

SCTP: Stream Control Transmission Protocol

An Analysis

Jorge Mena
UC Riverside
jmena@cs.ucr.edu

Ryan Rusich
UC Riverside
rusichr@cs.ucr.edu

January 5, 2006

Abstract

Currently the Internet infrastructure is dominated by the transport protocols TCP and UDP. While TCP provides strict sequencing and reliable delivery and UDP provides plain fast transmission, there exists applications that occupy the best out of the two worlds, such as telephony signaling. Moreover, since the emergence of the Internet, the industry is seeing converge the two largest networks in the world: the IP Network and the Public Switched Telephony Network, making the need of more suitable protocol more apparent. SCTP is a protocol that has been proposed by the IETF in 2000, and defined initially in the RFC 2960. With this situation in mind and after analyzing the current market conditions, the authors of this paper present an analysis of SCTP as a viable solution to the current situation and provide their own perspective and predictions about the future of SCTP.

1 Introduction

SCTP is transport layer protocol that has been proposed by IETF in 2000. The objective of the paper is to provide an analysis and a prediction on the future of SCTP in industry. We undertook a thorough investigation of the SCTP protocol. SCTP is currently supported by SS7 network that was introduced to exploit the packet switched signaling network the telecom companies use to route calls, perform minor book keeping tasks such as billing, as well as provide services like caller ID, call waiting, etc. Neither TCP nor UDP were ideal for the unique requirements SS7 placed on the underlying [25] IP network, so IETF's Signaling Transport (SIGTRAN) produced SCTP. This paper gives a detailed description of the standard, analyzes the protocol, isolates its advantages over competing transport protocols, and investigates through market research and forecasting whether or SCTO has a commercial future.

The rest of the paper is organized as follows: Section II provides a background information about SCTP; Section III presents a synthesis of SCTP; Section IV gives an analysis of who currently supports SCTP; Section V is our analysis; Section

VI is a prediction of the future of SCTP according to our own observations; Section VII concludes our paper; and Section VIII is an appendix that states the methodology used by Deloitte in their own predictions.

2 Background

The Stream Control Transmission Protocol (SCTP) is a connection-oriented protocol in nature [1] that seeks to combine the fast operation of UDP with the reliable, sequencing, and congestion control features of TCP. It is defined as a transport layer protocol by the IETF and accepted as a standard that works at the same level of TCP and UDP to provide transport services to its upper layers [2]. SCTP was conceived to address the needs of IP applications that require features that neither TCP nor UDP are able to provide.

In a scenario put in [3], the popularity of Internet Protocol (IP) networks is evident nowadays thanks to the Internet and the number of applications that are targeted to the end hosts. Moreover, new technologies like optical fibers or wireless standards such as GSM are making the Public Switched Telephone Network (PSTN) improve their services making it a viable media for aggregate services to its customers. This is causing the convergence of IP and PSTN into one network that supports these aggregate services, and SCTP is a serious candidate to satisfy the needs of telephony signaling and the reliability needed by the IP protocol.

2.1 Basic Features

SCTP, like TCP, is an end-to-end protocol that works in full duplex communication. It is also connection-oriented by using a 4-way hand shake mechanism to establish what is called an association between two communicating ends. It is message-oriented, which helps the protocol keep states during operation and react upon events occurring in the network, such as a new connection, disconnection, a failure, etc. Just like TCP, it provides message

bundling and message fragmentation to allow faster data transmission during initializations and the full utilization of the SCTP packets sent across the wire.

Another important feature is its reliable transmission feature that detects when data is discarded, reordered, duplicated or corrupted [2]. This makes it a desirable protocol that addresses the needs of many of the current applications existing in the IP networks. It also includes an exponential backoff mechanism similar to the one found in TCP in order to provide a rate adaptive service with the other end during data transmission. Sequencing is another feature that both TCP and SCTP offer as a mechanism to order the arrival of messages during a communication.

A distinguishing characteristic of SCTP is a solution to Denial of Service attacks, also known as blind attacks that TCP is vulnerable to. SCTP uses private session keys to digitally sign pairs (called cookies) of the session key and a hash of the control information obtained during the initialization phase of the protocol, using message authentication codes (MAC). This mechanism provides authentication and validation of the source [3]. Another feature is path selection and path monitoring. They refer to the ability of SCTP so choose a reliable communication path that is less probable to be subject to connectivity issues. Due to the Multi-homing feature presented next, SCTP is able to monitor alternative paths using special control messages, called HEARTBEATS, which are used when the primary path results broken.

2.2 Multi-Homing

Apart from the basic features that SCTP offers, what really separates this protocol from the others are the Multi-Homing and Multi-Streaming features. Multi-Homing refers to the ability of utilize multiple addresses for the same host in a network environment. In the case of an IP network, it is possible that one host owns two or more distinct public IP addresses from the same or distinct service providers, in the same way one can own two distinct telephone numbers (i.e., home phone, cell phone, and fax numbers) in PSTN.

SCTP uses its multi-homing nature to provide redundancy to the transport layer. During the association establishment, both end points will select one IP address from their pool of addresses and will designate this one as the "primary" path. The remaining pool is utilized as backup in case the connectivity of this primary path fails (a fail is dynamically detected using the special control message HEARTBEAT). Moreover, each end point will exchange their pool of IP addresses in order to let the other end know where this end point can be reached at aside from the primary advertised path. This situation creates at each end point a list of potential, hopefully distinct, paths that can be used to reach the other end point, which effectively adds redundancy to the protocol.

