



Computer Network Programming

SCTP Overview

Dr. Sam Hsu

Computer Science & Engineering

Florida Atlantic University



SCTP Overview

- Introduction
- Motivations
- Architectural & Functional Views
- Packet & Chunk Formats
- Key Terms
- Association Startup and Termination



SCTP

- **Stream Control Transmission Protocol**
- A connection-oriented Transport Protocol
 - Message-oriented
 - Reliable data transfer
 - Ordered delivery
 - Selective acknowledgement
 - Flow and congestion control
 - Multi-homing
 - Multi-streaming




Motivations

- To support the transport of PSTN signaling across the IP network.
- To build next generation commercial-grade communication infrastructures for various applications.


Architectural View of SCTP

- An SCTP association between two nodes

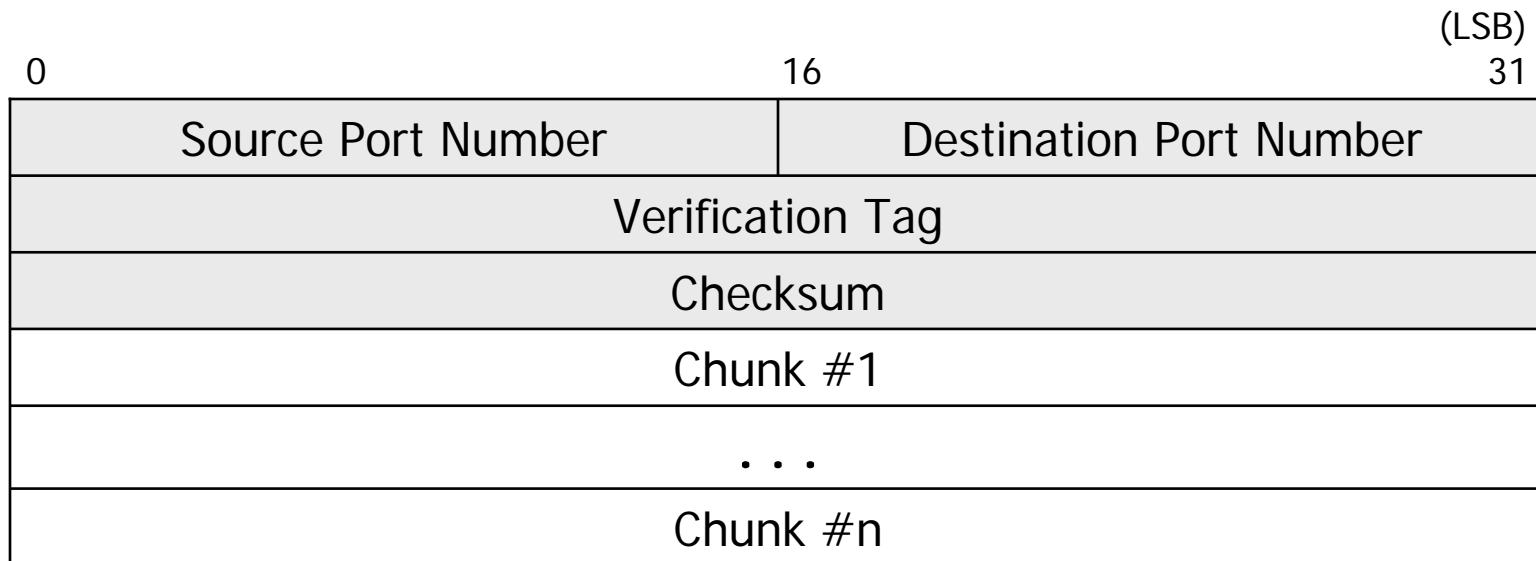




SCTP Packet Format



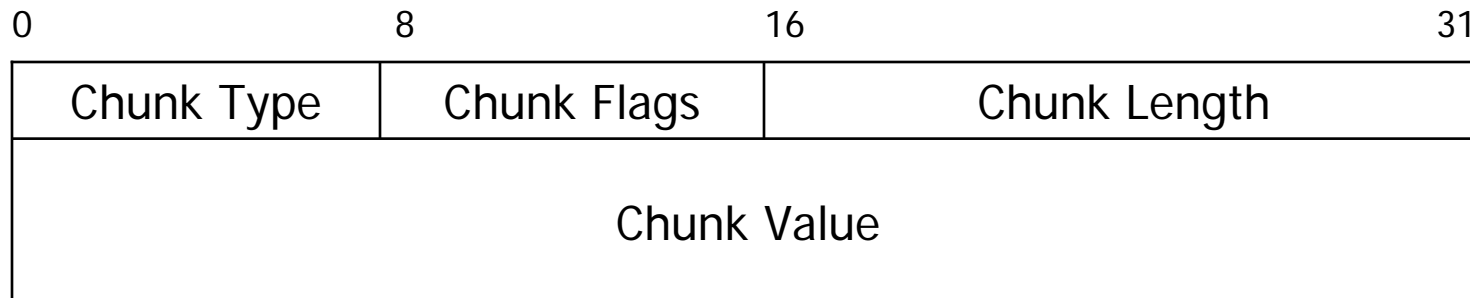
General format of an SCTP packet (max size = path MTU)



Format of SCTP common header (shaded area)



SCTP Chunk Format



General format of an SCTP control/data chunk

- *Type* is used to indicate the chunk type. There are 12 types of control and 1 type of data chunks defined so far.
- Bit values for *Flags* depend on chunk type.
- Chunk length is measured in bytes (8 bits) for all bytes in this chunk, excluding final zero-byte padding (to a multiple of 4 bytes) in general.
- *Value* is of variable length. Its usage and format depend on the chunk type.

