

Network Working Group
Request for Comments: 4553
Category: Standards Track

A. Vainshtein, Ed.
Axerra Networks
YJ. Stein, Ed.
RAD Data Communications
June 2006

Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes a pseudowire encapsulation for Time Division Multiplexing (TDM) bit-streams (T1, E1, T3, E3) that disregards any structure that may be imposed on these streams, in particular the structure imposed by the standard TDM framing.

Table of Contents

1. Introduction	3
2. Terminology and Reference Models	3
2.1. Terminology	3
2.2. Reference Models	4
3. Emulated Services	4
4. SAToP Encapsulation Layer	5
4.1. SAToP Packet Format	5
4.2. PSN and PW Demultiplexing Layer Headers	5
4.3. SAToP Header	6
4.3.1. Usage and Structure of the Control Word	8
4.3.2. Usage of RTP Header	9
5. SAToP Payload Layer	10
5.1. General Payloads	10
5.2. Octet-Aligned T1	11
6. SAToP Operation	12
6.1. Common Considerations	12
6.2. IWF Operation	12
6.2.1. PSN-Bound Direction	12
6.2.2. CE-Bound Direction	13
6.3. SAToP Defects	14
6.4. SAToP PW Performance Monitoring	15
7. Quality of Service (QoS) Issues	16
8. Congestion Control	16
9. Security Considerations	18
10. Applicability Statement	18
11. IANA Considerations	20
12. Acknowledgements	20
13. Co-Authors	20
14. Normative References	21
15. Informative References	22
Appendix A: Old Mode of SAToP Encapsulation over L2TPv3	24
Appendix B: Parameters That MUST Be Agreed upon during the PW Setup	24

1. Introduction

This document describes a method for encapsulating Time Division Multiplexing (TDM) bit-streams (T1, E1, T3, E3) as pseudowires over packet-switching networks (PSN). It addresses only structure-agnostic transport, i.e., the protocol completely disregards any structure that may possibly be imposed on these signals, in particular the structure imposed by standard TDM framing [G.704]. This emulation is referred to as "emulation of unstructured TDM circuits" in [RFC4197] and suits applications where the PEs have no need to interpret TDM data or to participate in the TDM signaling.

The SAToP solution presented in this document conforms to the PWE3 architecture described in [RFC3985] and satisfies both the relevant general requirements put forward in [RFC3916] and specific requirements for unstructured TDM signals presented in [RFC4197].

As with all PWs, SAToP PWs may be manually configured or set up using the PWE3 control protocol [RFC4447]. Extensions to the PWE3 control protocol required for setup and maintenance of SAToP pseudowires and allocations of code points used for this purpose are described in separate documents ([TDM-CONTROL] and [RFC4446], respectively).

2. Terminology and Reference Models

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.1. Terminology

The following acronyms used in this document are defined in [RFC3985] and [RFC4197]:

ATM	Asynchronous Transfer Mode
CE	Customer Edge
CES	Circuit Emulation Service
NSP	Native Service Processing
PE	Provider Edge
PDH	Plesiochronous Digital Hierarchy
PW	Pseudowire
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
TDM	Time Division Multiplexing

In addition, the following TDM-specific terms are needed:

- o Loss of Signal (LOS) - a condition of the TDM attachment circuit wherein the incoming signal cannot be detected. Criteria for entering and leaving the LOS condition can be found in [G.775].
- o Alarm Indication Signal (AIS) - a special bit pattern (e.g., as described in [G.775]) in the TDM bit stream that indicates presence of an upstream circuit outage. For E1, T1, and E3 circuits, the AIS pattern is a sequence of binary "1" values of appropriate duration (the "all ones" pattern), and hence it can be detected and generated by structure-agnostic means. The T3 AIS pattern requires T3 framing (see [G.704], Section 2.5.3.6.1) and hence can only be handled by a structure-aware NSP.

We also use the term Interworking Function (IWF) to describe the functional block that segments and encapsulates TDM into SAToP packets and that in the reverse direction decapsulates SAToP packets and reconstitutes TDM.

2.2. Reference Models

The generic models defined in Sections 4.1, 4.2, and 4.4 of [RFC3985] fully apply to SAToP.

The native service addressed in this document is a special case of the bit stream payload type defined in Section 3.3.3 of [RFC3985].

The Network Synchronization reference model and deployment scenarios for emulation of TDM services are described in [RFC4197], Section 4.3.

3. Emulated Services

This specification describes edge-to-edge emulation of the following TDM services described in [G.702]:

1. E1 (2048 kbit/s)
2. T1 (1544 kbit/s); this service is also known as DS1
3. E3 (34368 kbit/s)
4. T3 (44736 kbit/s); this service is also known as DS3

The protocol used for emulation of these services does not depend on the method in which attachment circuits are delivered to the PEs. For example, a T1 attachment circuit is treated in the same way

regardless of whether it is delivered to the PE on copper [G.703], multiplexed in a T3 circuit [T1.107], mapped into a virtual tributary of a SONET/SDH circuit [G.707], or carried over an ATM network using unstructured ATM Circuit Emulation Service (CES) [ATM-CES]. Termination of any specific "carrier layers" used between the PE and CE is performed by an appropriate NSP.

