

TAPS  
Internet-Draft  
Intended status: Informational  
Expires: September 9, 2017

M. Welzl  
University of Oslo  
M. Tuexen  
Muenster Univ. of Appl. Sciences  
N. Khademi  
University of Oslo  
March 8, 2017

On the Usage of Transport Features Provided by IETF Transport Protocols  
draft-ietf-taps-transport-usage-03

Abstract

This document describes how TCP, MPTCP, SCTP, UDP and UDP-Lite expose services to applications and how an application can configure and use the transport features that make up these services. It also discusses the service provided by the LEDBAT congestion control mechanism.

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 9, 2017.

Copyright Notice

Copyright (c) 2017 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. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Terminology . . . . .	3
2. Introduction . . . . .	3
3. Pass 1 . . . . .	5
3.1. Primitives Provided by TCP . . . . .	5
3.1.1. Excluded Primitives or Parameters . . . . .	8
3.2. Primitives Provided by MPTCP . . . . .	9
3.3. Primitives Provided by SCTP . . . . .	10
3.3.1. Excluded Primitives or Parameters . . . . .	17
3.4. Primitives Provided by UDP and UDP-Lite . . . . .	17
3.5. The service of LEDBAT . . . . .	17
4. Pass 2 . . . . .	18
4.1. CONNECTION Related Primitives . . . . .	19
4.2. DATA Transfer Related Primitives . . . . .	30
5. Pass 3 . . . . .	33
5.1. CONNECTION Related Transport Features . . . . .	33
5.2. DATA Transfer Related Transport Features . . . . .	39
5.2.1. Sending Data . . . . .	39
5.2.2. Receiving Data . . . . .	40
5.2.3. Errors . . . . .	40
6. Acknowledgements . . . . .	41
7. IANA Considerations . . . . .	41
8. Security Considerations . . . . .	41
9. References . . . . .	41
9.1. Normative References . . . . .	41
9.2. Informative References . . . . .	44
Appendix A. Overview of RFCs used as input for pass 1 . . . . .	45
Appendix B. How this document was developed . . . . .	45
Appendix C. Revision information . . . . .	47
Authors' Addresses . . . . .	47

## 1. Terminology

**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 Protocol Component:** an implementation of a Transport Feature within a protocol.

**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 of data across the network (this may also be an upper layer protocol or tunnel encapsulation).

**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.

**Primitive:** a function call that is used to locally communicate between an application and a transport endpoint and is related to one or more Transport Features.

**Parameter:** a value passed between an application and a transport protocol by a primitive.

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

**Transport Address:** the combination of an IP address, transport protocol and the port number used by the transport protocol.

## 2. Introduction

This document presents defined interactions between applications and the transport protocols TCP, MPTCP, SCTP, UDP and UDP-Lite as well as the LEDBAT congestion control mechanism in the form of primitives and Transport Features. Primitives can be invoked by an application or a transport protocol; the latter type is called an "event". The list of primitives and Transport Features in this document is strictly based on the parts of protocol specifications that describe what the protocol provides to an application using it and how the application interacts with it.

Parts of a protocol that are explicitly stated as optional to implement are not covered. Interactions between the application and a transport protocol that are not directly related to the operation of the protocol are also not covered. For example, [RFC6458] explains how an application can use socket options to indicate its interest in receiving certain notifications. However, for the purpose of identifying primitives and Transport Services, the ability to enable or disable the reception of notifications is irrelevant. Similarly, one-to-many style sockets described in [RFC6458] just affect the application programming style, not how the underlying protocol operates, and they are therefore not discussed here. The same is true for the ability to obtain the unchanged value of a parameter that an application has previously set (this is the case for the "get" in many get/set operations in [RFC6458]).

The document presents a three-pass process to arrive at a list of Transport Features. In the first pass, the relevant RFC text is discussed per protocol. In the second pass, this discussion is used to derive a list of primitives that are uniformly categorized across protocols. Here, an attempt is made to present or -- where text describing primitives does not yet exist -- construct primitives in a slightly generalized form to highlight similarities. This is, for example, achieved by renaming primitives of protocols or by avoiding a strict 1:1-mapping between the primitives in the protocol specification and primitives in the list. Finally, the third pass presents Transport Features based on pass 2, identifying which protocols implement them.

In the list resulting from the second pass, some Transport Features are missing because they are implicit in some protocols, and they only become explicit when we consider the superset of all features offered by all protocols. For example, TCP always carries out congestion control; we have to consider it together with a protocol like UDP (which does not have congestion control) before we can consider congestion control as a Transport Feature. The complete list of features across all protocols is therefore only available after pass 3.

This document discusses unicast transport protocols and a unicast congestion control mechanism. Transport protocols provide communication between processes that operate on network endpoints, which means that they allow for multiplexing of communication between the same IP addresses, and normally this multiplexing is achieved using port numbers. Port multiplexing is therefore assumed to be always provided and not discussed in this document.

Some protocols are connection-oriented. Connection-oriented protocols often use an initial call to a specific transport primitive

to open a connection before communication can progress, and require communication to be explicitly terminated by issuing another call to a transport primitive (usually called "close"). A "connection" is the common state that some transport primitives refer to, e.g., to adjust general configuration settings. Connection establishment, maintenance and termination are therefore used to categorize transport primitives of connection-oriented transport protocols in pass 2 and pass 3. For this purpose, UDP is assumed to be used with "connected" sockets, i.e. sockets that are bound to a specific pair of addresses and ports [FJ16].

### 3. Pass 1

This first iteration summarizes the relevant text parts of the RFCs describing the protocols, focusing on what each transport protocol provides to the application and how it is used (abstract API descriptions, where they are available).

#### 3.1. Primitives Provided by TCP

[RFC0793] states: "The Transmission Control Protocol (TCP) is intended for use as a highly reliable host-to-host protocol between hosts in packet-switched computer communication networks, and in interconnected systems of such networks". Section 3.8 in [RFC0793] further specifies the interaction with the application by listing several transport primitives. It is also assumed that an Operating System provides a means for TCP to asynchronously signal the application; the primitives representing such signals are called 'events' in this section. This section describes the relevant primitives.

open: this is either active or passive, to initiate a connection or listen for incoming connections. All other primitives are associated with a specific connection, which is assumed to first have been opened. An active open call contains a socket. A passive open call with a socket waits for a particular connection; alternatively, a passive open call can leave the socket unspecified to accept any incoming connection. A fully specified passive call can later be made active by calling 'send'. Optionally, a timeout can be specified, after which TCP will abort the connection if data has not been successfully delivered to the destination (else a default timeout value is used). [RFC1122] describes a procedure for aborting the connection that must be used to avoid excessive retransmissions, and states that an application must be able to control the threshold used to determine the condition for aborting -- and that this threshold may be measured in time units or as a count of retransmission.

This indicates that the timeout could also be specified as a count of retransmission.

Also optional, for multihomed hosts, the local IP address can be provided [RFC1122]. If it is not provided, a default choice will be made in case of active open calls. A passive open call will await incoming connection requests to all local addresses and then maintain usage of the local IP address where the incoming connection request has arrived. Finally, the 'options' parameter is explained in [RFC1122] to allow the application to specify IP options such as source route, record route, or timestamp. It is not stated on which segments of a connection these options should be applied, but probably all segments, as this is also stated in a specification given for the usage of source route (section 4.2.3.8 of [RFC1122]). Source route is the only non-optional IP option in this parameter, allowing an application to specify a source route when it actively opens a TCP connection.

send: this is the primitive that an application uses to give the local TCP transport endpoint a number of bytes that TCP should reliably send to the other side of the connection. The URGENT flag, if set, states that the data handed over by this send call is urgent and this urgency should be indicated to the receiving process in case the receiving application has not yet consumed all non-urgent data preceding it. An optional timeout parameter can be provided that updates the connection's timeout (see 'open').

receive: This primitive allocates a receiving buffer for a provided number of bytes. It returns the number of received bytes provided in the buffer when these bytes have been received and written into the buffer by TCP. The application is informed of urgent data via an URGENT flag: if it is on, there is urgent data. If it is off, there is no urgent data or this call to 'receive' has returned all the urgent data.

close: This primitive closes one side of a connection. It is semantically equivalent to "I have no more data to send" but does not mean "I will not receive any more", as the other side may still have data to send. This call reliably delivers any data that has already been given to TCP (and if that fails, 'close' becomes 'abort').

abort: This primitive causes all pending 'send' and 'receive' calls to be aborted. A TCP RESET message is sent to the TCP endpoint on the other side of the connection [RFC0793].

close event: TCP uses this primitive to inform an application that the application on the other side has called the 'close' primitive, so the local application can also issue a 'close' and terminate the connection gracefully. See [RFC0793], Section 3.5.

abort event: When TCP aborts a connection upon receiving a "Reset" from the peer, it "advises the user and goes to the CLOSED state." See [RFC0793], Section 3.4.

USER TIMEOUT event: This event, described in Section 3.9 of [RFC0793], is executed when the user timeout expires (see 'open'). All queues are flushed and the application is informed that the connection had to be aborted due to user timeout.

ERROR\_REPORT event: This event, described in Section 4.2.4.1 of [RFC1122], informs the application of "soft errors" that can be safely ignored [RFC5461], including the arrival of an ICMP error message or excessive retransmissions (reaching a threshold below the threshold where the connection is aborted).

Type-of-Service: Section 4.2.4.2 of [RFC1122] states that the application layer MUST be able to specify the Type-of-Service (TOS) for segments that are sent on a connection. The application should be able to change the TOS during the connection lifetime, and the TOS value should be passed to the IP layer unchanged. Since then the TOS field has been redefined. A part of the field has been assigned to ECN [RFC3168] and the six most significant bits have been assigned to carry the DiffServ CodePoint, DSField [RFC3260]. Staying with the intention behind the application's ability to specify the "Type of Service", this should probably be interpreted to mean the value in the DSField, which is the Differentiated Services Codepoint (DSCP).

Nagle: The Nagle algorithm, described in Section 4.2.3.4 of [RFC1122], delays sending data for some time to increase the likelihood of sending a full-sized segment. An application can disable the Nagle algorithm for an individual connection.

