

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 6 May 2021

H. Asai
Preferred Networks
2 November 2020

Separation of Data Path and Data Flow Sublayers in the Transport Layer
draft-asai-tsvwg-transport-review-00

Abstract

This document reviews the architectural design of the transport layer. In particular, this document separates the transport layer into two sublayers; the data path and the data flow layers. The data path layer provides functionality on the data path, such as connection handling, path quality and trajectory monitoring, waypoint management, and congestion control. The data flow layer provides additional functionality upon the data path layer, such as flow control for the receive buffer management, retransmission for reliable data delivery, and transport layer security. The data path layer multiplexes multiple data flow layer protocols and provides data path information to the data flow layer to control data transmissions, such as prioritization and inverse multiplexing for multipath protocols.

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 <https://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 6 May 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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. Introduction	3
2. Transport Layer Functionality Review	3
2.1. Conventional Layering	3
2.2. Data Path-aware Networking	4
2.3. Resource Management: Flow Control and Congestion Control	5
2.4. Multipath Protocols	6
2.5. Reliable Data Communication	6
2.6. Security	7
2.7. Summary	7
3. Specifications of Data Path Layer and Data Flow Layer	8
3.1. Data Path Layer	8
3.1.1. In-band trajectory monitoring	8
3.1.2. Waypoint management	9
3.1.3. Bidirectional connection establishment	9
3.1.4. Data path quality (congestion) monitoring and congestion control	10
3.1.5. Data flow multiplexing	10
3.1.6. Packet Duplication	10
3.2. Data Flow Layer	10
3.2.1. Retransmission for reliable data communication	11
3.2.2. Flow control for receive-buffer management	11
3.2.3. Flow prioritization	11
3.2.4. End-to-end security	11
3.2.5. Inverse multiplexing for multipath protocols	11
4. Use Cases	12
4.1. Multipath Transport Protocols	12
4.2. Congestion Control Acceleration	13
4.3. In-Network Computing	14
4.4. Flow Arbitration	15
5. Implementation Considerations	16
6. IANA Considerations	16
7. Security Considerations	16
8. Acknowledgements	17
9. References	17
9.1. Normative References	17
9.2. Informative References	19

Author's Address 19

1. Introduction

This document specifies two sublayers of the transport layer; the data path and the data flow layers. In this document, the transport layer's data path functionality, such as bidirectional connection handling, waypoint management, and congestion control, is separated from the data flow functionality, such as flow control for the receive buffer management, retransmission for reliable data delivery, and transport layer security. This document reviews the transport layer's functionality from the viewpoint of data paths and data flows in Section 2. It then specifies the data path layer and the data flow layer in Section 3. The data path and data flow layers provide the data path and the data flow functionality, respectively.

This document reviews the transport layer to clarify the transport layer's functionality and to invent data flow layer protocols for advanced Internet technologies, such as middleboxes and in-network computing. Hence, this document does not intend to obsolete the transport layer or violate the current Internet architecture.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Transport Layer Functionality Review

This section reviews the conventional layering and transport layer functionality. It then summarizes the transport layer functionality by distinguishing data paths and data flows.

2.1. Conventional Layering

RFC 1122 [RFC1122] defines three communication layers, link layer, Internet layer, transport layer, and the interfaces between these layers. Link-layer protocols provide hop-by-hop data communications. Internet layer protocols such as the Internet Protocol (IP), the Internet Control Message Protocol (ICMP), and the Internet Group Management Protocol (IGMP) provide fragmentation, hop-by-hop datagram forwarding, and end-to-end datagram delivery. RFC 1123 [RFC1123] defines the interface between the application layer and the transport layer.

The Internet design follows the end-to-end principle to keep the simplicity of the Internet layer. Thus, the main functionality of the Internet layer is to provide end-to-end reachability. It does not guarantee datagram integrity, which means that packet loss,

duplication, corruption, or reordering may occur. Over the Internet layer, the transport layer implements such functionality to provide datagram integrity on the end-hosts to achieve end-to-end communications.




