Deploying Services over IP NGN - SIGTRAN Design Considerations

What You Will Learn

This white paper will initially discuss the requirements to carry services as voice and signaling over an IP/Multiprotocol Label Switching (MPLS) network. A more detailed analysis will focus on design aspects of how to implement SIGTRAN over an IP network, with a final description of protocol optimization.

A conclusion summarizes the advantages of understanding and optimizing Stream Control Transmission Protocol (SCTP) over just the IP/MPLS core network.

Traditional Requirements to Carry Voice and Signaling over Core Packet Network (CPN)

Both mobile and fixed service providers have been migrating over the past decade to the IP next-generation network (NGN) in order to provide services such as triple play - voice, video, and Internet over the same network. The shrinking margins from the basic voice and Internet business has created an urgent need for new value-added services that will increase revenues and profits. Building an IP/MPLS core and converging all networks into a single one that makes use of both the new IP services and technologies together with the legacy ones will provide a major factor in getting savings.

Examples are the migration of voice and signaling services over an IP core.

Additional drivers for a converged IP packet core are:

- Reduced factor per transported bit
- Converged transport layer around IP/Ethernet
- Improved reliability and availability
- Flexible connectivity, easy to manage
- Rich service-aware functions with IP
- Reduced provisioning
- More capital expenditure (CapEx) and operating expenses (OpEx) efficient

When moving into a CPN core, the transport requirements in the network to deploy voice and signaling services are a primary issue. Operators have seen diverse values coming from different vendors over a period of years, as depicted in Figures 1, 2, and 3.

Figure 1. Transport Requirements for MSC Vendor 1

Transport Requirements for MSS

	CS User Plane	Signaling
Average delay (one way, lower 95%)	< 20 ms	< 20 ms
Maximum delay (one way)	< 100 ms	< 100 ms
Delay variation (jitter)	< 5 ms	N/A
Packet loss	< 10-4 (assuming voice concealment)	< 10-4
Max fail-over time	2 s (200 ms if end-user shall perceive no voice interruption)	N/A (SCTP path deversity assumed)

- Traffic Prioritization
- 99.999% Availability of the IP Backbone

Figure 2. Transport Requirement for Vendor 2

QoS Parameters	Delay (ms)		Jitter (ms)		Packet Loss Rate	
	Ideal	Max	Ideal	Max	Ideal	Max
Real-time service	<20	<40	<8	<15	<0.05%	<0.05%
Non real time	<20	<60	<8	<20	<0.05%	<1%

Figure 3. Transport Requirements for Vendor 3

QoS Parameters	Delay (ms)		Jitter (ms)		Packet Loss Rate	
	Ideal	Max	Ideal	Max	Ideal	Max
Iub real-time service	<5	<30	<5	<10	<0.01%	<0.01%
Iub non real time	<10	<100	<8	<20	<0.01%	<1%

There is a disparity of values that range from less than 5 msec to less than 20 msec in the transport network for real-time services. In order to put previous numbers into perspective, we need to consider:

- Each network is different in terms of span, numbers of elements, distance, physical connectivity, and other factors. We can't apply the same numbers to different networks. In reality, we measure these values (delay, jitter, packet loss) once the services such as voice and SIGTRAN have been deployed, finding average delay values that range from 50 msec to 5 msec depending on the network.
- We need to understand that mobile operators will incur interruptions in service at 100 msec with a hard handover (a handover where all the old radio links in the User Equipment (UE) are removed before the new radio links are established), so delay values of 5-20 msec are not relevant.

Another important requirement that has been requested from operators was to comply with the 50 msec protection time when designing a CPN network. This requirement comes from the 1988 ITU G.841, which governs SDH network protection architectures. The reason for this value lies in the use of voice channel banks in carrier networks in the early 1980s that could not tolerate failures that lasted longer than 200 ms to 300 ms. When failures reached that duration, a Carrier Group Alarm (CGA) would be activated, causing the channel bank to perform a "trunk conditioning" procedure that would terminate all connections carried over that given T3 or E3 trunk line. Carriers did not want to have calls dropped from a fiber trunk that was protected, so an outage budget was developed for this fault condition. Since the outage budget had to be less than 200 ms, and 50 ms protection time was already established as a de facto standard, 50 ms became forever associated with "voice service protection."