However, an important clarification about multi-homing in [2] is the following: SCTP does not do load sharing [and] multi-homing is used for redundancy purposes only [2]. The multiple interfaces are used mainly as backup during failures of the primary path, but

once this path is restored, normal traffic is once again moved to this path and the other secondary paths are released. In doing so, it guarantees more connectivity when the end point engages with multiple associations with multiple end points.

2.3 Multi-Streaming

In the SCTP context, multi-streaming refers to the parallel transmission of user data over the same association made between two end nodes. The term stream refers to one of these parallel "pipes" in the association and independently carries fragmented messages from one end to the other. This independence of transmission among the streams gives SCTP an advantage over TCP during the transmission of data and the ultimate cumulative throughput achieved.

In TCP, during the transmission of segments between the two end points, it is possible that the protocol wastes its bandwidth due to the strict sequence of message delivery [2] and thus decreases its cumulative throughput due to network failures. This problem is known as head-of-the-queue blocking. In SCTP, if a stream starts to fail due to message loss or network path failure, only this failing stream will block the delivery of its own sequential packets, while the remaining streams will continue to operate normally. This gives SCTP the ability to bypass unimportant messages over important ones by utilizing the knowledge of what the protocol transmits; contrary to TCP that only sees plain bytes that need to be sent over the wire [3].

In practice, the strict sequence delivery of segments might not be absolutely necessary [2]. In telephone communication, it is possible to tolerate this loss as long as the sequence of this stream is maintained; as for the content of the other streams is irrelevant for any other stream. Also online content benefits from this perspective of multiple streams. Important documents that have higher priority of delivery can be transmitted independently from other data such as images or other multimedia used in advertising, etc.

3 SCTP

The following is a synthesis of the protocol details initially described in [2], and later improved upon in [4, 5, 6, 7, 8, 9]. This section is intended to be informative and it includes in no way the personal perspectives of the authors.

3.1 Architecture

SCTP is defined to be a transport layer protocol by the IETF in RCF 2960 (figure 1). It sits in between the network layer and the upper layers that we will call application for short. User data is taken from the upper layers and received data is reassembled and returned to the user applications expecting this data. Inside the transport layer, the data is fragmented and sequenced before it is sent to the lower layers for its delivery. The received data from the lower layers is checked reassembled and checked for

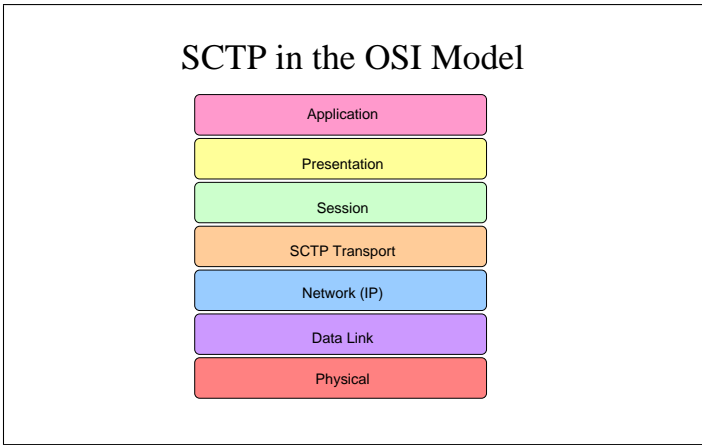


Figure 1: The OSI model and SCTP

validation before it is delivered to the upper layers.

The streams are the sequence of messages that need to be delivered to the users in the upper layers [1]. Since they carry independent data traffic, they can proceed in the event of a single stream failing for any reason. To do this, SCTP provides fragmentation mechanism that makes it able to accommodate large messages into the selected streams. These fragments are defined as chunks and there exists distinct kinds of chunks presented in [1].

Chunks can be batched into SCTP packets; this is done by a bundling function in the SCTP implementation. The other end will be in charge of interpreting this bundling and checking the sequence delivery of each stream. These SCTP packets, in turn, need to be validated by the protocol. This is done by using a Verification Tag field in the packet, which adds a CRC checksum to the packet [5].

For the delivery of these SCTP packets, it is necessary to use a primary path out of the pool of available IP/port pairs that are exchanged during the initialization phase. However, these paths also have to be validated periodically. SCTP uses heartbeat control (chunk) messages to achieve this goal. Both path management and packet verification are done at the same time [1].

For a list of key terms and other more detailed architecture conventions, check [1].

3.2 Packet Structure

Each SCTP packet contains both data and control chunks. The common header section includes information such as the source and destination port numbers to associate it with the application it belongs [3], the verification tag introduced before, and the checksum to test the integrity of the packet. The first chunks can be either control or data chunks (due to early start of data transfer during the initialization phase), but consecutive chunks are data. Finally, the number of chunks in a SCTP packet is determined during the initialization of the association in order to cooperate