Figure 1: Conventional layering of Internet architecture

The transport layer provides various functions over the IP. User Datagram Protocol (UDP) [RFC0768] and Transmission Control Protocol (TCP) [RFC0793] are the most commonly used transport layer protocols. UDP is a connectionless transport protocol with a minimum header. TCP implements flow control according to the receiver's buffer capacity, retransmission for reliable communication, congestion control by a packet delivery status such as packet loss and delay. The transport layer may implement an end-to-end security function, Transport Layer Security (TLS) [RFC8446].

Figure 1 illustrates the conventional layering of Internet architecture. A router forwards an IP datagram to a next-hop router corresponding to the destination address. In this way, the Internet layer protocol, IP, provides end-to-end reachability through hop-by-hop routing. The transport layer provides additional functions for end-to-end communications.

2.2. Data Path-aware Networking

As described above, the Internet layer's main functionality has been end-to-end reachability with hop-by-hop packet forwarding based on destination IP address. Therefore, it is unaware of data paths, i.e., trajectories of packets. Best-effort communications over ``dumb'' networks without the quality of service (QoS) have been the Internet principle [RFC2768]. This end-to-end principle of the Internet has achieved a scalable architecture. However, computer networks' advancement introduced QoS and middleboxes to achieve high-quality data communication and optimize communication over distributed and heterogeneous computer networks. These technologies require to be aware of data paths associated with data flows.

The Internet layer has been extended to support QoS. Differentiated Services, DiffServ [RFC2474], is one of the technologies to implement QoS. It enables autonomous and scalable service discrimination using an IP header field, differentiated services codepoint (DSCP), to implement QoS for a data path. Segment Routing [RFC8402] enables data path configuration for individual flows by leveraging the source routing paradigm. Thus, Segment Routing is another technology to treat QoS.

Middleboxes are more complex than QoS. They add various functions such as firewalls, TCP offloading, transcoding, and content caches to a data path. These functions require to be aware of waypoints to be activated. Routing technologies with packet classification using the IP and the transport protocol headers have collectively been leveraged to route traffic through the middleboxes as the IP layer is not aware of data paths. If a middlebox is transparent to end-hosts, it should be installed at the network gateway or redirected by policy-based routing or segment routing to activate the function. Otherwise, an end-host should specify the middlebox as a target end-host. In addition to these middleboxes, distributed computing technologies, such as in-network computing and multi-access edge computing (MEC), also require to be aware of data paths.

In summary, advanced computer networks require to treat data paths for QoS, middleboxes, and distributed computing. Packet classification for data path treatment may use the header information of transport layer protocols such as port numbers as well as IP header fields.

2.3. Resource Management: Flow Control and Congestion Control

As the Internet layer does not provide resource management functionality, transport layer protocols may implement it, such as flow control and congestion control. Both flow control and congestion control mechanisms control packet transmission for resource management. However, the target resource is different. They control packet transmission according to the receiver's buffer capacity and the network bandwidth capacity, respectively.

For example, in TCP's flow control, a receiver announces the remaining receive buffer size as the window size. Hence, flow control is not aware of network bandwidth capacity. On the other hand, congestion control is performed based on data communication quality information, such as packet loss and delay. The underlying Internet layer may provide congestion information by the Explicit Congestion Notification (ECN) [RFC3168]. In this manner, transport layer protocols control congestion depending on the data path's network resources. Therefore, congestion control is associated with a data path, while flow control is associated with the end-hosts.

As discussed above, congestion control should be performed on the associated data path. However, in current transport layer protocols such as TCP, Datagram Congestion Control Protocol (DCCP) [RFC4340], and Stream Control Transmission Protocol (SCTP) [RFC4960], an individual flow independently performs congestion control even if the same data path multiplexes multiple flows. Therefore, multiple flows cannot collectively perform congestion control for the data path. For example, when a data path multiplexes a TCP and a UDP flows, the TCP flow's congestion control may affect the other UDP flow's quality. Moreover, transport layer protocols utilize network resources on a fair basis. Thus, flow prioritization cannot be implemented at the transport layer, although a QoS technology such as DiffServ may be leveraged.

