

Network Working Group
Request for Comments: 4734
Obsoletes: 2833
Updates: 4733
Category: Standards Track

H. Schulzrinne
Columbia U.
T. Taylor
Nortel
December 2006

Definition of Events for Modem, Fax, and Text Telephony Signals

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 IETF Trust (2006).

Abstract

This memo updates RFC 4733 to add event codes for modem, fax, and text telephony signals when carried in the telephony event RTP payload. It supersedes the assignment of event codes for this purpose in RFC 2833, and therefore obsoletes that part of RFC 2833.

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Overview	3
2. Definitions of Events for Control of Data, Fax, and Text Telephony Sessions	5
2.1. V.8 bis Events	5
2.1.1. Handling of Congestion	9
2.2. V.21 Events	10
2.2.1. Handling of Congestion	11
2.3. V.8 Events	12
2.3.1. Handling of Congestion	15
2.4. V.25 Events	15
2.4.1. Handling of Congestion	17
2.5. V.32/V.32bis Events	18
2.5.1. Handling of Congestion	19
2.6. T.30 Events	19
2.6.1. Handling of Congestion	23
2.7. Events for Text Telephony	23
2.7.1. Signal Format Indicators for Text Telephony	23
2.7.2. Use of Events with V.18 Modems	27
2.8. A Generic Indicator	28
3. Strategies for Handling Fax and Modem Signals	29
4. Example of V.8 Negotiation	30
4.1. Simultaneous Transmission of Events and Retransmitted Events Using RFC 2198 Redundancy	35
4.2. Simultaneous Transmission of Events and Voice-Band Data Using RFC 2198 Redundancy	37
5. Security Considerations	39
6. IANA Considerations	40
7. Acknowledgements	42
8. References	43
8.1. Normative References	43
8.2. Informative References	44

1. Introduction

1.1. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [1].

In addition to those defined for specific events, this document uses the following abbreviations:

Fax facsimile

HDLC High-level Data Link Control

PSTN Public Switched (circuit) Telephone Network

1.2. Overview

This document extends the set of telephony events defined within the framework of RFC 4733 [5] to include the control events and tones that can appear on a subscriber line serving a fax machine, a modem, or a text telephony device. The events are organized into several groups, corresponding to the ITU-T Recommendation in which they are defined. Their purpose is to support negotiation, start-up and takedown of fax, modem, or text telephony sessions and transitions between operating modes. The actual fax, modem, and text payload is typically carried by other payload types (e.g., V.150.1 [32] modem relay, voice-band data as formalized in ITU-T Rec. V.152 [33], Clearmode [17] for digital data, T.38 [21] for fax, or RFC 4103 [18] for character-mode text).

NOTE: implementers SHOULD NOT rely on the descriptions of the various modem protocols described below without consulting the original references (generally ITU-T Recommendations). The descriptions are provided in this document to give a context for the use of the events defined here. They frequently omit important details needed for implementation.

The typical application of these events is to allow the Internet to serve as a bridge between terminals operating on the PSTN. This application is characterized as follows:

- o each gateway will act both as sender and as receiver;
- o time constraints apply to the exchange of signals, making the early identification and reporting of events desirable so that receiver playout can proceed in a timely fashion;

- o the receiver must play out events in their proper order;
- o transfer of the events must be reliable. Applications will vary in their ability to recover from missing events.

In some cases, an implementation may simply ignore certain events, such as fax tones, that do not make sense in a particular environment. Section 2.4.1 of RFC 4733 [5] specifies how an implementation can use the Session Description Protocol (SDP) "fmt" parameter within an SDP description [4] to indicate which events it is prepared to handle.

Regardless of which events they support, implementations **MUST** be prepared to send and receive data signals using payload types other than telephone-event, simultaneously with the use of the latter. This is discussed further in Section 3.

In many cases, continuity of playout is critical. In principle, this is achieved through buffering at the receiving end. It is generally desirable to minimize such buffering to reduce round-trip response times. Maintenance of a constant packetization interval at the sending end while reporting events is helpful for this purpose.

A further word on time constraints is in order. Time constraints governing the duration of tones do not pose a problem when using the telephone-event payload type: the payload specifies the duration and the receiving gateway can play out the tones accordingly. Problems occur when time constraints are specified for the duration of silence between tones. A silent period of "at least x ms" is not a problem -- event notifications can be received late, but they can still be played out at their specified durations.

The problem occurs if silence must last for a specific duration or at most some specific period. The most general constraint of the latter type has to do with the operation of echo suppressors (ITU-T Rec. G.164 [6]) and echo cancellers (ITU-T Rec. G.165 [7]). These devices may re-activate after as little as 100 ms of no signal on the line. As a result, in any situation where echo suppressors or cancellers must be disabled for signalling to work, tone events must be reported quickly enough to ensure that these devices do not become re-enabled.

2. Definitions of Events for Control of Data, Fax, and Text Telephony Sessions

2.1. V.8 bis Events

Recommendation V.8 bis [10] is a general procedure for two endpoints to establish each other's capabilities and to transition between different operating modes, both at call startup and after the call has been established. It supports many of the same terminals as V.8 [9] (Section 2.3 below), but allows more detailed parameter negotiation. It lacks support for some of the older V-series modems defined in V.8, but adds capabilities for simultaneous or alternating voice and data, H.324 [20] multilink, and T.120 [23] conferencing.

Following V.8 bis capability negotiations, if the terminals have negotiated a modem-based operating mode, they initiate the actual modem session using either V.8, a truncated version of V.8 (preferred), or V.25 start-up. V.25 is described in Section 2.4.

V.8 bis distinguishes between "signals" and "messages". The V.8 bis signals -- ESi/ESr, MRe/MRd, and CRe/CRd -- consist of tones, as described in the next few paragraphs. The V.8 bis messages -- MS, CL, CLR, ACK(1), ACK(2), NAK(1), NAK(2), NACK(3), and NACK(4) -- consist of sequences of bits transported over V.21 [12] modulation.

Signals are intended to be comprehensible at the receiver even in the presence of voice content. They consist of two tone segments. The first segment consists of a dual-frequency tone held for 400 ms, and has the function of preparing the receiver and any in-line echo suppressor or canceller for what follows. The specific frequencies depend only on whether the signal is from the initiator or the responder in a transaction. When using the telephone-event payload, the V8bISeg and V8bRSeg events in Table 1 represent the first segment of any V.8 bis signal in the initiating and responding case, respectively.

The complete V.8 bis strategy for dealing with echo suppressors or cancellers is described in Rec. V.8 bis Appendix III. The only silent period constraints imposed are of the "at least" type, posing no difficulties for the use of the telephone-event payload.