4. SAToP Encapsulation Layer

4.1. SAToP Packet Format

The basic format of SAToP packets is shown in Figure 1 below.

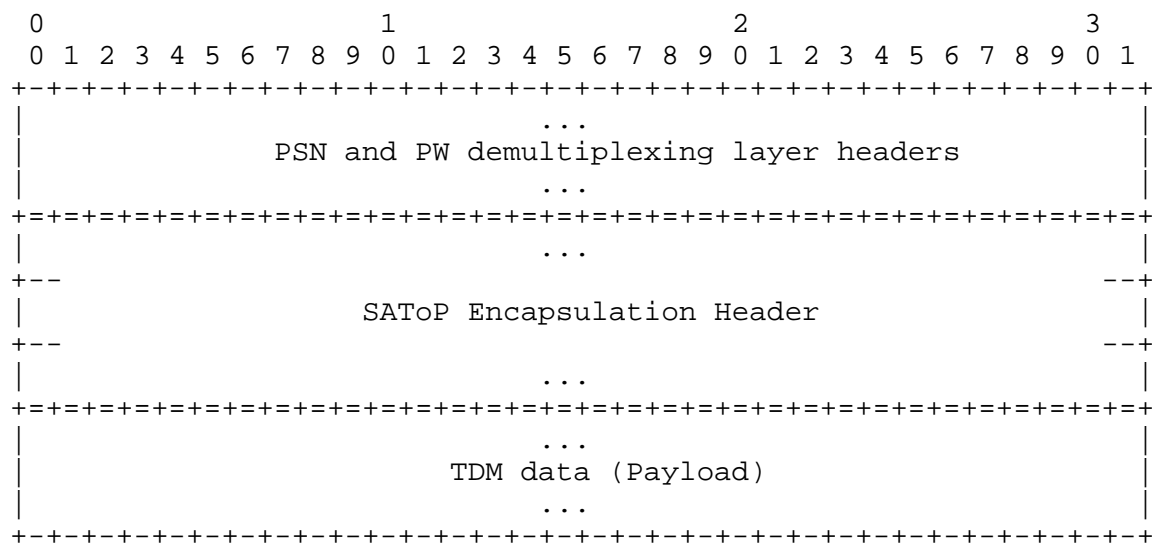


Figure 1. Basic SAToP Packet Format

4.2. PSN and PW Demultiplexing Layer Headers

Both UDP and L2TPv3 [RFC3931] can provide the PW demultiplexing mechanisms for SAToP PWs over an IPv4/IPv6 PSN. The PW label provides the demultiplexing function for an MPLS PSN as described in Section 5.4.2 of [RFC3985].

The total size of a SAToP packet for a specific PW MUST NOT exceed path MTU between the pair of PEs terminating this PW. SAToP implementations using IPv4 PSN MUST mark the IPv4 datagrams they generate as "Don't Fragment" [RFC791] (see also [PWE3-FRAG]).

4.3. SAToP Header

The SAToP header MUST contain the SAToP Control Word (4 bytes) and MAY also contain a fixed RTP header [RFC3550]. If the RTP header is included in the SAToP header, it MUST immediately follow the SAToP control word in all cases except UDP multiplexing, where it MUST precede it (see Figures 2a, 2b, and 2c below).

Note: Such an arrangement complies with the traditional usage of RTP for the IPv4/IPv6 PSN with UDP multiplexing while making SAToP PWs Equal Cost Multi-Path (ECMP)-safe for the MPLS PSN by providing for PW-IP packet discrimination (see [RFC3985], Section 5.4.3). Furthermore, it facilitates seamless stitching of L2TPv3-based and MPLS-based segments of SAToP PWs (see [PWE3-MS]).




