

PWE3 Working Group
Internet Draft
[draft-anavi-tdmoip-03.txt](#)
Expires: August 2002

Yaakov (Jonathan) Stein
Motty (Mordechai) Anavi
Ronen Shaashoua
RAD Data Communications

February 2002

TDM over IP

[draft-anavi-tdmoip-03.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

TDM over IP

February, 2002

Abstract

This document describes methods for transporting time division multiplexed (TDM) digital voice and data signals over Pseudo Wires. It is a revision of the document [draft-anavi-tdmoip-02](#).

Conventions used in this document

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 [RFC 2119](#).

Table of Contents

1.	Introduction	2
2.	TDMoIP Encapsulation	3
3.	Encapsulation Details for Specific PSNs	6
4.	TDMoIP Payload types	10
5.	AAL1 Format Payload (FORMID=001XX)	10
6.	AAL2 Format Payload (FORMID=01000)	14
7.	HDLC Format Payload (FORMID=10000)	15
8.	OAM Signaling	16
9.	General Requirements	19
10.	Security Considerations	21
11.	IANA Considerations	21
12.	Applicability Statement	21
13.	References	22
14.	Acknowledgments	23
15.	Contact Information	23

[1.](#) Introduction

Telephony traffic is conventionally carried over connection-oriented synchronous or plesiosynchronous networks (which will be loosely called TDM networks herein). With the proliferation of

packet-switched networks (PSNs), telephony carriers desire integration of TDM services into a unified PSN infrastructure. This integration requires emulation of TDM circuits within the PSN, a function that can be carried out using Pseudo Wires (PWs), as described in the PWE3 requirements [PWE-REQ] and framework [PWE-FR] documents. This emulation must ensure QoS and voice quality similar to those of existing circuit-based networks as well as preserving signaling features.


In this document we describe a protocol for tunneling TDM circuits through PSNs using PWs. This protocol can support various TDM traffic types, including n*64K, unstructured T1/E1, structured T1/E1 with and without CAS signaling, T3/E3, and TDM in AAL1 and AAL2 networks. The precise requirements of the emulation for each of these types are described in the Applicability Statement, below.

The protocol as herein described is agnostic to the underlying PSN, which may be IPv4, IPv6, MPLS, L2TPv3, Ethernet, etc. Implementation specifics for particular PSNs are discussed in [section 3](#). Although the protocol should be more generally called TDMoPW and its specific implementations TDMoIP, TDMoMPLS, etc. we will use the nomenclature TDMoIP for reasons of consistency with previous versions of this draft.

[2](#). TDMoIP Encapsulation

2.1 Layering Model

The protocol-layering model used by TDMoIP is shown in the figure, where the order is that of the physical packet, so that higher layers appear lower in the diagram.



+-----+

2.2 PSN-specific layers

The PSN-specific layers contain all necessary infrastructure, and may consist of UDP/IP, MPLS, L2TPv3 over IP, layer 2 Ethernet, etc. The PSN is assumed to be reliable enough and of sufficient bandwidth to enable transport of the required TDM data.

TDMoIP edge devices typically handle more than one circuit bundle at a time. A circuit bundle is defined as a stream of bits that have originated from the same physical interface or from interfaces that share a common clock, which are transmitted from a single TDMoIP source device to a single TDMoIP destination device. For example, bundles may comprise some number of 64 Kbps timeslots originating from a single E1, or an entire T3 or E3. Circuit bundles are single direction streams, but are frequently coupled with bundles in the opposite direction to enable full duplex communications. More than one bundle can be transmitted between two edge devices, as is the case when the PW limits the bundle's packet length.

Stein et al.

[PAGE 3]


TDM over IP

February, 2002

If a TDMoIP edge device is required to handle multiple circuit bundles, then it is the responsibility of the PSN-specific layers to provide a circuit bundle identifier (CBID) in order to enable differentiation between these circuits. Furthermore, These layers will be more fully discussed in [section 3](#).

2.3 RTP

If timing needs to be delivered over the PSN RTP MUST be used. When the TDMoIP edges have sufficiently accurate local clocks its use is optional. If RTP is used, the header of following figure, as defined in [[RTP](#)] SHALL appear.



RTV (2 bits) the RTP Version number SHALL be set to RTV=010

P (1 bit) the RTP padding indicator SHALL be set to P=0

X (1 bit) the RTP extension SHALL be set to X=0

CC (4 bits) the RTP CSRC count SHALL be set to CC=0000

M (1 bit) the RTP marker indicator SHALL be set to M=0

PT (7 bits) the RTP Payload Type

RTP Sequence Number (16 bits) is defined separately for each circuit bundle and increments by one for each TDMoIP packet sent for that circuit bundle. It MAY be used by the receiver to detect packet loss and to restore packet sequence. The initial value of the sequence number SHOULD be random (unpredictable) for security purposes.

RTP Timestamp (32 bits) The RTP timestamp indicates the precise sampling instant of the first octet in the TDM payload. It is derived from a clock that increments monotonically and linearly in time to allow synchronization and jitter calculations, and that has sufficient resolution.

SSRC Identifier (32 bits) the RTP synchronization source identifier uniquely identifies the circuit bundle's timing source. It is chosen randomly for independent timing source as described in [RTP].

Stein et al.

[PAGE 4]


TDM over IP

February, 2002

The main difficulty with the use of RTP is its large overhead. For this reason TDMoIP allows the RTP header to be omitted when timing information need not be transferred over the network.

2.4 TDMoIP Control Word

The 32-bit control word **MUST** appear in every TDMoIP packet. Its format is given in the following figure.



FORMID Format identifier (5 bits)

The following values are presently defined:

00100 AAL1 unstructured

00101 AAL1 structured

00110 AAL1 structured w/ CAS

01000 AAL2

10000 HDLC

The payload format for each of these cases will be described in sections [4](#), [5](#), and [6](#) below.

L Loss of Sync failure (1 bit) The L bit being set indicates that the source has detected or has been informed of a TDM physical layer fault impacting the data to be transmitted. This bit can be used to indicate Physical layer LOS that should trigger AIS generation at the far end.