2.4. Multipath Protocols

Multipath TCP [RFC8684] and SCTP utilize multiple data paths over multiple endpoint addresses. As these multipath protocols are unaware of data paths, they distinguish the data paths by endpoint IP addresses. Accordingly, multiple flows of these protocols may use an identical data path without recognizing it.

Multipath protocols are also responsible for inverse multiplexing to split a data stream into multiple data paths. This inverse multiplexing is independent of data paths except for congestion control.

2.5. Reliable Data Communication

Transport layer protocols may implement retransmission and reordering functions to recover lost or reordered datagrams for reliable data communication. This functionality is independent of data paths, and consequently, should be implemented over data paths. Instead, ``smart'' end-hosts or middleboxes may implement it.

An end-host may transmit a duplicate packet for improving reliability [RFC8655]. This functionality is also independent of data paths. However, a router in a data path may be capable of duplicating a packet because the duplication process does not require significant computing resources, unlike retransmission.

2.6. Security

The transport layer may implement an end-to-end security function, such as Transport Layer Security (TLS) [RFC8446] Datagram Transport Layer Security (DTLS) [RFC6347]. As TLS does not encrypt the underlying transport protocol header, middleboxes such as TCP offloading can still work. However, some protocols implementing a transport layer protocol over DTLS, such as SCTP over DTLS used in WebRTC Data Channel [I-D.ietf-rtcweb-data-channel], encrypt the transport layer header, and consequently, have difficulty cooperating with these middleboxes.

2.7. Summary

As described above, the transport layer functionality is summarized by distinguishing data paths and data flows as follows:

The following functionality is categorized as data path functions.

- * In-band trajectory monitoring
- * Waypoint management
- * Bidirectional connection establishment
- * Data path quality (congestion) monitoring and congestion control
- * Data flow multiplexing
- * Packet duplication

The following functionality is categorized as data flow functions.

- * Retransmission for reliable data communication
- * Flow control for receive-buffer management
- * Flow prioritization
- * End-to-end security
- * Inverse multiplexing for multipath protocols

One approach to implement this functionality is prepending, appending, or XOR-ing device identifiers like in-band network telemetry or in-situ OAM [I-D.ietf-ippm-ioam-data] by data path layer protocol's routers. Figure 2 shows an example of in-band trajectory monitoring. A router that handles a data path layer protocol appends the router identifier to the header. In case the upper-layer protocol does not require the path information but path change events, the router can apply an XOR operation with a hashed identifier to a fixed-length field. Other alternative approaches may be used.

The latter approach separates the control plane and the data plane. The control plane manages the data path, and the data plane monitors if packets pass through the dedicated data path. The data path control may leverage underlay protocols such as Segment Routing.

3.1.2. Waypoint management

The data path layer may support waypoint management functionality. The in-band trajectory monitoring functionality described above does not provide the functionality to designate the data path's waypoints. However, some middleboxes such as firewalls and in-network computing require to designate waypoints to activate the in-network functions.

A data path layer protocol may define a specification to control the waypoints of a data path. There are two approaches to designate waypoints over the Internet layer; 1) setting a waypoint as the destination IP address and 2) using routing technologies such as source routing and policy-based routing. The former approach includes Network Address Translation (NAT) [RFC3022] and proxy services. The latter approach may leverage underlay protocols to designate waypoints, such as Segment Routing and IPv6 Type 2 Routing Header [RFC6275]. Other alternative approaches may be used in a data path layer protocol.

3.1.3. Bidirectional connection establishment

The data path layer should support both unidirectional and bidirectional paths. A bidirectional path may be symmetric or asymmetric. As the characteristics of unidirectional and bidirectional paths are different, some data path layer protocols may only support either unidirectional or bidirectional paths.