Figure 2a. SAToP Packet Format for an IPv4/IPv6 PSN with
UDP PW Demultiplexing

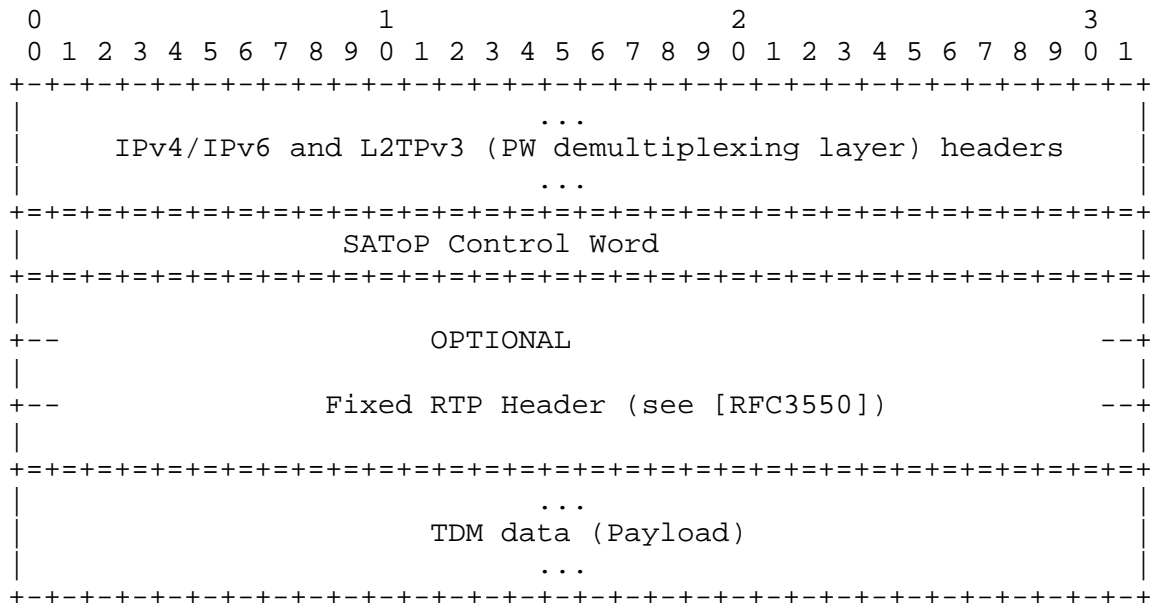


Figure 2b. SAToP Packet Format for an IPv4/IPv6 PSN with L2TPv3 PW Demultiplexing

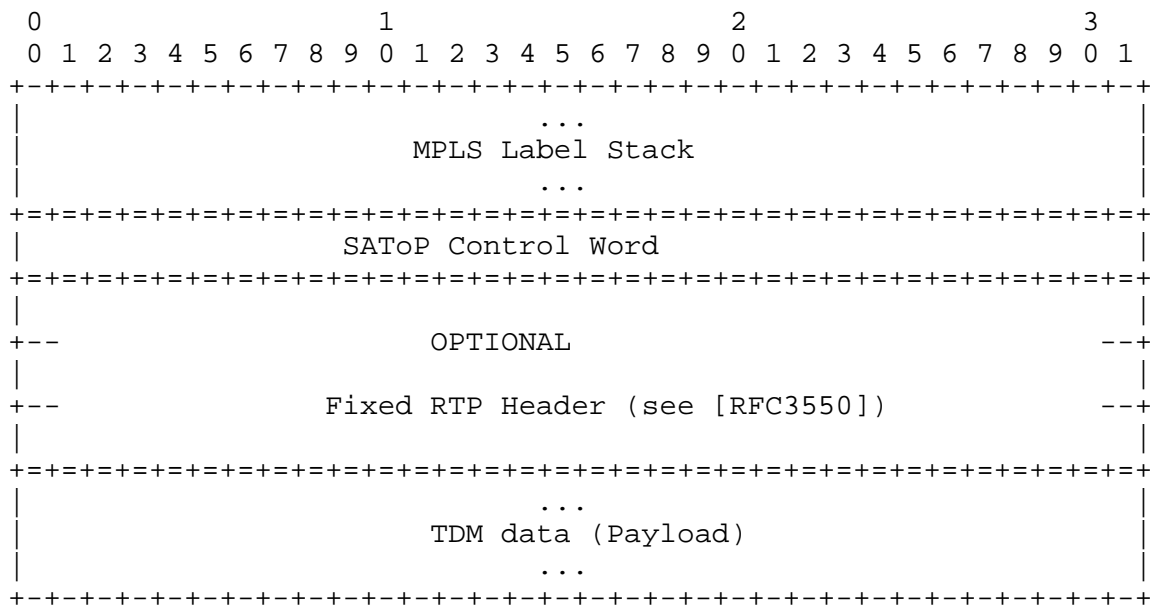


Figure 2c. SAToP Packet Format for an MPLS PSN

4.3.1. Usage and Structure of the Control Word

Usage of the SAToP control word allows:

1. Detection of packet loss or misordering
2. Differentiation between the PSN and attachment circuit problems as causes for the outage of the emulated service
3. PSN bandwidth conservation by not transferring invalid data (AIS)
4. Signaling of faults detected at the PW egress to the PW ingress.

The structure of the SAToP Control Word is shown in Figure 3 below.




Figure 3. Structure of the SAToP Control Word

The use of Bits 0 to 3 is described in [RFC4385]. These bits MUST be set to zero unless they are being used to indicate the start of an Associated Channel Header (ACH). An ACH is needed if the state of the SAToP PW is being monitored using Virtual Circuit Connectivity Verification [PWE3-VCCV].

L - If set, indicates that TDM data carried in the payload is invalid due to an attachment circuit fault. When the L bit is set the payload MAY be omitted in order to conserve bandwidth. The CE-bound IWF MUST play out an appropriate amount of filler data regardless of the payload size. Once set, if the fault is rectified, the L bit MUST be cleared.

Note: This document does not specify which TDM fault conditions are treated as invalidating the data carried in the SAToP packets. Possible examples include, but are not limited to LOS and AIS.

R - If set by the PSN-bound IWF, indicates that its local CE-bound IWF is in the packet loss state, i.e., has lost a preconfigured number of consecutive packets. The R bit MUST be cleared by the PSN-bound IWF once its local CE-bound IWF has exited the packet loss state, i.e., has received a preconfigured number of consecutive packets.

RSV and FRG (bits 6 to 9) - MUST be set to 0 by the PSN-bound IWF and MUST be ignored by the CE-bound IWF. RSV is reserved. FRG is fragmentation; see [PWE3-FRAG].