User Timeout Option: The User Timeout Option (UTO) [RFC5482] allows one end of a TCP connection to advertise its current user timeout value so that the other end of the TCP connection can adapt its own user timeout accordingly. In addition to the configurable value of the User Timeout (see 'send'), [RFC5482] introduces three per-connection state variables that an application can adjust to control the operation of the User Timeout Option (UTO): ADV\_UTO is the value of the UTO advertised to the remote TCP peer (default: system-wide default user timeout); ENABLED (default false) is a boolean-type flag that controls whether the UTO option is enabled

for a connection. This applies to both sending and receiving. CHANGEABLE is a boolean-type flag (default true) that controls whether the user timeout may be changed based on a UTO option received from the other end of the connection. CHANGEABLE becomes false when an application explicitly sets the user timeout (see 'send').

Fast Open: TCP Fast Open (TFO) [RFC7413] allows to immediately hand over a message from the active open to the passive open side of a TCP connection together with the first message establishment packet (the SYN). This can be useful for applications that are sensitive to TCP's connection setup delay. TCP implementations MUST NOT use TFO by default, but only use TFO if requested explicitly by the application on a per-service-port basis. To benefit from TFO, the first application data unit (e.g., an HTTP request) needs to be no more than TCP's maximum segment size (minus options used in the SYN). For the active open side, [RFC7413] recommends changing or replacing the connect() call in order to support a user data buffer argument. For the passive open side, the application needs to enable the reception of Fast Open requests, e.g. via a new TCP\_FASTOPEN setsockopt() socket option before listen(). The receiving application must be prepared to accept duplicates of the TFO message, as the first data written to a socket can be delivered more than once to the application on the remote host.

#### 3.1.1. Excluded Primitives or Parameters

The 'open' primitive specified in [RFC0793] can be handed optional Precedence or security/compartiment information according to [RFC0793], but this was not included here because it is mostly irrelevant today, as explained in [RFC7414].

The 'status' primitive was not included because [RFC0793] describes this primitive as "implementation dependent" and states that it "could be excluded without adverse effect". Moreover, while a data block containing specific information is described, it is also stated that not all of this information may always be available. The 'send' primitive described in [RFC0793] includes an optional PUSH flag which, if set, requires data to be promptly transmitted to the receiver without delay; the 'receive' primitive described in [RFC0793] can (under some conditions) yield the status of the PUSH flag. Because PUSH functionality is made optional to implement for both the 'send' and 'receive' primitives in [RFC1122], this functionality is not included here. [RFC1122] also introduces keep-alives to TCP, but these are optional to implement and hence not considered here. [RFC1122] describes that "some TCP implementations

have included a FLUSH call", indicating that this call is also optional to implement. It is therefore not considered here.

### 3.2. Primitives Provided by MPTCP

Multipath TCP (MPTCP) is an extension to TCP that allows the use of multiple paths for a single data-stream. It achieves this by creating different so-called TCP subflows for each of the interfaces and scheduling the traffic across these TCP subflows. The service provided by MPTCP is described in [RFC6182] "Multipath TCP MUST follow the same service model as TCP [RFC0793]: in-order, reliable, and byte-oriented delivery. Furthermore, a Multipath TCP connection SHOULD provide the application with no worse throughput or resilience than it would expect from running a single TCP connection over any one of its available paths."

Further, [RFC6182] states constraints on the API exposed by MPTCP: "A multipath-capable equivalent of TCP MUST retain some level of backward compatibility with existing TCP APIs, so that existing applications can use the newer merely by upgrading the operating systems of the end hosts." As such, the primitives provided by MPTCP are equivalent to the ones provided by TCP. Nevertheless, [RFC6824] and [RFC6897] clarify some parts of TCP's primitives with respect to MPTCP and add some extensions for better control on MPTCP's subflows. Hereafter is a list of the clarifications and extensions the above cited RFCs provide to TCP's primitives.

open: [RFC6897] states "An application should be able to request to turn on or turn off the usage of MPTCP.". The RFC states that this functionality can be provided through a socket-option called TCP\_MULTIPATH\_ENABLE. Further, [RFC6897] says that MPTCP must be disabled in case the application is binding to a specific address.

send/receive: [RFC6824] states that the sending and receiving of data does not require any changes to the application when MPTCP is being used. The MPTCP-layer will "take one input data stream from an application, and split it into one or more subflows, with sufficient control information to allow it to be reassembled and delivered reliably and in order to the recipient application." The use of the Urgent-Pointer is special in MPTCP and [RFC6824] says "a TCP subflow MUST NOT use the Urgent Pointer to interrupt an existing mapping."

address and subflow management: MPTCP uses different addresses and allows a host to announce these addresses as part of the protocol. [RFC6897] says "An application should be able to restrict MPTCP to binding to a given set of addresses." and thus allows applications to limit the set of addresses that are being used by MPTCP.

Further, "An application should be able to obtain information on the pairs of addresses used by the MPTCP subflows."

### 3.3. Primitives Provided by SCTP

Section 1.1 of [RFC4960] lists limitations of TCP that SCTP removes. Three of the four mentioned limitations directly translate into Transport Features that are visible to an application using SCTP: 1) it allows for preservation of message delineations; 2) these messages, while reliably transferred, do not require to be in order unless the application wants it; 3) multi-homing is supported. In SCTP, connections are called "associations" and they can be between not only two (as in TCP) but multiple addresses at each endpoint.

Section 10 of [RFC4960] further specifies the interaction with the application (which RFC [RFC4960] calls the "Upper Layer Protocol" (ULP)). It is assumed that the Operating System provides a means for SCTP to asynchronously signal the application; the primitives representing such signals are called 'events' in this section. Here, we describe the relevant primitives. In addition to the abstract API described in Section 10 of [RFC4960], an extension to the socket API is described in [RFC6458], covering the functionality of the base protocol specified in [RFC4960] and its extensions specified in [RFC3758], [RFC4895], and [RFC5061]. For the protocol extensions specified in [RFC6525], [RFC6951], [RFC7053], [RFC7496], [RFC7829] and [I-D.ietf-tsvwg-sctp-ndata], the corresponding extensions of the socket API are specified in these protocol specifications. The functionality exposed to the ULP through this socket API is considered here in addition to the abstract API specified in Section 10 of [RFC4960].

[RFC4960] contains a "SETPROTOCOLPARAMETERS" primitive that allows to adjust elements of a parameter list; it is stated that SCTP implementations "may allow ULP to customize some of these protocol parameters", indicating that none of the elements of this parameter list are mandatory to make ULP-configurable. Thus, we only consider the parameters in [RFC4960] that are also covered in one of the other RFCs listed above, which leads us to exclude the parameters RTO.Alpha, RTO.Beta and HB.Max.Burst. For clarity, we also replace "SETPROTOCOLPARAMETERS" itself with primitives that adjust parameters or groups of parameters which fit together.

**Initialize:** Initialize, described in [RFC4960], creates a local SCTP instance that it binds to a set of local addresses (and, if provided, port number). Initialize needs to be called only once per set of local addresses. [RFC6458] also describes a number of per-association initialization parameters that can be used when an

association is created, but before it is connected (via the primitive 'Associate' below): the maximum number of inbound streams the application is prepared to support, the maximum number of attempts to be made when sending the INIT (the first message of association establishment), and the maximum retransmission timeout (RTO) value to use when attempting an INIT. At this point, before connecting, an application can also enable UDP encapsulation by configuring the remote UDP encapsulation port number [RFC6951].

**Associate:** This creates an association (the SCTP equivalent of a connection) that connects the local SCTP instance and a remote SCTP instance. To identify the remote endpoint, it can be given one or multiple (using connectx as described in section 9.9 of [RFC6458]) sockets. Most primitives are associated with a specific association, which is assumed to first have been created. Associate can return a list of destination transport addresses so that multiple paths can later be used. One of the returned sockets will be selected by the local endpoint as default primary path for sending SCTP packets to this peer, but this choice can be changed by the application using the list of destination addresses. Associate is also given the number of outgoing streams to request and optionally returns the number of negotiated outgoing streams. An optional parameter of 32 bits, the adaptation layer indication, can be provided, as specified in [RFC5061]. If the extension specified in [RFC4895] is used, the chunk types required to be sent authenticated by the peer can be provided. [RFC6458] describes a 'SCTP\_CANT\_STR\_ASSOC' notification that is used to inform the application of a failure to create an association. [RFC6458] describes how an application could use sendto() or sendmsg() to implicitly setup an association, thereby handing over a message that SCTP might send during the association setup phase. Note that this mechanism is different from TCP's TFO mechanism: the message would arrive only once, after at least one RTT, as it is sent together with the third message exchanged during association setup, the COOKIE-ECHO chunk).

**Send:** This sends a message of a certain length in bytes over an association. A number can be provided to later refer to the correct message when reporting an error, and a stream id is provided to specify the stream to be used inside an association (we consider this as a mandatory parameter here for simplicity: if not provided, the stream id defaults to 0). A condition to abandon the message can be specified (for example limiting the number of retransmissions or the lifetime of the user message). This allows to control the partial reliability extension specified in [RFC3758] and [RFC7496]. An optional maximum life time can specify the time after which the message should be discarded

rather than sent. A choice (advisory, i.e. not guaranteed) of the preferred path can be made by providing a socket, and the message can be delivered out-of-order if the unordered flag is set. An advisory flag indicates that the peer should not delay the acknowledgement of the user message provided by making use of the I-bit specified in [RFC7053]. Another advisory flag indicates whether the application prefers to avoid bundling user data with other outbound DATA chunks (i.e., in the same packet). A payload protocol-id can be provided to pass a value that indicates the type of payload protocol data to the peer. If the extension specified in [RFC4895] is used, the key identifier used for authenticating the DATA chunks can be provided.

**Receive:** Messages are received from an association, and optionally a stream within the association, with their size returned. The application is notified of the availability of data via a DATA ARRIVE notification. If the sender has included a payload protocol-id, this value is also returned. If the received message is only a partial delivery of a whole message, a partial flag will indicate so, in which case the stream id and a stream sequence number are provided to the application. A delivery number lets the application detect reordering.