A data path protocol supporting bidirectional data paths should implement a handshake mechanism to establish a bidirectional connection, such as a 3-way handshake of TCP and a 4-way handshake of SCTP. It may provide a 0-RTT connection establishment feature such as TLS 1.3 [RFC8446] and TCP Fast Open [RFC7413].

3.1.4. Data path quality (congestion) monitoring and congestion control

The data path layer should implement congestion control. However, Data path layer protocols supporting only unidirectional paths may not implement congestion control because it cannot implement congestion or data path quality reporting.

Congestion control is performed based on data path quality or congestion reporting from the receiver end-host or intermediate nodes that process a data path layer protocol. Data path quality or congestion signals may be packet losses, delay, or ECN, as used in many transport layer protocols. Other signals or metrics, such as in-band network telemetry information, may be used.

3.1.5. Data flow multiplexing

The data path layer enables us to collectively handle different characteristics (e.g., service level requirements) of transport layer protocols such as stream and datagram protocols.

3.1.6. Packet Duplication

On a lossy data path, a data path layer protocol may implement packet duplication for improving reliability. However, the data path layer should not provide retransmission because it requires significant resources.

3.2. Data Flow Layer

The data flow layer may implement various functions for end-to-end data communication over the data path layer. Data flow layer protocols may be stream, datagram, or message-oriented protocols. Middleboxes, MEC nodes, or in-network computing nodes may process data flow layer protocols. Therefore, a node processing data flow layer protocols should be as ``smart'' as end-hosts.

A data flow layer protocol may implement a retransmission function for reliable data communication, a flow control mechanism for receive-buffer management, flow prioritization for effective network resource utilization, a security extension such as TLS and DTLS, and a multipath transport protocol using multiple bidirectional paths over the data path layer.

3.2.1. Retransmission for reliable data communication

A data flow layer protocol may provide retransmission for reliable data communication. To this end, a node implementing the data flow layer protocol must equip a transmit buffer for retransmission to retransmit lost corrupted packets. It must also equip a receive buffer to reassemble a stream and a message from a sequence of packets in a stream and message-oriented data flow layer protocols. The receive buffer also handles packet reordering and duplication.

3.2.2. Flow control for receive-buffer management

Flow control is necessary for receive-buffer management. Therefore, a data flow layer protocol may implement the flow control mechanism. A receiver node advertises the available receive-buffer size in the data flow layer, like TCP's window size, when implementing the flow control mechanism. A sender node controls the transmission so that transmitted data do not exceed the receive-buffer size.

3.2.3. Flow prioritization

The prioritization of data flows multiplexed in a data path layer protocol may be implemented on the data flow layer using the data path layer information, such as monitored data path quality or congestion. A data flow layer protocol may provide a service code point or priority so that nodes processing the data flow protocol discriminate a flow.

3.2.4. End-to-end security

A data flow layer protocol may implement the end-to-end security function. TLS and DTLS can be used for stream and datagram protocols, respectively.

3.2.5. Inverse multiplexing for multipath protocols

A multipath data flow protocol may leverage multiple bidirectional data paths established by data path layer protocols. The data flow layer is responsible for the inverse multiplexing functionality. Therefore, the multipath data flow protocol divides a stream, a message, or a datagram into these multiple data paths. Note that the data path layer performs congestion control for each data path, while the data flow layer performs flow control.

4. Use Cases

This document does not specify any data path protocols or data flow protocols, but the data path and data flow layers' architectural design. For better understanding, this section describes the use cases of the data path and the data flow layers. In the use cases, a data path router (DPR) and a data flow layer node (DFN) denote a router that processes a data path layer protocol and a node that processes a data flow layer protocol, respectively.

4.1. Multipath Transport Protocols

Hosts A and B communicate with each other over multiple data paths; Path R1-DPR3-DPR2 and Path DPR4-DPR2. R1 is an IP router. DPR2, DPR3, and DPR4 are data path routers that treat data path layer protocols.