R Receive failure (1 bit) The R bit being set indicates that the source is not receiving packets at its TDMoIP receive port.

Z (1 bit) The Z bit indicates an extended header and MUST be set to zero at present.

Length (8 bits) is used to indicate the use of padding to meet minimum transmission unit requirements of the PSN. It MUST be used if the total packet length is less than 64 bytes, and MUST be set to zero if this length exceeds 255 bytes.

Sequence number (16 bits) The TDMoIP sequence number MUST be present when the RTP header is not used and fulfills the same function as the RTP sequence number. In addition, since the basic clock rate for each circuit bundle is constant, the sequence number may be used as an approximate timestamp. The initial value of the sequence number SHOULD be random (unpredictable) for security purposes, and the value is incremented modulo 2^{16} separately for each circuit bundle.

[3](#). Encapsulation Details for Specific PSNs

3.1 UDP/IP

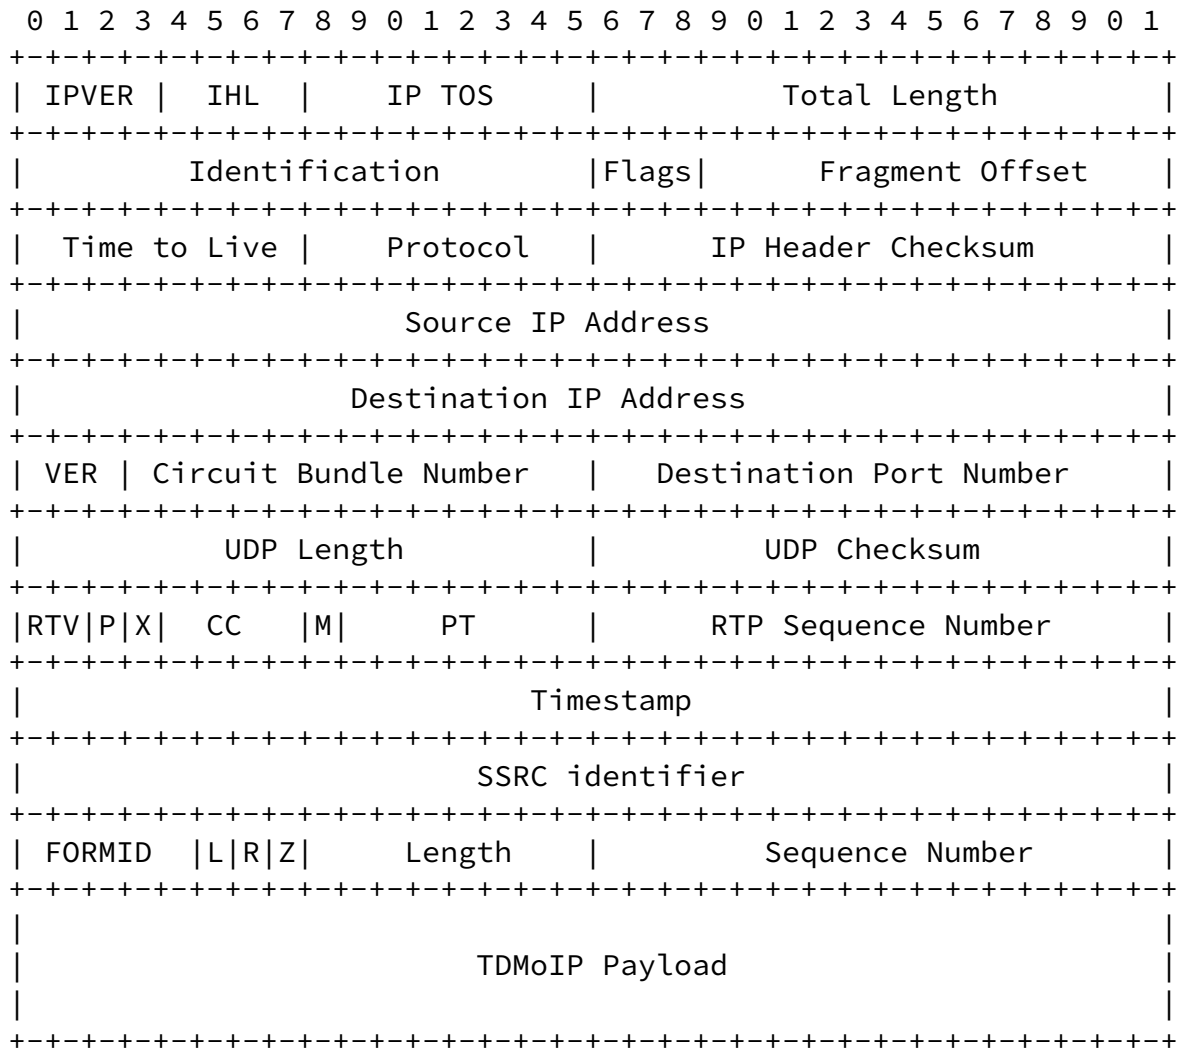
The UDP/IP header as described in [[UDP](#)] and [IP] is prefixed to the TDMoIP data. The frame structure is as follows:

0

1

2

3



The first five rows are the IP header, the sixth and seventh rows are the UDP header. Rows 8 through 10 are the optional RTP header. Row 11 is the TDMoIP control word.

IPVER (4 bits) is the IP version number, for IPv4 IPVER=4.

IHL (4 bits) is the length in 32-bit words of the IP header, IHL=5.

IP TOS (8 bits) is the IP type of service.

Total Length (16 bits) is the length in octets of header and data.

Identification (16 bits) is the IP fragmentation identification field.

Flags (3 bits) are the IP control flags and MUST be set to Flags=010 to avoid fragmentation.

Fragment Offset (13 bits) indicates where in the datagram the fragment belongs and is not used here.

Time to Live (8 bits) is the IP time to live field. Datagrams with zero in this field are to be discarded.

Protocol (8 bits) MUST be set to 0x11 for UDP.