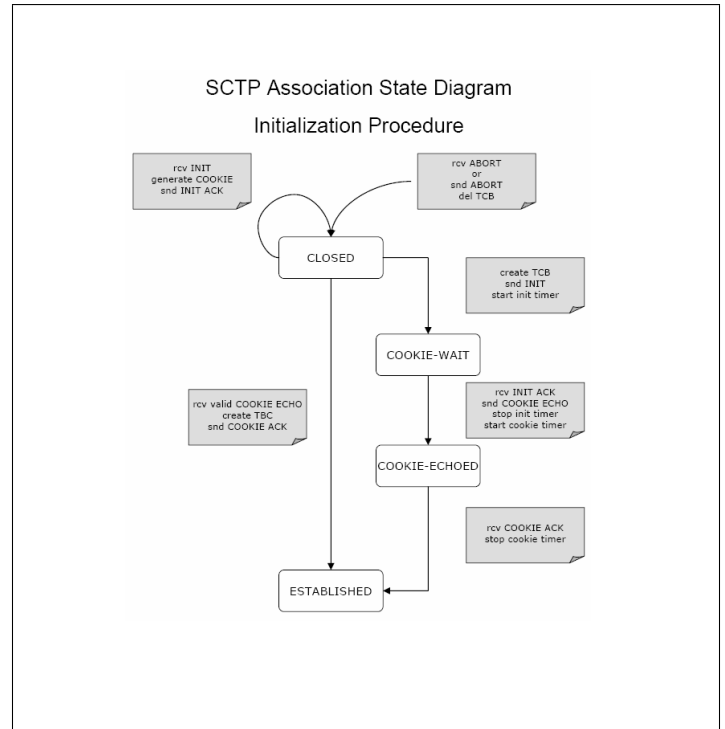


Figure 2: Initialization Phase

with each other end in the communication.

3.3 Chunks

The chunk is the minimal data unit that can be transmitted in SCTP. Contrary to TCP that transmits bytes, chunks in SCTP can be used to control the association between nodes, test the validity of the announce paths, and provide mechanisms for diverse network events, such as failure, disconnection, or abort (half-open connections).

3.4 Initialization Phase

A connection in SCTP is termed "association," and it is much more than the simple three-way hand-shake in TCP. SCTP uses a four-way hand-shake mechanism that seeks to eliminate the SYN in TCP that can initiate a ultimate result in the Denial of Service on the host. A cookie mechanism is used to achieve this purpose. For illustrative purposes, the connecting end will be called A, and the receiving connection will be called B from now on.

The state diagram in figure 2 depicts graphically the initialization phase. All the ends start at the CLOSED state and are listening for network events. There are several events that lead to this state. First, and the less obvious, is the receipt of an ABORT cookie from an end node or the sent of an ABORT due to errors passed in its parameters. If a current association has been previously established, the presence of an ABORT message will lead to this

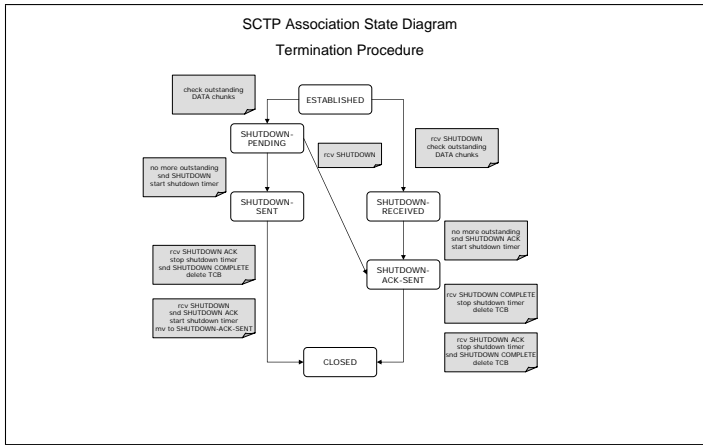


Figure 3: Termination Phase

state during the initialization.

In the case of node A trying to connect with node B, the receipt of an INIT cookie will put the end node B in this state as well. This is because SCTP provides the cookie mechanism to validate the authenticity of its peer end node, A. At this point, B does not allocate any resources for the contacting node, but creates a cookie message that is composed of a hash of the TCB information together with lifetime and a signature for authentication. This cookie is packaged in a chunk message and labeled INIT ACK, and is replied back to A, and B stays in the CLOSED state. On the other end, A has created a TCB entry for the association before sending its INIT chunk; however, as soon as it sends this message, it initialized its init timer and moves to the state COOKIE-WAIT. When A receives a cookie from B, A validates the cookie and sends a COOKIE ECHO chunk back to B; then, it stops its init timer and starts its cookie timer. Now A moves to the COOKIE-ECHOED state. Node B is expecting this COOKIE ECHO message and as soon as it receives it, it creates a TCP entry (allocates resources) and replies back to A with a COOKIE ACK message, and B moves from the CLOSED state into the ESTABLISHED state. When A received the COOKIE ACK, it stops its cookie timer and moves into the ESTABLISHED state as well. At this point, both A and B nodes have establish an association and they can start exchanging DATA chunks among themselves. It is important to notice that it is here where both A and B exchange parameters such as MTU to avoid overfilling of the other nodes buffer, IP address lists for multi-homing, multi-stream parameters.

3.5 Graceful Termination

SCTP provides graceful termination in all its associations made. However, there are occasions when this is not achieved. For example, TCP supports half-open nodes with no problem, but this is not the case of SCTP. In the event an association is not terminated appropriately, ABORT messages are used to cleanly

terminate the rest of the association that stills alive. On the other hand, we have the normal termination of an association that is started with SHUTDOWN messages.

Assuming nodes A and B have established an association, and node A desires to terminate the connection, the protocol uses the flow diagram presented in Figure 3. At the beginning of this phase, both A and B are in the ESTABLISHED state, but A, at some point in time A moves into the SHUTDOWN-PENDING state after it checks its outstanding DATA packets that it needs to transmit before announcing the shutdown. After this happens, A sends a SHUTDOWN chunk message to the other end in order to initiate the process of association termination with B and starts its shutdown timer; A moves now to the SHUTDOWN-SENT state. On the other end, B received the request to shutdown and immediately checks its outstanding DATA chunks; B moves to the SHUTDOWN-RECEIVED state and continues sending its outstanding DATA chunks. At some point B terminates sending its chunks and replies back to A with a SHUTDOWN ACK message and it also starts its shutdown timer. B moves to the SHUTDOWN-ACK-SENT state.