The second segment follows immediately after the first, and is a single tone held for 100 ms. The frequency used indicates the specific signal of the six signals defined. When using the telephone-event payload, the second segment of a V.8 bis signal is represented by the applicable event: CRdSeg, CReSeg, MRdSeg, MReSeg,

ESiSeg, or ESrSeg, as defined in Table 1. ESiSeg and ESrSeg use the same frequencies as V.21 low and high channel '1' bits, respectively (see Table 2), and are therefore assigned the same event codes.

V.8 bis messages use V.21 [12] frequency-shift signalling to transfer message content. V.21 is described in the next section. V.8 bis uses V.21 in half-duplex mode at 300 bits/s, with the lower channel assigned to the initiator and the upper channel to the responder.

Each V.8 bis message is preceded by a 100-ms preamble of continuous V.21 marking frequency except if it was immediately preceded by an ESi or ESr signal (the second segment of which is that same V.21 marking frequency). The sender SHALL NOT report this preamble tone using the ESiSeg or ESrSeg events; these are to be used only for the V.8 bis signals to which they pertain.

Spelling this out, continuous V.21 marking tone immediately following V8bISeg and V8bRSeg is reported as ESiSeg or ESrSeg, respectively. Continuous V.21 marking tone occurring in any other context, and particularly after CRdSeg, CReSeg, MRdSeg, or MReSeg, is reported by other means such as a different payload type or using the V.21 '1' bit events defined in Section 2.2.

No events are defined for V.8 bis messages, but a brief description follows.

- o the V.8 bis CL message describes the sending terminal's capabilities;
- o the CLR message also describes capabilities, but indicates that the sender wants to receive a CL in return;
- o the MS establishes a particular operating mode;
- o the ACK and NAK messages are used to terminate the message transactions.

The V.8 bis messages are organized as a sequence of octets. The first two to five octets are HDLC flags (0x7E). Then comes a message type identifier (four bits), a V.8 bis version identifier (four bits), zero to two more octets of identifying information, followed by zero or more information field parameters in the form of bit maps. An individual bit map is one to five octets in length. Up to 64 octets of non-standard information may also be present. The information fields are followed by a checksum and one to three HDLC flags. Because of limits on the size of any one information field, V.8 bis defines segmentation procedures. Excess data is sent in an additional message, but only after prompting from the receiving end.

Applications supporting V.8 bis signalling using the telephone-event payload MAY transfer V.8 bis messages in the form of sequences of bits, using the V.21 bit events defined in the next section. If they do so, the transmitted information MUST include the complete contents of the message: the initial HDLC flags, the information field, the checksum, and the terminating HDLC flags.

Transmission MUST also include the extra '0' bits added according to the procedures of Rec. V.8 bis, clause 7.2.8, to prevent false recognition of HDLC flags at the receiver. Implementers should note that these extra '0' bits mean that in general V.8 bis messages as transmitted on the wire will not come out to an even multiple of octets. Sending implementations MAY choose to vary the packetization interval to include exactly one octet of information plus any extra '0' bits inserted into that octet; the resulting variation will be insignificant compared with the amount of buffering required to guard against network delays in delivery of packets to the receiver (see below).

One reason for reporting the V.21 bits exactly as presented on the wire is to match the corresponding content if it is also carried by other means, such as voice-band data.

The power levels of the V.8 bis and V.21 signals are subject to national regulation. Thus, it seems suitable to model V.8 bis events as tones for which the volumes SHOULD be specified by the sender. If the receiver is rendering the V.8 bis tones as audio content for onward transmission, the receiver MAY use the volumes contained in the event reports, or MAY modify the volumes to match downstream national requirements.

Table 1 summarizes the event codes defined for V.8 bis signalling in this document. The individual events are described following the table. Each event begins when the beginning of the tone segment is detected and ends when the tone is no longer detected.

Event	Freq. (Hz)	Dur. (ms)	Event Code	Type	Volume?
ESiSeg	980	100	38	tone	yes
ESrSeg	1650	100	40	tone	yes
CRdSeg	1900	100	23	tone	yes
CReSeg	400	100	24	tone	yes
MRdSeg	1150	100	25	tone	yes
MReSeg	650	100	26	tone	yes
V8bISeg	1375 + 2002	400	28	tone	yes
V8bRSeg	1529 + 2225	400	29	tone	yes

Table 1: Events for V.8 bis Signals

ESiSeg:

The second segment of a V.8 bis initiating Escape Signal (ESi). The complete ESi signal is represented by events V8bISeg followed by ESiSeg. ESi will be followed by an MS, CL, or CLR message from the same terminal. A 1.5-s silent interval may come between the ESi signal and the transmission of the MS, CL, or CLR message to accommodate network echo suppressors.

ESrSeg:

The second segment of a V.8 bis responding Escape Signal (ESr). The complete ESr signal is represented by events V8bRSeg followed by ESrSeg. ESr is always sent by the calling terminal in response to an MRe or CRe from an automatic answering station. It will be followed by an MS, CL, or CLR message. The ESr signal turns off any announcement being generated by the automatic answering station.

CRdSeg:

The second segment of a V.8 bis Capabilities Request signal (CRd). The first segment of the CRd signal is represented either by V8bISeg or V8bRSeg, depending on context. The other end will return a capabilities list (CL or CLR message).

CReSeg:

The second segment of a V.8 bis Capabilities Request signal (CRe) initiated by an automatic answering terminal. The complete CRe signal is represented by events V8bISeg followed by CReSeg. The calling terminal will respond with a CRd signal or a CL or CLR message.

MRdSeg:

The second segment of a V.8 bis Mode Request signal (MRd). The first segment of the MRd signal is represented either by V8bISeg or V8bRSeg, depending on context. The other end will return a CRd signal or an MS message.

MReSeg:

The second segment of a V.8 bis Mode Request signal (MRe) initiated by an automatic answering terminal. The complete MRe signal is represented by events V8bISeg followed by MReSeg. The calling terminal will respond with an MRd or CRd signal or an MS message.

V8bISeg:

The first segment of an initiating V.8 bis signal, which may be one of ESi, CRd, CRe, MRd, or MRe.

V8bRSeg:

The first segment of a responding V.8 bis signal, which may be one of ESr, CRd, or MRd.

2.1.1.1. Handling of Congestion

V.8 bis implementations are unlikely to tolerate gaps or extensions in playout times due to congestion-caused packet delay. At a minimum, the current transaction is liable to be reset when these defects in playout occur. As a result, careful management of the playout buffer is required at the receiver to increase robustness in the face of possible lost or delayed packets. The playout algorithm should also be such as not to cause event playout to exceed the nominal duration of the event.

