

414 Precis  
SCTP: State of the Art in Research,  
Products and Technical Challenges

Date : 16/05/06  
Matthew Kiernan  
Andrew Watson

SCTP is a reliable transport protocol initially designed for signaling transport. It is attractive due to several features such as multistreaming and multihoming, and support for timely, reliable message based communication. The paper describes several of these features, and explains some situations where SCTP can be used.

PSTNs are being integrated with IP networks. In order to accomplish this successfully, special signaling protocols are necessary such as SS7. Packets in an SS7 format need to be transported reliably and quickly. Because UDP is not connection orientated, it cannot be used for signaling. TCP is more suitable as it is connection orientated, yet it has several deficiencies. Strict byte order can cause head of line blocking. TCP is stream rather than message based and does not support multihoming, therefore it is less reliable. It is also vulnerable to DoS attacks.

SCTP was introduced in 2000 by the IETF. Later it was realized that it could be used for more than just signaling. It followed in the footsteps of TCP by using window based congestion control and several other features.

The multihoming functionality of SCTP allows an association to be set up with multiple IP addresses at each endpoint. An association is similar to a TCP connection. This means that redundancy can be incorporated into the network so that switch over can occur between networks links to prevent lengthy interruption to data transfer. One of the IP addresses on this association can be designated the primary, the other the secondary. Data retransmission can be sent through the secondary link to better the chances of delivery. This is useful for SS7 signaling which needs to be reliable.

The multistreaming capability of SCTP means that various sections of data can be split in terms of ordering. If packets from another stream are received out of order or delayed, another stream may be free, preventing blocking of data. This lessens the Head of Line blocking problem TCP faces. Only individually affected streams have the HOL effect, while others streams continue on passing packets up to higher layers.

An example of multistreaming can be given using a HTML page with 5 objects. TCP would create 5 separate connections for each of the objects. SCTP uses multistreaming to send each object in a different stream, while opening only one SCTP association.

Congestion control is based on TCP's rate adaptive window based scheme. It provides reliability, by detecting lost/corrupt packets and resends them. SCTP differs from TCP in that it does not have a fast recovery phase. SCTP achieves fast recovery by using a method of acknowledging packets known as SACK (Selective Acknowledgement). In Contrast to TCP, SCTP regards SACK as compulsory. This allows a more reliable reaction to multiple packets lost in a single window. This avoids the time consuming slow start stage after multiple segment losses.

SCTP also differs from TCP in that during a slow start, the window is increased by the number of acknowledged bytes. TCP advances this window by the number of ACK segments received. In TCP 'delayed ACKs' mean the smaller number of ACKs returns slows the window growth rate.

During the congestion avoidance stage of SCTP, the window size can only be increased when the full window is used. This prevents the window size creeping up when only a small amount of data is sent. If the window size did creep up, a large burst of data could travel through, and may cause congestion for others in the network.

SCTP waits for 4 rather than 3 duplicate ACKS for a fast transmission.

The transport protocol might be carrying sensitive billing or critical signaling messages so security measures have been taken to ensure the safe and authenticated transport of this data through the network.

To prevent DOS attacks that TCP suffers from SCTP does not store state on the server about each pending connection. A four way handshake is implemented, rather than a three way with the TCP protocol. Only on the 3<sup>rd</sup> part of the handshake is state stored from the response. A cookie is then sent. This cookie includes information about IP addresses, stream numbers, TTL and a signature. This means data is stored on the network or client, preventing server resources being uselessly consumed and failure occurring during a DoS attack.

IPSec or Transport layer security is used if necessary to protect the payload. IPSec works with both IP4 and 6 and works by cryptography at the network layer. IPSec offers both Authenticating header and Encapsulating security payload – which ensures secrecy as well as integrity and authentication. Transport layer security can provide server and client authentication and encrypted messages.

### **Differences between TCP and SCTP**

SCTP does not allow a situation where one endpoint has finished sending, but is expecting further data from the other end. This is because SCTP uses a three way Hand shake for shut down, rather than four.

Some other features include

- Optionally supports unordered delivery. This provides better performance for multimedia applications.
- Message Based – rather than stream like TCP. This supports a system based around the communication of signaling messages.
- Keep alive messages are mandatory, instead of relying on implementation like TCP. This can be useful when up to date information about the state of distributed processes is required.

### **Current research in the area of SCTP**

SCTP traffic can interoperate alongside TCP without noticeable deterioration of service offered by either method due to the similarities between congestion control in TCP and SCTP, that is, network resources are shared fairly.

Multihoming support in SCTP allows a greater degree of reliability and faster failure recovery. It is not recommended to use the redundant paths for load balancing, as this can cause byte reordering in some scenarios. The redundant paths are useful for retransmission of lost segments or if the primary path fails.

An experiment was carried out to determine the benefits of multistreaming. The results show that multistreaming results in a slower degradation of throughput as the error rate increases. Multistreaming reduces the buffer requirements at the receiver.

Because SIP is usually carried by UTP, SCTP's out of order service has been investigated for carrying SIP. An IETF draft suggests SCTP could carry traffic when only partially reliable transport is needed.

Although SCTP was designed for wireline networks, its potential use in wireless networks was also investigated. SCTP suffers 'Go Back N' problems after a delay spike just as in TCP. The use of selective acknowledgements can prevent pointless fast retransmissions.

It has been shown that using SCTP in a mobile IP environment was found to be more effective in operation than TCP Reno and TCP SACK. There is a possibility to use multihoming to reduce the load on a home agent after a handover. The association on the mobile host can store a permanent home address as well as a care of address. The care of address can be updated to the current location of the mobile on the network.

Transport Layer Seamless Handover (TraSH) has been proposed as a new mobile handover method. It uses the multihoming feature to add the new IP address from the new domain which eliminates the triangle routing problem.

SCTP products are already available. Some Unix operating system such as NetBSD as well as Linux have implemented support into the kernel for SCTP. Most SCTP commercial products are put in place for signaling purposes

### **Implementation considerations**

Socket API needs to be made. Needs to be consistent and support TCP and UDP style interfaces. Special new API's for some applications are required to utilize the fullest possible power of SCTP. Notably, the use of stream connections needs to be emulated.

The IETF has provided an implementing guide in RFC 2960. SCTP is relatively new, and still requires some work in meeting requirements of SS7 and other signaling protocols. More work needed to solve problems with delay spikes.

Dynamic address configuration allows IP addresses to be added to an association without interrupting data transfer. There are security issues with this.

### **Conclusion**

The Stream Control protocol is being standardized as a reliable transport protocol to address limitations in other protocols such as TCP with dealing in signaling transport.