**Shutdown:** This primitive gracefully closes an association, reliably delivering any data that has already been handed over to SCTP. A parameter lets the application control whether further receive or send operations or both are disabled when the call is issued. A return code informs about success or failure of this procedure.

**Abort:** This ungracefully closes an association, by discarding any locally queued data and informing the peer that the association was aborted. Optionally, an abort reason to be passed to the peer may be provided by the application. A return code informs about success or failure of this procedure.

**Change Heartbeat / Request Heartbeat:** This allows the application to enable/disable heartbeats and optionally specify a heartbeat frequency as well as requesting a single heartbeat to be carried out upon a function call, with a notification about success or failure of transmitting the HEARTBEAT chunk to the destination.

**Configure Max. Retransmissions of an Association:** The parameter Association.Max.Retrans in [RFC4960], called sasoc\_maxrxt in [RFC6458], allows to configure the number of unsuccessful retransmissions after which an entire association is considered as failed (which should invoke a COMMUNICATION LOST notification).

**Set Primary:** This allows to set a new primary default path for an association by providing a socket. Optionally, a default source address to be used in IP datagrams can be provided.

**Change Local Address / Set Peer Primary:** This allows an endpoint to add/remove local addresses to/from an association. In addition, the peer can be given a hint which address to use as the primary address. This is provided by the protocol extension defined in [RFC5061].

**Configure Path Switchover:** [RFC4960] contains a primitive called SET FAILURE THRESHOLD. This configures the parameter "Path.Max.Retrans", which determines after how many retransmissions a particular transport address is considered as unreachable. If there are more transport addresses available in an association, reaching this limit will invoke a path switchover. [RFC7829] extends this method with a concept of "Potentially Failed" (PF) paths. When a path is in PF state, SCTP will not entirely give up sending on that path, but it will preferably send data on other active paths if such paths are available. Entering the PF state is done upon exceeding a configured maximum number of retransmissions. Thus, for all paths where this mechanism is used, there are two configurable error thresholds: one to decide that a path is in PF state, and one to decide that the transport address is unreachable.

**Set / Get Authentication Parameters:** This allows an endpoint to add/remove key material to/from an association. In addition, the chunk types being authenticated can be queried. This is provided by the protocol extension defined in [RFC4895].

**Add / Reset Streams, Reset Association:** This allows an endpoint to add streams to an existing association or to reset them individually. Additionally, the association can be reset. This is provided by the protocol extension defined in [RFC6525].

**Status:** The 'Status' primitive returns a data block with information about a specified association, containing: association connection state; destination transport address list; destination transport address reachability states; current local and peer receiver window sizes; 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 [RFC4960] and MTU per path [RFC6458].

**Enable / Disable Interleaving:** This allows to enable or disable the negotiation of user message interleaving support for future associations. For existing associations it is possible to query whether user message interleaving support was negotiated or not on a particular association [I-D.ietf-tsvwg-sctp-ndata].

**Set Stream Scheduler:** This allows to select a stream scheduler per association, with a choice of: First Come First Serve, Round Robin, Round Robin per Packet, Priority Based, Fair Bandwidth, Weighted Fair Queuing. How these schedulers operate is described in detail in [I-D.ietf-tsvwg-sctp-ndata].

**Configure Stream Scheduler:** This allows to change a parameter per stream for the schedulers: a priority value for the Priority Based scheduler and a weight for the Weighted Fair Queuing scheduler.

**Enable/disable NODELAY:** This turns on/off any Nagle-like algorithm for an association [RFC6458].

**Configure send buffer size:** This controls the amount of data SCTP may have waiting in internal buffers to be sent or retransmitted [RFC6458].

**Configure receive buffer size:** This sets the receive buffer size in octets, thereby controlling the receiver window for an association [RFC6458].

**Configure message fragmentation:** If a user message causes an SCTP packet to exceed the maximum fragmentation size (which can be provided by the application, and is otherwise the PMTU size), then the message will be fragmented by SCTP. Disabling message fragmentation will produce an error instead of fragmenting the message [RFC6458].

**Configure Path MTU Discovery:** Section 8.1.12 of [RFC6458] explains how Path MTU Discovery can be enabled or disabled per peer address of an association. When it is enabled, the current Path MTU value can be obtained. When it is disabled, the Path MTU to be used can be controlled by the application.

**Configure delayed SACK timer:** The time before sending a SACK can be adjusted; delaying SACKs can be disabled; the number of packets that must be received before a SACK is sent without waiting for the delay timer to expire can be configured [RFC6458].

Set Cookie life value: The Cookie life value can be adjusted as explained in Section 8.1.2 of [RFC6458]. "Valid.Cookie.Life" is also one of the parameters listed as potentially adjustable with SETPROTOCOLPARAMETERS in [RFC4960].

Set maximum burst: The maximum burst of packets that can be emitted by a particular association (default 4, and values above 4 are optional to implement) can be adjusted as explained in Section 8.1.2 of [RFC6458]. "Max.Burst" is also one of the parameters listed as potentially adjustable with SETPROTOCOLPARAMETERS in [RFC4960].

Configure RTO calculation: [RFC4960] lists the following adjustable parameters: RTO.Initial; RTO.Min; RTO.Max; RTO.Alpha; RTO.Beta. Only the initial, minimum and maximum RTO are also described as configurable [RFC6458].

Set DSCP value: Section 8.1.12 of [RFC6458] explains how to set the DSCP value per peer address of an association.

Set IPv6 flow label: Section 8.1.12 of [RFC6458] explains how to set the flow label field per peer address of an association.

Set Partial Delivery Point: This allows to specify the size of a message where partial delivery will be invoked. Setting this to a lower value will cause partial deliveries to happen more often [RFC6458].

COMMUNICATION UP notification: When a lost communication to an endpoint is restored or when SCTP becomes ready to send or receive user messages, this notification informs the application process about the affected association, the type of event that has occurred, the complete set of sockets of the peer, the maximum number of allowed streams and the inbound stream count (the number of streams the peer endpoint has requested). If interleaving is supported by both endpoints, this information is also included in this notification.

RESTART notification: When SCTP has detected that the peer has restarted, this notification is passed to the upper layer [RFC6458].

DATA ARRIVE notification: When a message is ready to be retrieved via the Receive primitive, the application is informed by this notification.

SEND FAILURE notification / Receive Unsent Message / Receive Unacknowledged Message: When a message cannot be delivered via an association, the sender can be informed about it and learn whether the message has just not been acknowledged or (e.g. in case of lifetime expiry) if it has not even been sent. This can also inform the sender that a part of the message has been successfully delivered.

NETWORK STATUS CHANGE notification: The NETWORK STATUS CHANGE notification informs the application about a socket becoming active/inactive [RFC4960] or "Potentially Failed" [RFC7829].

COMMUNICATION LOST notification: When SCTP loses communication to an endpoint (e.g. via Heartbeats or excessive retransmission) or detects an abort, this notification informs the application process of the affected association and the type of event (failure OR termination in response to a shutdown or abort request).

SHUTDOWN COMPLETE notification: When SCTP completes the shutdown procedures, this notification is passed to the upper layer, informing it about the affected association.

AUTHENTICATION notification: When SCTP wants to notify the upper layer regarding the key management related to the extension defined in [RFC4895], this notification is passed to the upper layer.

ADAPTATION LAYER INDICATION notification: When SCTP completes the association setup and the peer provided an adaptation layer indication, this is passed to the upper layer. This extension is defined in [RFC5061] and [RFC6458].

STREAM RESET notification: When SCTP completes the procedure for resetting streams as specified in [RFC6525], this notification is passed to the upper layer, informing it about the result.

ASSOCIATION RESET notification: When SCTP completes the association reset procedure as specified in [RFC6525], this notification is passed to the upper layer, informing it about the result.

STREAM CHANGE notification: When SCTP completes the procedure used to increase the number of streams as specified in [RFC6525], this notification is passed to the upper layer, informing it about the result.

SENDER DRY notification: When SCTP has no more user data to send or retransmit on a particular association, this notification is passed to the upper layer [RFC6458].

PARTIAL DELIVERY ABORTED notification: When a receiver has begun to receive parts of a user message but the delivery of this message is then aborted, this notification is passed to the upper layer (section 6.1.7 of [RFC6458]).

### 3.3.1. Excluded Primitives or Parameters

The 'Receive' primitive can return certain additional information, but this is optional to implement and therefore not considered. With a COMMUNICATION LOST notification, some more information may optionally be passed to the application (e.g., identification to retrieve unsent and unacknowledged data). SCTP "can invoke" a COMMUNICATION ERROR notification and "may send" a RESTART notification, making these two notifications optional to implement. The list provided under 'Status' includes "etc", indicating that more information could be provided. The primitive 'Get SRTT Report' returns information that is included in the information that 'Status' provides and is therefore not discussed. The 'Destroy SCTP Instance' API function was excluded: it erases the SCTP instance that was created by 'Initialize', but is not a Primitive as defined in this document because it does not relate to a Transport Feature. The SHUTDOWN EVENT described in Section 6.1 of [RFC6458] informs an application that the peer has sent a SHUTDOWN, and hence no further data should be sent on this socket. However, if an application would try to send data on the socket, it would get an error message anyway; thus, this event is classified as "just affecting the application programming style, not how the underlying protocol operates" and not included here.

### 3.4. Primitives Provided by UDP and UDP-Lite

The primitives provided by UDP and UDP-Lite are described in [FJ16].

### 3.5. The service of LEDBAT

The service of the Low Extra Delay Background Transport (LEDBAT) congestion control mechanism is described in the abstract of [RFC6817] as follows: "LEDBAT is designed for use by background bulk-transfer applications to be no more aggressive than standard TCP congestion control (as specified in RFC 5681) and to yield in the presence of competing flows, thus limiting interference with the network performance of competing flows."

LEDBAT does not have any primitives, as LEDBAT is not a transport protocol. [RFC6817] states: "LEDBAT can be used as part of a transport protocol or as part of an application, as long as the data transmission mechanisms are capable of carrying timestamps and acknowledging data frequently. LEDBAT can be used with TCP, Stream Control Transmission Protocol (SCTP), and Datagram Congestion Control Protocol (DCCP), with appropriate extensions where necessary; and it can be used with proprietary application protocols, such as those built on top of UDP for peer-to-peer (P2P) applications." At the time of writing, the appropriate extensions for TCP, SCTP or DCCP do not exist.