V.8 bis does not appear to offer opportunities for dynamic adaptation to congestion through manipulation of the packetization interval.

2.2. V.21 Events

V.21 [12] is a modem protocol offering data transmission at a maximum rate of 300 bits/s. Two channels are defined, supporting full duplex data transmission if required. The low channel uses frequencies 980 Hz for '1' (mark) and 1180 Hz for '0' (space); the high channel uses frequencies 1650 Hz for '1' and 1850 Hz for '0'. The modem can operate synchronously or asynchronously.

V.21 is used by other protocols (e.g., V.8 bis, V.18, T.30) for transmission of control data, and is also used in its own right between text terminals. The V.21 events are summarized in Table 2.

Sending implementations SHOULD report a completed event for every bit transmitted (i.e., rather than at transitions between '0' and '1'). Bit events are assumed to begin and end with the clock interval for the event, neglecting the rise and fall times between bit transitions. Thus, it is important for a gateway to determine the actual bit rate in use before beginning to report V.21 events.

Sometimes determination of the bit rate is not immediately possible, as in the case of the 100-ms training signal at V.21 mark frequency used before V.8 bis messages. Transmission of a single longer-duration V.21 event is reasonable under these circumstances and should not cause any difficulties at the receiving end.

Implementations SHOULD pack multiple events into one packet, using the procedures of Section 2.5.1.5 of RFC 4733 [5]. Eight to ten bits is a reasonable packetization interval.

Reliable transmission of V.21 events is important, to prevent data corruption. Reporting an event per bit rather than per transition increases reporting redundancy and thus reporting reliability, since each event completion is transmitted three times as described in Section 2.5.1.4 of RFC 4733 [5]. To reduce the number of packets required for reporting, implementations SHOULD carry the retransmitted events using RFC 2198 [2] redundancy encoding. This is illustrated in the example in Section 4.1.

The time to transmit one V.21 bit at the nominal rate of 300 bits/s is 3.33 ms, or 26.67 timestamp units at the default 8000-Hz sampling rate for the telephone-event payload type. Because this duration is not an integral number of timestamp units, accurate reporting of the beginning of the event and the event duration is impossible. Sending gateways SHOULD round V.21 event starting times to the nearest whole timestamp unit.

When sending multiple consecutive V.21 events in a succession of packets, the sending gateway MUST ensure that individual event durations reported do not cause the last event of one packet to overlap with the first event of the next, taking into account the respective initial event timestamps. To accomplish this, the sending gateway MUST derive the individual event durations as the succession of differences between the event starting times (so that, at 8000 Hz, every third event has reported duration 26 units, the remainder 27 units).

Where a receiving gateway recognizes that a packet reports a consecutive series of V.21 bit events, it SHOULD play them out at a uniform rate despite the possible one-timestamp-unit discrepancies in their reported spacing and duration.

Event	Frequency (Hz)	Event Code	Type	Volume?
V.21 channel 1, '0' bit	1180	37	tone	yes
V.21 channel 1, '1' bit	980	38	tone	yes
V.21 channel 2, '0' bit	1850	39	tone	yes
V.21 channel 2, '1' bit	1650	40	tone	yes

Table 2: Events for V.21 Signals

Implementations that choose to transmit V.21 content using a different payload type may wish to use one of the indicator events defined in Table 7 to alert the receiver to the nature of the content. It is not expected that an implementation will send both one of these indicator events and the V.21 bit events defined above for the same content.

2.2.1. Handling of Congestion

The duration of V.21 bits cannot be extended from its nominal value (which depends on the transmission rate). The playout algorithm at the receiver should take this constraint into account when compensating for the delay or loss of packets due to congestion.

Other congestion-related considerations depend on the specific application for which the V.21 bit events are being used.

2.3. V.8 Events

V.8 [9] is an older general negotiation and control protocol, supporting startup for the following terminals: H.324 [20] multimedia, V.18 [11] text, T.101 [22] videotext, T.30 [8] send or receive fax, and a long list of V-series modems including V.34 [28], V.90 [29], V.91 [30], and V.92 [31]. In contrast to V.8 bis [10], in V.8 only the calling terminal can determine the operating mode.

V.8 does not use the same terminology as V.8 bis. Rather, it defines four signals that consist of bits transferred by V.21 [12] at 300 bits/s: the call indicator signal (CI), the call menu signal (CM), the CM terminator (CJ), and the joint menu signal (JM). In addition, it uses tones defined in V.25 [13] and T.30 [8] (described below), and one tone (ANSam) defined in V.8 itself. The calling terminal sends using the V.21 low channel; the answering terminal uses the high channel.

The basic protocol sequence is subject to a number of variations to accommodate different terminal types. A pure V.8 sequence is as follows:

1. After an initial period of silence, the calling terminal transmits the V.8 CI signal. It repeats CI at least three times, continuing with occasional pauses until it detects ANSam tone. The CI indicates whether the calling terminal wants to function as H.324, V.18, T.30 send, T.30 receive, or a V-series modem.
2. The answering terminal transmits ANSam after detecting CI. ANSam will disable any G.164 [6] echo suppressors on the circuit after 400 ms and any G.165 [7] echo cancellers after one second of ANSam payout.
3. On detecting ANSam, the calling terminal pauses at least half a second, then begins transmitting CM to indicate detailed capabilities within the chosen mode.
4. After detecting at least two identical sequences of CM, the answering terminal begins to transmit JM, indicating its own capabilities (or offering an alternative terminal type if it cannot support the one requested).

5. After detecting at least two identical sequences of JM, the calling terminal completes the current octet of CM, then transmits CJ to acknowledge the JM signal. It pauses exactly 75 ms, then starts operating in the selected mode.
6. The answering terminal transmits JM until it has detected CJ. At that point, it stops transmitting JM immediately, pauses exactly 75 ms, then starts operating in the selected mode.

The CI, CM, and JM signals all consist of a fixed sequence of ten '1' bits followed by a signal-dependent pattern of ten synchronization bits, followed by one or more octets of variable information. Each octet is preceded by a '0' start bit and followed by a '1' stop bit. The combination of the synchronization pattern and V.21 channel uniquely identifies the message type. The CJ signal consists of three successive octets of all zeros with stop and start bits but without the preceding '1's and synchronizing pattern of the other signals.

Applications MAY report each instance of a CM, JM, and CJ signal, respectively, as a series of V.21 bit events (Section 2.2), or may use another payload type to carry this information. Applications supporting V.8 signalling using the telephone-event payload MAY report the synchronization part of the CI signal (ten '1's followed by '00000 00001') both as a series of V.21 bit events and, when it has been recognized, as a single CI event.