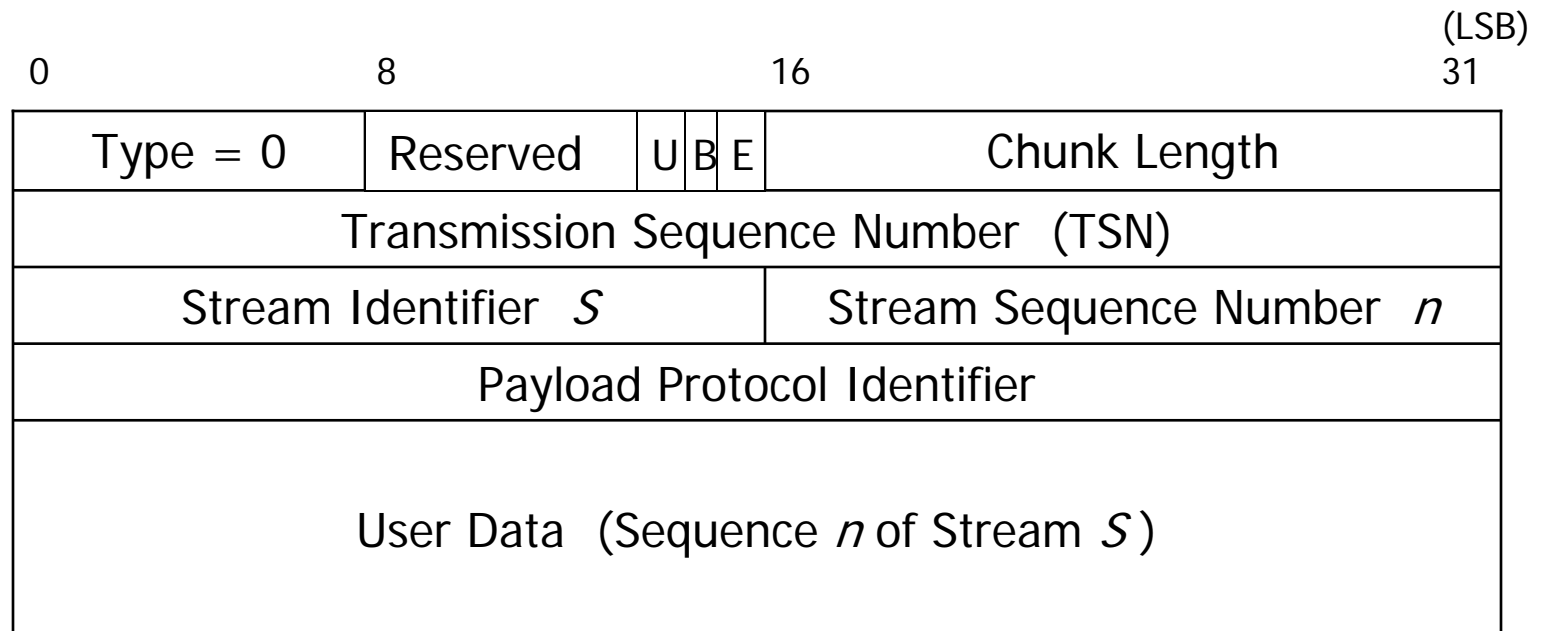


Some Chunk Type Values

Value	Chunk Type
0	Payload Data (DATA)
1	Initiation (INIT)
2	Initiation Acknowledgement (INIT ACK)
3	Selective Acknowledgement (SACK)
4	Heartbeat Request (HEARTBEAT)
5	Heartbeat Acknowledgement (HEARTBEAT ACK)
6	Abort (ABORT)
7	Shutdown (SHUTDOWN)
8	Shutdown Acknowledgement (SHUTDOWN ACK)
9	Operation Error (ERROR)
10	State COOKIE (COOKIE ECHO)
11	Cookie Acknowledgement (COOKIE ACK)
14	Shutdown Complete (SHUTDOWN COMPLETE)




Data Chunk Format



Format of an SCTP data chunk



Data Chunk Flag Bits



U,B,E bits:

- U: If set, indicates an unordered DATA chunk, and therefore, no Stream Sequence Number is assigned to this chunk.
- B: If set, indicates the first fragment of a user message.
- E: If set, indicates the last fragment of a user message.



Control Chunk Types

- 12 control chunk types defined so far
 - INIT, INIT ACK
 - COOKIE ECHO, COOKIE ACK
 - HEARTBEAT, HEARTBEAT ACK
 - SACK
 - SHUTDOWN, SHUTDOWN ACK,
SHUTDOWN COMPLETE
 - ABORT, ERROR



Some Key Terms (1/3)

- Message: User message
 - Data submitted to SCTP by the upper layer protocol.
 - There are ordered and un-ordered messages.
- Packet: The unit of data delivery across the interface between SCTP and the underlying connectionless packet network (e.g., IP).
 - A packet is composed of a common header and chunks.
- Chunk: A unit of information within an SCTP packet, consisting of a chunk header and chunk-specific content.
- Stream: A unidirectional logical channel established from one to another associated SCTP endpoints.
 - All messages within a stream are delivered in sequence, except for those submitted to the unordered delivery service.



Some Key Terms (2/3)

- Association: A protocol relationship between two SCTP endpoints.
 - An association can be uniquely identified by the transport addresses used by the endpoints in the association.
 - Two SCTP endpoints **MUST NOT** have more than one SCTP association between them at any given time.
- Endpoint: The logical sender/receiver of SCTP packets.
- Transport Address: In SCTP, a transport address is defined by the combination of an IP address and an SCTP port number.
 - A transport address is unique to an SCTP endpoint.
 - In a multi-homed SCTP host, all transport addresses must use the same port number, but may contain multiple IP addresses.

