WIKIPEDIA

Transport layer

In computer networking, the **transport layer** is a conceptual division of methods in the layered architecture of protocols in the network stack in the Internet protocol suite and the OSI model. The protocols of this layer provide host-to-host communication services for applications.[1]:§1.1.3 It provides services such as connection-oriented communication, reliability, flow control, and multiplexing.

The details of implementation and semantics of the transport layer of the Internet protocol suite,[1] which is the foundation of the Internet, and the OSI model of general networking are different. The protocols in use today in this layer for the Internet all originated in the development of TCP/IP. In the OSI model the transport layer is often referred to as **Layer 4**, or **L4**,[2] while numbered layers are not used in TCP/IP.

The best-known transport protocol of the Internet protocol suite is the Transmission Control Protocol (TCP). It is used for connection-oriented transmissions, whereas the connectionless User Datagram Protocol (UDP) is used for simpler messaging transmissions. TCP is the more complex protocol, due to its stateful design incorporating reliable transmission and data stream services. Together, TCP and UDP comprise essentially all traffic on the Internet and are the only protocols implemented in every major operating system. Additional transport layer protocols that have been defined and implemented include the Datagram Congestion Control Protocol (DCCP) and the Stream Control Transmission Protocol (SCTP).

Contents

Services

Analysis

Protocols

- Comparison of transport layer protocols
- Comparison of OSI transport protocols

References

Services

Transport layer services are conveyed to an application via a programming interface to the transport layer protocols. The services may include the following features:

- Connection-oriented communication: It is normally easier for an application to interpret a connection as a data stream rather than having to deal with the underlying connection-less models, such as the datagram model of the User Datagram Protocol (UDP) and of the Internet Protocol (IP).
- Same order delivery: The network layer doesn't generally guarantee that packets of data will arrive in the same order that they were sent, but often this is a desirable feature. This is usually done through the use of segment numbering, with the receiver passing them to the application in order. This can cause head-of-line blocking.
- Reliability: Packets may be lost during transport due to network congestion and errors. By means of an error detection code, such as a checksum, the transport protocol may check that

the data is not corrupted, and verify correct receipt by sending an ACK or NACK message to the sender. Automatic repeat request schemes may be used to retransmit lost or corrupted data.

- Flow control: The rate of data transmission between two nodes must sometimes be managed to prevent a fast sender from transmitting more data than can be supported by the receiving data buffer, causing a buffer overrun. This can also be used to improve efficiency by reducing buffer underrun.
- Congestion avoidance: Congestion control can control traffic entry into a telecommunications network, so as to avoid congestive collapse by attempting to avoid oversubscription of any of the processing or link capabilities of the intermediate nodes and networks and taking resource reducing steps, such as reducing the rate of sending packets. For example, automatic repeat requests may keep the network in a congested state; this situation can be avoided by adding congestion avoidance to the flow control, including slow-start. This keeps the bandwidth consumption at a low level in the beginning of the transmission, or after packet retransmission.
- Multiplexing: Ports can provide multiple endpoints on a single node. For example, the name on a postal address is a kind of multiplexing, and distinguishes between different recipients of the same location. Computer applications will each listen for information on their own ports, which enables the use of more than one network service at the same time. It is part of the transport layer in the TCP/IP model, but of the session layer in the OSI model.

Analysis

The transport layer is responsible for delivering data to the appropriate application process on the host computers. This involves statistical multiplexing of data from different application processes, i.e. forming data segments, and adding source and destination port numbers in the header of each transport layer data segment. Together with the source and destination IP address, the port numbers constitute a network socket, i.e. an identification address of the process-to-process communication. In the OSI model, this function is supported by the session layer.

Some transport layer protocols, for example TCP, but not UDP, support virtual circuits, i.e. provide connection-oriented communication over an underlying packet-oriented datagram network. A byte-stream is delivered while hiding the packet mode communication for the application processes. This involves connection establishment, dividing of the data stream into packets called segments, segment numbering and reordering of out-of order data.

Finally, some transport layer protocols, for example TCP, but not UDP, provide end-to-end reliable communication, i.e. error recovery by means of error detecting code and automatic repeat request (ARQ) protocol. The ARQ protocol also provides flow control, which may be combined with congestion avoidance.

UDP is a very simple protocol, and does not provide virtual circuits, nor reliable communication, delegating these functions to the application program. UDP packets are called datagrams, rather than segments.

TCP is used for many protocols, including HTTP web browsing and email transfer. UDP may be used for multicasting and broadcasting, since retransmissions are not possible to a large amount of hosts. UDP typically gives higher throughput and shorter latency, and is therefore often used for real-time multimedia communication where packet loss occasionally can be accepted, for example IP-TV and IP-telephony, and for online computer games.

Many non-IP-based networks, such as X.25, Frame Relay and ATM, implement the connection-oriented communication at the network or data link layer rather than the transport layer. In X.25, in telephone network modems and in wireless communication systems, reliable node-to-node communication is implemented at lower protocol layers.

The OSI connection-mode transport layer protocol specification defines five classes of transport protocols: *TP0*, providing the least error recovery, to *TP4*, which is designed for less reliable networks.

Protocols

This list shows some protocols that are commonly placed in the transport layers of the Internet protocol suite, the OSI protocol suite, NetWare's IPX/SPX, AppleTalk, and Fibre Channel.