But by the time the requirement was actually adopted, newer channel banks, as well as digital telephone switches with T1/E1 ports, implemented a CGA timer of 2 seconds. So the 50 ms protection time was adopted to protect a small and diminishing fraction of the digital transmission plant.

Nowadays, the voice channel banks that prompted the original fast restoration requirements are not being used, so modern network applications don't need 50 ms recovery.

Real Requirements to Carry Signaling over CPN

By converging network services into a single infrastructure, multiple requirements emerge and new properties have to be considered. The services that require tight service-level agreements (SLAs) are voice and signaling, and we are going to study the real requirements for SIGTRAN.

SIGTRAN is the name, derived from *signaling transport*, of the former IETF working group that produced specifications for a family of protocols that provide reliable datagram service and user layer adaptations for Signaling System 7 (SS7) and ISDN communication protocols. The SIGTRAN protocols are an extension of the SS7 protocol

family. It supports the same application and call management paradigms as SS7 but uses an IP transport called Stream Control Transmission Protocol. Indeed, the most significant protocol defined by the SIGTRAN group is SCTP, which is used to carry public switched telephone network (PSTN) signaling over IP.

SCTP is an IP transport protocol, similar to Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) and provides transport layer functionality to many Internet-based applications. It has been designed to overcome the problems found in TCP when dealing with sensitive signaling traffic.

When deploying SIGTRAN over an IP/MPLS core, we have to change the paradigm, the focus has to be set in the proper design of SCTP parameters to adapt to the actual CPN network, not requiring unrealistic transport values to the network, as we have seen in Figures 1, 2, and 3. As a simile, we could think of the network as a bottle that will contain SCTP, the liquid. The liquid will adapt to any container shape, provided we have properly designed the liquid parameters.

The purpose of this paper is not to cover SCTP in detail but to indicate the tools available to design a proper SCTP over CPN network.

In order to carry SIGTRAN over CPN, we will need to comply with the following requirements:

- SCTP to work in multihomed topology: Each endpoint will use multiple IP addresses and/or multiple network interfaces.
- SIGTRAN path diversity at both the physical and logical layers: Underlying IP network will provide two totally diverse paths. Any two SCTP associations between two signaling endpoints (SEPs) will have no common resources on the IP/MPLS path, irrespective of the Mobile Vendor software release.
- Optimize SCTP SIGTRAN timers to adapt to the specific IP/MPLS network.

As we can see, we don't have an exact transport requirement in terms of delay or jitter in the CPN to carry SIGTRAN, but we will need to provide resiliency and fine-tune SCTP.

Design Recommendations for SIGTRAN

SIGTRAN is transparent of the underlying MPLS design due to the SCTP multihoming built-in protocol, but in order to provide resiliency, we need to assure physical path diversity for multihoming SCTP associations.

Architectural choices depend on the redundancy options and complexity that we want to configure in the nodes that will transport SIGTRAN traffic.

In Figure 4, we have a simple design for SIGTRAN, where two VLANs and two VPNs are used between the SEPs. The CPN network will have to provide path diversity avoiding a single point of failure.

Figure 4. SIGTRAN Architecture

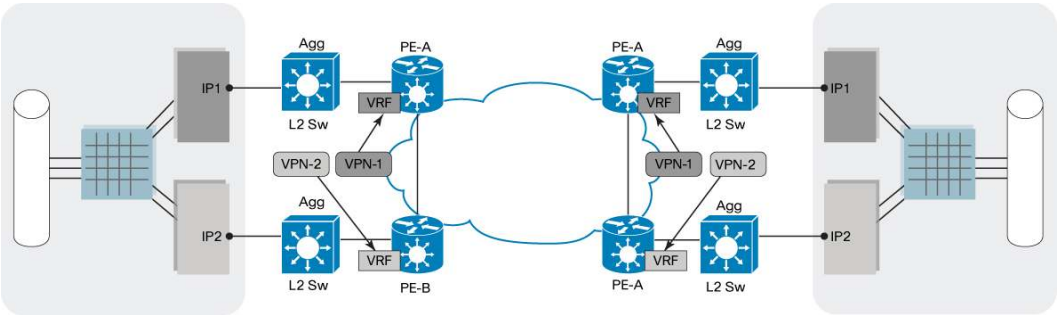

In Figure 5, we have additional redundancy options with Hot Standby Router Protocol (HSRP) implemented in the provider edge nodes.