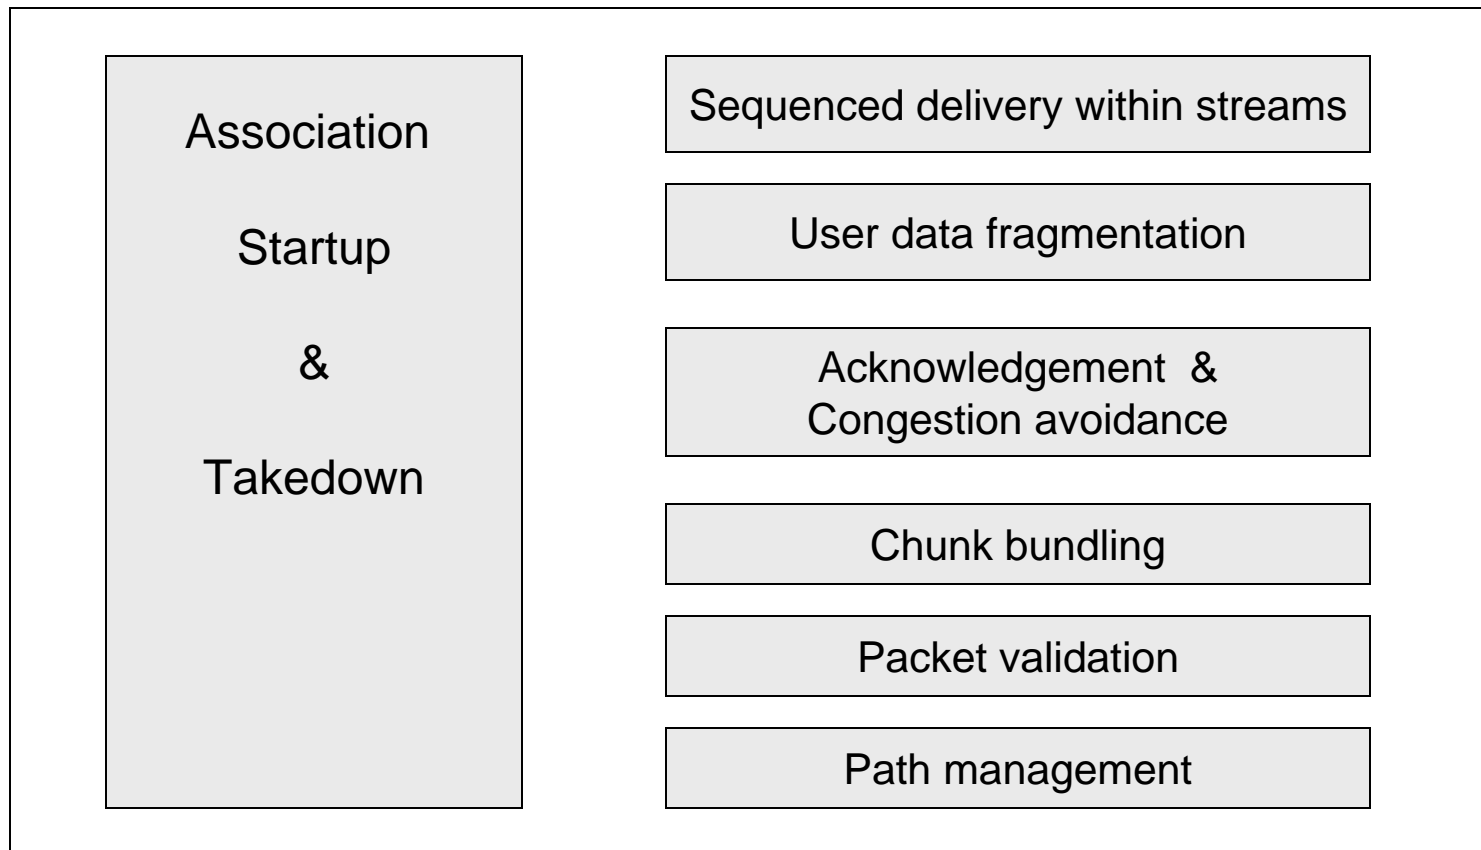


Some Key Terms (3/3)

- Transition Sequence Number: A 32-bit number used internally by SCTP.
 - One TSN is attached to each chunk containing user data to permit the receiving SCTP endpoint to acknowledge its receipt and detect duplicate deliveries.
- Verification Tag: A randomly generated 32-bit number.
 - It provides a key that allows a receiver to verify that the SCTP packet belongs to the current association, and is not an old or stale packet from a previous association.