A number of configurable parameters exist in the LEDBAT specification: TARGET, which is the queuing delay target at which LEDBAT tries to operate, must be set to 100ms or less. ALLOWED\_INCREASE (should be 1, must be greater than 0) limits the speed at which LEDBAT increases its rate. GAIN, which MUST be set to 1 or less to avoid a faster ramp-up than TCP Reno, determines how quickly the sender responds to changes in queuing delay. Implementations may divide GAIN into two parameters, one for increase and a possibly larger one for decrease. We call these parameters GAIN\_INC and GAIN\_DEC here. BASE\_HISTORY is the size of the list of measured base delays, and SHOULD be 10. This list can be filtered using a FILTER() function which is not prescribed in [RFC6817], yielding a list of size CURRENT\_FILTER. The initial and minimum congestion windows, INIT\_CWND and MIN\_CWND, should both be 2.

Regarding which of these parameters should be under control of an application, the possible range goes from exposing nothing on the one hand, to considering everything that is not fully prescribed with a MUST in [RFC6817] as a parameter on the other hand. Function implementations are not provided as a parameter to any of the transport protocols discussed here, and hence we do not regard the FILTER() function as a parameter. However, to avoid unnecessarily limiting future implementations, we consider all other parameters above as tunable parameters that a TAPS system should expose.

#### 4. Pass 2

This pass categorizes the primitives from pass 1 based on whether they relate to a connection or to data transmission. Primitives are presented following the nomenclature "CATEGORY.[SUBCATEGORY].PRIMITIVENAME.PROTOCOL". The CATEGORY can be CONNECTION or DATA. Within the CONNECTION category, ESTABLISHMENT, AVAILABILITY, MAINTENANCE and TERMINATION subcategories can be considered. The DATA category does not have any SUBCATEGORY. The PROTOCOL name "UDP(-Lite)" is used when primitives are equivalent for

UDP and UDP-Lite; the PROTOCOL name "TCP" refers to both TCP and MPTCP. We present "connection" as a general protocol-independent concept and use it to refer to, e.g., TCP connections (identifiable by a unique pair of IP addresses and TCP port numbers), SCTP associations (identifiable by multiple IP address and port number pairs), as well UDP and UDP-Lite connections (identifiable by a unique socket pair).

Some minor details are omitted for the sake of generalization -- e.g., SCTP's 'close' [RFC4960] returns success or failure, and lets the application control whether further receive or send operations or both are disabled [RFC6458]. This is not described in the same way for TCP in [RFC0793], but these details play no significant role for the primitives provided by either TCP or SCTP (for the sake of being generic, it could be assumed that both receive and send operations are disabled in both cases).

The TCP 'send' and 'receive' primitives include usage of an "URGENT" mechanism. This mechanism is required to implement the "synch signal" used by telnet [RFC0854], but SHOULD NOT be used by new applications [RFC6093]. Because pass 2 is meant as a basis for the creation of TAPS systems, the "URGENT" mechanism is excluded. This also concerns the notification "Urgent pointer advance" in the ERROR\_REPORT described in Section 4.2.4.1 of [RFC1122].

Since LEDBAT is a congestion control mechanism and not a protocol, it is not currently defined when to enable / disable or configure the mechanism. For instance, it could be a one-time choice upon connection establishment or when listening for incoming connections, in which case it should be categorized under CONNECTION.ESTABLISHMENT or CONNECTION.AVAILABILITY, respectively. To avoid unnecessarily limiting future implementations, it was decided to place it under CONNECTION.MAINTENANCE, with all parameters that are described in [RFC6817] made configurable.

#### 4.1. CONNECTION Related Primitives

##### ESTABLISHMENT:

Active creation of a connection from one transport endpoint to one or more transport endpoints.

Interfaces to UDP and UDP-Lite allow both connection-oriented and connection-less usage of the API . [RFC8085]

##### o CONNECT.TCP:

Pass 1 primitive / event: 'open' (active) or 'open' (passive) with socket, followed by 'send'

Parameters: 1 local IP address (optional); 1 destination transport address (for active open; else the socket and the local IP address

of the succeeding incoming connection request will be maintained); timeout (optional); options (optional); user message (optional)  
Comments: If the local IP address is not provided, a default choice will automatically be made. The timeout can also be a retransmission count. The options are IP options to be used on all segments of the connection. At least the Source Route option is mandatory for TCP to provide. The user message may be transmitted to the peer application immediately upon reception of the TCP SYN packet. To benefit from the lower latency this provides as part of the experimental TFO mechanism, its length must be at most the TCP's maximum segment size (minus TCP options used in the SYN). The message may also be delivered more than once to the application on the remote host.

- o CONNECT.SCTP:  
Pass 1 primitive / event: 'initialize', followed by 'enable / disable interleaving' (optional), followed by 'associate'  
Parameters: list of local SCTP port number / IP address pairs (initialize); one or several sockets (identifying the peer); outbound stream count; maximum allowed inbound stream count; adaptation layer indication (optional); chunk types required to be authenticated (optional); request interleaving on/off; maximum number of INIT attempts (optional); maximum init. RTO for INIT (optional); user message (optional); remote UDP port number (optional)  
Returns: socket list or failure  
Comments: 'initialize' needs to be called only once per list of local SCTP port number / IP address pairs. One socket will automatically be chosen; it can later be changed in MAINTENANCE. The user message may be transmitted to the peer application immediately upon reception of the packet containing the COOKIE-ECHO chunk. To benefit from the lower latency this provides, its length must be limited such that it fits into the packet containing the COOKIE-ECHO chunk. If a remote UDP port number is provided, SCTP packets will be encapsulated in UDP.
- o CONNECT.MPTCP:  
This is similar to CONNECT.TCP except for one additional boolean parameter that allows to enable or disable MPTCP for a particular connection or socket (default: enabled).
- o CONNECT.UDP(-Lite):  
Pass 1 primitive / event: 'connect' followed by 'send'.  
Parameters: 1 local IP address (default (ANY), or specified); 1 destination transport address; 1 local port (default (OS chooses), or specified); 1 destination port (default (OS chooses), or specified).  
Comments: Associates a transport address creating a UDP(-Lite)

socket connection. This can be called again with a new transport address to create a new connection. The CONNECT function allows an application to receive errors from messages sent to a transport address.

AVAILABILITY:

Preparing to receive incoming connection requests.

- o LISTEN.TCP:  
Pass 1 primitive / event: 'open' (passive)  
Parameters: 1 local IP address (optional); 1 socket (optional);  
timeout (optional); buffer to receive a user message (optional)  
Comments: if the socket and/or local IP address is provided, this  
waits for incoming connections from only and/or to only the  
provided address. Else this waits for incoming connections  
without this / these constraint(s). ESTABLISHMENT can later be  
performed with 'send'. If a buffer is provided to receive a user  
message, a user message can be received from a TFO-enabled sender  
before TCP's connection handshake is completed. This message may  
arrive multiple times.
- o LISTEN.SCTP:  
Pass 1 primitive / event: 'initialize', followed by 'COMMUNICATION  
UP' or 'RESTART' notification and possibly 'ADAPTATION LAYER'  
notification  
Parameters: list of local SCTP port number / IP address pairs  
(initialize)  
Returns: socket list; outbound stream count; inbound stream count;  
adaptation layer indication; chunks required to be authenticated;  
interleaving supported on both sides yes/no  
Comments: initialize needs to be called only once per list of  
local SCTP port number / IP address pairs. COMMUNICATION UP can  
also follow a COMMUNICATION LOST notification, indicating that the  
lost communication is restored. If the peer has provided an  
adaptation layer indication, an 'ADAPTATION LAYER' notification is  
issued.
- o LISTEN.MPTCP:  
This is similar to LISTEN.TCP except for one additional boolean  
parameter that allows to enable or disable MPTCP for a particular  
connection or socket (default: enabled).
- o LISTEN.UDP(-Lite):  
Pass 1 primitive / event: 'receive'.  
Parameters: 1 local IP address (default (ANY), or specified); 1  
destination transport address; local port (default (OS chooses),  
or specified); destination port (default (OS chooses), or

specified).

Comments: The receive function registers the application to listen for incoming UDP(-Lite) datagrams at an endpoint.

#### MAINTENANCE:

Adjustments made to an open connection, or notifications about it. These are out-of-band messages to the protocol that can be issued at any time, at least after a connection has been established and before it has been terminated (with one exception: `CHANGE_TIMEOUT.TCP` can only be issued for an open connection when `DATA.SEND.TCP` is called). In some cases, these primitives can also be immediately issued during `ESTABLISHMENT` or `AVAILABILITY`, without waiting for the connection to be opened (e.g. `CHANGE_TIMEOUT.TCP` can be done using TCP's 'open' primitive). For UDP and UDP-Lite, these functions may establish a setting per connection, but may also be changed per datagram message.

#### o `CHANGE_TIMEOUT.TCP`:

Pass 1 primitive / event: 'open' or 'send' combined with unspecified control of per-connection state variables  
Parameters: timeout value (optional); `ADV_UTO` (optional); boolean `UTO_ENABLED` (optional, default false); boolean `CHANGEABLE` (optional, default true)

Comments: when sending data, an application can adjust the connection's timeout value (time after which the connection will be aborted if data could not be delivered). If `UTO_ENABLED` is true, the user timeout value (or, if provided, the value `ADV_UTO`) will be advertised for the TCP on the other side of the connection to adapt its own user timeout accordingly. `UTO_ENABLED` controls whether the UTO option is enabled for a connection. This applies to both sending and receiving. `CHANGEABLE` controls whether the user timeout may be changed based on a UTO option received from the other end of the connection; it becomes false when 'timeout value' is used.