It is possible at this point in the flow diagram that A receives a SHUTDOWN message from B, before A sends the SHUTDOWN and with A in the SHUTDOWN-PENDING state. In this case, B will become the initiator of the shutdown and A will respond to A requests.

Following the normal procedure, when A receives the SHUTDOWN-ACK message, it stops its timer, sends back to B another message called SHUTDOWN COMPLETE, and deletes its TCB entry from its table. At this point A considers its association with B terminated and it moves to the CLOSED state. On the other hand, B receives the SHUTDOWN COMPLETE message from A and stops its timer. B removes the TCB entry for A from its table and finally B moves to the CLOSED state and terminates the association with A.

In the case where A and B changed their roles, A would receive a SHUTDOWN message from B (the opposite case). A would reply with a SHUTDOWN ACK, start its timer, and move to the SHUTDOWN-ACK-SENT. From this point, A would follow the same procedure as B in the previous case and vice versa.

There are special cases when in the middle of the flow diagram, a node receives an unexpected packet. These are special circumstances that are left to the interested reader to further investigate in the RFCs.

4 Who supports SCTP

SS7 - Signal System 7 As described in [23], in the earliest days of telephony, if one wanted to make phone calls, they bought a pair of handsets connect by a wire, and they laid that wire between themselves and someone that they intended to have conversations with. To initiate a call, they could yell into their handset and hope the person on the other end would hear them through their

handset and pick up. Calls between another third party and the handset owner would require a second set of handsets connected by a wire. Three way calling was nonexistent.

Enter the telephone company. Now there were centralized switchboards where a caller need only connect to the centralized hub and a call could be set up by allotting a circuit to two individuals who needed to talk to one another. Eventually dialing of numbers was introduced, under which numbers of clicks (interruptions) on the line allowed the user to directly dial the number they wished to call. These clicks were signals that allowed one to set-up a call.

As discussed in [17], the huge inefficiency in this system was that in order to connect, a caller would dial a number, that number would then be taken by the phone company and compared against a routing table, and be routed to the next switch, which would compare its routing table to the number and route the call to the next switch, and so on until a connection across the voice line circuit could be established. If the callee was in fact already speaking with somebody on a separate connection, or had their phone off the hook, then the voice line would be wasted just trying to set up a call. This was called In Channel Signaling, in that all setup and signaling for a call traveled on the same phone line as the voice. Call waiting was nonexistent.

Enter the digital age. Phone companies realizing that they would be wasting valuable resources trying to accommodate high voice traffic began using a separate digital line to transfer all of the signaling for calls. Under this arrangement, one could try and place a call to a busy phone, and could get a response that the line was busy without setting up a circuit for the call. The loss of one voice line, that could transport thousands of signals while not tying up voice circuits was called Common Channel Signaling (CCS).

Phone calls in the form of voice are circuit switched, signaling messages are packet switched. Immediately the phone companies began to use this packet switched network to provide services such as call waiting, caller ID, busy signal call back as well as sending information such as start time and end time for calls for billing purposes. Other services such as credit card calling cards became a reality. In short SS7 is a packet switched network that sits logically on top of the circuit switched

SS7 is the standardization of the digital signaling network phone companies use. It is the only network that currently supports SCTP. It was the result of Consultive Committee for International Telephone and Telegraphs (CCITT) recommendation on a standard. The 7 represents version seven of the standard though not all were physically implemented.

5 Analysis

SCTP is indeed a robust protocol that adds additional features to the users of the transport layer protocols than those provided by others, such as TCP and UDP. Although they are conceived in distinct time frames, a direct comparison is fair and desirable at the same time, since both of them are globally used and dominant. That has been already discussed in the introduction but in this

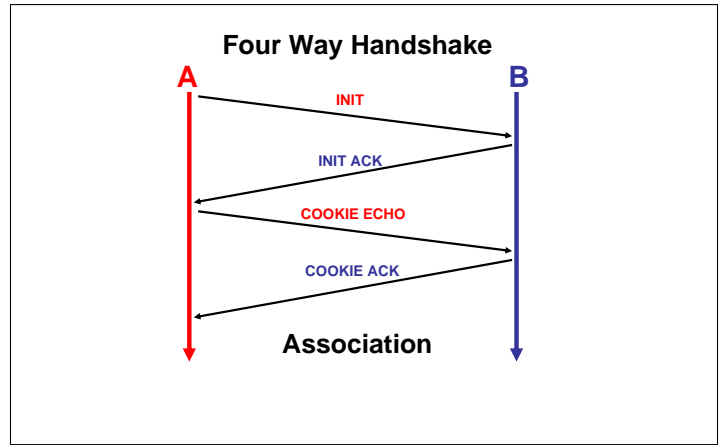


Figure 4: 4-way Handshake

section, we touch upon some of those arguments by presenting the advantages and disadvantages of the protocol and prepare the reader for a discussion of our take on the future of this protocol.

5.1 Advantages

We identified the following advantages in SCTP.