Note that the CI event covers only the synchronization part of the CI signal. The remaining call function octet and its start and stop bits need to be transmitted also, either as a series of V.21 bit events or in some other payload format. Presumably, the calling end gateway will use the same format for the CM and CJ signals.

The overlapping nature of V.8 signalling means that there is no risk of silence exceeding 100 ms once ANSam has disabled any echo control circuitry. However, the 75-ms pause before entering operation in the selected data mode will require both the calling and the answering gateways to recognize the completion of CJ, so they can change from playout of telephone-event to playout of the data-bearing payload after the 75-ms period.

Event	Frequency (Hz)	Event Code	Type	Volume?
ANSam	2100 x 15	34	tone	yes
/ANSam	2100 x 15 phase rev.	35	tone	yes
CI	(V.21 bits)	53	tone	yes

Table 3: Events for V.8 Signals

ANSam:

The modified answer tone ANSam consists of a sinewave signal at 2100 Hz, amplitude-modulated by a sine wave at 15 Hz. The beginning of the event is at the beginning of the tone. The end of the event is at the sooner of the ending of the tone or the occurrence of a phase reversal (marking the beginning of a /ANSam event). Phase reversals are used to disable echo cancellation; if they are being applied, they occur at 450-ms intervals.

An ANSam event packet SHOULD NOT be sent until it is possible to discriminate between an ANSam event and an ANS event (see V.25 events, below).

The modulated envelope for the ANSam tone ranges in amplitude between 0.8 and 1.2 times its average amplitude. The average transmitted power is governed by national regulations. Thus, it makes sense to indicate the volume of the signal.

/ANSam:

/ANSam reports the same physical signal as ANSam, but is reported following the first phase reversal in that signal. It begins with the phase reversal and ends at the end of the tone. The receiver of /ANSam MUST reverse the phase of the tone at the beginning of playout of /ANSam and every 450 ms thereafter until the end of the tone is reached.

CI:

CI reports the occurrence of the V.21 bit pattern '11111 11111 00000 00001' indicating the beginning of a V.8 CI signal. The event begins at the beginning of the first bit and ends at the end of the last one. This event MUST NOT be reported except in a context where a V.8 CI signal might be expected (i.e., at the calling end during call setup). Note that if the calling modem

sends the CI signal at all, it will typically repeat the signal several times.

It is expected that the CI event will be most useful when the modem content is being transmitted primarily using another payload type. The event acts as a commentary on that content, allowing the receiver to recognize that V.8 signalling is in progress.

2.3.1. Handling of Congestion

The tolerances built into V.8 suggest that it may be mostly robust in the face of packet losses or delays. Payout of ANSam and /ANSam can be extended for multiple packetization periods without harm, provided that phase reversals occur on schedule at 450-ms intervals during payout of the latter.

To increase robustness of transmission of the V.21-based signals, sending applications using the V.21 events SHOULD include an integral number of octets, including start and stop bits, in each packet. The presence of start and stop bits provides some hope that receiving implementations can withstand unavoidable gaps in payout between octets. When a message is being repeated (as is possible for CI, CM, and JM), an even stronger robustness measure would be for the receiver to retain a copy of the message when it is first received, and when a packet is delayed or lost to continue playing out the current message instance and commence a new repetition as if packets had continued to arrive on schedule.

2.4. V.25 Events

V.25 [13] is a start-up protocol predating V.8 [9] and V.8 bis [10]. It specifies the exchange of two tone signals: CT and ANS.

CT (calling tone) consists of a series of interrupted bursts of 1300-Hz tone, on for a duration of not less than 0.5 s and not more than 0.7 s and off for a duration of not less than 1.5 s and not more than 2.0 s. [13]. Modems not starting with the V.8 CI signal often use this tone.

ANS (Answer tone) is a 2100-Hz tone used to disable echo suppression for data transmission [13], [8]. For fax machines, Recommendation T.30 [8] refers to this tone as called terminal identification (CED) answer tone. ANS differs from V.8 ANSam in that, unlike the latter, it has constant amplitude.

V.25 specifically includes procedures for disabling echo suppressors as defined by ITU-T Rec. G.164 [6]. However, G.164 echo suppressors have now for the most part been replaced by G.165 [7] echo

cancellers, which require phase reversals in the disabling tone (see ANSam above). As a result, Recommendation V.25 was modified in July 2001 to say that phase reversal in the ANS tone is required if echo cancellers are to be disabled.

One possible V.25 sequence is as follows:

1. The calling terminal starts generating CT as soon as the call is connected.
2. The called terminal waits in silence for 1.8 to 2.5 s after answer, then begins to transmit ANS continuously. If echo cancellers are on the line, the phase of the ANS signal is reversed every 450 ms. ANS will not reach the calling terminal until the echo control equipment has been disabled. Since this takes about a second, it can only happen in the gap between one burst of CT and the next.
3. Following detection of ANS, the calling terminal may stop generating CT immediately or wait until the end of the current burst to stop. In any event, it must wait at least 400 ms (at least 1 s if phase reversal of ANS is being used to disable echo cancellers) after stopping CT before it can generate the calling station response tone. This tone is modem-specific, not specified in V.25.
4. The called terminal plays out ANS for 2.6 to 4.0 seconds or until it has detected calling station response for 100 ms. It waits 55-95 ms (nominal 75 ms) in silence. (Note that the upper limit of 95 ms is rather close to the point at which echo control may reestablish itself.) If the reason for ANS termination was timeout rather than detection of calling station response, the called terminal begins to play out ANS again to maintain disabling of echo control until the calling station responds.

The events defined for V.25 signalling are shown in Table 4.

Event	Frequency (Hz)	Event Code	Type	Volume?
Answer tone (ANS)	2100	32	tone	yes
/ANS	2100 ph. rev.	33	tone	yes
CT	1300	49	tone	yes

Table 4: Events for V.25 Signals

ANS:

The beginning of the event is at the beginning of the 2100-Hz tone. The end of the event is at the sooner of the ending of the tone or the occurrence of a phase reversal (marking the beginning of a /ANS event).

An initial ANS event packet SHOULD NOT be sent until it is possible to discriminate between an ANS event and an ANSam event (see V.8 events, above).

/ANS:

/ANS reports the same physical signal as ANS, but is reported following the first phase reversal in that signal. It begins with the phase reversal and ends at the end of the tone. The receiver of /ANS MUST reverse the phase of the tone at the beginning of playout of /ANS and every 450 ms thereafter until the end of the tone is reached.