Figure 3: An example of multipath data communication

Figure 3 shows an example of multipath data communication over two data paths. The end-hosts A and B establish two bidirectional data paths; A-R1-DPR3-R2-B and A-DPR4-DPR2-B, and they use a data flow layer protocol over these data paths for end-to-end communication. The data path layer protocols monitor the data path quality and perform congestion control for each data path. The data flow layer protocol performs flow control and inverse multiplexing into these data paths.

Multipath transport protocols may leverage in-band trajectory monitoring to detect a shared waypoint. We assume that the data path routers manipulate the header of a data path layer protocol. They add the router identifier to the data path layer protocol's header to report the trajectory to the end-hosts. When sending packets from end-host A to end-host B over these paths, each packet reports trajectory A-DPR3-DPR2-B or A-DPR4-DPR2-B. In this way, end-host B recognizes two different data paths and detects a shared data path router, DPR2, on the multiple paths, potentially a single point of failure for end-to-end communication.

4.2. Congestion Control Acceleration

Some TCP congestion control acceleration functions proxy TCP sessions to shorten the perspective latency. The acceleration functions acknowledge receipt of packets. Using a loss-based congestion control algorithm, both congestion control and flow control use the acknowledgments (ACKs). For congestion control, missing ACKs report packet losses, and consequently, they present the data path quality to control the transmission rate to avoid congestion. For flow control, ACKs report successful data delivery. Then, the sender can release part of the transmit buffer. Otherwise, the sender retransmits the lost data.

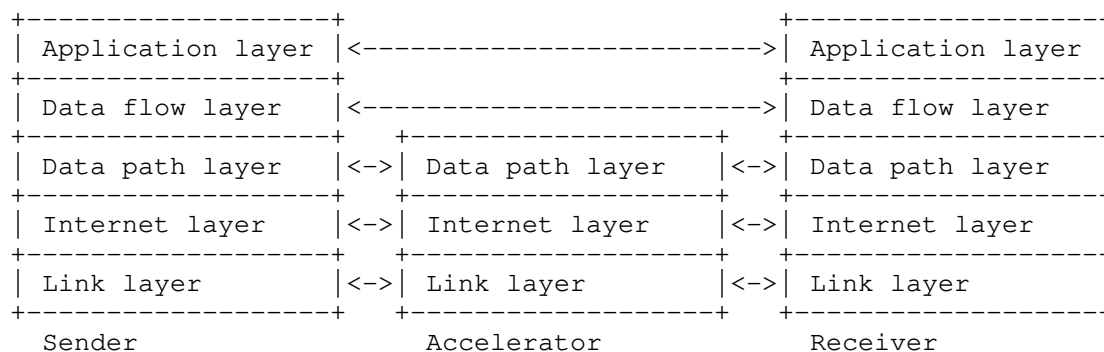


Figure 4: An example of a communication model for congestion control acceleration

Figure 4 illustrates an example of a communication model for congestion control acceleration using the data path layer. The accelerator or the receiver reports the data path quality to the sender on the data path layer to perform congestion control. There are various ways to measure data path quality. For example, data path routers may use the buffer capacity in the same way as ECN. If a data path router can detect packet losses for a connection of data path layer protocols, the packet losses or statistics may be used to report the data path quality. The data path quality report by the accelerator enables the fast response of congestion control algorithms. The data flow layer is responsible for retransmission for reliable communication and flow control for receive-buffer management.

4.3. In-Network Computing

In-network computing [I-D.kunze-coin-industrial-use-cases] is a novel distributed computing paradigm. It requires to be aware of the data path's waypoints because computing components are placed in between end-hosts. The message-oriented protocol may adopt the pub/sub communication model, widely used in machine-to-machine communication.

The data path layer establishes a data path while designating the waypoints to perform in-network computing. The data flow layer over the data path implements a message-oriented protocol or a stream protocol to feed data into in-network computing nodes. The message-oriented protocol may adopt the pub/sub communication model, widely used in machine-to-machine communication.