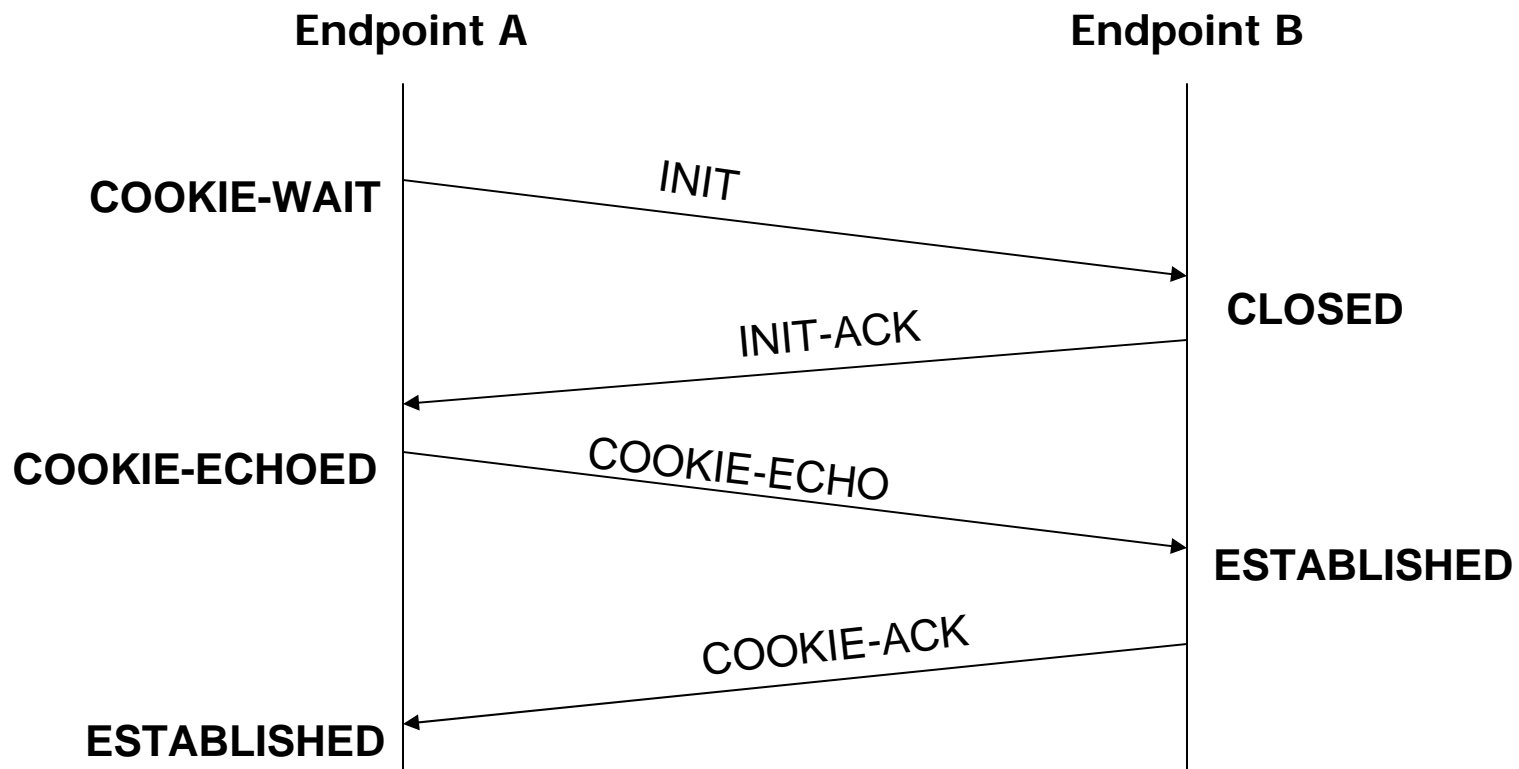



Functional View of SCTP



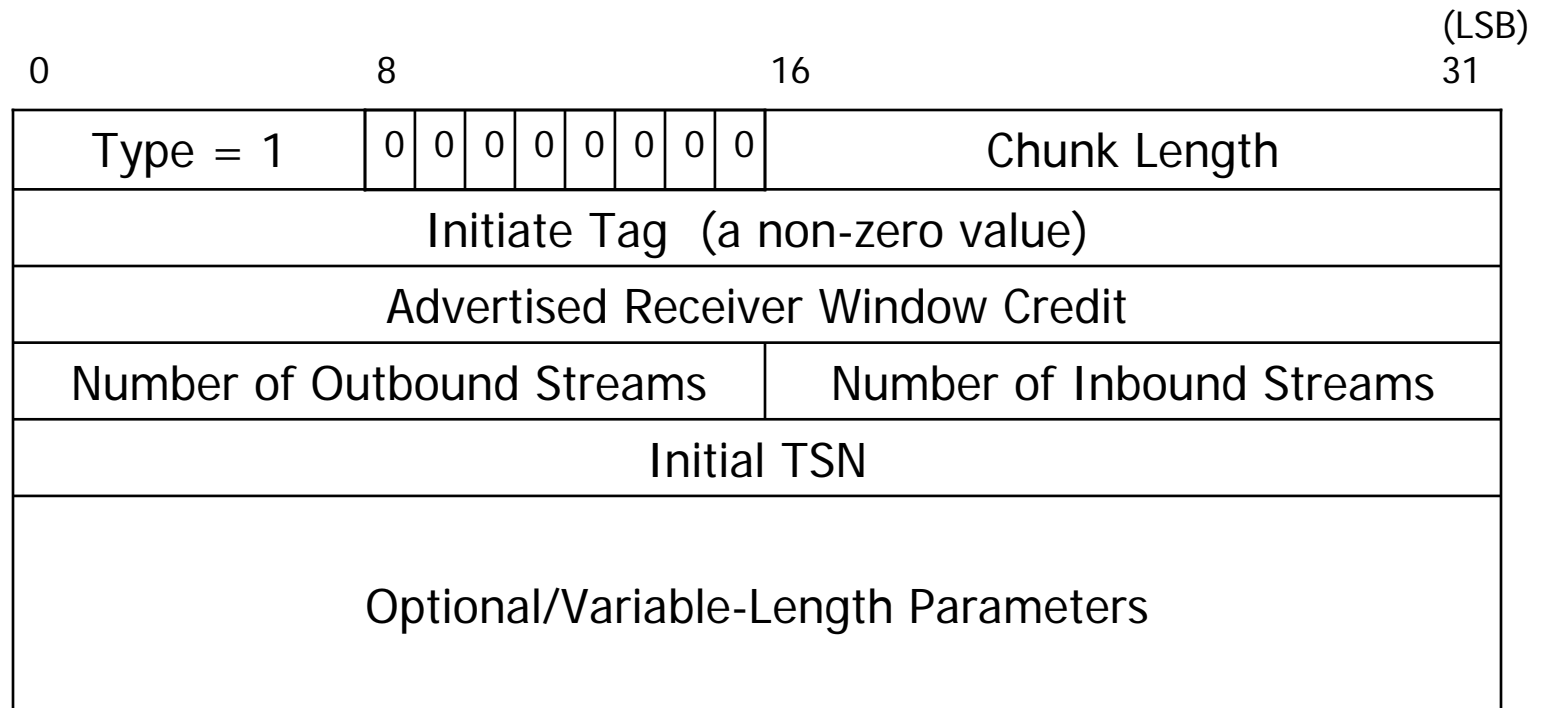
SCTP Association Startup

- Four-way Handshake





Init Chunk Format



Format of INIT control chunk



Initial Startup (1/6)

- An association is initiated by an SCTP endpoint.
 - It is a 4-way handshake.
 - The last two legs are allowed to carry user data for fast setup.
- Two random numbers are generated.
 - One for initiate tag
 - A non-zero unsigned integer.
 - One for initial TSN
 - An unsigned integer.
- Max receiving buffer size and numbers of inbound and outbound streams.
- Optional parameters
 - IPv4/IPv6 addresses
 - Additional services



Initial Startup (2/6)

- Once an INIT chunk is sent, the sender will start a timer.
 - It is named the *T1-init* timer.
 - A second INIT chunk, a duplicate of the previous one, will be sent after the timer expires.
 - The timer value doubles.
 - It is up to the local setup regarding the number of retransmissions to be allowed.
- The sender will enter the COOKIE-WAIT state after the first INIT chunk is sent.
- In the meantime, a TCB (Transmission Control Block) storing information about the new association is created.