CT:

The beginning of the CT event is at the beginning of an individual burst of the 1300-Hz tone. The end of the event is at the end of that tone burst. The gateway at the calling end SHOULD use a packetization interval smaller than the nominal duration of a CT burst, to ensure that CT playout at the called end precedes the sending of ANS from that end.

2.4.1. Handling of Congestion

The V.25 sequence appears to be robust in the face of lost or delayed packets, provided that the receiver continues to play out any tone it is in the process of playing until more packets are received. The receiver must play out the phase transitions for /ANS on schedule, at 450-ms intervals, even if updates of the /ANS event have been delayed. It also appears to be possible for the sender to temporarily increase the packetization interval to reduce packet volumes when congestion is encountered. The one risk is that extended playout proceeds past the actual end of the tone (as determined retroactively), and the receiver is forced to continue imposing an additional playout buffering lag in order to meet the constraint on maximum duration of the nominal 75-ms silent period following tone playout.

2.5. V.32/V.32bis Events

ITU-T Recommendation V.32 [14] is a modem using phase-shift keying with quadrature amplitude modification. It operates on a carrier at 1800 Hz, modulated at 2400 symbols/s. The basic data rates for V.32 are 4800 and 9600 bits/s. V.32bis [15] extends the data rates up to 14,400 bits/s. Most or all existing deployments are V.32bis, typically in support of point-of-sale terminals and the like.

One reason V.32bis is still used is because of its relatively rapid start-up sequence, particularly on leased lines. Operating over the public telephone network, the start-up begins as follows:

- a. the answering end begins with the V.25 answering procedure (1.8 to 2.5 s of silence followed by continuous ANS tone to a maximum of 3.3 s, with possible phase reversals to disable echo cancelling equipment);
- b. the calling end waits in silence until it has detected ANS for 1 s;
- c. the calling end begins to transmit a V.32/V.32bis pattern designated AA, i.e., a series of '0000' bit sequences transmitted at 4800 bits/s;
- d. upon detecting the AA pattern for at least 100 ms, the called modem is silent for 75 +/- 20 ms, then responds with an AC pattern, which is a series of '0011' bit sequences transmitted at 4800 bits/s.

The difference in leased line operation is that the calling modem starts the session by sending AA. After that, the called modem responds with AC, and the rest of the sequence is unchanged.

In support of V.32/V.32bis operation, Table 5 defines two events, V32AA and V32AC.

Event	Bit Pattern	Event Code	Type	Volume?
V32AA	b'0000' repeated	63	tone	yes
V32AC	b'0011' repeated	27	tone	yes

Table 5: Events for V.32/V.32bis Signals

V32AA:

Indicates that the AA calling pattern of a V.32/V.32bis terminal has been detected.

V32AC:

Indicates that the AC answering pattern of a V.32/V.32bis terminal has been detected.

Each of these two events begins at the beginning of its pattern, and ends nominally when the pattern stops being received. Following the sending of either of these events the session may continue using V.150.1 modem relay [32] or Clearmode [17] as negotiated or configured in advance. To help make the transition as quickly as possible, the V32AA or V32AC event SHOULD be reported as soon as the corresponding pattern is detected. It seems likely that the implementation will be transmitting the event reports simultaneously with the same data in an alternate form, typically using RFC 2198 [2] redundancy.

2.5.1. Handling of Congestion

The primary issue raised by congestion is the loss or undue delay of the initial report. Once the receiver is aware that an AA or AC pattern has been detected, further reports are of no interest. The actual duration of the AC pattern may be as short as 27 ms. On this basis, the appropriate sender behavior may be to send at least three packets reporting the event using normal event updates and end of event retransmission behavior and a fairly short packetization interval (20-30 ms).

2.6. T.30 Events

ITU-T Recommendation T.30 [8] defines the procedures used by Group III fax terminals. The pre-message procedures for which the events of this section are defined are used to identify terminal capabilities at each end and negotiate operating mode. Post-message procedures are also included, to handle cases such as multiple document transmission. Fax terminals support a wide variety of protocol stacks, so T.30 has a number of options for control protocols and sequences.

T.30 defines two tone signals used at the beginning of a call. The CNG signal is sent by the calling terminal. It is a pure 1100-Hz tone played in bursts: 0.5 s on, 3 s off. It continues until timeout

or until the calling terminal detects a response. Its primary purpose is to let human operators at the called end know that a fax terminal has been activated at the calling end.

The called terminal waits in silence for at least 200 ms. It then may return CED tone (which is physically identical to V.25 ANS), or else V.8 ANSam if it has V.8 capability. If called and calling terminals both support V.8, the called terminal will detect CI or more likely CM in response to its ANSam and will continue with V.8 negotiation. Otherwise, the called terminal stops transmitting CED after 2.6 to 4 seconds, waits 75 +/- 20 ms in silence, then enters the T.30 negotiation phase.

In the T.30 negotiation phase the terminals exchange binary messages using V.21 signals, high channel frequencies only, at 300 bits/s. Each message is preceded by a one-second (nominal) preamble consisting entirely of HDLC flag octets (0x7E). This flag has the function of preparing echo control equipment for the message that follows.

The pre-transfer messages exchanged using the V.21 coding are:

Digital Identification Signal (DIS):

Characterizes the standard ITU-T capabilities of the called terminal. This is always the first message sent.

Digital Transmit Command (DTC):

A possible response to the DIS signal by the calling terminal. It requests the called terminal to be the transmitter of the fax content.

Digital Command Signal (DCS):

A command message sent by the transmitting terminal to indicate the options to be used in the transmission and request that the other end prepare to receive fax content. This is sent by the calling end if it will transmit, or by the called end in response to a DTC from the calling end. It is followed by a training signal, also sent by the transmitting terminal.

Confirmation To Receive (CFR):

A digital response confirming that the entire pre-message procedure including training has been completed and the message transmissions may commence.

Each message may consist of multiple frames bounded by HDLC flags. The messages are organized as a series of octets, but like V.8 bis, T.30 calls for the insertion of extra '0' bits to prevent spurious recognition of HDLC flags.

T.30 also provides for the transmission of control messages after document transmission has completed (e.g., to support transmission of multiple documents). The transition to and from the modem used for document transmission (V.17 [24], V.27ter [26], V.29 [27], V.34 [28]) is preceded by 75 ms (nominal) of silence).