IP Header Checksum (16 bits) is a checksum on the IP header.
Source IP Address (32 bits) is the IP address of the source.

Destination IP Address (32 bits) is the IP address of the destination.

VER (3 bits) is the TDMoIP version number. The original version (VER=000) was experimental and should no longer be used. Presently VER=001 when RTP is not used, and VER=011 when RTP is used.

Circuit Bundle Number (13 bits) This field is usually dedicated to the Source Port Number, but here identifies the unique data stream emanating from a given trunk and sharing a common destination. This nonstandard use of a UDP port number is similar to RTP/RTCP's use of port numbers to uniquely identify sessions, and the common practice (sanctioned in H.225) of randomly allocating port numbers for VoIP sessions. Here placing the circuit bundle identifier in the UDP header rather than the application area enables fast switching. The available circuit bundle numbers are 1-8063; 0 is invalid; 8191 (1FFF) is used for OAM control messages (see [section 8](#)); and the 127 ports 8064-8190 are reserved for advanced applications.

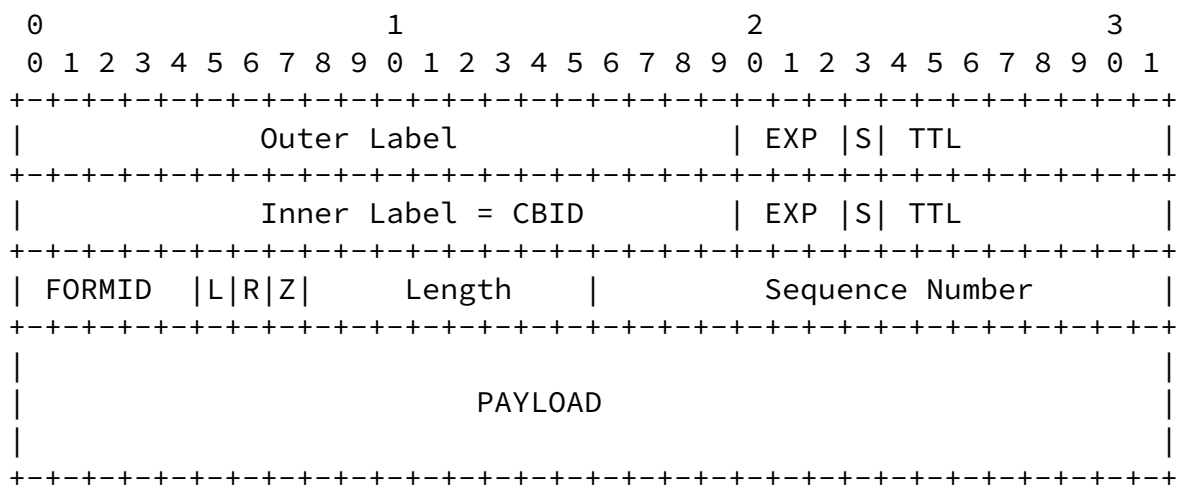
Destination Port Number (16 bits) MUST be set to 0x085E (2142), the user port number which has been assigned by the Internet Assigned Numbers Authority (IANA) to TDMoIP.

UDP Length (16 bits) is the length in octets of UDP header and data.

UDP Checksum (16 bits) is the checksum of UDP/IP header and data. If not computed it must be set to zero.

3.2 MPLS

The MPLS header as described in [\[MPLS\]](#) is prefixed to the TDMoIP data. The frame structure (as seen at the edges) is as follows:



The first two rows are the MPLS header; the third is the control word.

Outer Label (20 bits) is the MPLS label which identifies the MPLS LSP used to tunnel the TDM packets through the MPLS network. The label number can be assigned either by manual provisioning or via the MPLS control protocol.

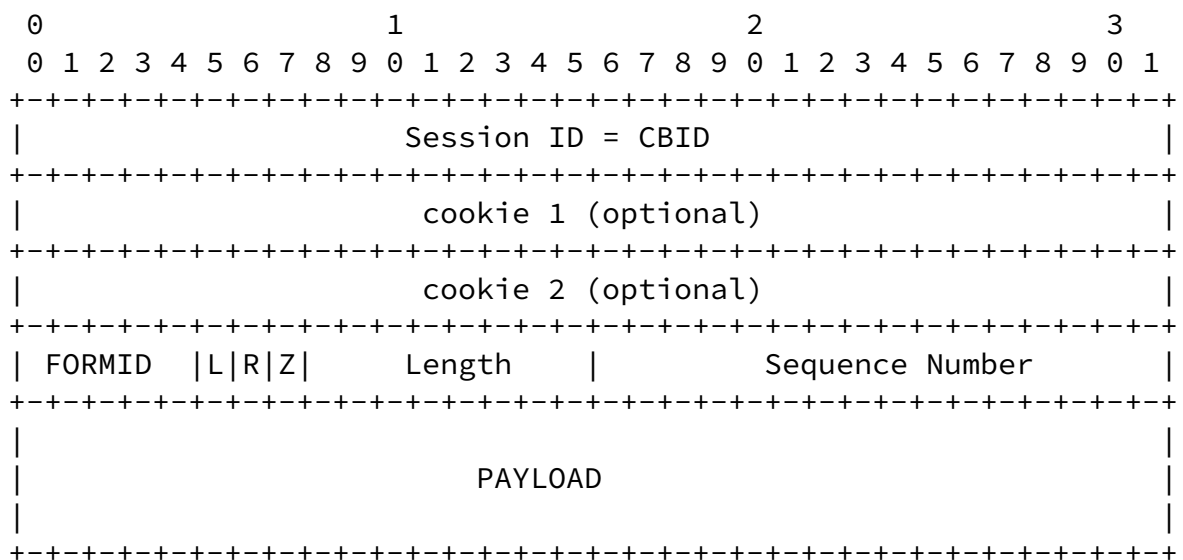
EXP (3 bits) experimental field
 S (1 bit) stacking bit
 TTL (8 bits) MPLS Time to live

Inner Label (20 bits) the MPLS inner label, contains the circuit bundle identifier used to multiplex multiple circuit bundles within the same tunnel. Valid values are as in [subsection 3.1](#) supra. For label stack usage see [\[MPLS\]](#).