#### o `CHANGE_TIMEOUT.SCTP`:

Pass 1 primitive / event: 'Change HeartBeat' combined with 'Configure Max. Retransmissions of an Association'  
Parameters: 'Change HeartBeat': heartbeat frequency; 'Configure Max. Retransmissions of an Association': `Association.Max.Retrans`  
Comments: Change Heartbeat can enable / disable heartbeats in SCTP as well as change their frequency. The parameter `Association.Max.Retrans` defines after how many unsuccessful transmissions of any packets (including heartbeats) the association will be terminated; thus these two primitives / parameters together can yield a similar behavior for SCTP associations as `CHANGE_TIMEOUT.TCP` does for TCP connections.

- `DISABLE_NAGLE.TCP`:  
Pass 1 primitive / event: not specified  
Parameters: one boolean value  
Comments: the Nagle algorithm delays data transmission to increase the chance to send a full-sized segment. An application must be able to disable this algorithm for a connection.
- `DISABLE_NAGLE.SCTP`:  
Pass 1 primitive / event: 'Enable/disable NODELAY'  
Parameters: one boolean value  
Comments: Nagle-like algorithms delay data transmission to increase the chance to send a full-sized packet.
- `REQUEST_HEARTBEAT.SCTP`:  
Pass 1 primitive / event: 'Request HeartBeat'  
Parameters: socket  
Returns: success or failure  
Comments: requests an immediate heartbeat on a path, returning success or failure.
- `ADD_PATH.MPTCP`:  
Pass 1 primitive / event: not specified  
Parameters: local IP address and optionally the local port number  
Comments: the application specifies the local IP address and port number that must be used for a new subflow.
- `ADD_PATH.SCTP`:  
Pass 1 primitive / event: Change Local Address / Set Peer Primary  
Parameters: local IP address
- `REM_PATH.MPTCP`:  
Pass 1 primitive / event: not specified  
Parameters: local IP address, local port number, remote IP address, remote port number  
Comments: the application removes the subflow specified by the IP/port-pair. The MPTCP implementation must trigger a removal of the subflow that belongs to this IP/port-pair.
- `REM_PATH.SCTP`:  
Pass 1 primitive / event: 'Change Local Address / Set Peer Primary'  
Parameters: local IP address
- `SET_PRIMARY.SCTP`:  
Pass 1 primitive / event: 'Set Primary'  
Parameters: socket  
Returns: result of attempting this operation  
Comments: update the current primary address to be used, based on

the set of available sockets of the association.

- o SET\_PEER\_PRIMARY.SCTP:  
Pass 1 primitive / event: 'Change Local Address / Set Peer Primary'  
Parameters: local IP address  
Comments: this is only advisory for the peer.
- o CONFIG\_SWITCHOVER.SCTP:  
Pass 1 primitive / event: 'Configure Path Switchover'  
Parameters: primary max retrans (no. of retransmissions after which a path is considered inactive), PF max retrans (no. of retransmissions after which a path is considered to be "Potentially Failed", and others will be preferably used) (optional)
- o STATUS.SCTP:  
Pass 1 primitive / event: 'Status', 'Enable / Disable Interleaving' and 'NETWORK STATUS CHANGE notification'.  
Returns: data block with information about a specified association, containing: association connection state; destination transport address list; destination transport address reachability states; current local and peer receiver window sizes; 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.  
Comments: The NETWORK STATUS CHANGE notification informs the application about a socket becoming active/inactive; this only affects the programming style, as the same information is also available via 'Status'.
- o STATUS.MPTCP:  
Pass 1 primitive / event: not specified  
Returns: list of pairs of tuples of IP address and TCP port number of each subflow. The first of the pair is the local IP and port number, while the second is the remote IP and port number.
- o SET\_DSCP.TCP:  
Pass 1 primitive / event: not specified  
Parameters: DSCP value  
Comments: this allows an application to change the DSCP value for outgoing segments. For TCP this was originally specified for the TOS field [RFC1122], which is here interpreted to refer to the DSField [RFC3260].

- SET\_DSCP.SCTP:  
Pass 1 primitive / event: 'Set DSCP value'  
Parameters: DSCP value  
Comments: this allows an application to change the DSCP value for outgoing packets on a path.
- SET\_DSCP.UDP(-Lite):  
Pass 1 primitive / event: 'SET\_DSCP'  
Parameter: DSCP value  
Comments: This allows an application to change the DSCP value for outgoing UDP(-Lite) datagrams. [RFC7657] and [RFC8085] provide current guidance on using this value with UDP.
- ERROR.TCP:  
Pass 1 primitive / event: 'ERROR\_REPORT'  
Returns: reason (encoding not specified); subreason (encoding not specified)  
Comments: soft errors that can be ignored without harm by many applications; an application should be able to disable these notifications. The reported conditions include at least: ICMP error message arrived; Excessive Retransmissions.
- ERROR.UDP(-Lite):  
Pass 1 primitive / event: 'ERROR\_REPORT'.  
Returns: Error report  
Comments: This returns soft errors that may be ignored without harm by many applications; An application must connect to be able receive these notifications.
- SET\_AUTH.SCTP:  
Pass 1 primitive / event: 'Set / Get Authentication Parameters'  
Parameters: key\_id, key, hmac\_id
- GET\_AUTH.SCTP:  
Pass 1 primitive / event: 'Set / Get Authentication Parameters'  
Parameters: key\_id, chunk\_list
- RESET\_STREAM.SCTP:  
Pass 1 primitive / event: 'Add / Reset Streams, Reset Association'  
Parameters: sid, direction
- RESET\_STREAM-EVENT.SCTP:  
Pass 1 primitive / event: 'STREAM RESET notification'  
Parameters: information about the result of RESET\_STREAM.SCTP.  
Comments: This is issued when the procedure for resetting streams has completed.

- RESET\_ASSOC.SCTP:  
Pass 1 primitive / event: 'Add / Reset Streams, Reset Association'  
Parameters: information related to the extension defined in [RFC3260].
- RESET\_ASSOC-EVENT.SCTP:  
Pass 1 primitive / event: 'ASSOCIATION RESET notification'  
Parameters: information about the result of RESET\_ASSOC.SCTP.  
Comments: This is issued when the procedure for resetting an association has completed.
- ADD\_STREAM.SCTP:  
Pass 1 primitive / event: 'Add / Reset Streams, Reset Association'  
Parameters: number if outgoing and incoming streams to be added
- ADD\_STREAM-EVENT.SCTP:  
Pass 1 primitive / event: 'STREAM CHANGE notification'  
Parameters: information about the result of ADD\_STREAM.SCTP.  
Comments: This is issued when the procedure for adding a stream has completed.
- SET\_STREAM\_SCHEDULER.SCTP:  
Pass 1 primitive / event: 'Set Stream Scheduler'  
Parameters: scheduler identifier  
Comments: choice of First Come First Serve, Round Robin, Round Robin per Packet, Priority Based, Fair Bandwidth, Weighted Fair Queuing.
- CONFIGURE\_STREAM\_SCHEDULER.SCTP:  
Pass 1 primitive / event: 'Configure Stream Scheduler'  
Parameters: priority  
Comments: the priority value only applies when Priority Based or Weighted Fair Queuing scheduling is chosen with SET\_STREAM\_SCHEDULER.SCTP. The meaning of the parameter differs between these two schedulers but in both cases it realizes some form of prioritization regarding how bandwidth is divided among streams.
- SET\_FLOWLABEL.SCTP:  
Pass 1 primitive / event: 'Set IPv6 flow label'  
Parameters: flow label  
Comments: this allows an application to change the IPv6 header's flow label field for outgoing packets on a path.
- AUTHENTICATION\_NOTIFICATION-EVENT.SCTP:  
Pass 1 primitive / event: 'AUTHENTICATION notification'  
Returns: information regarding key management.

- CONFIG\_SEND\_BUFFER.SCTP:  
Pass 1 primitive / event: 'Configure send buffer size'  
Parameters: size value in octets
- CONFIG\_RECEIVE\_BUFFER.SCTP:  
Pass 1 primitive / event: 'Configure receive buffer size'  
Parameters: size value in octets  
Comments: this controls the receiver window.
- CONFIG\_FRAGMENTATION.SCTP:  
Pass 1 primitive / event: 'Configure message fragmentation'  
Parameters: one boolean value (enable/disable), maximum fragmentation size (optional; default: PMTU)  
Comments: if fragmentation is enabled, messages exceeding the maximum fragmentation size will be fragmented. If fragmentation is disabled, trying to send a message that exceeds the maximum fragmentation size will produce an error.
- CONFIG\_PMTUD.SCTP:  
Pass 1 primitive / event: 'Configure PathMTU Discovery'  
Parameters: one boolean value (PMTUD on/off), PMTU value (optional)  
Returns: PMTU value  
Comments: This returns a meaningful PMTU value when PMTUD is enabled (the boolean is true), and the PMTU value can be set if PMTUD is disabled (the boolean is false)
- CONFIG\_DELAYED\_SACK.SCTP:  
Pass 1 primitive / event: 'Configure delayed SACK timer'  
Parameters: one boolean value (delayed SACK on/off), timer value (optional), number of packets to wait for (default 2)  
Comments: If delayed SACK is enabled, SCTP will send a SACK upon either receiving the provided number of packets or when the timer expires, whatever occurs first.
- CONFIG\_RTO.SCTP:  
Pass 1 primitive / event: 'Configure RTO calculation'  
Parameters: init (optional), min (optional), max (optional)  
Comments: This adjusts the initial, minimum and maximum RTO values.
- SET\_COOKIE\_LIFE.SCTP:  
Pass 1 primitive / event: 'Set Cookie life value'  
Parameters: cookie life value
- SET\_MAX\_BURST.SCTP:  
Pass 1 primitive / event: 'Set maximum burst'  
Parameters: max burst value

Comments: not all implementations allow values above the default of 4.