Applications supporting T.30 signalling using the telephone-event payload MAY report the preamble preceding each message both as a series of V.21 bit events and, when it has been recognized, as a single V.21 preamble event. The T.30 control message following the preamble MAY be reported in the form of a sequence of V.21 bit events or using some other payload type. If transmitted as bit events, the transmitted information MUST include the complete contents of the message: the initial HDLC flags, the information field, the checksum, the terminating HDLC flags, and the extra '0' bits added to prevent false recognition of HDLC flags at the receiver. Implementers should note that these extra '0' bits mean that in general T.30 messages as transmitted on the wire will not come out to an even multiple of octets.

The training signal sent by the transmitting terminal after DCS consists of a steady string of V.21 high channel zeros (1850-Hz tone) for 1.5 s. Since the bit rate (nominally 300 bits/s) should have been clearly established when processing the preceding signalling, it is natural that if the telephony-event payload type is being used, this training signal will also be sent as a series of V.21 bit events at that bit rate. However, if the sending gateway is capable of recognizing the transition from the end of the DCS to the start of training, it MAY report the training signal as a single extended V.21 (high channel) '0' event.

The events defined for T.30 signalling are shown in Table 6. The CED and /CED events represent exactly the same tone signals as V.25 ANS and /ANS, and are given the same codepoints; they are reproduced here only for convenience.

Event	Frequency (Hz)	Event Code	Type	Volume?
CED (Called tone)	2100	32	tone	yes
/CED	2100 ph. rev.	33	tone	yes
CNG (Calling tone)	1100	36	tone	yes
V.21 preamble flag	(V.21 bits)	54	tone	yes

Table 6: Events for T.30 Signals

CED:

The beginning of the event is at the beginning of the 2100-Hz tone. The end of the event is at the sooner of the ending of the tone or the occurrence of a phase reversal (marking the beginning of a /CED event).

An initial CED event packet SHOULD NOT be sent until it is possible to discriminate between a CED event and an ANSam event (see V.8 events, above).

/CED:

/CED reports the same physical signal as CED, but is reported following the first phase reversal in that signal. It begins with the phase reversal and ends at the end of the tone. The receiver of /CED MUST reverse the phase of the tone at the beginning of playout of /CED and every 450 ms thereafter until the end of the tone is reached.

CNG:

The beginning of the CNG event is at the beginning of an individual burst of the 1100-Hz tone. The end of the event is at the end of that tone burst.

V.21 preamble flag:

This event begins with the first V.21 bits transmitted after a period of silence. It ends when a pattern of V.21 bits other than an HDLC flag is observed. This means that the V.21 preamble event absorbs the initial HDLC flags of the following message.

It is expected that the V.21 preamble flag event will be most useful when the modem content is being transmitted primarily using another payload type. The event acts as a commentary on that content, allowing the receiver to prepare itself to transition to fax mode.

2.6.1. Handling of Congestion

T.30 appears to be an intermediate case in terms of its vulnerability to congestion. Tone playout in the face of packet delay or loss is subject to the same considerations as for V.25 (see Section 2.4.1). Similarly, the receiver may extend playout of the preamble event while waiting for further reports. However, gaps or extended playout of the V.21 sequences are not feasible. This means, as with V.8 bis, that the receiver must manage its playout buffer appropriately to increase robustness in the face of congestion.

2.7. Events for Text Telephony

2.7.1. Signal Format Indicators for Text Telephony

Legacy text telephony uses a wide variety of terminals, with different standards favored in different parts of the world. Going forward, the vision is that new terminals will work directly into the packet network and be based on RFC 4103 [18] packetization of character data. In anticipation of this migration, it is RECOMMENDED that text carried in the PSTN by legacy modem protocols be converted to RFC 4103 packets at the sending gateway.

During a transitional period, however, gateways of a lesser capability may be able to recognize the nature of incoming content, but may only be able to encode it as voice-band data on the packet side. In such circumstances, it will help to optimize processing of the signal at the receiving end if that end receives an indication of the nature of the voice-encoded data signals. The events defined in this section provide such indications, and MAY be used in conjunction with ITU-T Recommendation V.152 [33], as one example, to carry the content as voice-band data.

Implementers should take note of an additional class of text terminals not considered in the events below. These terminals use dual tone multi-frequency (DTMF) tones to encode and exchange signals. This application is described in RFC 4733 [5], Section 3.1, in conjunction with the registration of DTMF events.

The events shown in Table 7 correspond to signals coming from the following modem types:

- o Baudot [34], a five bit character encoding nominally operating at 45.45 or 50 bits/s with frequencies 1800 Hz = '0', 1400 Hz = '1';
- o EDT, which is V.21 [12] operating at 110 bits/s in half-duplex mode (lower channel only); characters are 7-bit IA5 plus initial start bit, trailing parity bit, and two stop bits;
- o Bell 103 mode (documented in Recommendation V.18 Annex D), which is structurally similar to V.21, but uses different frequencies: lower channel, 1070 Hz = '0', 1270 Hz = '1'; upper channel, 2025 Hz = '0', 2225 Hz = '1'; characters are US ASCII framed by one start bit, one trailing parity bit, and one stop bit;
- o V.23 [25] based videotex, in Minitel and Prestel versions. V.23 offers a forward channel operating at 1200 bits/s if possible (2100 Hz = '0', 1300 Hz = '1') or otherwise at 600 bits/s (1700 Hz = '0', 1300 Hz = '1'), and a 75 bits/s backward channel, which is transmitting 390 Hz (continuous '1's) except when '0' is to be transmitted (450 Hz);
- o a non-V.18 text terminal using V.21 [12] at 300 bits/s. Characters are 7-bit national (e.g., US ASCII) with a start bit, parity, and one stop bit.

Event	Bit Rate bits/s	Frequency (Hz)	Event Code	Type	Volume?
ANS2225	N/A	2225	52	tone	yes
V21L110	110	980/1180	55	other	no
V21L300	300	980/1180	30	other	no
V21H300	300	1650/1850	31	other	no
B103L300	300	1070/1270	56	other	no
V23Main	600/1200	1700-2100/1300	57	other	no
V23Back	75	450/390	58	other	no
Baud4545	45.45	1800/1400	59	other	no
Baud50	50	1800/1400	60	other	no
XCIMark	1200	2100/1300	62	tone	yes

Table 7: Indicators for Text Telephony

ANS2225:

indicates that a 2225-Hz answer tone has been detected. This is a pure tone with no amplitude modulation and no semantics attached to phase reversals, if there are any. The sender SHOULD report the beginning of the event when the tone is detected. The sender MAY send updates as the tone continues, and MUST report the end of the event when the tone ceases. The tone concerned is generated by a Bell 103-type modem in answer mode. This event MUST NOT be reported outside of the startup context (i.e., on the answering side at the beginning of a call).