EXP (3 bits) experimental field
 S (1 bit) stacking bit must equal 1 to indicate bottom of stack
 TTL (8 bits) MPLS Time to live

3.3 L2TPv3

If L2TP is used over UDP the frame described in [subsection 3.1](#) is to be used. For use of L2TP over IPv4 without UDP the header defined in subsection 4.1.1.1 of [[L2TPv3](#)] is prefixed to the TDMoIP data.



Session ID (32 bits) is the locally significant L2TP session identifier, and contains the circuit bundle identifier used to multiplex multiple circuit bundles within the same tunnel. Valid values are as in [subsection 3.1](#) supra.

Cookie (32 or 64 bits) is an optional field that contains a randomly selected value that can be used to validate association of the received frame with the expected circuit bundle.

3.4 Ethernet

In many applications the TDMoIP frame as described in the previous

subsections will be further encapsulated in an Ethernet frame by prefixing the Ethernet preamble, destination and source MAC addresses, optional VLAN header, etc and appending the four octet frame check sequence after the TDMoIP frame.

A TDMoIP implementation MUST be able to receive both industry standard (DIX) Ethernet and IEEE 802.3 CSMA/CD frames and SHOULD transmit Ethernet frames.

Ethernet encapsulation introduces restrictions on both minimum and maximum packet size. Whenever the entire TDMoIP packet is less than 64 bytes, zero padding is introduced and the true length indicated by using the Length field in the control word. In order to avoid fragmentation the TDMoIP frame must be restricted to the maximum payload size. For example, the length of the Ethernet payload for a non-RTP AAL2 adapted E1 trunk with 31 channels is $8 \times 4 + 31 \times 47 = 1489$ octets. This falls below the maximal permitted payload size of 1500 bytes.

The direct use of layer 2 Ethernet frames without higher layers (such as IP or MPLS) is for further study.

Stein et al.

[PAGE 9]

TDM over IP

February, 2002

4. TDMoIP Payload types

TDMoIP is a trunking application, i.e. it transports entire trunks potentially containing multiple voice or data streams. Trunking can be carried out at two levels - circuit emulation and loop emulation. In circuit emulation the entire TDM trunk is transferred across the network as a whole without attempting to separate it into individual timeslots, while in loop emulation the individual timeslots are identified and transported, albeit while preserving the trunk integrity.

Presently we define three different payload types, namely AAL1, AAL2, and HDLC. AAL1 is used for circuit emulation, while AAL2 is used for loop emulation. AAL1 is thus best for unstructured trunks, or for structured trunks with relatively constant usage. AAL2 is used to conserve bandwidth for structured trunks in which usage is highly variable. The HDLC mode is to enable efficient transport of trunk associated CCS signaling.

The AAL family of protocols is a natural choice for trunking applications. Although originally developed to adapt various types

of application data to the rigid format of ATM, the mechanisms are general solutions to the problem of transporting constant or variable bandwidth data streams over a byte-oriented packet network.

In addition, since the AAL mechanisms are extensively used within and on the edge of the telephony system, they were specifically designed to efficiently treat audio, non-audio data and telephony signaling.

Finally, simple service interworking with legacy TDM networks is a major design goal of TDMoIP. Hence, payload types should be chosen in order to simplify interworking with the existing infrastructure, including AAL1 and AAL2 networks.

5. AAL1 Format Payload (FORMID=001XX)

For the prevalent case for which the timeslot allocation is static and no activity detection is performed, the payload can be efficiently encoded using constant bit rate AAL1 adaptation. The AAL1 format is described in [[AAL1](#)] and its use for circuit emulation over ATM in [[CES](#)].

In AAL1 mode the TDMoIP payload consists of between one and thirty 48-octet subframes. The number of subframes, which can be inferred by the receiving side from the total packet length as specified in the PSN header, is pre-configured and typically chosen according to latency and bandwidth constraints. Using a single subframe reduces latency to a minimum, but incurs the highest overhead,

Stein et al.

[PAGE 10]

TDM over IP

February, 2002

while using, for example, eight subframes reduces the overhead percentage while increasing the latency by a factor of eight.

```
+-----+-----+
|control word |48-octet subframe|
+-----+-----+
```

Single TDMoIP-AAL1 subframe per TDMoIP frame

```
+-----+-----+ +-----+
|control word |48-octet subframe|---|48-octet subframe|
+-----+-----+ +-----+
```

Multiple TDMoIP-AAL1 subframes per TDMoIP frame

The first octet of each 48-octet AAL1 subframe consists of an error protected three-bit sequence number.

```

 1 2 3 4 5 6 7 8
+---+---+---+---+---+---+
|C| SN | CRC |P| 47 octets of payload
+---+---+---+---+---+---+
```

where

C (1 bit) convergence sublayer indication, its use here is limited to indication of the existence of a pointer (see below)

C=0 means no pointer, C=1 means a pointer is present.

SN (3 bits) The AAL1 sequence number increments from subframe to subframe.

CRC (3 bits) is a 3 bit error cyclic redundancy code on C and SN

P (1 bit) even byte parity

As can be readily inferred this octet can only take on eight different values, and incrementing the sequence number forms an eight subframe sequence number cycle, the importance of which will become clear shortly.

The structure of the remaining 47 octets in the TDMoIP-AAL1 subframe depends on the subframe type, of which there are three, corresponding to the three types of AAL1 circuit emulation service defined in [\[CES\]](#). These are known as namely unstructured circuit emulation, structured circuit emulation and structured circuit emulation with CAS.