Four-way Hand-shake One of the most noticeable problems found in TCP is its vulnerability to Denial of Service attacks (DoS), also known as blind attacks [10, 11, 12, 13]. SCTP provides a mechanism to authenticate the initiator of an association with the server by the use of a cookie mechanism in a four-way hand-shake. By sending a challenge to the requester of the communication, the server will not allocate any resources at this point, but effectively will transmit information that is necessary to establish the association. This exchange of information helps both ends of the association to build their TCB entries after they have been validated. However, not until the server has received a response to its cookie challenge will it assign any resources devoted to the other end. Thus, if an attacker decides to flood the server with INIT messages, the server will respond with INIT ACK messages to the source of the INIT message. This, of course, can be interpreted as an attack on the source of the INIT message, but the protocol specifies the behavior of SCTP in case of receiving messages unexpectedly. In this particular case, the node receiving the INIT ACK messages will establish an association with the sender of the message, but once it is established, the INIT ACK messages are dropped and ignored by the other end.

Parallel Association Integration Parallel association refers to the attempt of a node to establish more than one association with the same node. This is possible since the protocol assumes

that the nodes are multi-homed. Because it is desirable to have multiple connections between nodes, it is not desirable to have multiple associations, though, due to the introduction of unnecessary additional overhead. For this reason SCTP is a multi-stream protocol, and uses these streams to handle the multiple connections. We believe that an abstraction such as the association that controls all the connections with the end nodes gives SCTP much better control of who connects to the host and how are the resources handled.

Multi-streams Introduced in the previous paragraph, the use of multi-streams is one of the best additions of SCTP to the transport layer. We mentioned before how there exists many applications that do not necessitate the strict sequencing that TCP has during transmission of data. SCTP has not that type of sequencing since it uses multiple streams to transmit data to the requester. Because these streams are independent, and because SCTP is a message-oriented protocol, SCTP has a comparative advantage over TCP by resolving the issue of packet loss. If a stream stops transmitting data due to network problems, the other streams will not be obstructed by this stream and will continue to deliver messages without any problems.

An important observation is the following. Overall, it is expected that SCTP improves its throughput compared with the one obtained in TCP. For example, assuming that there is only one stream between nodes A and B, and this stream starts to drop messages. In this situation, SCTP would degrade at least to how TCP would behave, but the throughput would still be better in SCTP because in case of a packet dropped, TCP retransmits packets that have already received by the other end, and SCTP does not, thus wasting bandwidth.

Redundancy SCTP provides redundancy by the use making the end nodes multi-homing. Each node that participates in an association keeps a list of IP addresses and port numbers and exchanges it with its other end nodes. In case of a failure in its primary path, one of these secondary paths is selected and used without the user knowing that the failure has occurred. This is indeed a desirable feature in any system.

Reachability Monitoring SCTP provides for the support for continuous monitoring of reachability. Through the mechanism of the heartbeat chunk, connectivity is constantly checked to determine whether or not a particular IP to IP connection is available. This becomes important in that a failure of connectivity is immediately detected and the appropriate action can be taken, usually a closing of the association. An added benefit of the heartbeat chunk is that under the condition that not only a primary IP mapping has been established at association initialization, but rather a set of backup IPs has been exchanges

between end nodes as described in the SCTP protocol, a failure to receive a heartbeat ACK can initiate and immediate transfer to a secondary IP to IP connection between end nodes, exploiting the redundancy built into this protocol.

Notorious Network Failures This continuous monitoring of reachability is a hallmark of SCTP and non-existent in the TCP protocol. One might ask why this is of tremendous importance. Consider two notorious network failures that occurred in 1991, and the economic costs as well as increased danger caused. On September 17, 1991 [14] AT&T switching centers in New York City ran out of power. Neighboring airports LaGuardia, Kennedy, and Newark lost all voice and data communications. Five hundred flights were cancelled and another 500 were delayed. Eighty-five thousand passengers were affected by the outage included the Chairman of the Federal Communications Commission. Five million calls were blocked. Ultimately, the Wall Street Journal reported on page C18 of the October 1, 1991 issue that "AT&T Tells FCC a Lapse In Procedure Led to Outage." AT&T revealed that standard procedure called for a supervisor to select a technician to inspect each of the Thomas St. facility's power plants when AT&T switched to its own electrical power from the grid operated by New York utility [Con Ed]. Instead, the supervisor took his technicians to a class on a new power alarm system, leaving the plant unsupervised. To add insult to injury, AT&T found that the plant in question had a faulty alarm system [24].

Earlier that same year, June 26, a circuit board in call-routing equipment facility in Baltimore failed, sounding an alarm that technicians, at the time, did not consider to be overly ominous. The system was supposed to recover gracefully via a built-in software solution. The software turned out to have a minor bug, three bad bits to be exact, and nearly 6.5 million customers in three East Coast States and Washington D.C, became unable to complete local calls [15]. The board failure and underlying software bug caused an SS7 computer to attempt to automatically reconfigure itself to defend against traffic disruption. In the process it[SS7 computer] flooded itself with so many internal messages, it rapidly used up any resources that might have been used to route calls.

There is another less insipient risk to connectivity, a physical disruption or severing of a link that is in the ground. The dangers are considerable, natural disasters like earthquakes or hurricanes; back hoes severing lines. All of these physical disruptions as well as the two examples above illustrate the problem of reachability and the desire for a protocol that can recover gracefully and in a timely manner to the particular outage.

Security SCTP provides for protection against blind denial-of-service attacks. During the handshaking set-up of a connection in a TCP environment, the initial SYN packet that gets sent to an

end point to initiate the handshake process, elicits a SYN-ACK packet. This allows of blind denial of service (DoS) attacks under which a malicious sender uses an invalid sender IP in their SYN packet and floods a server with SYN packets. The attacker has no access of the traffic to and from the target node [RFC2960]. The target nodes kernel resources quickly become depleted with the chore of replying to these farcical SYN requests. TCP has no intrinsic defense to this type of attack, rather the network administrator must take independent precautionary measures like configuring the network to drop repeated and stale queue requests, not responding to unexpected packets, giving priority to completing processes underway over new requests, or performing a reset if traffic to a particular node rises above a certain threshold.