V21L110:

indicates that the sender has detected V.21 modulation operating in the lower channel at 110 bits/s. Note that it may take some time to distinguish between 300 bits/s and 110 bits/s operation. It is expected that implementations will not transmit both this event and individual V.21 bit events for the same content.

V21L300:

indicates that the sender has detected V.21 modulation operating in the lower channel at 300 bits/s. Note that it may take some time to distinguish between 300 bits/s and 110 bits/s operation. It is expected that implementations will not transmit both this event and individual V.21 bit events for the same content.

V21H300:

indicates that the sender has detected V.21 modulation operating in the upper channel at 300 bits/s. It is expected that implementations will not transmit both this event and individual V.21 bit events for the same content.

B103L300:

indicates that the sending device has detected Bell 103 class modulation operating in the low channel at 300 bits/s.

V23Main:

indicates that the sending device has detected V.23 modulation operating in the high-speed channel. As described below, this indicator may alternate with the XCIMark indication.

V23Back:

indicates that the sending device has detected V.23 modulation operating in the 75 bit/s back-channel.

Baud4545:

indicates that the sending device has detected Baudot modulation operating at 45.45 bits/s.

Baud50:

indicates that the sending device has detected Baudot modulation operating at 50 bits/s.

XCIMark:

Indicates that the sending device has detected the specific bit pattern (0) 1111 1111(1)(0)1111 1111(1) sent at 1200 bits/s using V.23 upper-channel modulation, following a period of V.23 main channel "mark" (1300 Hz).

It is assumed in all cases that the event reports described here are being transmitted in addition to another media encoding, typically G.711 [19] voice-band data, reporting the same information. A natural method to do this is to combine the voice-band data with event reports in an RFC 2198 [2] redundancy payload.

The handling of ANS2225 has been indicated above. Since it is a specific tone, it can be handled like any other tone event.

For all of the other indicators, the sender SHOULD generate an initial event report as soon as the nature of the audio content has been recognized. For reliability, the initial event report SHOULD be retransmitted twice at short intervals. (20 ms is a suggested value, although the packetization period of the associated media may be sufficient.) The sender MAY continue to send additional reports of the same indicator event, although these have little value once the receiver has adjusted itself to the type of content it is receiving.

If the nature of the content changes (e.g., because it is coming from a V.18 terminal in the probing stage), the sender MUST send an event report for the new content type as soon as it is recognized. If the sender has been sending updates for the previous indicator, it SHOULD report the end of that previous indicator event along with the beginning of the new one.

2.7.1.1. Handling of Congestion

In the face of packet loss or delay, it is appropriate for the receiver to continue to play out the ANS2225 event until further packets are received. For the other events, the issue is loss of the initial event report rather than maintenance of playout continuity. The advice on retransmission of these other events already given above is sufficient to deal with packet loss or delay due to congestion.

2.7.2. Use of Events with V.18 Modems

ITU-T Recommendation V.18 [11] defines a terminal for text conversation, possibly in combination with voice. V.18 is intended to interoperate with a variety of legacy text terminals, so its start-up sequence can consist of a series of stimuli designed to determine what is at the other end. Two V.18 terminals talking to each other will use V.8 to negotiate startup and continue at the physical level with V.21 at 300 bits/s carrying 7-bit characters bounded by start and stop bits.

The V.18 terminal is also designed to interoperate with the text modems listed in the previous sub-section. The startup sequences for all these different terminal types are naturally quite different. The V.18 initial startup sequence specifically addresses itself to V.8-capable terminals and V.21 terminals and, by the combination of signals, to V.23 videotex terminals. During the initial startup sequence, the V.18 terminal listens for frequency responses characterizing the other terminal types. If it does not make contact in the preliminary step, it probes for each type specifically. By the nature of the application, V.18 has been designed to provide an extremely robust startup capability.

The handling of the V.18 XCI signal is a specific case of the procedures described in the previous section. XCI is a signal transmitted in high-band V.23 modulation to stimulate V.23 terminals to respond and to allow detection of V.18 capabilities in a DCE. The 3-second XCI signal uses the V.23 upper channel having periods of "mark" (i.e., 1300 Hz) alternating with the XCIMark pattern. The full definition is found in V.18, Section 3.13. The sender SHOULD indicate V23Main during the transmission of the "mark" portion of XCI, and change the indication to XCIMark when that pattern is detected.

2.8. A Generic Indicator

Numerous proprietary modem protocols exist, as well as standardized protocols not identified above. Table 8 defines a single indicator event that may be used to identify modem content when a more specific event is unavailable. Typically, this would be sent in combination with another payload type, for example, voice-band data as specified by ITU-T Recommendation V.152 [33].

As with the indicators in the previous section, the sender SHOULD generate an initial event report as soon as the nature of the audio content has been recognized. For reliability, the initial event report SHOULD be retransmitted twice at short intervals. (20 ms is a suggested value, although the packetization period of the associated media may be sufficient.) The sender MAY continue to send additional reports of the VBDGen event, although these have little value once the receiver has adjusted itself to the type of content it is receiving.

Event	Bit Rate bits/s	Frequency (Hz)	Event Code	Type	Volume?
VBDGen	Variable	Variable	61	other	no

Table 8: Generic Modem Signal Indicator

VBDGen:

indicates that the sender has detected tone patterns indicating the operation of some form of modem. This indicator SHOULD NOT be sent if a more specific event is available.

3. Strategies for Handling Fax and Modem Signals

As described in Section 1.2, the typical data application involves a pair of gateways interposed between two terminals, where the terminals are in the PSTN. The gateways are likely to be serving a mixture of voice and data traffic, and need to adopt payload types appropriate to the media flows as they occur. If voice compression is in use for voice calls, this means that the gateways need the flexibility to switch to other payload types when data streams are recognized.

Within the established IETF framework, this implies that the gateways must negotiate the potential payloads (voice, telephone-event, tones, voice-band data, T.38 fax [21], and possibly RFC 4103 [18] text and Clearmode [17] octet streams) as separate payload types. From a timing point of view, this is most easily done at the beginning of a call, but results in an over-allocation of resources at the gateways and in the intervening network.

One alternative is to use named events to buy time while out-of-band signals are exchanged to update to the new payload type applicable to the session. Thanks to the events defined in this document, this is a viable approach for sessions beginning with V.8, V.8 bis, T.30, or V.25 control sequences.

Named data-related events also allow gateways to optimize their operation when data signals are received in a relatively general form. One example is the use of V.8-related events to deduce that the voice-band data being sent in a G.711 payload comes from a higher-speed modem and therefore requires disabling of echo cancellers.