App., DF, DP, IP, and Link denote the application layer, the data flow layer, the data path layer, the IP, and the link layer, respectively. H1 and H2 are end-hosts, C1 and C2 are in-network computing nodes, P1 is a data path router, and R1 is an IP router.

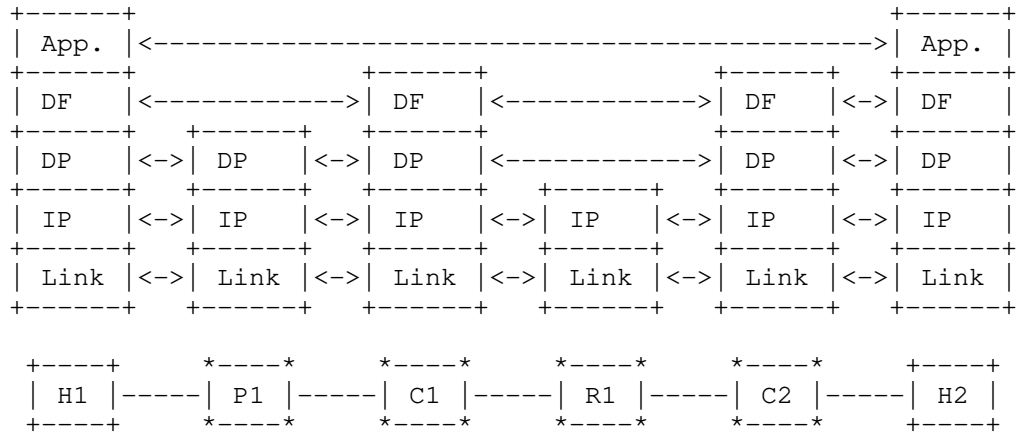


Figure 5: An example of communication model of in-network computing

Figure 5 illustrates an example of the communication model of in-network computing. A data path is established between H1 and H2. A data path layer protocol designates C1 and C2 as the waypoints. Over the established data path, H1 transmits messages to H2 using a data flow layer protocol. C1 and C2 process the messages according to the programs running on C1 and C2.

4.4. Flow Arbitration

The separation of the data path and the data flow layer enables flow-level prioritization. As the data path layer is responsible for congestion control, the multiplexed data flows over a data path can collectively perform resource arbitration for the data path bandwidth capacity.

For example, a data path multiplexes two data flows; one of them requires a constant data rate, and the other is tolerant of delayed data delivery. Flow control can prioritize the former data flow and decrease the transmission rate when the data path layer detects congestion. This flow arbitration is crucial in coexisting deadline-based and best-effort flows.

H1, H2, H3, and H4 are end-hosts. R1, R2, and R3 are IP routers. P1 is a data path router that support the data path quality monitoring function.




Figure 6: An example of flow arbitration over multiple data paths

Moreover, the data path layer may support data path quality monitoring for multiple data paths. If data path routers monitor multiple data paths' bandwidth capacity, they can report the estimated capacity to arbitrate multiple flows over the data paths. Figure 6 shows an example flow arbitration over multiple data paths. Suppose data flows X and Y are over the data paths H1-R1-P1-R2-H3 and H2-R1-P1-R3-H4, respectively; the data path router P1 can monitor the data path quality such as packet losses for these data paths. If the bandwidth utilization reaches the capacity, the data path router requests resource arbitration. A flow may reduce the data rate using the flow control mechanism in the data flow layer if its priority is not high. Thus, these two flows can autonomously perform the flow arbitration. This flow arbitration may fail and perform best-effort communication, such in case both flows are the high priority and do not reduce the data rate.

5. Implementation Considerations

Although this document does not specify the data path layer and the data flow layer protocols, it discusses the implementation of these protocols' specifications. The data path layer may be implemented over UDP for interoperability with the current Internet operations because other protocols than ICMP, TCP, and UDP may be filtered on the Internet. Data path layer protocols must implement the maximum transmission unit (MTU) discovery [RFC8899].