LEN (bits 10 to 15) - MAY be used to carry the length of the SAToP packet (defined as the size of the SAToP header + the payload size) if it is less than 64 bytes, and MUST be set to zero otherwise. When the LEN field is set to 0, the preconfigured size of the SAToP packet payload MUST be assumed to be as described in Section 5.1, and if the actual packet size is inconsistent with this length, the packet MUST be considered malformed.

Sequence number - used to provide the common PW sequencing function as well as detection of lost packets. It MUST be generated in accordance with the rules defined in Section 5.1 of [RFC3550] for the RTP sequence number:

- o Its space is a 16-bit unsigned circular space
- o Its initial value SHOULD be random (unpredictable).

It MUST be incremented with each SAToP data packet sent in the specific PW.

4.3.2. Usage of RTP Header

When RTP is used, the following fields of the fixed RTP header (see [RFC3550], Section 5.1) MUST be set to zero: P (padding), X (header extension), CC (CSRC count), and M (marker).

The PT (payload type) field is used as follows:

1. One PT value MUST be allocated from the range of dynamic values (see [RTP-TYPES]) for each direction of the PW. The same PT value MAY be reused for both directions of the PW and also reused between different PWs.
2. The PSN-bound IWF MUST set the PT field in the RTP header to the allocated value.
3. The CE-bound IWF MAY use the received value to detect malformed packets.

The sequence number MUST be the same as the sequence number in the SAToP control word.

The RTP timestamps are used for carrying timing information over the network. Their values are generated in accordance with the rules established in [RFC3550].

The frequency of the clock used for generating timestamps MUST be an integer multiple of 8 kHz. All implementations of SAToP MUST support the 8 kHz clock. Other multiples of 8 kHz MAY be used.

The SSRC (synchronization source) value in the RTP header MAY be used for detection of misconnections, i.e., incorrect interconnection of attachment circuits.

Timestamp generation MAY be used in the following modes:

1. Absolute mode: The PSN-bound IWF sets timestamps using the clock recovered from the incoming TDM attachment circuit. As a consequence, the timestamps are closely correlated with the sequence numbers. All SAToP implementations that support usage of the RTP header MUST support this mode.
2. Differential mode: Both IWFs have access to a common high-quality timing source, and this source is used for timestamp generation. Support of this mode is OPTIONAL.

Usage of the fixed RTP header in a SAToP PW and all the options associated with its usage (the timestamping clock frequency, the timestamping mode, selected PT and SSRC values) MUST be agreed upon between the two SAToP IWFs during PW setup as described in [TDM-CONTROL]. Other, RTP-specific methods (e.g., see [RFC3551]) MUST NOT be used.

5. SAToP Payload Layer

5.1. General Payloads

In order to facilitate handling of packet loss in the PSN, all packets belonging to a given SAToP PW are REQUIRED to carry a fixed number of bytes filled with TDM data received from the attachment circuit. The packet payload size MUST be defined during the PW setup, MUST be the same for both directions of the PW, and MUST remain unchanged for the lifetime of the PW.

The CE-bound and PSN-bound IWFs MUST agree on SAToP packet payload size during PW setup (default payload size values defined below guarantee that such an agreement is always possible). The SAToP packet payload size can be exchanged over the PWE3 control protocol ([TDM-CONTROL]) by using the Circuit Emulation over Packet (CEP)/TDM Payload Bytes sub-TLV of the Interface Parameters TLV ([RFC4446]).

SAToP uses the following ordering for packetization of the TDM data:

- o The order of the payload bytes corresponds to their order on the attachment circuit.
- o Consecutive bits coming from the attachment circuit fill each payload byte starting from most significant bit to least significant.

All SAToP implementations MUST be capable of supporting the following payload sizes:

- o E1 - 256 bytes
- o T1 - 192 bytes
- o E3 and T3 - 1024 bytes.

Notes:

1. Whatever the selected payload size, SAToP does not assume alignment to any underlying structure imposed by TDM framing (byte, frame, or multiframe alignment).
2. When the L bit in the SAToP control word is set, SAToP packets MAY omit invalid TDM data in order to conserve PSN bandwidth.
3. Payload sizes that are multiples of 47 bytes MAY be used in conjunction with unstructured ATM-CES [ATM-CES].

5.2. Octet-Aligned T1

An unstructured T1 attachment circuit is sometimes provided already padded to an integer number of bytes, as described in Annex B of [G.802]. This occurs when the T1 is de-mapped from a SONET/SDH virtual tributary/container, or when it is de-framed by a dual-mode E1/T1 framer.

In order to facilitate operation in such cases, SAToP defines a special "octet-aligned T1" transport mode. In this mode, the SAToP payload consists of a number of 25-byte subframes, each subframe carrying 193 bits of TDM data and 7 bits of padding. This mode is depicted in Figure 4 below.




Figure 4. SAToP Payload Format for Octet-Aligned T1 Transport

Notes:

1. No alignment with the framing structure that may be imposed on the T1 bit-stream is implied.
2. An additional advantage of the octet-aligned T1 transport mode is the ability to select the SAToP packetization latency as an arbitrary integer multiple of 125 microseconds.