- o SET\_PARTIAL\_DELIVERY\_POINT.SCTP:  
Pass 1 primitive / event: 'Set Partial Delivery Point'  
Parameters: partial delivery point (integer)  
Comments: this parameter must be smaller or equal to the socket receive buffer size.
- o CHECKSUM.UDP:  
Pass 1 primitive / event: 'DISABLE\_CHECKSUM'.  
Parameters: 0 when no checksum is used at sender, 1 for checksum at sender (default)
- o CHECKSUM\_REQUIRED.UDP:  
Pass 1 primitive / event: 'REQUIRE\_CHECKSUM'.  
Parameter: 0 when checksum is required at receiver, 1 to allow zero checksum at receiver (default)
- o SET\_CHECKSUM\_COVERAGE.UDP-Lite:  
Pass 1 primitive / event: 'SET\_CHECKSUM\_COVERAGE'  
Parameters: Coverage length at sender (default maximum coverage)
- o SET\_MIN\_CHECKSUM\_COVERAGE.UDP-Lite:  
Pass 1 primitive / event: 'SET\_MIN\_COVERAGE'.  
Parameter: Coverage length at receiver (default minimum coverage)
- o SET\_DF.UDP(-Lite):  
Pass 1 primitive event: 'SET\_DF'.  
Parameter: 0 when DF is not set (default), 1 when DF is set
- o SET\_TTL.UDP(-Lite) (IPV6\_UNICAST\_HOPS):  
Pass 1 primitive / event: 'SET\_TTL' and 'SET\_IPV6\_UNICAST\_HOPS'  
Parameters: IPv4 TTL value or IPv6 Hop Count value  
Comments: This allows an application to change the IPv4 TTL of IPv6 Hop count value for outgoing UDP(-Lite) datagrams.
- o GET\_TTL.UDP(-Lite) (IPV6\_UNICAST\_HOPS):  
Pass 1 primitive / event: 'GET\_TTL' and 'GET\_IPV6\_UNICAST\_HOPS'  
Returns: IPv4 TTL value or IPv6 Hop Count value  
Comments: This allows an application to read the the IPv4 TTL of IPv6 Hop count value from a received UDP(-Lite) datagram.
- o SET\_ECN.UDP(-Lite):  
Pass 1 primitive / event: 'SET\_ECN'  
Parameters: ECN value  
Comments: This allows a UDP(-Lite) application to set the ECN codepoint field for outgoing UDP(-Lite) datagrams.

- GET\_ECN.UDP(-Lite):  
Pass 1 primitive / event: 'GET\_ECN'  
Parameters: ECN value  
Comments: This allows a UDP(-Lite) application to read the ECN codepoint field from a received UDP(-Lite) datagram.
- SET\_IP\_OPTIONS.UDP(-Lite):  
Pass 1 primitive / event: 'SET\_IP\_OPTIONS'  
Parameters: options  
Comments: This allows a UDP(-Lite) application to set IP Options for outgoing UDP(-Lite) datagrams. These options can at least be the Source Route, Record Route, and Time Stamp option.
- GET\_IP\_OPTIONS.UDP(-Lite):  
Pass 1 primitive / event: 'GET\_IP\_OPTIONS'  
Returns: options  
Comments: This allows a UDP(-Lite) application to receive any IP options that are contained in a received UDP(-Lite) datagram.
- CONFIGURE.LEDBAT:  
Pass 1 primitive / event: N/A  
Parameters: enable (boolean), TARGET, ALLOWED\_INCREASE, GAIN\_INC, GAIN\_DEC, BASE\_HISTORY, CURRENT\_FILTER, INIT\_CWND, MIN\_CWND  
Comments: enable is a newly invented parameter that enables or disables the whole LEDBAT service.

**TERMINATION:**

Gracefully or forcefully closing a connection, or being informed about this event happening.

- CLOSE.TCP:  
Pass 1 primitive / event: 'close'  
Comments: this terminates the sending side of a connection after reliably delivering all remaining data.
- CLOSE.SCTP:  
Pass 1 primitive / event: 'Shutdown'  
Comments: this terminates a connection after reliably delivering all remaining data.
- CLOSE.UDP(-Lite):  
Pass 1 primitive event: 'CLOSE'  
Comments: No further UDP(-Lite) datagrams are sent/received on this connection.

- ABORT.TCP:  
Pass 1 primitive / event: 'abort'  
Comments: this terminates a connection without delivering remaining data and sends an error message to the other side.
- ABORT.SCTP:  
Pass 1 primitive / event: 'abort'  
Parameters: abort reason to be given to the peer (optional)  
Comments: this terminates a connection without delivering remaining data and sends an error message to the other side.
- TIMEOUT.TCP:  
Pass 1 primitive / event: 'USER TIMEOUT' event  
Comments: the application is informed that the connection is aborted. This event is executed on expiration of the timeout set in CONNECTION.ESTABLISHMENT.CONNECT.TCP (possibly adjusted in CONNECTION.MAINTENANCE.CHANGE\_TIMEOUT.TCP).
- TIMEOUT.SCTP:  
Pass 1 primitive / event: 'COMMUNICATION LOST' event  
Comments: the application is informed that the connection is aborted. this event is executed on expiration of the timeout that should be enabled by default (see beginning of section 8.3 in [RFC4960]) and was possibly adjusted in CONNECTION.MAINTENANCE.CHANGE\_TIMEOUT.SCTP.
- ABORT-EVENT.TCP:  
Pass 1 primitive / event: not specified.
- ABORT-EVENT.SCTP:  
Pass 1 primitive / event: 'COMMUNICATION LOST' event  
Returns: abort reason from the peer (if available)  
Comments: the application is informed that the other side has aborted the connection using CONNECTION.TERMINATION.ABORT.SCTP.
- CLOSE-EVENT.TCP:  
Pass 1 primitive / event: not specified.
- CLOSE-EVENT.SCTP:  
Pass 1 primitive / event: 'SHUTDOWN COMPLETE' event  
Comments: the application is informed that CONNECTION.TERMINATION.CLOSE.SCTP was successfully completed.

#### 4.2. DATA Transfer Related Primitives

All primitives in this section refer to an existing connection, i.e. a connection that was either established or made available for

receiving data (although this is optional for the primitives of UDP(-Lite)). In addition to the listed parameters, all sending primitives contain a reference to a data block and all receiving primitives contain a reference to available buffer space for the data. Note that CONNECT.TCP and LISTEN.TCP in the ESTABLISHMENT and AVAILABILITY category also allow to transfer data (an optional user message) before the connection is fully established.

- o SEND.TCP:  
Pass 1 primitive / event: 'send'  
Parameters: timeout (optional)  
Comments: this gives TCP a data block for reliable transmission to the TCP on the other side of the connection. The timeout can be configured with this call whenever data are sent (see also CONNECTION.MAINTENANCE.CHANGE-TIMEOUT.TCP).
- o SEND.SCTP:  
Pass 1 primitive / event: 'Send'  
Parameters: stream number; context (optional); socket (optional); unordered flag (optional); no-bundle flag (optional); payload protocol-id (optional); pr-policy (optional) pr-value (optional); sack-immediately flag (optional); key-id (optional)  
Comments: this gives SCTP a data block for transmission to the SCTP on the other side of the connection (SCTP association). The 'stream number' denotes the stream to be used. The 'context' number can later be used to refer to the correct message when an error is reported. The 'socket' can be used to state which path should be preferred, if there are multiple paths available (see also CONNECTION.MAINTENANCE.SETPRIMARY.SCTP). The data block can be delivered out-of-order if the 'unordered flag' is set. The 'no-bundle flag' can be set to indicate a preference to avoid bundling. The 'payload protocol-id' is a number that will, if provided, be handed over to the receiving application. Using pr-policy and pr-value the level of reliability can be controlled. The 'sack-immediately' flag can be used to indicate that the peer should not delay the sending of a SACK corresponding to the provided user message. If specified, the provided key-id is used for authenticating the user message.
- o SEND.UDP(-Lite):  
Pass 1 primitive / event: 'SEND'  
Parameters: IP Address and Port Number of the destination endpoint (optional if connected).  
Comments: This provides a message for unreliable transmission using UDP(-Lite) to the specified transport address. IP address and Port may be omitted for connected UDP(-Lite) sockets. All CONNECTION.MAINTENANCE.SET\_\*.UDP(-Lite) primitives apply per message sent.

- RECEIVE.TCP:  
Pass 1 primitive / event: 'receive'.
- RECEIVE.SCTP:  
Pass 1 primitive / event: 'DATA ARRIVE' notification, followed by 'Receive'  
Parameters: stream number (optional)  
Returns: stream sequence number (optional), partial flag (optional)  
Comments: if the 'stream number' is provided, the call to receive only receives data on one particular stream. If a partial message arrives, this is indicated by the 'partial flag', and then the 'stream sequence number' must be provided such that an application can restore the correct order of data blocks that comprise an entire message. Additionally, a delivery number lets the application detect reordering.
- RECEIVE.UDP(-Lite):  
Pass 1 primitive / event: 'RECEIVE',  
Parameters: Buffer for received datagram.  
Comments: All CONNECTION.MAINTENANCE.GET\_\*.UDP(-Lite) primitives apply per message received.
- SENDFAILURE-EVENT.SCTP:  
Pass 1 primitive / event: 'SEND FAILURE' notification, optionally followed by 'Receive Unsent Message' or 'Receive Unacknowledged Message'  
Returns: cause code; context; unsent or unacknowledged message (optional)  
Comments: 'cause code' indicates the reason of the failure, and 'context' is the context number if such a number has been provided in DATA.SEND.SCTP, for later use with 'Receive Unsent Message' or 'Receive Unacknowledged Message', respectively. These primitives can be used to retrieve the unsent or unacknowledged message (or part of the message, in case a part was delivered) if desired.
- SEND\_FAILURE.UDP(-Lite):  
Pass 1 primitive / event: 'SEND'  
Comments: This may be used to probe for the effective PMTU when using in combination with the 'MAINTENANCE.SET\_DF' primitive.
- SENDER\_DRY-EVENT.SCTP:  
Pass 1 primitive / event: 'SENDER DRY' notification  
Comments: This informs the application that the stack has no more user data to send.

- o PARTIAL\_DELIVERY\_ABORTED-EVENT.SCTP:  
Pass 1 primitive / event: 'PARTIAL DELIVERY ABORTED' notification  
Comments: This informs the receiver of a partial message that the further delivery of the message has been aborted.

## 5. Pass 3

This section presents the superset of all Transport Features in all protocols that were discussed in the preceding sections, based on the list of primitives in pass 2 but also on text in pass 1 to include features that can be configured in one protocol and are static properties in another (congestion control, for example). Again, some minor details are omitted for the sake of generalization -- e.g., TCP may provide various different IP options, but only source route is mandatory to implement, and this detail is not visible in the Pass 3 feature "Specify IP Options".