The simplest subframe is the unstructured one which is used for transparent transfer of whole trunks (T1,E1,T3,E3). The 47 octets after the sequence number octet contain 376 bits from the TDM bit stream. No frame synchronization is supplied or implied, and framing is the sole responsibility of the end-user equipment. Hence the unstructured mode can be used for leased lines which carry data rather than N*64 Kbps timeslots, and even for trunks

with nonstandard frame synchronization. For the T1 case the raw frame consists of 193 bits, and hence 1 183/193 T1 frames fit into each TDMoIP-AAL1 subframe. The E1 frame consists of 256 bits, and so 1 15/32 E1 frames fit into each subframe.

When the TDM trunk is segmented into timeslots according to [G704], and it is desired to transport $N \times 64$ Kbps circuit where N is only a fraction of the full E1 or T1, it is more efficient to use one of the structured AAL1 circuit emulation services. Structured AAL1 views the data not merely as a bit stream, but as a circuit bundle of timeslots. Furthermore, when CAS signaling is used it can be formatted such that it can be readily detected and manipulated.

In the structured circuit emulation mode without CAS, N octets from the N timeslots to be transported are first arranged in order of timeslot number. Thus if timeslots 2, 3, 5, 7 and 11 are to be transported the corresponding five octets are placed in the subframe immediately after the sequence number octet. This placement is repeated until all 47 octets in the subframe are taken;

octet	1	2	3	4	5	6	7	8	9	10	---	41	42	43	44	45	46	47
timeslot	2	3	5	7	11	2	3	5	7	11	---	2	3	5	7	11	2	3

the next subframe commences where the present subframe left off

octet	1	2	3	4	5	6	7	8	9	10	---	41	42	43	44	45	46	47
timeslot	5	7	11	2	3	5	7	11	2	3	---	5	7	11	2	3	5	7

and so forth. The set of timeslots 2,3,5,7,11 is called a structure and the point where one structure ends and the next commences is a structure boundary.

The problem with this arrangement is the lack of explicit indication of the octet identities. As can be seen in the above example, each TDMoIP-AAL1 subframe starts with a different timeslot, so a single lost frame will result in misidentifying timeslots from that point onwards, without possibility of recovery. The solution to this deficiency is the periodic introduction of a pointer to the next structure boundary. This pointer need not be used too frequently, as the timeslot identification are uniquely inferable unless frames are lost.

The particular method used in AAL1 is to insert a pointer once every sequence number cycle of length eight subframes. The pointer is seven bits and protected by an even parity MSB, and so occupies a single octet. Since seven bits are sufficient to represent offsets larger than 47, we can limit the placement of the pointer octet to subframes with even sequence number. Unlike usual TDMoIP-AAL1 subframes with 47 octets available for payload, subframes which contain a pointer, called P-format subframes, have the following format.

```

+--+-----+-----+--+-----+-----+
|C| SN  | CRC |P|E|   pointer   | 46 octets of payload
+--+-----+-----+--+-----+-----+

```

where

C (1 bit) convergence sublayer indication, C=1 for P-format subframes

SN (3 bits) is an even AAL1 sequence number,

CRC (3 bits) is a 3 bit error cyclic redundancy code on C and SN

P (1 bit) even byte parity LSB for sequence number octet

E (1 bit) even byte parity MSB for pointer octet
 pointer (7 bits) pointer to next structure boundary

Since P-format subframes have 46 octets of payload and the next subframe has 47 octets, viewed as a single entity the pointer needs to indicate one of 93 octets. If P=0 it is understood that the structure commences with the following octet (i.e. the first octet in the payload belongs to the lowest numbered timeslot). P=93 means that the last octet of the second subframe is the final octet of the structure, and the following subframe commences with a new structure. The special value P=127 indicates that there is no structure boundary to be indicated (needed when extremely large structures are being transported).

The P-format subframe is always placed at the first possible position in the sequence number cycle that a structure boundary occurs, and can only occur once per cycle.

The only difference between the structured circuit emulation format and structured circuit emulation with CAS is the definition of the structure. Whereas in structured circuit emulation the structure is composed of the N timeslots, in structured circuit emulation with CAS the structure encompasses the super-frame consisting of multiple repetitions of the N timeslots and then the CAS signaling bits. The CAS bits are tightly packed into octets and the final octet is padded with zeros if required.

For example, for E1 trunks the CAS signaling bits are updated once per superframe of 16 frames. Hence the structure for N*64 derived from an E1 with CAS signaling consists of 16 repetitions of N

octets, followed by N sets of the four ABCD bits, and finally four zero bits if N is odd. For example, the structure for timeslots 2,3 and 5 will be as follows

```

2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 2 3 5
2 3 5 2 3 5 2 3 5 2 3 5 2 3 5 [ABCD2 ABCD3] [ABCD5 0000]

```

Similarly for T1 ESF trunks the superframe is 24 frames, and the structure consists of 24 repetitions of N octets, followed by the ABCD bits as before. For the T1 case the signaling bits will in general appear twice, in their regular (bit-robbed) positions and at the end of the structure.

6. AAL2 Format Payload (FORMID=01000)

When timeslot allocation is dynamic or activity detection is performed, the payload can be more efficiently encoded using variable bit rate AAL2 adaptation.

The variable bit rate AAL2 format is described in [[AAL2](#)] and its use for loop emulation over ATM is explained in [[LES](#)]. The AAL2 is subdivided into the Common Part Sublayer (CPS) and the Service Specific Convergence Sublayer (SSCS); for PWs we need only the CPS.

The basic AAL2-CPS packet is as follows:

```

|   Octet 1   |   Octet 2   |   Octet 3   | Octets 4-47
| 0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7 0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   CID   |   LI   |   UUI   |   HEC   |   PAYLOAD
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---

```

CID (8 bits) channel identifier is a unique identifier for the bundle. The values below 8 are reserved and so there are 248 possible channels. The mapping of CID values to trunk timeslots is outside the scope of the TDMoIP protocol and must be configured during installation or via network management.

LI (6 bits) length indicator is one less than the length of the payload in octets. (Note that the payload is limited to 64 octets.)

UUI (5 bits) user-to-user indication is the higher layer (application) identifier and counter. For voice data the UUI will always be in the range 0-15, and SHOULD be incremented modulo 16 each time a channel buffer is sent. The receiver MAY monitor this sequence. UUI is set to 24 for signaling packets.