Similarly in a blind masquerade attacks are defended against since the four-way handshake requires that a challenge be met before an association can be met. TCP again has no built in defense mechanism due to the fact that only a two way handshake solidifies a connection. UDP has no defense against either of these types of attacks. In fact UDP being a connectionless transport protocol enables any packet to be sent to any host at any time.

Security Limitations of SCTP. SCTP is superior to TCP due primarily to the cookie based four way handshake process, but it is not impervious to all forms of attack. [RFC2960] Should a man in the middle attack be instantiated under which the attacker is able to not only intercept packets received in association, but alter those packets, this attack would not be stopped by SCTP. This is due to the fact that the INIT ACK packet sends sufficient information for the hijacker actually hijack the association. SCTP also has not defense for attacks that come from outside the SCTP node, or within the connection itself, such as an insider attacks.

5.2 Disadvantages

The following paragraphs discuss some of the disadvantages that we believe will deteriorate the opportunities of SCTP to become a used standard.

Multi-homing The whole idea of multi-homing, besides redundancy, is that the node can do load sharing or load balancing; however, this is not possible with SCTP, as stated in the definition in [1]. Aside from that, SCTP also assumes that the hosts will have a descent number of IP addresses that identify it in the network. In the current situation that exists now, this assumptions turns out to be an expensive one for the customer that does not have or is not willing to pay for the royalties that this implicates. If this protocol is to substitute, or even compete, against TCP, it must provide the same functionality and flexibility to its users in any environment. While this is possible in the telecommunication

sector where large Autonomous Systems would find it desirable in telephony signaling, SCPT would not be suitable in the much simpler home user computer. Nowadays, service providers charge for the rent of a public IP address that can be used to route traffic. It is not plausible to expect customers to assimilate an extra expense for the use of a protocol that will use additional IP addresses as backups of one.

Moreover, considering the current IPv4 pool of addresses as a scarce resource, the protocol would simple deplete the existing IP addresses that are not utilized. One can, however, argue that the IPv6 standard will not suffer of this problem, since it provides the same number of IP addresses as IPv4 squared for use. However, the current infrastructure does not support this standard at the moment and it cannot be considered as an option.

NAT Problem Network Address Translation allows the use of a single IP address to represent multiple hosts connected to one router or firewall. This introduces problems with the protocol and its assumptions. SCTP assumes that the IP addresses that are exchanged during the initialization phase are routable ones, but the protocol does not mention the problems that are originated when the participating hosts are connected in a private network, using unroutable (private) IP addresses. Even worse, it does not detect its existence until they are received in the other end.

During the initialization, the end host, multi-homed with many private IP addresses, will compile a list of its available addresses that would be sent to the other end host. Notice that in this case, we are assuming that the gateway supports SCTP as well and that NAT supports SCTP traffic in the same way. The gateway will normally use its public IP address and log the host that owns this packet in case or response. The SCTP packet will be routed in the global network until it arrives to its destination. This receiving host will find no problem in using the public IP address that belongs to the gateway and it will proceed to evaluate the list of IP addresses passed to end node by the contacting end. At this point, it is undefined of what happens to the protocol. However, let us assume that private IP addresses are not used by SCTP. In this case, all the private addresses would be dropped and the redundancy will potentially be limited to the use of the primary path. However, this assumption is not stated, thus it is possible that the protocol actually validates the use of private IP addresses. If this is the case, then if the end host is behind a private network and "somehow" this SCTP packet got routed to this node, then it would mistakenly try to establish an association with a distinct host connected to its private network.

Also, knowing that NAT is used as a way to share a connection among multiple hosts in a private network, it would only be possible to establish a single association among all the hosts in two distinct private networks that use NAT because of the Parallel Association Integration feature of SCTP described above. This leads to another undefined situation of SCTP where we have streams carrying data that belongs to distinct hosts.

Host Name Resolution SCTP specifies that it is possible to use host names as members of the IP address list that is used to provide redundancy of the association. It also acknowledges that the resolution of host names could be a lengthy process and it could be considered as a vulnerability of the end host that could be exploited by attackers. The way SCTP resolves this is by not resolving the host name at this point but to continue with the association establishment. However, no data can be transmitted until this resolution is complete. If the resolution is not successful, an ABORT message is sent and the connection is terminated.

Assuming that this issue can be discarded as a potential Denial of Service attack, it does introduce a similar problem stated above with the NATs. If the host names of the end host reside in a private network, two cases are possible. First, the host name resolves to the Public IP Address of the gateway, which becomes the same problem as described above. Second, protocol is subject of DNS misconfigurations, which can resolve the host name to an address that is not routable, thus, introducing the potential issue of establishing an association is a host that is not the one that initiated the association establishment.

Dynamic IP Addressing SCTP assumes that an IP address names a host, but this situation gets interesting when we find the use of Dynamic IP addresses by service providers. Every now and then a customer gets assigned a Public IP address dynamically by the routers of the ISP. Assuming that an association is established, there are no guarantees of when the IP address of one host would change; if this happens, the parameters (and their TCBS) would be corrupted and the association will pass to an undefined state. In the best case, both nodes participating in the association would try to send message chunks to the new owners of the IP addresses dynamically allocated and eventually drop their association. However, this opens up vulnerability for an attacker to pick up the association and masquerade as the previous node.