- ATP, AppleTalk Transaction Protocol
- CUDP, Cyclic UDP[3]
- DCCP, Datagram Congestion Control Protocol
- FCP, Fibre Channel Protocol
- IL, IL Protocol
- MPTCP, Multipath TCP
- RDP, Reliable Data Protocol
- RUDP, Reliable User Datagram Protocol
- SCTP, Stream Control Transmission Protocol
- SPX, Sequenced Packet Exchange
- SST, Structured Stream Transport
- TCP, Transmission Control Protocol
- UDP, User Datagram Protocol
- UDP-Lite
- µTP, Micro Transport Protocol

Comparison of transport layer protocols

Feature	UDP	UDP-Lite	TCP	Multipath TCP	SCTP	DCCP	RUDP[a]
Packet header size	8 bytes	8 bytes	20–60 bytes	50–90 bytes	12 bytes[b]	12 or 16 bytes	14+ bytes
Typical data-packet overhead	8 bytes	8 bytes	20 bytes	?? bytes	44–48+ bytes[c]	12 or 16 bytes	14 bytes
Transport-layer packet entity	Datagram	Datagram	Segment	Segment	Datagram	Datagram	Datagram
Connection-oriented	No	No	Yes	Yes	Yes	Yes	Yes
Reliable transport	No	No	Yes	Yes	Yes	No	Yes
Unreliable transport	Yes	Yes	No	No	Yes	Yes	Yes
Preserve message boundary	Yes	Yes	No	No	Yes	Yes	Yes
Delivery	Unordered	Unordered	Ordered	Ordered	Ordered / Unordered	Unordered	Unordered
Data checksum	Optional	Yes	Yes	Yes	Yes	Yes	Optional
Checksum size	16 bits	16 bits	16 bits	16 bits	32 bits	16 bits	16 bits
Partial checksum	No	Yes	No	No	No	Yes	No
Path MTU	No	No	Yes	Yes	Yes	Yes	?
Flow control	No	No	Yes	Yes	Yes	No	Yes
Congestion control	No	No	Yes	Yes	Yes	Yes	?
Explicit Congestion Notification	No	No	Yes	Yes	Yes	Yes	?
Multiple streams	No	No	No	No	Yes	No	No
Multi-homing	No	No	No	Yes	Yes	No	No
Bundling / Nagle	No	No	Yes	Yes	Yes	No	?

a. RUDP is not officially standardized. There have been no standard-related developments since 1999.

b. Excluding data chunk headers and overhead chunks. Without embedded chunks, an SCTP packet is essentially useless.

c. Counted as follows: 12 bytes SCTP header + 16 bytes DATA chunk header or 20 bytes I-DATA chunk header + 16+ bytes SACK chunk. Additional non-data chunks (e.g. AUTH) and/or headers for additional data chunks, which might easily increase the overhead with 50 bytes or more, not counted.

Comparison of OSI transport protocols

ISO/IEC 8073/ITU-T Recommendation X.224, "Information Technology - Open Systems Interconnection - Protocol for providing the connection-mode transport service", defines five classes of connection-mode transport protocols designated class 0 (TP0) to class 4 (TP4). Class 0 contains no error recovery, and was designed for use on network layers that provide error-free connections. Class 4 is closest to TCP, although TCP contains functions, such as the graceful close, which OSI assigns to the session layer. All OSI connection-mode protocol classes provide expedited data and preservation of record boundaries. Detailed characteristics of the classes are shown in the following table:[4]

Service	TP0	TP1	TP2	TP3	TP4
Connection-oriented network	Yes	Yes	Yes	Yes	Yes
Connectionless network	No	No	No	No	Yes
Concatenation and separation	No	Yes	Yes	Yes	Yes
Segmentation and reassembly	Yes	Yes	Yes	Yes	Yes
Error recovery	No	Yes	No	Yes	Yes
Reinitiate connection (if an excessive number of PDUs are unacknowledged)	No	Yes	No	Yes	No
Multiplexing and demultiplexing over a single virtual circuit	No	No	Yes	Yes	Yes
Explicit flow control	No	No	Yes	Yes	Yes
Retransmission on timeout	No	No	No	No	Yes
Reliable Transport Service	No	Yes	No	Yes	Yes

There is also a connectionless transport protocol, specified by ISO/IEC 8602/ITU-T Recommendation X.234.[5]

References

1. R. Braden, ed. (October 1989). *Requirements for Internet Hosts – Communication Layers* (https://tools.ietf.org/html/rfc1122). doi:10.17487/RFC1122 (https://doi.org/10.17487%2FRFC1122). RFC 1122 (https://tools.ietf.org/html/rfc1122).
2. "Introducing the Internet Protocol Suite" (https://docs.oracle.com/cd/E19455-01/806-0916/6ja85398m/index.html). *System Administration Guide, Volume 3*.
3. Brian C. Smith, *Cyclic-UDP: A Priority-Driven Best-Effort Protocol* (https://www.cs.cornell.edu/zeno/Papers/cyclicudp.pdf) (PDF), retrieved February 23, 2020
4. "ITU-T Recommendation X.224 (11/1995) ISO/IEC 8073" (http://www.itu.int/rec/T-REC-X.224-199511-I/en/). *Itu.int*. Retrieved January 17, 2017.
5. "ITU-T Recommendation X.234 (07/1994) ISO/IEC 8602" (http://www.itu.int/rec/T-REC-X.234-199407-I/en/). *Itu.int*. Retrieved January 17, 2017.