HEC (5 bits) the header error control

Payload - voice

A block of length indicated by LI of voice samples are placed as-is into the AAL2 packet.

Payload - CAS signaling

For CAS signaling the payload is formatted as a type 3 packet (in the notation of [AAL2]) in order to ensure error protection. The signaling is sent with the same CID as the corresponding voice channel. Signaling is sent whenever the state of the ABCD bits changes, and is sent with triple redundancy, i.e. sent three times spaced 5 milliseconds apart. In addition, the entire set of the signaling bits is sent periodically to ensure reliability.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|RED|          timestamp          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| RES | ABCD |   type   | CRC
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
      CRC (cont) |      PAD      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

RED (2 bits) is the triple redundancy counter. For the first packet it takes the value 00, for the second 01 and for the third 10. RED=11 means non-redundant information and is used for periodic refresh of the CAS information.

Timestamp (14 bits)

The timestamp is the same for all three redundant transmissions.

RES (4 bits) is reserved and MUST be set to zero

ABCD (4 bits) are the CAS signaling bits

type (6 bits) for CAS signaling this is 000011

CRC-10 (10 bits) is a 10 bit CRC error detection code

PAD (8 bits) is set to zero.

7. HDLC Format Payload (FORMID=10000)

The motivation for handling HDLC in TDMoIP is to efficiently transport CCS (common channel signaling such as SS7) which is embedded in the TDM stream. This mechanism is not intended for general HDLC payloads.

In order to transport HDLC the sender monitors flags until a frame is detected. The contents of the frame are collected and the FCS tested. If the FCS is incorrect the frame is discarded, otherwise the frame is sent (without the FCS, flags or transparency zero-insertions) as-is in the payload. When an HDLCoIP frame is received its FCS is calculated, and the original HDLC frame reconstituted.

Stein et al.

[PAGE 15]

TDM over IP

February, 2002

This format assumes that the HDLC messages are shorter than the maximum packet size and hence avoid fragmentation.

8. OAM Signaling

Since the TDMoIP PW is not absolutely reliable, and hence requires a signaling mechanism to provide feedback regarding problems in the communications environment. In addition, such signaling can be used to collect statistics relating to the performance of the underlying PSN [[IPPM](#)].

If the underlying PSN has adequate signaling mechanisms then these are to be used. If not, the ICMP-like procedures detailed below SHOULD be followed.

All TDMoIP OAM signaling messages MUST use CBID 8191 (1FFF). All PSN layer parameters (for example, IP addresses, TOS, and VLAN ID) MUST remain those of the circuit bundle being investigated.

8.1 Connectivity-Check Messages

In most conventional IP applications a server sends some finite amount of information over the network upon explicit request from a client. With TDMoIP the source sends a continuous stream of packets towards the destination without knowing whether the destination device is ready to accept them, leading to flooding of the PSN.

The problem may occur when an edge device fails or is disconnected from the PSN, or the PW is broken. After an aging time the destination edge disappears from the routing tables, and intermediate routers may flood the network with the TDMoIP packets in an attempt to find a new path.

The solution to this problem is to significantly reduce the number of TDMoIP packets transmitted per second when PW failure is detected, and to return to full rate only when the PW is restored. The detection of failure and restoration and is made possible by the periodic exchange of one-way connectivity-check messages, as defined in [[CONNECT](#)].

Connectivity is tested by periodically sending OAM messages from the source edge to the destination edge, and having the destination reply to each message. The format of connectivity-check messages is given in [subsection 8.3](#) infra.

The connectivity check mechanism can also be useful during setup and configuration. Without OAM signaling one must ensure that the destination edge is ready to receive packets before starting to send them. Since TDMoIP edge devices usually operate full-duplex,

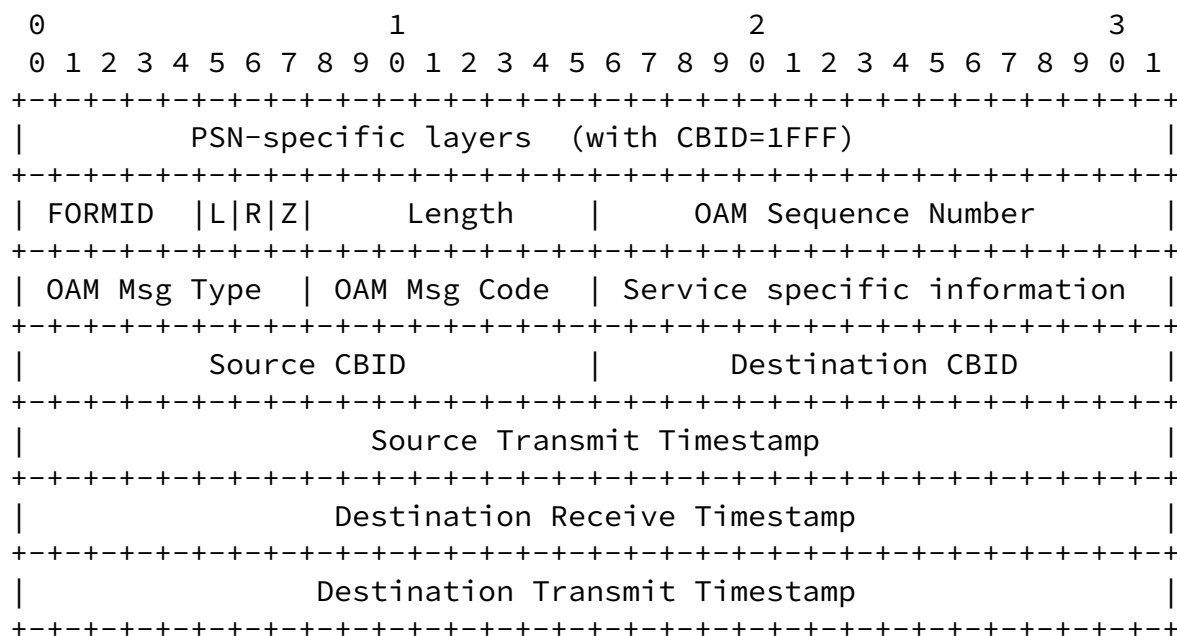
both edges must be set up and properly configured simultaneously if flooding is to be avoided. By using the connectivity mechanism a configured edge device waits until it can detect its destination before transmitting at full rate. In addition, errors in configuration can be readily discovered by using the service specific field.