Initial Startup (3/6)

- Upon receiving an INIT chunk successfully, the remote endpoint will acknowledge the receipt of it with the INIT ACK chunk.
 - It has the similar format as the INIT chunk, with chunk type value set to 2.
- A mandatory *State Cookie* parameter must be included in the optional parameter part.
 - While the format of the cookie is up to the sender, the following are considered crucial:
 - Information related to the corresponding INIT chunk.
 - Timestamp of cookie creation.
 - Some security parameters.
- At this point, the receiving side does not keep any information or knowledge about this association.
 - It still remains in the CLOSED state, no TCB created.



Initial Startup (4/6)

- At the initiating sender side, things to be done upon receiving the INIT ACK chunk include:
 - Find the association.
 - Stop the timer.
 - Update the address list.
 - Assemble a COOKIE ECHO chunk.
 - The chunk type value is 10 (0x0a).
 - The received cookie data are simply echoed back.
 - Start a new timer *T1-cookie* after sending out the new chunk.
 - Enter the COOKIE-ECHOED state.
- Data, if present, can be bundled into this chunk for an early start.



Initial Startup (5/6)

- After receiving the COOKIE ECHO chunk, the receiving SCTP endpoint will perform tasks such as:
 - Authenticate the cookie received.
 - Make sure the cookie has not expired.
 - Created an association TCB based on the information carried in the cookie.
 - Assemble a COOKIE ACK chunk to acknowledge the successful receipt of the COOKIE ECHO chunk.
 - Chunk type value = 0x0b
 - Data bundled in the chunk received will also be acknowledged if no error.
 - Enter the ESTABLISHED state.
- Should an error be detected in any stage of the above process, an ERROR chunk will be sent back.
 - The received chunk is thus dropped.



Initial Startup (6/6)

- The last leg in this 4-way handshake goes back to the initiating SCTP endpoint.
- Once the COOKIE ACK chunk is received, it will enter the ESTABLISHED state.
 - An association between two SCTP endpoints is thus established.