Primary Path selection It not clear on how this path is selected by the host initially; even less on what parameters are used to select the remaining backup paths. If this is left to the implementation details, then the security of the protocol could become compromised by making it dependent on whoever implements the protocol. The use of random seems to be a very nave way of selection, since it does not take into considerations how routing policies among Autonomous Systems are handled, for example.

Required Modification of the Current Infrastructure: In order to support SCTP, it is necessary to investigate how many changes are necessary to make to the current infrastructure. To start, it is

necessary to know how an SCTP packet will be handled when routers look inside the IP packet payload (i.e., when routers do packet filtering using Access Control Lists). This imposes a big problem if SCTP does not provide a mechanism that closely imitates the behavior of TCP, otherwise, it would be required to change the operating system of all routers in the network, which make it prohibitively costly. On top of that, end users will need to upgrade their system to include an implementation of the protocol, which will add another cost to the users. Until all these issues with the routers are resolved with considerably costs, the need to change routes will be an obstacle to the global adoption of SCTP.

Ordered and unordered data delivery on a per-stream basis

TCP also allows for strict ordered data delivery, but suffers from one key drawback. It cannot proceed with a transmission (single stream) if a particular sequenced packet has not yet arrived and been acknowledged. This phenomenon is referred to as head of the line blocking. In TCP, only a single stream is allowed, which also complicates the identification of beginning and ends packet transfers, demanding that this information be transmitted in the individual packets as well. The largest setback to the head of the line blocking is the inability to continue to transfer data, until the lost packet has been retransmitted and acknowledged as received. SCTP contains all the ordered data delivery capability of TCP with an added advantage. Multi-streaming ensures that if a head of the line blocking situation arises in an individual stream, the hold up of transmission until the lost packet is retransmitted and the streams sequence preserved has no effect on all other streams, that currently comprise the association. As an added flexibility provided feature, individual streams can be designated as unordered in which case no blocking will occur.

UDP does not provide for any order preserving data delivery. It is a connectionless best effort transport protocol

6 Future of SCTP

6.1 Industry & Market Trends

Unlike the eventual displacement of CRT monitors by the more progressive flat screen LCDs and plasma screens voice transmission is not going to go away. People will continue to demand the ability to transfer their voices over existing and ensuing network connections. Deloitte & Touche LLPs Technology, Media and Telecommunications (TMT) industry group, a paid market consultant, in their January 2005 report on Top Trends for Telecom in 2005, reported that there would be nearly two billion mobile phone subscriptions by years end, with strongest growth coming from developing nations in South America and Asia where land line infrastructure is weak and cost prohibitive to implement. Secondly they predict on average that over 80% of all rev-

venues from these subscriptions will come from voice. A more ambitious, yet dubious prediction is that of 100+% penetration as many customers take second subscription for data or personal use. A reasonable conclusion to draw is that voice and data will share a common network, which will interface with the more heterogeneous Internet.

A second more poignant prediction Deloitte makes is that of continued preference on PSTN (Public Switched Telephone Network) service due to superior call quality and reliability, but continued growth in Voice over IP (VoIP). Admittedly a majority of calls will still originate from and terminate to PSTNs. However some companies are predicted to transfer internal communication over to VoIP, while continuing to use PSTN as a portal to the outside world. In general VoIP is expected to only grow its market share. Deloitte's Prediction Methodology is included at the conclusion of this paper as Appendix.

6.2 Plausability

The bottom line is money. Nothing in business is done without it in mind. Large telecom companies like Verizon and SBC are currently reconstituting Ma Bell. The irony is that the U.S. government broke up AT&T into regional pieces in the late 1980s because of unfair business practices and other anti-trust concerns. However in an evolving global market, these corporations need leverage and power. Edward Whitaker, chief executive of SBC has publicly stated his intention to reform AT&T to its former glory. We begin to get an idea of the costs by looking at changes that must be made in the hardware and software that will come to support SCTP.

One logistical problem that arises with the possibility for the adoption of SCTP is the network routers. All network routers would need to be reconfigured to handle the protocol. A logical time for doing so would be when IPv6 actually came on line as the wholesale change could be implemented suddenly rather than by piecemeal. This would entail the owners of those routers, mostly telecom businesses to fit the bill. This outlaying of capital would only be undertaken if it resulted in one of two conditions; future profits, or increased market share. If the implementation of this change could somehow hinder competition, by weakening opponents technology then it could occur. We would argue that a better transfer protocol can bring better reliability and service. When one considers reliability and service the benchmark is the five nines (99.999 call success rate)[Wireline Predictions] offered by PSTN providers like AT&T.

A second consideration is the updating of entire Network Address Translators (NATs) with the ability to look inside chunks to translate the entire list of IP addresses that provide for the multi-homing of SCTP. Without this ability, the association degrades back to a single IP to IP connection and the multi-homing feature is completely eliminated. Whether or not the backup IPs are inside the packet, the NATs would need, via software to look inside the packet. In an IPv6 world NATs would not be necessary

since there should be an abundance of IP addresses. This assumes of course that IP addresses are not wasted unnecessarily.

VoIP does not need SCTP to work. VoIP might benefit from the protocol with added reliability and increased throughput. VoIP is a catalyst for another reason. VoIP is an irritant to the big telecom companies. First they [19] can charge far less because they send the call in packets along with all other packet based traffic. This allows for multiplexing and does not demand a dedicated circuit like a PSTN call. This has drawn the ire of SBC CEO Ed Whitaker. Here is a segment of an interview published in Business Week in November of 2005, slightly more than a month before this writing. The question goes to the heart of what telecom feels its obligations are to providing cheap and easy access to their resources.

Q: How concerned are you about Internet upstarts like Google, MSN, Vonage, and others?