Support of the octet-aligned T1 transport mode is OPTIONAL. An octet-aligned T1 SAToP PW is not interoperable with a T1 SAToP PW that carries a non-aligned bit-stream, as described in the previous section.

Implementations supporting octet-aligned T1 transport mode MUST be capable of supporting a payload size of 200 bytes (i.e., a payload of eight 25-byte subframes) corresponding to precisely 1 millisecond of TDM data.

6. SAToP Operation

6.1. Common Considerations

Edge-to-edge emulation of a TDM service using SAToP is only possible when the two PW attachment circuits are of the same type (T1, E1, T3, E3). The service type is exchanged at PW setup as described in [RFC4447].

6.2. IWF Operation

6.2.1. PSN-Bound Direction

Once the PW is set up, the PSN-bound SAToP IWF operates as follows:

TDM data is packetized using the configured number of payload bytes per packet.

Sequence numbers, flags, and timestamps (if the RTP header is used) are inserted in the SAToP headers.

SAToP, PW demultiplexing layer, and PSN headers are prepended to the packetized service data.

The resulting packets are transmitted over the PSN.

6.2.2. CE-Bound Direction

The CE-bound SAToP IWF SHOULD include a jitter buffer where the payload of the received SAToP packets is stored prior to play-out to the local TDM attachment circuit. The size of this buffer SHOULD be locally configurable to allow accommodation to the PSN-specific packet delay variation.

The CE-bound SAToP IWF SHOULD use the sequence number in the control word for detection of lost and misordered packets. If the RTP header is used, the RTP sequence numbers MAY be used for the same purposes.

Note: With SAToP, a valid sequence number can be always found in bits 16 - 31 of the first 32-bit word immediately following the PW demultiplexing header regardless of the specific PSN type, multiplexing method, usage or non-usage of the RTP header, etc. This approach simplifies implementations supporting multiple encapsulation types as well as implementation of multi-segment (MS) PWs using different encapsulation types in different segments.

The CE-bound SAToP IWF MAY reorder misordered packets. Misordered packets that cannot be reordered MUST be discarded and treated as lost.

The payload of the received SAToP packets marked with the L bit set SHOULD be replaced by the equivalent amount of the "all ones" pattern even if it has not been omitted.

The payload of each lost SAToP packet MUST be replaced with the equivalent amount of the replacement data. The contents of the replacement data are implementation-specific and MAY be locally configurable. By default, all SAToP implementations MUST support generation of the "all ones" pattern as the replacement data. Before a PW has been set up and after a PW has been torn down, the IWF MUST play out the "all ones" pattern to its TDM attachment circuit.

Once the PW has been set up, the CE-bound IWF begins to receive SAToP packets and to store their payload in the jitter buffer but continues to play out the "all ones" pattern to its TDM attachment circuit. This intermediate state persists until a preconfigured amount of TDM

data (usually half of the jitter buffer) has been received in consecutive SAToP packets or until a preconfigured intermediate state timer (started when the PW setup is completed) expires.

Once the preconfigured amount of the TDM data has been received, the CE-bound SAToP IWF enters its normal operation state where it continues to receive SAToP packets and to store their payload in the jitter buffer while playing out the contents of the jitter buffer in accordance with the required clock. In this state, the CE-bound IWF performs clock recovery, MAY monitor PW defects, and MAY collect PW performance monitoring data.

If the CE-bound SAToP IWF detects loss of a preconfigured number of consecutive packets or if the intermediate state timer expires before the required amount of TDM data has been received, it enters its packet loss state. While in this state, the local PSN-bound SAToP IWF SHOULD mark every packet it transmits with the R bit set. The CE-bound SAToP IWF leaves this state and transitions to the normal one once a preconfigured number of consecutive valid SAToP packets have been received. (Successfully reordered packets contribute to the count of consecutive packets.)

The CE-bound SAToP IWF MUST provide an indication of TDM data validity to the CE. This can be done by transporting or by generating the native AIS indication. As mentioned above, T3 AIS cannot be detected or generated by structure-agnostic means, and hence a structure-aware NSP MUST be used when generating a valid AIS pattern.

6.3. SAToP Defects

In addition to the packet loss state of the CE-bound SAToP IWF defined above, it MAY detect the following defects:

- o Stray packets
- o Malformed packets
- o Excessive packet loss rate
- o Buffer overrun
- o Remote packet loss

Corresponding to each defect is a defect state of the IWF, a detection criterion that triggers transition from the normal operation state to the appropriate defect state, and an alarm that MAY be reported to the management system and thereafter cleared. Alarms are only reported when the defect state persists for a preconfigured amount of time (typically 2.5 seconds) and MUST be

cleared after the corresponding defect is undetected for a second preconfigured amount of time (typically 10 seconds). The trigger and release times for the various alarms may be independent.

Stray packets MAY be detected by the PSN and PW demultiplexing layers. When RTP is used, the SSRC field in the RTP header MAY be used for this purpose as well. Stray packets MUST be discarded by the CE-bound IWF, and their detection MUST NOT affect mechanisms for detection of packet loss.

Malformed packets are detected by mismatch between the expected packet size (taking the value of the L bit into account) and the actual packet size inferred from the PSN and PW demultiplexing layers. When RTP is used, lack of correspondence between the PT value and that allocated for this direction of the PW MAY also be used for this purpose. Malformed in-order packets MUST be discarded by the CE-bound IWF and replacement data generated as with lost packets.