### 5.1. CONNECTION Related Transport Features

#### ESTABLISHMENT:

Active creation of a connection from one transport endpoint to one or more transport endpoints.

- o Connect  
Protocols: TCP, SCTP, UDP(-Lite)
- o Specify which IP Options must always be used  
Protocols: TCP
- o Request multiple streams  
Protocols: SCTP
- o Limit the number of inbound streams  
Protocols: SCTP
- o Specify number of attempts and/or timeout for the first establishment message  
Protocols: TCP, SCTP
- o Obtain multiple sockets  
Protocols: SCTP
- o Disable MPTCP  
Protocols: MPTCP

- o Specify which chunk types must always be authenticated  
Protocols: SCTP  
Comments: DATA, ACK etc. are different 'chunks' in SCTP; one or more chunks may be included in a single packet.
- o Indicate an Adaptation Layer (via an adaptation code point)  
Protocols: SCTP
- o Request to negotiate interleaving of user messages  
Protocols: SCTP
- o Hand over a message to transfer (possibly multiple times) before connection establishment  
Protocols: TCP
- o Hand over a message to transfer during connection establishment  
Protocols: SCTP
- o Enable UDP encapsulation with a specified remote UDP port number  
Protocols: SCTP

**AVAILABILITY:**

Preparing to receive incoming connection requests.

- o Listen, 1 specified local interface  
Protocols: TCP, SCTP, UDP(-Lite)
- o Listen, N specified local interfaces  
Protocols: SCTP, UDP(-Lite)
- o Listen, all local interfaces  
Protocols: TCP, SCTP, UDP(-Lite)
- o Obtain requested number of streams  
Protocols: SCTP
- o Limit the number of inbound streams  
Protocols: SCTP
- o Specify which IP Options must always be used  
Protocols: TCP
- o Disable MPTCP  
Protocols: MPTCP

- o Specify which chunk types must always be authenticated  
Protocols: SCTP  
Comments: DATA, ACK etc. are different 'chunks' in SCTP; one or more chunks may be included in a single packet.
- o Indicate an Adaptation Layer (via an adaptation code point)  
Protocols: SCTP

**MAINTENANCE:**

Adjustments made to an open connection, or notifications about it.

- o Change timeout for aborting connection (using retransmit limit or time value)  
Protocols: TCP, SCTP
- o Suggest timeout to the peer  
Protocols: TCP
- o Disable Nagle algorithm  
Protocols: TCP, SCTP
- o Request an immediate heartbeat, returning success/failure  
Protocols: SCTP
- o Notification of Excessive Retransmissions (early warning below abortion threshold)  
Protocols: TCP
- o Add path  
Protocols: MPTCP, SCTP  
MPTCP Parameters: source-IP; source-Port; destination-IP; destination-Port  
SCTP Parameters: local IP address
- o Remove path  
Protocols: MPTCP, SCTP  
MPTCP Parameters: source-IP; source-Port; destination-IP; destination-Port  
SCTP Parameters: local IP address
- o Set primary path  
Protocols: SCTP
- o Suggest primary path to the peer  
Protocols: SCTP

- o Configure Path Switchover  
Protocols: SCTP
- o 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 sizes; 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)
- o Specify DSCP field  
Protocols: TCP, SCTP, UDP(-Lite)
- o Notification of ICMP error message arrival  
Protocols: TCP, UDP(-Lite)
- o Change authentication parameters  
Protocols: SCTP
- o Obtain authentication information  
Protocols: SCTP
- o Reset Stream  
Protocols: SCTP
- o Notification of Stream Reset  
Protocols: STCP
- o Reset Association  
Protocols: SCTP
- o Notification of Association Reset  
Protocols: STCP
- o Add Streams  
Protocols: SCTP
- o Notification of Added Stream  
Protocols: STCP

- o Choose a scheduler to operate between streams of an association  
Protocols: SCTP
- o Configure priority or weight for a scheduler  
Protocols: SCTP
- o Specify IPv6 flow label field  
Protocols: SCTP
- o Configure send buffer size  
Protocols: SCTP
- o Configure receive buffer (and rwnd) size  
Protocols: SCTP
- o Configure message fragmentation  
Protocols: SCTP
- o Configure PMTUD  
Protocols: SCTP
- o Configure delayed SACK timer  
Protocols: SCTP
- o Set Cookie life value  
Protocols: SCTP
- o Set maximum burst  
Protocols: SCTP
- o Configure size where messages are broken up for partial delivery  
Protocols: SCTP
- o Disable checksum when sending  
Protocols: UDP
- o Disable checksum requirement when receiving  
Protocols: UDP
- o Specify checksum coverage used by the sender  
Protocols: UDP-Lite
- o Specify minimum checksum coverage required by receiver  
Protocols: UDP-Lite
- o Specify DF field  
Protocols: UDP(-Lite)

- o Specify TTL/Hop count field  
Protocols: UDP(-Lite)
- o Obtain TTL/Hop count field  
Protocols: UDP(-Lite)
- o Specify ECN field  
Protocols: UDP(-Lite)
- o Obtain ECN field  
Protocols: UDP(-Lite)
- o Specify IP Options  
Protocols: UDP(-Lite)
- o Obtain IP Options  
Protocols: UDP(-Lite)
- o Enable and configure "Low Extra Delay Background Transfer"  
Protocols: A protocol implementing the LEDBAT congestion control mechanism

**TERMINATION:**

Gracefully or forcefully closing a connection, or being informed about this event happening.

- o Close after reliably delivering all remaining data, causing an event informing the application on the other side  
Protocols: TCP, SCTP  
Comments: A TCP endpoint locally only closes the connection for sending; it may still receive data afterwards.
- o Abort without delivering remaining data, causing an event informing the application on the other side  
Protocols: TCP, SCTP  
Comments: In SCTP a reason can optionally be given by the application on the aborting side, which can then be received by the application on the other side.
- o Timeout event when data could not be delivered for too long  
Protocols: TCP, SCTP  
Comments: the timeout is configured with CONNECTION.MAINTENANCE "Change timeout for aborting connection (using retransmit limit or time value)".

## 5.2. DATA Transfer Related Transport Features

All features in this section refer to an existing connection, i.e. a connection that was either established or made available for receiving data. Note that TCP allows to transfer data (a single optional user message, possibly arriving multiple times) before the connection is fully established. Reliable data transfer entails delay -- e.g. for the sender to wait until it can transmit data, or due to retransmission in case of packet loss.

### 5.2.1. Sending Data

All features in this section are provided by DATA.SEND from pass 2. DATA.SEND is given a data block from the application, which we here call a "message" if the beginning and end of the data block can be identified at the receiver, and "data" otherwise.

- o Reliably transfer data, with congestion control  
Protocols: TCP
- o Reliably transfer a message, with congestion control  
Protocols: SCTP
- o Unreliably transfer a message, with congestion control  
Protocols: SCTP
- o Unreliably transfer a message, without congestion control  
Protocols: UDP(-Lite)
- o Configurable Message Reliability  
Protocols: SCTP
- o Choice of stream  
Protocols: SCTP
- o Choice of path (destination address)  
Protocols: SCTP
- o Choice between unordered (potentially faster) or ordered delivery of messages  
Protocols: SCTP
- o Request not to bundle messages  
Protocols: SCTP
- o Specifying a "payload protocol-id" (handed over as such by the receiver)  
Protocols: SCTP

- o Specifying a key id to be used to authenticate a message  
Protocols: SCTP
- o Request not to delay the acknowledgement (SACK) of a message  
Protocols: SCTP

### 5.2.2. Receiving Data

All features in this section are provided by DATA.RECEIVE from pass 2. DATA.RECEIVE fills a buffer provided by the application, with what we here call a "message" if the beginning and end of the data block can be identified at the receiver, and "data" otherwise.

- o Receive data (with no message delineation)  
Protocols: TCP
- o Receive a message  
Protocols: SCTP, UDP(-Lite)
- o Choice of stream to receive from  
Protocols: SCTP
- o Information about partial message arrival  
Protocols: SCTP  
Comments: In SCTP, partial messages are combined with a stream sequence number so that the application can restore the correct order of data blocks an entire message consists of.
- o Obtain a message delivery number  
Protocols: SCTP  
Comments: This number can let applications detect and, if desired, correct reordering.

### 5.2.3. Errors

This section describes sending failures that are associated with a specific call to DATA.SEND from pass 2.

- o Notification of an unsent (part of a) message  
Protocols: SCTP, UDP(-Lite)
- o Notification of an unacknowledged (part of a) message  
Protocols: SCTP

- o Notification that the stack has no more user data to send  
Protocols: SCTP
- o Notification to a receiver that a partial message delivery has  
been aborted  
Protocols: SCTP

## 6. Acknowledgements

The authors would like to thank (in alphabetical order) Bob Briscoe, Gorry Fairhurst, David Hayes, Tom Jones, Karen Nielsen, Joe Touch and Brian Trammell for providing valuable feedback on this document. We especially thank Christoph Paasch for providing input related to Multipath TCP. 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).

## 7. IANA Considerations

XX RFC ED - PLEASE REMOVE THIS SECTION XXX

This memo includes no request to IANA.

## 8. 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 these features on its own. Therefore, these features are not considered in this document, with the exception of native authentication capabilities of SCTP for which the security considerations in [RFC4895] apply.

## 9. References

### 9.1. Normative References

- [FJ16] Fairhurst, G. and T. Jones, "Features of the User Datagram Protocol (UDP) and Lightweight UDP (UDP-Lite) Transport Protocols", draft-ietf-taps-transport-usage-udp-00 (work in progress), November 2016.