Figure 5. SIGTRAN Architecture with Redundancy Options

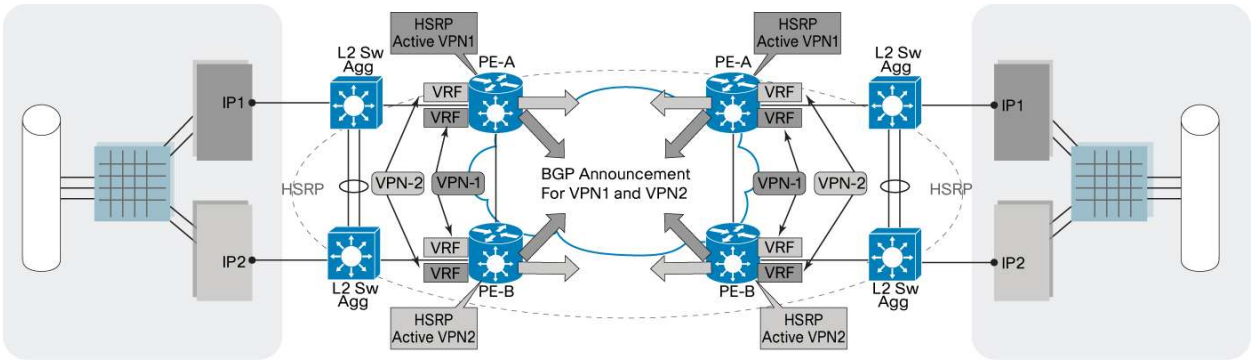

Understanding SCTP Parameters

SCTP provides several protocol parameters that can be customized by the upper layer protocol. These protocol parameters can be customized to control and influence SCTP performance behavior. Different network designs and implementations pose their own unique performance requirements. It is not possible to provide customized protocol parameters that are suitable for all implementations. The tuning information is provided as a guide for understanding what the SCTP parameters are and how they affect the various SCTP algorithms.

Connection Establishment

The protocol parameters assoc-retransmit, init-retransmit, and init-timeout can be customized to control connection establishment. During SCTP association initialization, sometimes packet retransmissions occur. The first initialization packet timeout occurs after 1 second. When initialization packet retransmissions occur, the timeout value is doubled for each retransmission. The maximum timeout value is bound by the init-timeout parameter. The init-timeout parameter is used to control the time between initialization packet retries. As a general rule, init-timeout should be configured to reflect the round-trip time (RTT) for packets to traverse the network. An init-timeout value that is too small can cause excessive retries of initialization packets. Large init-timeout values can increase connection establishment times.

The number of retries allowed for connection establishment packets is controlled by the init-retransmit protocol parameter. When selecting the number of retries, the number of attempts should take into account varying network conditions that may prevent initialization packets from traversing the network.

The defaults used by M3UA/M2PA are recommendations from RFC 4960.

SCTP Multihoming

A key feature of SCTP is multihoming. An SCTP endpoint is considered multihomed if more than one IP address can be used as a destination to reach that endpoint. Upon failure of the primary destination address, SCTP switches to an alternate address.

It is important to note that each vendor has a different implementation of the SCTP multihoming algorithm (as per RFC 4960 6.4.1, "*Rules for picking the most divergent source-destination pair are an implementation decision*"), and it is key to understand each vendor algorithm to provide a proper parameter optimization.

In a Cisco Signaling Gateway: IP Transfer Point (ITP), the configuration of a multihomed endpoint, the first remote IP address specified on the peer link, is defined as the primary address. If the primary address is determined to be unreachable, SCTP multihoming switches to one of the alternate addresses specified on the peer link. SCTP will monitor the reachability of the failed destination address. Upon notification that reachability is reestablished to the primary address, M3UA/M2PA directs SCTP to switch back to the primary address.

SCTP sends the data to the primary address. If timeout occurs, it resends the data to an alternate address. All new data will still be sent to the primary address. If a timeout occurs again, it resends the data to the alternate address. It continues this process until path-retransmit is reached on the primary address. Once path-retransmit is reached on the primary address, it marks the primary address unreachable and sends all data to the alternate address. It begins heartbeats on the primary address. When the heartbeat is successful, it marks the primary address reachable and starts sending data to the primary address again.