A: How do you think they're going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes?

Which is in stark contrast to his statement about SBC's \$16 Billion acquisition of AT&T, "The commission vote demonstrates recognition that the merger of SBC and AT&T will enhance competition, help bring new technologies to market faster, and provide real benefits to consumers and businesses."

7 Conclusion

Telecom will probably move towards more utilization of the SS7 network, by pushing for enhanced handsets similar to the mobile ones people use today. This will include ringtones and wallpapers, and other packet transferred data files. Currently however home telephone handsets spend 1% on R&D that mobile phone manufacturers do.[TMT trends]. The implication is that large telecom companies will continue to extract as much profit as they can from any of their infrastructurally mature capital resources. This does not ensure that they will push the state of the art or foster blossoming technologies.

The network developing community would enjoy being able to implement every outstanding technology they invent. Sadly this is not the case. Profit margins drive what rolls out. Mr. Whitaker started with his phone company as a linesman, one who lays lines for calls. Before SS7, to handle the growing call demand, more lines were simply laid. He is not a Computer Scientist and cannot be expected to lead his corporation in a manner that is consistent with championing novel approaches. SCTP cannot be established without telecom's support due to their overwhelming control of the network topology. Possibly in another iteration of

the Internet it will happen, in the near future probably not.

8 Appendix

From [20], Deloitte, Inc.

8.1 Predictions Methodology

These predictions have been compiled by Deloitte Research (a part of Deloitte Services LP) on behalf of Deloitte & Touches Technology, Media and Telecommunications (TMT) Group. The major inputs used in writing the predictions were: input from the 5,000 strong TMT team around the world; discussions with leading industry and financial analysts; interaction and conversations with clients from the telecommunications and related sectors. These predictions do not claim to be fully comprehensive, but rather provide a commentary on major industry trends and developments.

8.2 About Deloitte & Touches Technology, Media & Telecommunications (TMT) Group

The TMT Group is composed of service professionals who have a wealth of experience serving technology, media and telecommunications companies throughout the world in areas including cable, communications providers, computers and peripherals, entertainment, media and publishing, networking, semiconductors, software, wireless, and related industries. These specialists understand the challenges that these companies face throughout all stages of their business growth cycle and are committed to helping them succeed. Deloitte & Touche is a leader in providing strategic, financial and operational assistance to its technology, media and telecommunications clients.

Bibliography

- [1] RFC 2960 <http://www.ietf.org/rfc/rfc2960.txt>
- [2] RFC 3286 <http://www.ietf.org/rfc/rfc3286.txt>
- [3] Stream Control Transmission Protocol (SCTP), The International Engineering Consortium <http://www.iec.org>
- [4] RFC 3257 <http://www.ietf.org/rfc/rfc3257.txt>
- [5] RFC 3309 <http://www.ietf.org/rfc/rfc3309.txt>
- [6] RFC 3436 <http://www.ietf.org/rfc/rfc3436.txt>
- [7] RFC 3554 <http://www.ietf.org/rfc/rfc3554.txt>
- [8] RFC 3758 <http://www.ietf.org/rfc/rfc3758.txt>
- [9] RFC 3873 <http://www.ietf.org/rfc/rfc3873.txt>
- [10] Denial of Service Attacks, CERT Coordination Center http://www.cert.org/tech_tips/denial_of_service.html
- [11] Distributed Denial of Service Attacks, Dave Dittrich <http://staff.washington.edu/dittrich/misc/ddos/>
- [12] L. Garber, Denial of Service, Rip the Internet, IEEE Computer, 2000, www.ieeexplore.ieee.org
- [13] F. Lau, S.H. Rubin, M. H. Smith, L. Trajkovic. Distributed DoS, Proc IEEE Int Conf Syst Man Cybern, 2000. www.ieeexplore.ieee.org
- [14] Wicker, S. B. Complexity and Constraint in Complex Adaptive Networks. School of Electrical and Computer Engineering, Cornell University.
- [15] Burgess John. A Little 'Bit' Goes a Long Way in Fouling Up Complex Networks, International Herald Tribune, 1991, http://www.iht.com/bin/print_ipub.php?file=/articles/1991/10/08/soft.php
- [16] The Risks Digest, Email Thread, <http://catless.ncl.ac.uk/Risks/12.43.html>
- [17] SS8 Networks, SS7 Tutorial, <http://www.ss7.com/History.pdf>
- [18] Reimer J. Verizon and SBC purchases cleared by FCC. ArsTechnica, 2005, <http://arstechnica.com/news.ars/post/20051031-5505.html>
- [19] Business Week Online, At SBC, It's All About "Scale and Scope", November 2005, http://www.businessweek.com/@n34h*IUQu7KtOwgA/magazine/content/05_45/b3958092.htm
- [20] Deloitte, TMT Trends: Predictions, 2005, A Focus on the wireline sector. Internet PDF, 2005 http://www.deloitte.com/dtt/cda/doc/content/Wireline_FINAL_01FEB05_LR_FA_LOC KED%28%29.pdf
- [21] Stewart, R., Xie, Q., Stream Control Transmission Protocol (SCTP): A Reference Guide, October, 2001
- [22] SCTP, <http://www.sctp.org/>
- [23] Arias Rodríguez, Iván, Stream Control Transmission Protocol, The design of a new reliable transport protocol for IP networks.
- [24] Scecof, Mark, AT&T tells FCC a lapse in Procedure led to Outage, Wall Street Journal, October, 1991.
- [25] Stafford, M. Signaling & Switching for Packet Telephony. Artech House Telecommunications Library, 2004