Excessive packet loss rate is detected by computing the average packet loss rate over a configurable amount of times and comparing it with a preconfigured threshold.

Buffer overrun is detected in the normal operation state when the jitter buffer of the CE-bound IWF cannot accommodate newly arrived SAToP packets.

Remote packet loss is indicated by reception of packets with their R bit set.

6.4. SAToP PW Performance Monitoring

Performance monitoring (PM) parameters are routinely collected for TDM services and provide an important maintenance mechanism in TDM networks. The ability to collect compatible PM parameters for SAToP PWs enhances their maintenance capabilities.

Collection of the SAToP PW performance monitoring parameters is OPTIONAL and, if implemented, is only performed after the CE-bound IWF has exited its intermediate state.

SAToP defines error events, errored blocks, and defects as follows:

- o A SAToP error event is defined as insertion of a single replacement packet into the jitter buffer (replacement of payload of SAToP packets with the L bit set is not considered insertion of a replacement packet).

- o A SAToP errored data block is defined as a block of data played out to the TDM attachment circuit and of a size defined in accordance with the [G.826] rules for the corresponding TDM service that has experienced at least one SAToP error event.
- o A SAToP defect is defined as the packet loss state of the CE-bound SAToP IWF.

The SAToP PW PM parameters (Errored, Severely Errored, and Unavailable Seconds) are derived from these definitions in accordance with [G.826].

7. Quality of Service (QoS) Issues

SAToP SHOULD employ existing QoS capabilities of the underlying PSN.

If the PSN providing connectivity between PE devices is Diffserv-enabled and provides a PDB [RFC3086] that guarantees low jitter and low loss, the SAToP PW SHOULD use this PDB in compliance with the admission and allocation rules the PSN has put in place for that PDB (e.g., marking packets as directed by the PSN).

If the PSN is Intserv-enabled, then GS (Guaranteed Service) [RFC2212] with the appropriate bandwidth reservation SHOULD be used in order to provide a bandwidth guarantee equal or greater than that of the aggregate TDM traffic.

8. Congestion Control

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion. SAToP PWs represent inelastic constant bit-rate (CBR) flows and cannot respond to congestion in a TCP-friendly manner prescribed by [RFC2914], although the percentage of total bandwidth they consume remains constant.

Unless appropriate precautions are taken, undiminished demand of bandwidth by SAToP PWs can contribute to network congestion that may impact network control protocols.

Whenever possible, SAToP PWs SHOULD be carried across traffic-engineered PSNs that provide either bandwidth reservation and admission control or forwarding prioritization and boundary traffic conditioning mechanisms. IntServ-enabled domains supporting Guaranteed Service (GS) [RFC2212] and DiffServ-enabled domains [RFC2475] supporting Expedited Forwarding (EF) [RFC3246] provide examples of such PSNs. Such mechanisms will negate, to some degree, the effect of the SAToP PWs on the neighboring streams. In order to facilitate boundary traffic conditioning of SAToP traffic over IP

PSNs, the SAToP IP packets SHOULD NOT use the DiffServ Code Point (DSCP) value reserved for the Default Per-Hop Behavior (PHB) [RFC2474].

If SAToP PWs run over a PSN providing best-effort service, they SHOULD monitor packet loss in order to detect "severe congestion". If such a condition is detected, a SAToP PW SHOULD shut down bi-directionally for some period of time as described in Section 6.5 of [RFC3985].

Note that:

1. The SAToP IWF can inherently provide packet loss measurement since the expected rate of arrival of SAToP packets is fixed and known
2. The results of the SAToP packet loss measurement may not be a reliable indication of presence or absence of severe congestion if the PSN provides enhanced delivery. For example:
 - a) If SAToP traffic takes precedence over non-SAToP traffic, severe congestion can develop without significant SAToP packet loss.
 - b) If non-SAToP traffic takes precedence over SAToP traffic, SAToP may experience substantial packet loss due to a short-term burst of high-priority traffic.
3. The TDM services emulated by the SAToP PWs have high availability objectives (see [G.826]) that MUST be taken into account when deciding on temporary shutdown of SAToP PWs.

This specification does not define the exact criteria for detecting "severe congestion" using the SAToP packet loss rate or the specific methods for bi-directional shutdown the SAToP PWs (when such severe congestion has been detected) and their subsequent re-start after a suitable delay. This is left for further study. However, the following considerations may be used as guidelines for implementing the SAToP severe congestion shutdown mechanism:

1. SAToP Performance Monitoring techniques (see Section 6.4) provide entry and exit criteria for the SAToP PW "Unavailable" state that make it closely correlated with the "Unavailable" state of the emulated TDM circuit as specified in [G.826]. Using the same criteria for "severe congestion" detection may decrease the risk of shutting down the SAToP PW while the emulated TDM circuit is still considered available by the CE.

2. If the SAToP PW has been set up using either PWE3 control protocol [RFC4447] or L2TPv3 [RFC3931], the regular PW teardown procedures of these protocols SHOULD be used.
3. If one of the SAToP PW end points stops transmission of packets for a sufficiently long period, its peer (observing 100% packet loss) will necessarily detect "severe congestion" and also stop transmission, thus achieving bi-directional PW shutdown.