8.2 Performance Measurements

In addition to one way connectivity, the OAM signaling mechanism can be used to request and report on various PSN metrics, such as one way delay, round trip delay, packet delay variation, etc. It

can also be used for remote diagnostics, and for unsolicited reporting of potential problems (e.g. dying gasp messages).

8.3 The format of an OAM message packet is depicted in the following figure. Note that PSN-specific layers are identical to those used to carry the TDMoIP data, with the exception that CBID = 1FFF is used instead of the usual circuit bundle identifier.



FORMID and L, R, and Z are identical to those used for the circuit bundle being tested.

Length is the length in bytes of the OAM message packet.

OAM Sequence Number (16 bits) is used to uniquely identify the message. Its value is unrelated to the sequence number of the TDMoIP data packets for the circuit bundle in question. It is incremented in query messages, and replicated without change in replies.

OAM Msg Type (8 bits) indicates the function of the message. At present the following are defined:

0 for one way connectivity query message

8 for one way connectivity reply message.

OAM Msg Code (8 bits) is used to carry information related to the message, and its interpretation depends on the message type. For type 0 (connectivity query) messages the following codes are defined:

- 0 validate connection.
- 1 do not validate connection

for type 8 (connectivity reply) messages the available codes are:

- 0 -
 - acknowledge valid query
- 1 -
 - invalid query (configuration mismatch)

Service specific information (16 bits) is a field that can be used to exchange configuration information between edge devices. If it is not used this field MUST contain zero. Its interpretation depends on the FORMID field. At present the following is defined for AAL1 payloads.

0	1								
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5								
+---+---+---+---+---+---+---+---+---+---+									
Number of TSs					Number of SFs				
+---+---+---+---+---+---+---+---+---+---+									

Number of TSs (8 bits) is the number of timeslots being transported, e.g. 24 for full T1.

Number of SFs (8 bits) is the number of 48-octet AAL1 subframes per packet, e.g. 8 when packing 8 subframes per packet.

Source CBID (16 bits) uniquely identifies the circuit bundle as labeled by the source edge.

Destination CBID (16 bits) uniquely identifies the circuit bundle as labeled by the destination edge.

Source Transmit Timestamp (32 bits) represents the time the source edge transmitted the query message in units of 100 microseconds. This field and the following ones only appear if delay is being measured.

Destination Receive Timestamp (32 bits) represents the time the destination edge received the query message in units of 100 microseconds

Destination Transmit Timestamp (32 bits) represents the time the destination edge transmitted the reply message in units of 100 microseconds.

9. General Requirements

9.1 Quality of Service

TDMoIP does not provide mechanisms to ensure timely delivery or provide other quality-of-service guarantees; hence it is required that the lower-layer services do so. Layer 2 priority can be bestowed upon a TDMoIP stream by using the VLAN priority field, layer 3 priority is controllable by using TOS. Switches and routers which the TDMoIP stream must traverse should be configured to respect these priorities.

TDMoIP assumes a relatively benevolent environment. On the IP side this means a network with prioritization and sufficient bandwidth (sufficient bandwidth can be guaranteed by undersubscription and/or traffic engineering), low probability of bit error, packet order interchange or lost packets. In particular, use of TDMoIP over the public Internet is not presently envisaged.

9.2 Timing

TDM networks are inherently synchronous; somewhere in the network there will always be at least one extremely accurate primary reference clock, with long-term accuracy of one part in 10^{11} . This node, whose accuracy is called "stratum 1", provides reference timing to secondary nodes with lower "stratum 2" accuracy, and these in turn provide reference clock to "stratum 3" nodes. This hierarchy of time synchronization is essential for the proper functioning of the network as a whole; for details [[G823](#),[G824](#)]. The use of time standards less accurate than stratum 4 is NOT RECOMMENDED as it may result in service impairments.

Packets in IP networks reach their destination with delay that has a random component, known as jitter. When emulating TDM on a PSN, it is possible to overcome this randomness by using a "jitter buffer" on all incoming data, assuming the proper time reference is available. The problem is that the original TDM time reference information is not disseminated through the PSN.

In broadest terms there are two methods of overcoming this difficulty; in one the timing information is provided by some means independent of the PSN, while in the other the timing must be transferred over the PSN.

For example, if the entire TDM infrastructure (or at least major

portions of it) is replaced by TDMoIP timing information MUST be delivered over the IP network, and the reconstructed TDM stream SHOULD conform to ITU-T recommendations [G823] for E1 and [G824] for T1 trunks. IP networks can disseminate accurate timing information through NTP [NTP], and synchronous operation of a real-time stream can be emulated using RTP as detailed in subsection 2.3 supra. RTP appends a sequence number and a timestamp to each packet, enabling the reconstitution of timing at the expense of considerable overhead.

However, TDMoIP is frequently used in a "toll-bypass" scenario, where an IP link connects two existing TDM networks. In such a case both TDMoIP devices SHALL receive accurate timing from the TDM networks to which they connect, and MUST use this local timing when outputting to the TDM network. There is no need to carry timing over the IP network and the overhead associated with RTP can be avoided.

9.3 Jitter and Packet Loss

In order to compensate for packet delay variation that exists in any IP network a jitter buffer SHALL be provided. The length of this buffer SHOULD be configurable and MAY be dynamic (i.e. grow and shrink in length according to the statistics of the delay variation).

