RFC7305, July 2014, <http://www.rfc-editor.org/info/rfc7305>.

[RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <http://www.rfc-editor.org/info/rfc7413>.

[WWDC2015] Lakhera, P. and S. Cheshire, "Your App and Next Generation Networks", Apple Worldwide Developers Conference 2015, San Francisco, USA, June 2015, <https://developer.apple.com/videos/wwdc/2015/?id=719>.

Appendix A. Revision information

XXX RFC-Ed please remove this section prior to publication.

-02: implementation suggestions added, discussion section added, terminology extended, DELETED category removed, various other fixes; list of Transport Features adjusted to -01 version of [TAPS2] except

that MPTCP is not included.

-03: updated to be consistent with -02 version of [TAPS2].

-04: updated to be consistent with -03 version of [TAPS2]. Reorganized document, rewrote intro and conclusion, and made a first stab at creating a real "minimal set".

Authors' Addresses

Stein Gjessing
University of Oslo
PO Box 1080 Blindern
Oslo, N-0316
Norway

Phone: +47 22 85 24 44
Email: steing@ifi.uio.no

Michael Welzl
University of Oslo
PO Box 1080 Blindern
Oslo, N-0316
Norway

Phone: +47 22 85 24 20
Email: michawe@ifi.uio.no