9. Security Considerations

SAToP does not enhance or detract from the security performance of the underlying PSN; rather, it relies upon the PSN mechanisms for encryption, integrity, and authentication whenever required.

SAToP PWs share susceptibility to a number of pseudowire-layer attacks and will use whatever mechanisms for confidentiality, integrity, and authentication are developed for general PWs. These methods are beyond the scope of this document.

Although SAToP PWs MAY employ an RTP header when explicit transfer of timing information is required, SRTP (see [RFC3711]) mechanisms are NOT RECOMMENDED as a substitute for PW layer security.

Misconnection detection capabilities of SAToP increase its resilience to misconfiguration and some types of denial-of-service (DoS) attacks.

Random initialization of sequence numbers, in both the control word and the optional RTP header, makes known-plaintext attacks on encrypted SAToP PWs more difficult. Encryption of PWs is beyond the scope of this document.

10. Applicability Statement

SAToP is an encapsulation layer intended for carrying TDM circuits (E1/T1/E3/T3) over PSN in a structure-agnostic fashion.

SAToP fully complies with the principle of minimal intervention, thus minimizing overhead and computational power required for encapsulation.

SAToP provides sequencing and synchronization functions needed for emulation of TDM bit-streams, including detection of lost or misordered packets and appropriate compensation.

TDM bit-streams carried over SAToP PWs may experience delays exceeding those typical of native TDM networks. These delays include the SAToP packetization delay, edge-to-edge delay of the underlying PSN, and the delay added by the jitter buffer. It is recommended to estimate both delay and delay variation prior to setup of a SAToP PW.

SAToP carries TDM streams over PSN in their entirety, including any TDM signaling contained within the data. Consequently, the emulated TDM services are sensitive to the PSN packet loss. Appropriate generation of replacement data can be used to prevent shutting down the CE TDM interface due to occasional packet loss. Other effects of packet loss on this interface (e.g., errored blocks) cannot be prevented.

Note: Structure-aware TDM emulation (see [CESoPSN] or [TDMoIP]) completely hides effects of the PSN packet loss on the CE TDM interface (because framing and Cyclic Redundancy Checks (CRCs) are generated locally) and allows usage of application-specific packet loss concealment methods to minimize effects on the applications using the emulated TDM service.

SAToP can be used in conjunction with various network synchronization scenarios (see [RFC4197]) and clock recovery techniques. The quality of the TDM clock recovered by the SAToP IWF may be implementation-specific. The quality may be improved by using RTP if a common clock is available at both ends of the SAToP PW.

SAToP provides for effective fault isolation by carrying the local attachment circuit failure indications.

The option not to carry invalid TDM data enables PSN bandwidth conservation.

SAToP allows collection of TDM-like faults and performance monitoring parameters and hence emulates 'classic' carrier services of TDM.

SAToP provides for a carrier-independent ability to detect misconnections and malformed packets. This feature increases resilience of the emulated service to misconfiguration and DoS attacks.

Being a constant bit rate (CBR) service, SAToP cannot provide TCP-friendly behavior under network congestion.

Faithfulness of a SAToP PW may be increased by exploiting QoS features of the underlying PSN.

SAToP does not provide any mechanisms for protection against PSN outages, and hence its resilience to such outages is limited. However, lost-packet replacement and packet reordering mechanisms increase resilience of the emulated service to fast PSN rerouting events.

11. IANA Considerations

Allocation of PW Types for the corresponding SAToP PWs is defined in [RFC4446].

12. Acknowledgements

We acknowledge the work of Gil Biran and Hugo Silberman who implemented TDM transport over IP in 1998.

We would like to thank Alik Shimelmits for many productive discussions and Ron Insler for his assistance in deploying TDM over PSN.

We express deep gratitude to Stephen Casner who has reviewed in detail one of the predecessors of this document and provided valuable feedback regarding various aspects of RTP usage, and to Kathleen Nichols who has provided the current text of the QoS section considering Diffserv-enabled PSN.

We thank William Bartholomay, Robert Biksner, Stewart Bryant, Rao Cherukuri, Ron Cohen, Alex Conta, Shahram Davari, Tom Johnson, Sim Narasimha, Yaron Raz, and Maximilian Riegel for their valuable feedback.

13. Co-Authors

The following are co-authors of this document:

Motty Anavi	RAD Data Communications
Tim Frost	Zarlink Semiconductors
Eduard Metz	TNO Telecom
Prayson Pate	Overture Networks
Akiva Sadoski	
Israel Sasson	Axerra Networks
Ronen Shashoua	RAD Data Communications

14. Normative References

- [G.702] ITU-T Recommendation G.702 (11/88) - Digital Hierarchy Bit Rates.
- [G.703] ITU-T Recommendation G.703 (10/98) - Physical/Electrical Characteristics of Hierarchical Digital Interfaces.
- [G.704] ITU-T Recommendation G.704 (10/98) - Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 Kbit/s hierarchical levels.
- [G.707] ITU-T Recommendation G.707 (03/96) - Network Node Interface for the Synchronous Digital Hierarchy (SDH).
- [G.775] ITU-T Recommendation G.775 (10/98) - Loss of Signal (LOS), Alarm Indication Signal (AIS) and Remote Defect Indication (RDI) Defect Detection and Clearance Criteria for PDH Signals.
- [G.802] ITU-T Recommendation G.802 (11/88) - Interworking between Networks Based on Different Digital Hierarchies and Speech Encoding Laws.
- [G.826] ITU-T Recommendation G.826 (02/99) - Error performance parameters and objectives for international, constant bit rate digital paths at or above the primary rate.
- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.

- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC 3086, April 2001.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RTP-TYPES] RTP PARAMETERS, <<http://www.iana.org/assignments/rtp-parameters>>.
- [T1.107] American National Standard for Telecommunications - Digital Hierarchy - Format Specifications, ANSI T1.107-1988.

15. Informative References

- [ATM-CES] ATM forum specification af-vtoa-0078 (CES 2.0) Circuit Emulation Service Interoperability Specification Ver. 2.0.
- [CESoPSN] Vainshtein, A., Ed., Sasson, I., Metz, E., Frost, T., and P. Pate, "TDM Circuit Emulation Service over Packet Switched Network (CESoPSN)", Work in Progress, November 2005.
- [PWE3-MS] Martini, L., Metz, C., Nadeau, T., Duckett, M., and F. Balus, "Segmented Pseudo Wire", Work in Progress, March 2006.

- [PWE3-FRAG] Malis, A. and M. Townsley, "PWE3 Fragmentation and Reassembly", Work in Progress, November 2005.
- [PWE3-VCCV] Nadeau, T. and R. Aggarwal, "Pseudo Wire Virtual Circuit Connectivity", Work in Progress, August 2005.
- [RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, September 1997.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4197] Riegel, M., "Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks", RFC 4197, October 2005.
- [TDM-CONTROL] Vainshtein, A. and Y. Stein, "Control Protocol Extensions for Setup of TDM Pseudowires", Work in Progress, July 2005.
- [TDMoIP] Stein, Y., "TDMoIP", Work in Progress, February 2005.

Appendix A: Old Mode of SAToP Encapsulation over L2TPv3

Previous versions of this specification defined a SAToP PW encapsulation over L2TPv3, which differs from that described in Section 4.3 and Figure 2b. In these versions, the RTP header, if used, precedes the SAToP control word.

Existing implementations of the old encapsulation mode MUST be distinguished from the encapsulations conforming to this specification via the SAToP PW setup.

Appendix B: Parameters That MUST Be Agreed upon during the PW Setup

The following parameters of the SAToP IWF MUST be agreed upon between the peer IWFs during the PW setup. Such an agreement can be reached via manual configuration or via one of the PW setup protocols:

1. Type of the Attachment Circuit (AC)

As mentioned in Section 3, SAToP supports the following AC types:

- i) E1 (2048 kbit/s)
- ii) T1 (1544 kbit/s); this service is also known as DS1
- iii) E3 (34368 kbit/s)
- iv) T3 (44736 kbit/s); this service is also known as DS3

SAToP PWs cannot be established between ACs of different types.

2. Usage of octet-aligned mode for T1

- a) This OPTIONAL mode of emulating T1 bit-streams with SAToP PWs is described in Section 5.2.
- b) Both sides MUST agree on using this mode for a SAToP PW to be operational.

3. Payload size, i.e., the amount of valid TDM data in a SAToP packet

- a) As mentioned in Section 5.1:
 - i) The same payload size MUST be used in both directions of the SAToP PW.
 - ii) The payload size cannot be changed once the PW has been set up.
- b) In most cases, any mutually agreed upon value can be used. However, if octet-aligned T1 encapsulation mode is used, the payload size MUST be an integral multiple of 25, and it expresses the amount of valid TDM data including padding.

4. Usage of the RTP header in the encapsulation

- a) Both sides MUST agree on using RTP header in the SAToP PW.
- b) In the case of a SAToP PW over L2TPv3 using the RTP header, both sides MUST agree on usage of the "old mode" described in Appendix A.

5. RTP-dependent parameters. The following parameters MUST be agreed upon if usage of the RTP header for the SAToP PW has been agreed upon.

- a) Timestamping mode (absolute or differential); this mode MAY be different for the two directions of the PW, but the receiver and transmitter MUST agree on the timestamping mode for each direction of the PW
- b) Timestamping clock frequency:
 - i) The timestamping frequency MUST be an integral multiple of 8 kHz.
 - ii) The timestamping frequency MAY be different for the two directions of the PW, but the receiver and transmitter MUST agree on the timestamping mode for each direction of the PW.
- c) RTP Payload Type (PT) value; any dynamically assigned value can be used with SAToP PWs.
- d) Synchronization Source (SSRC) value; the transmitter MUST agree to send the SSRC value requested by the receiver.

Editors' Addresses

Alexander ("Sasha") Vainshtein
Axerra Networks
24 Raoul Wallenberg St.,
Tel Aviv 69719, Israel

EMail: sasha@axerra.com

Yaakov (Jonathan) Stein
RAD Data Communications
24 Raoul Wallenberg St., Bldg C
Tel Aviv 69719, Israel

EMail: yaakov_s@rad.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