LH: Link-layer protocol header, IPH: IP header, UDPH: UDP header, DPH: Data path layer protocol header, DFH: Data flow layer protocol header




Figure 7: An implementation of data path layer and data flow layer protocols over UDP

Figure 7 illustrates an example packet format of data path layer and data flow layer protocols over UDP. In this figure, the UDP header is not intended to provide the transport layer's functionality but is introduced for interoperability with existing middleboxes on the Internet.

A data path layer protocol may use hop-by-hop UDP connections to designate waypoints like an overlay network. Otherwise, data path routers require to leverage routing technologies such as Segment Routing and policy-based routing to support waypoint management.

6. IANA Considerations

This document does not have IANA Considerations.

7. Security Considerations

End-to-end encryption becomes critical to Internet protocols for privacy and security. The data path layer must be visible to intermediate nodes, such as routers and middleboxes, by design to process and manipulate a data path layer protocol. Therefore, this document proposes to implement transport layer security at the data flow layer. It exposes data path information. Therefore, data path layer protocols must not include privacy-sensitive information in the protocol header.

Data path layer protocols may not be visible to intermediate nodes if an underlay protocol encrypts the payload. For example, IPsec [RFC6071] encrypts the payload of IP datagrams. In such cases, intermediate nodes cannot process or manipulate data path layer protocols.

8. Acknowledgements

We thank Kenichi Murata, Ryokichi Onishi, Yusuke Doi, and Masahiro Ishiyama for their comments and review of this document.

9. References

9.1. Normative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://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, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC1123] Braden, R., Ed., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, DOI 10.17487/RFC1123, October 1989, <<https://www.rfc-editor.org/info/rfc1123>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.

- [RFC2768] Aiken, B., Strassner, J., Carpenter, B., Foster, I., Lynch, C., Mambretti, J., Moore, R., and B. Teitelbaum, "Network Policy and Services: A Report of a Workshop on Middleware", RFC 2768, DOI 10.17487/RFC2768, February 2000, <<https://www.rfc-editor.org/info/rfc2768>>.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <<https://www.rfc-editor.org/info/rfc3022>>.
- [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, <<https://www.rfc-editor.org/info/rfc3168>>.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, DOI 10.17487/RFC4340, March 2006, <<https://www.rfc-editor.org/info/rfc4340>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.
- [RFC6071] Frankel, S. and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", RFC 6071, DOI 10.17487/RFC6071, February 2011, <<https://www.rfc-editor.org/info/rfc6071>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC8684] Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 8684, DOI 10.17487/RFC8684, March 2020, <<https://www.rfc-editor.org/info/rfc8684>>.
- [RFC8899] Fairhurst, G., Jones, T., Tüxen, M., Rüngeler, I., and T. Völker, "Packetization Layer Path MTU Discovery for Datagram Transports", RFC 8899, DOI 10.17487/RFC8899, September 2020, <<https://www.rfc-editor.org/info/rfc8899>>.

9.2. Informative References

- [I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-10, 13 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-ippm-ioam-data-10.txt>>.
- [I-D.ietf-rtcweb-data-channel]
Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data Channels", Work in Progress, Internet-Draft, draft-ietf-rtcweb-data-channel-13, 4 January 2015, <<http://www.ietf.org/internet-drafts/draft-ietf-rtcweb-data-channel-13.txt>>.
- [I-D.kunze-coin-industrial-use-cases]
Kunze, I. and K. Wehrle, "Industrial Use Cases for In-Network Computing", Work in Progress, Internet-Draft, draft-kunze-coin-industrial-use-cases-03, 8 September 2020, <<http://www.ietf.org/internet-drafts/draft-kunze-coin-industrial-use-cases-03.txt>>.

Author's Address

Hirochika Asai
Preferred Networks, Inc.
1-6-1 Otemachi, Chiyoda, Tokyo
100-0004
Japan

Email: panda@wide.ad.jp