The protocol parameters path-retransmit and retransmit-timeout can be customized to control how long SCTP waits before switching to an alternate address. The path-retransmit parameter controls the number of times that SCTP attempts to retransmit a packet before declaring the destination address unreachable.

The retransmit-timeout parameter is used to determine whether a packet must be retransmitted. If an acknowledgement is not received by the time the retransmission timer expires, all packets that have been transmitted, but not acknowledged are retransmitted.

SCTP Optimization for the CPN Network

Operators have different options to adapt SCTP to the IP/MPLS network. The following highlights the main considerations for optimizing SCTP parameters:

- If we have a different vendor in each SEP, it is highly recommended to match the SCTP parameters in both sides to avoid unnecessary SCTP retransmissions.
- As the SCTP multihoming vendor implementation is open, the behavior upon a failure could be different in a scenario where the CPN network provides two paths to each SEP or four paths as illustrated in the Figure 6.

Figure 6. SCTP Multihoming Paths between Two SEPs

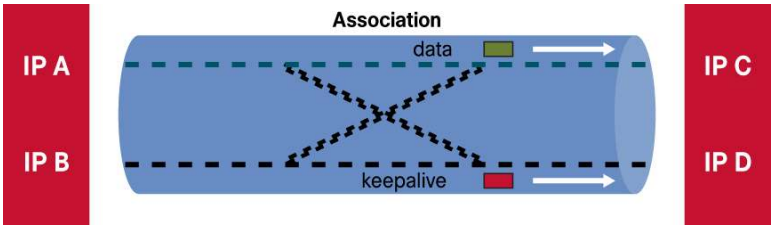

For simplicity purposes, many operators implement in the network two paths instead of the four paths possible, so the heartbeat (HB) mechanism keeps tracks of the two possible paths A-C and B-D, marking the paths A-D and B-C as unreachable in the case of Cisco ITP. If the other SEP SCTP implementation follows another mechanism, the SCTP parameters have to be adjusted accordingly.

- By lowering the minimum retransmission timeout (RTOmin) parameter, the failover times as well the maximum message delays can be further reduced. However, with very low values of RTOmin, associations may become more susceptible to early, unwanted retransmission timer timeouts, and thus retransmissions. With very low Path Retransmission Limit (PRL) values, this may even result in use of the secondary path before any actual failure has occurred, so it is generally not recommended to lower the RTOmin parameter below RTOmin_recommended = 2 * RTT. These spurious timeouts must also be avoided since they have a negative effect on the protocol throughput.
- RTOmax can be different in each customer, as it is tunable to adapt to each IP/MPLS network. The rule of thumb is: Verify the RTT in the worst-case scenario, in the suboptimal path. Make the RTO significantly bigger than the experienced RTT.

- The trade-off of bounding the maximum RTO close to the minimum RTO is the frequency of retransmissions versus increasing transmission delays for packets on the transmit queue. We have normally deployed values of RTO max < 2 x RTOmin or equal values for both, taking into account the compromise between allowing a long delay and having responsive switching to an alternate IP address.
- Another aspect is the HB-TIMEOUT and RTOmax. If the heartbeat interval is too close to the RTO values, that is, MAX_RTO = 2000, HB_TIMEOUT = 2000, this means the time we wait for appropriate heartbeat response (RTO) is too close to the heartbeat interval itself, hb_timeout + RTO ms. We don't want to hit the max_path_retrans too quickly and mark paths inactive by sending heartbeats too frequently and marking them unsuccessful within RTO just because the stack is busy processing other high-priority messages. So, the heartbeat interval has to be set to a value a lot greater than the max_rto.

Conclusion

Operators have been experiencing problems with the migration of services into the CPN network, correlated with the increase of voice, signaling, and data traffic. A proper design has to be implemented to focus not only on the CPN transport core but more importantly on each of the service protocols that will be carried over the IP/MPLS network.

We have seen that when we optimize each protocol at the correct level, we can get better design and avoid complexity in the CPN network. A proper parameter setting of SCTP provides a resilient network with the means to avoid congestion and drops.

For More Information

The paper "Understanding Congestion and SCTP Multihoming with Cisco ITP and Mobile Vendor MSC" shows in detail congestion, SCTP drop analysis, troubleshooting, and fine-tuning over CPNs.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C07-651864-00 04/11