

SCTP Multihoming

An essential property of SCTP is its support of multi-homed nodes, i.e. nodes which can be reached under several IP addresses. If the SCTP nodes and the according IP network are configured in such a way that traffic from one node to another travels on physically different paths if different destination IP addresses are used, associations become tolerant against physical network failures and other problems of that kind.

Address Management at Association Setup

If a client is multi-homed, it informs the server about all its IP addresses with the INIT chunk's address parameters. Thereby, the client is only required to know one IP address of the server because the server provides all its IP addresses to the client in the INIT-ACK chunk. SCTP is able to handle IP version 4 and IP version 6 addresses (even mixed). An SCTP instance regards each IP address of its peer as one "transmission path" towards this endpoint.

If no explicit IP addresses are contained in the INIT or INIT-ACK chunk, the source IP address of the IP packet which carries the SCTP datagram is used. This eases application of SCTP when Network Address Translation, (NAT), e.g. at the edge of large private IP networks, is involved. To facilitate this further, an additional optional feature has been introduced into the [RFC2960](#) which allows the usage of host names in addition to or instead of IP addresses.

Path and Peer Monitoring

An SCTP instance monitors all transmission paths to the peer instance of an association. To this end, HEARTBEAT chunks are sent over all paths which are currently not used for the transmission of data chunks. Each HEARTBEAT chunk has to be acknowledged by a HEARTBEAT-ACK chunk.

Each path is assigned a state: it is either active or inactive. A path is *active* if it has been used in the recent past to transmit an (arbitrary) SCTP datagram which has been acknowledged by the peer. If transmissions on a certain path seem to fail repeatedly, the path is regarded as *inactive*.

The number of events where heartbeats were not acknowledged within a certain time, or retransmission events occurred is counted on a per association basis, and if a certain limit is exceeded (the value of which may be configurable), the peer endpoint is considered unreachable, and the association will be terminated.

Path Selection

At the set-up of an SCTP association, one of the IP addresses from the returned list is selected as initial *primary path*. Data chunks are transmitted over this primary transmission path by default. For retransmissions however, another active path may be selected, if one is available. To support the measurement of round trip delays, SACK chunks should be sent to the source address of the IP packet which carried the data chunk that triggered the SACK.

The users of SCTP are informed about the status (state and measurements) of a transmission path on request or when a transmission path changes its state. They may then instruct the local SCTP instance to use a new primary path.