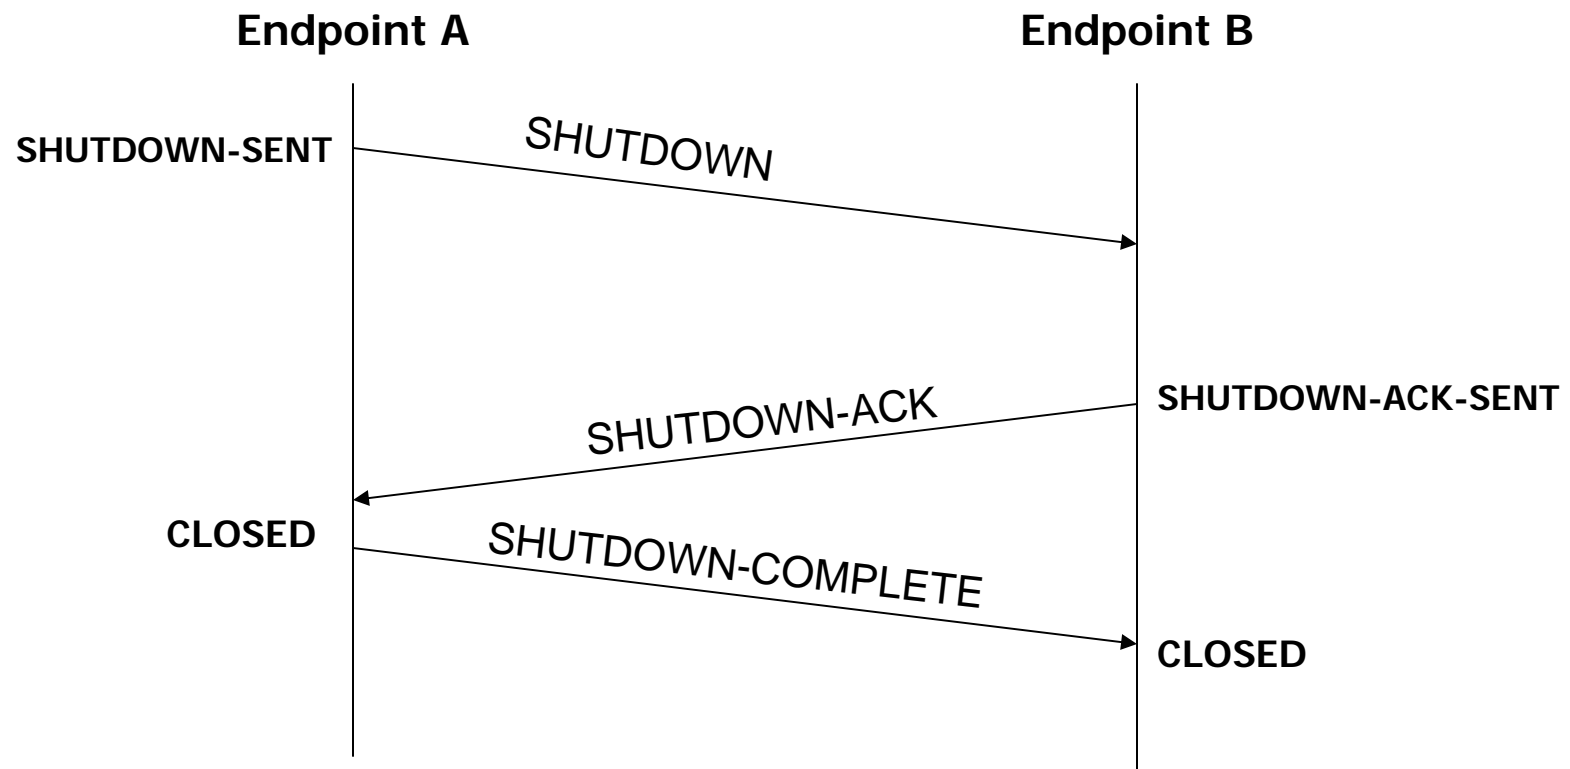


Final Takedown

- It is a 3-way handshake.
- SCTP will gracefully close an association.
- It will stop receiving any data from the upper layer.
- Any locally queued user data will be delivered to the remote peer.
- The association will be terminated only after the peer acknowledges all the SCTP packets sent.
- It will then delete its TCB and enter the CLOSED state.

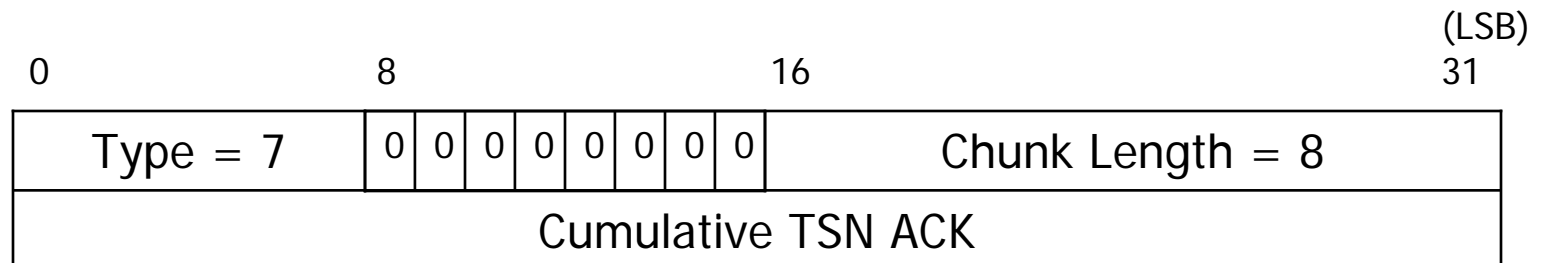
SCTP Association Termination

- Three-way handshake





Shutdown Chunk Format



Format of SHUTDOWN control chunk

- *Flags* are set to all zeros on transmit and ignored on receipt.



Reading Assignment

- Read Chapters 9-10, 23.



For More Information

- RFC 2960 – Stream Control Transmission Protocol
October 2000.
- RFC 3309 – Stream Control Transmission Protocol
(SCTP) Checksum Change, September 2002.
- RFC 4460 – Stream Control Transmission Protocol
(SCTP) Specification Errata and Issues , April 2006.
- <http://www.faqs.org/rfcs/rfc-index.html>
- SCTP for beginners, A. Jungmaier,
http://tdrwww.exp-math.uni-essen.de/inhalt/forschung/sctp_fb/index.html