- [I-D.ietf-tsvwg-sctp-ndata]  
Stewart, R., Tuexen, M., Loreto, S., and R. Seggelmann,  
"Stream Schedulers and User Message Interleaving for the  
Stream Control Transmission Protocol",  
draft-ietf-tsvwg-sctp-ndata-08 (work in progress),  
October 2016.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7,  
RFC 793, DOI 10.17487/RFC0793, September 1981,  
<<http://www.rfc-editor.org/info/rfc793>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts -  
Communication Layers", STD 3, RFC 1122, DOI 10.17487/  
RFC1122, October 1989,  
<<http://www.rfc-editor.org/info/rfc1122>>.
- [RFC3758] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P.  
Conrad, "Stream Control Transmission Protocol (SCTP)  
Partial Reliability Extension", RFC 3758, DOI 10.17487/  
RFC3758, May 2004,  
<<http://www.rfc-editor.org/info/rfc3758>>.
- [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>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol",  
RFC 4960, DOI 10.17487/RFC4960, September 2007,  
<<http://www.rfc-editor.org/info/rfc4960>>.
- [RFC5061] Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M.  
Kozuka, "Stream Control Transmission Protocol (SCTP)  
Dynamic Address Reconfiguration", RFC 5061, DOI 10.17487/  
RFC5061, September 2007,  
<<http://www.rfc-editor.org/info/rfc5061>>.
- [RFC5482] Eggert, L. and F. Gont, "TCP User Timeout Option",  
RFC 5482, DOI 10.17487/RFC5482, March 2009,  
<<http://www.rfc-editor.org/info/rfc5482>>.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J.  
Iyengar, "Architectural Guidelines for Multipath TCP  
Development", RFC 6182, DOI 10.17487/RFC6182, March 2011,  
<<http://www.rfc-editor.org/info/rfc6182>>.
- [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>>.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, DOI 10.17487/RFC6817, December 2012, <<http://www.rfc-editor.org/info/rfc6817>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.
- [RFC6897] Scharf, M. and A. Ford, "Multipath TCP (MPTCP) Application Interface Considerations", RFC 6897, DOI 10.17487/RFC6897, March 2013, <<http://www.rfc-editor.org/info/rfc6897>>.
- [RFC6951] Tuexen, M. and R. Stewart, "UDP Encapsulation of Stream Control Transmission Protocol (SCTP) Packets for End-Host to End-Host Communication", RFC 6951, DOI 10.17487/RFC6951, May 2013, <<http://www.rfc-editor.org/info/rfc6951>>.
- [RFC7053] Tuexen, M., Ruengeler, I., and R. Stewart, "SACK-IMMEDIATELY Extension for the Stream Control Transmission Protocol", RFC 7053, DOI 10.17487/RFC7053, November 2013, <<http://www.rfc-editor.org/info/rfc7053>>.
- [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>>.
- [RFC7496] Tuexen, M., Seggelmann, R., Stewart, R., and S. Loreto, "Additional Policies for the Partially Reliable Stream Control Transmission Protocol Extension", RFC 7496, DOI 10.17487/RFC7496, April 2015, <<http://www.rfc-editor.org/info/rfc7496>>.
- [RFC7829] Nishida, Y., Natarajan, P., Caro, A., Amer, P., and K. Nielsen, "SCTP-PF: A Quick Failover Algorithm for the Stream Control Transmission Protocol", RFC 7829, DOI 10.17487/RFC7829, April 2016,

<<http://www.rfc-editor.org/info/rfc7829>>.

- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<http://www.rfc-editor.org/info/rfc8085>>.

## 9.2. Informative References

- [RFC0854] Postel, J. and J. Reynolds, "Telnet Protocol Specification", STD 8, RFC 854, DOI 10.17487/RFC0854, May 1983, <<http://www.rfc-editor.org/info/rfc854>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC3260] Grossman, D., "New Terminology and Clarifications for Diffserv", RFC 3260, DOI 10.17487/RFC3260, April 2002, <<http://www.rfc-editor.org/info/rfc3260>>.
- [RFC5461] Gont, F., "TCP's Reaction to Soft Errors", RFC 5461, DOI 10.17487/RFC5461, February 2009, <<http://www.rfc-editor.org/info/rfc5461>>.
- [RFC6093] Gont, F. and A. Yourtchenko, "On the Implementation of the TCP Urgent Mechanism", RFC 6093, DOI 10.17487/RFC6093, January 2011, <<http://www.rfc-editor.org/info/rfc6093>>.
- [RFC7414] Duke, M., Braden, R., Eddy, W., Blanton, E., and A. Zimmermann, "A Roadmap for Transmission Control Protocol (TCP) Specification Documents", RFC 7414, DOI 10.17487/RFC7414, February 2015, <<http://www.rfc-editor.org/info/rfc7414>>.
- [RFC7657] Black, D., Ed. and P. Jones, "Differentiated Services (Diffserv) and Real-Time Communication", RFC 7657, DOI 10.17487/RFC7657, November 2015, <<http://www.rfc-editor.org/info/rfc7657>>.
- [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>>.

#### Appendix A. Overview of RFCs used as input for pass 1

TCP: [RFC0793], [RFC1122], [RFC5482], [RFC7413]  
MPTCP: [RFC6182], [RFC6824], [RFC6897]  
SCTP: RFCs without a socket API specification: [RFC3758], [RFC4895],  
[RFC4960], [RFC5061]. RFCs that include a socket API  
specification: [RFC6458], [RFC6525], [RFC6951], [RFC7053],  
[RFC7496] [RFC7829].  
UDP(-Lite): See [FJ16]  
LEDBAT: [RFC6817].

#### Appendix B. How this document was developed

This section gives an overview of the method that was used to develop this document. It was given to contributors for guidance, and it can be helpful for future updates or extensions.

This document is only concerned with Transport Features that are explicitly exposed to applications via primitives. It also strictly follows RFC text: if a feature is truly relevant for an application, the RFCs should say so, and they should describe how to use and configure it. Thus, the approach followed for developing this document was to identify the right RFCs, then analyze and process their text.

Primitives that MAY be implemented by a transport protocol were excluded. To be included, the minimum requirement level for a primitive to be implemented by a protocol was SHOULD. Where [RFC2119]-style requirements levels are not used, primitives were excluded when they are described in conjunction with statements like, e.g.: "some implementations also provide" or "an implementation may also". Excluded primitives or parameters were briefly described in a dedicated subsection.

Pass 1: This began by identifying text that talks about primitives. An API specification, abstract or not, obviously describes primitives -- but we are not *\*only\** interested in API specifications. The text describing the 'send' primitive in the API specified in [RFC0793], for instance, does not say that data transfer is reliable. TCP's reliability is clear, however, from this text in Section 1 of [RFC0793]: "The Transmission Control Protocol (TCP) is intended for use as a highly reliable host-to-host protocol between hosts in packet-switched computer communication networks, and in

interconnected systems of such networks."

Some text for pass 1 subsections was developed copy+pasting all the relevant text parts from the relevant RFCs, then adjusting terminology to match the terminology in Section 1 and adjusting (shortening!) phrasing to match the general style of the document. An effort was made to formulate everything as a primitive description such that the primitive descriptions became as complete as possible (e.g., the "SEND.TCP" primitive in pass 2 is explicitly described as reliably transferring data); text that is relevant for the primitives presented in this pass but still does not fit directly under any primitive was used in a subsection's introduction.

Pass 2: The main goal of this pass is unification of primitives. As input, only text from pass 1 was used (no exterior sources). The list in pass 2 is not arranged by protocol ("first protocol X, here are all the primitives; then protocol Y, here are all the primitives, ..") but by primitive ("primitive A, implemented this way in protocol X, this way in protocol Y, ..."). It was a goal to obtain as many similar pass 2 primitives as possible. For instance, this was sometimes achieved by not always maintaining a 1:1 mapping between pass 1 and pass 2 primitives, renaming primitives etc. For every new primitive, the already existing primitives were considered to try to make them as coherent as possible.

For each primitive, the following style was used:

- o PRIMITIVENAME.PROTOCOL:  
Pass 1 primitive / event:  
Parameters:  
Returns:  
Comments:

The entries "Parameters", "Returns" and "Comments" were skipped when a primitive had no parameters, no described return value or no comments seemed necessary, respectively. Optional parameters are followed by "(optional)". When a default value is known, this was also provided.

Pass 3: the main point of this pass is to identify transport protocol features that are the result of static properties of protocols, for which all protocols have to be listed together; this is then the final list of all available Transport Features. This list was primarily based on text from pass 2, with additional input from pass 1 (but no external sources).

## Appendix C. Revision information

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

-00 (from draft-welzl-taps-transport): this now covers TCP based on all TCP RFCs (this means: if you know of something in any TCP RFC that you think should be addressed, please speak up!) as well as SCTP, exclusively based on [RFC4960]. We decided to also incorporate [RFC6458] for SCTP, but this hasn't happened yet. Terminology made in line with [RFC8095]. Addressed comments by Karen Nielsen and Gorry Fairhurst; various other fixes. Appendices (TCP overview and how-to-contribute) added.

-01: this now also covers MPTCP based on [RFC6182], [RFC6824] and [RFC6897].

-02: included UDP, UDP-Lite, and all extensions of SCTPs. This includes fixing the [RFC6458] omission from -00.

-03: wrote security considerations. The "how to contribute" section was updated to reflect how the document WAS created, not how it SHOULD BE created; it also no longer wrongly says that Experimental RFCs are excluded. Included LEDBAT. Changed abstract and intro to reflect which protocols/mechanisms are covered (TCP, MPTCP, SCTP, UDP, UDP-Lite, LEDBAT) instead of talking about "transport protocols". Interleaving and stream scheduling added (draft-ietf-tsvwg-sctp-ndata). TFO added. "Set protocol parameters" in SCTP replaced with per-parameter (or parameter group) primitives. More primitives added, mostly previously overlooked ones from [RFC6458]. Updated terminology (s/transport service feature/transport feature) in line with an update of [RFC8095]. Made sequence of transport features / primitives more logical. Combined MPTCP's add/rem subflow with SCTP's add/remove local address.

## Authors' Addresses

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

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

Michael Tuexen  
Muenster University of Applied Sciences  
Stegerwaldstrasse 39  
Steinfurt 48565  
Germany

Email: [tuexen@fh-muenster.de](mailto:tuexen@fh-muenster.de)

Naeem Khademi  
University of Oslo  
PO Box 1080 Blindern  
Oslo, N-0316  
Norway

Email: [naeemk@ifi.uio.no](mailto:naeemk@ifi.uio.no)