In order to handle (infrequent) packet loss and misordering a packet order integrity mechanism SHALL be provided. This mechanism SHALL track the serial numbers of packets in the jitter buffer and MUST take appropriate action when faults are detected. When missing packet(s) are detected the mechanism SHALL output interpolation packet(s) in order to retain TDM timing. Packets with incorrect serial numbers or other detectable header errors MAY be discarded. Packets arriving in incorrect order SHOULD be swapped. Whenever possible, interpolation packets SHOULD ensure that proper synchronization bits are sent to the TDM network.

9.4 Overhead vs. Latency Trade-off

TDMoIP is designed to be parsimonious in bandwidth. To assist in achieving this goal we allow merging multiple subframes into a single packet in order to incur a single header. For example, for AAL1 payloads, there can be between N subframes per packet, where N is a configurable parameter. For structured T1/E1 when N=1 each

timeslot appears only once per packet, and hence packets are transmitted every 125 microsecond (on the average); when N=8 this becomes 1 millisecond.

While higher values of N reduce overhead, they increase the amount of time that passes between ingress of a TDM sample and its transmission over the PSN.

This buffering delay must be added to the network propagation delay and all other delays the packet experiences. Since the amount of latency that can be considered acceptable is dependent upon the application, it is essential that there be a method for tuning this trade-off between efficiency and latency.

10. Security Considerations

TDMoIP does not enhance or detract from the security performance of the underlying PSN, rather it relies upon the PSN's mechanisms for encryption, integrity, and authentication when these are required.

TDMoIP does not provide protection against malicious users utilizing snooping or packet injection during setup or operation.

Circuit bundle identifiers SHOULD be selected in an unpredictable manner rather than sequentially or otherwise in order to deter session hijacking. When using L2TP randomly selected cookies MAY be used to validate circuit bundle origin. Sequence numbers SHOULD be randomly initialized in order to increase the difficulty of decrypting based on packet headers.

11. IANA Considerations

All IANA considerations of the underlying PSN (e.g. UDP or L2TP IP types) MUST be abided.

When used with UDP/IP the destination port number MUST be set to 0x085E (2142), the user port number which has been assigned by the

to TDMoIP.

12. Applicability Statement

TDMoIP is designed to transport TDM services over a PSN. The TDM traffic may be any of the following types:

- n*64K
- unstructured T1
- unstructured E1
- structured T1 without CAS
- structured E1 without CAS
- structured T1 with CAS
- structured E1 with CAS
- T3
- E3
- AAL1
- AAL2

Both data and clock information must be transferred edge to edge. Some degradation in the stratum level is to be expected.

When present, CAS signaling is transparently transferred edge to edge. The precise phase of signaling bits inside multiframes is not always retained.

Trunk associated CCS signaling is transferred edge to edge. For HDLC type payloads the precise number of idle flags is not preserved.

Standard TDM alarms, such as those described in [subsection 2.4](#), are both transferred edge to edge and generated when required.

Latency and hence round trip delay may increased when using TDMoIP. When required echo cancellation must be provided. Voice quality will be similar to those of existing circuit-based networks, but may suffer minor degradation due to increased delay and jitter.

SONET/SDH rings are considered very reliable because of their capability to rapidly recovery from a failure or fiber cut. TDMoIP does not inherently have this capability, but the underlying PSN

may be required to redirect traffic to an alternative trunk within a few milliseconds.

When transporting SS7 signaling, the availability and dependability performance objectives of [Q766] will not be met unless further mechanisms are put in place.

13. References

- [UDP] [RFC 768](#) (STD0006) User Datagram Protocol (UDP)
- [Ipv4] [RFC 791](#) (STD0005) Internet Protocol (IP)
- [NTP] [RFC 1305](#) Network Time Protocol (NTP) (Version 3)
- [RTP] [RFC 1889](#) RTP: A Transport Protocol for Real-Time Applications
- [IPPM] [RFC 2330](#) Framework for IP Performance Metrics
- [CONNECT] [RFC 2678](#) IPPM Metrics for Measuring Connectivity
- [DELAY] [RFC 2679](#) A One-way Delay Metric for IPPM
- [MPLS] [RFC 3032](#) MPLS Label Stack encoding
- [L2TPv3] [draft-ietf-l2tpext-l2tp-base-01.txt](#)
Layer Two Tunneling Protocol (L2TP)
- [G704] ITU-T Recommendation G.704 (10/98)
Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44736 Kbit/s hierarchical levels

[G823] ITU-T Recommendation G.823 (03/00)
The control of jitter and wander within digital networks which are based on the 2048 Kbit/s hierarchy

[G824] ITU-T Recommendation G.824 (03/00)
The control of jitter and wander within digital networks which are based on the 1544 Kbit/s hierarchy

[Q766] ITU-T Recommendation Q.766 (03/93)
Performance Objectives in the Integrated Services Digital Network
Application (SS7 ISUP)

[AAL1] ITU-T Recommendation I.363.1 (08/96)
B-ISDN ATM Adaptation Layer (AAL) specification: Type 1

[AAL2] ITU-T Recommendation I.363.2 (11/00)
B-ISDN ATM Adaptation Layer (AAL) specification: Type 2

[CES] ATM forum specification atm-vtoa-0078 (CES 2.0)
Circuit Emulation Service Interoperability Specification Ver. 2.0

[LES] ATM forum specification atm-vmoa-0145 (LES)
Voice and Multimedia over ATM - Loop Emulation Service Using AAL2

14. Acknowledgments

The authors would like to thank Hugo Silberman, Shimon HaLevy, Ron Insler, and Eitan Schwartz of RAD Data Communications as well as HannsJuerген Schwarzbauer and Maximilian Riegel of Siemens for their valuable contributions to the technology described herein.

15. Contact Information

Yaakov (Jonathan) Stein
RAD Data Communications
24 Raoul Wallenburg St., Bldg C
Tel-Aviv 69719
ISRAEL
Phone: +972 3 645-5389
Email: yaakov_s@rad.co.il

Motty (Mordechai) Anavi
RAD Data Communications
900 Corporate Drive,
Mahwah, NJ 07430
USA
Phone: +1 201 529-1100 Ext. 213
Email: motty@radusa.com

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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."