All of the control procedures described in the sub-sections of Section 2 eventually give way to data content. As mentioned above, this content will be carried by other payload types. Receiving gateways MUST be prepared to switch to the other payload type within the time constraints associated with the respective applications. (For several of the procedures documented above, the sender provides 75 ms of silence between the initial control signalling and the sending of data content.) In some cases (V.8 bis [10], T.30 [8]), further control signalling may happen after the call has been established.

A possible strategy is to send both the telephone-event and the data payload in an RFC 2198 [2] redundancy arrangement. The receiving gateway then propagates the data payload whenever no event is in progress. For this to work, the data payload and events (when present) MUST cover exactly the same content over the same time period; otherwise, spurious events will be detected downstream. An example of this mode of operation is shown below.

Note that there are a number of cases where no control sequence will precede the data content. This is true, for example, for a number of legacy text terminal types. In such instances, the events defined in Section 2.7 in particular MAY be sent to help the remote gateway optimize its handling of the alternative payload.

4. Example of V.8 Negotiation

This section presents an example of the use of the event codes defined in Section 2. The basic scenario is the startup sequence for duplex V.34 modem operation. It is assumed that once the initial V.8 sequence is complete, the gateways will enter into voice-band data operation using G.711 encoding to transmit the modem signals. The basic packet sequence is indicated in Table 9. Sample packets are then shown in detail for two variants on event transmission strategy:

- o simultaneous transmission of events and retransmitted events using RFC 2198 [2] redundancy;
- o simultaneous transmission of events, retransmitted events, and voice-band data covering the same content using RFC 2198 redundancy.

For simplicity and semi-realism, the times shown for the example scenario assume a fixed lag at each gateway of 20 ms between the packet side of the gateway and the local user equipment and vice versa (i.e., minimum of 40 ms between packet received and packet sent specifically in response to the received packet). A propagation delay of 5 ms is assumed between gateways. It is assumed that the

event packetization interval is 30 ms, a reasonable compromise between packet volume and buffering delay, particularly for V.21 events.

At the basic V.8 protocol level, the table assumes that the answering modem waits 0.2 s (200 ms) from the beginning of the call to start transmitting ANSam. The calling modem waits 1 s (1000 ms) from the time it begins to receive ANSam until it begins to send the V.8 CM signal. Both modems wait 75 ms from the time they finish sending and receiving CJ, respectively, until they begin sending V.34 modem signals.

Time (ms)	Event
220.0	The called gateway detects the start of ANSam from its end.
250.0	The called gateway sends out the first ANSam event packet. M bit is set, timestamp is $ts_0 + 1760$ (where ts_0 is the timestamp value at the start of the call). The initial ANSam event continues until a phase shift is detected at 670.0 ms (see below). Up to this time, the called gateway sends out further ANSam event updates, with the same initial timestamp, M bit off, and cumulative duration increasing by 240 units each time.
255.0	The calling gateway receives the first ANSam event report and begins playout of ANSam tone at its end.
275.0	The calling terminal receives the beginning of ANSam tone and starts its timer. It will begin sending the CM signal 1 s later (at 1275.0 ms into the call).
670.0	The called gateway detects a phase shift in the incoming signal, marking a change from ANSam to /ANSam. This happens to coincide with the end of a packetization interval. For the sake of the example, assume that the called gateway does not detect this in time for the event report it sends out.

- 700.0 | The called gateway issues its next-scheduled event report packet, indicating an initial report for /ANSam (M bit set, timestamp $ts_0 + 5360$, duration 240 timestamp units). The packet also carries the first retransmission of the final ANSam report, total duration 3600 units, this time with the E bit set.
- 1295.0 | The calling gateway begins to receive the CM signal from the calling modem.
- 1325.0 | The calling gateway sends a packet containing the first 9 bits of the CM signal.
- 1445.0 | The calling gateway sends out a packet containing the last 4 bits of the first CM signal, plus the first 5 bits of the next repetition of that signal. CM bits will continue to be transmitted from the calling gateway until 2015.0 ms (see below), for a total of 24 packets. (The final packet also carries the beginning of the CJ signal.)
- 1596.7 | The called gateway completes playout of the final bit of the second occurrence of the CM signal.
- 1636.7 | The called gateway detects end of /ANSam (and beginning of JM) from the called modem. The next packet is not yet due to go out.
- 1660.0 | The called gateway sends out a packet combining the final /ANSam event report (E bit set and total duration 533 timestamp units) with the first 7 bits of the JM signal. The M bit for the packet is set and the packet timestamp is $ts_0 + 12560$ (the start of the now-discontinued /ANSam event).
- 1690.0 | The called gateway sends out a packet containing the next nine bits of JM signal. The M bit is set and the timestamp is $ts_0 + 13280$ (beginning of the first bit in the packet). JM will continue to be transmitted until 2170.0 ms (see below), for a total of 18 packets (plus two for final retransmissions).
- 1938.3 | The calling gateway completes playout of the final packet of the second occurrence of the JM signal.
- 1995.0 | The calling gateway begins to receive the initial bits of the CJ signal.

- 2015.0 The calling gateway sends a packet containing the final 3 bits of the first decad of a CM signal and first 6 bits of a CJ signal.
- 2095.0 The calling gateway receives the last bit of the CJ signal. A period of silence lasting 75-ms begins at the called end. It is not yet time to send out an event report.
- 2105.0 The calling gateway sends out a packet containing the final 6 bits of the CJ signal.
- 2130.0 The called gateway finishes playing out the last CJ signal bit sent to it.
- 2135.0 The calling gateway sends a packet containing no new events, but retransmissions of the last 15 bits of the CJ signal (in two generations).
- 2165.0 The calling gateway sends out a packet containing no new events, but retransmissions of the final 6 bits of the CJ signal.
- 2170.0 The called gateway sends out the last packet containing bits of the JM signal (except for retransmissions). Note that according to the V.8 specification these bits do not in general complete a JM signal or even an "octet" of that signal (although they happen to do so in this example). A 75 ms period of silence begins at the called end.
- 2170.0 The calling gateway begins to receive V.34 signalling from the called modem.
- 2175.0 The calling gateway finishes playing out the last JM signal bit sent to it.
- 2195.0 The calling gateway sends out a first packet of V.34 signalling as voice-band data (PCMU). Timestamp is $ts_0 + 17360$ and M bit is set to indicate the beginning of content after silence. The packet contains 200 8-bit samples. Packetization interval is shown here as continuing to be 30 ms. It could be less, but MUST NOT be more because that would make the silent period too long.

2200.0	The called gateway sends a packet containing no new events, but retransmissions of the last 18 bits of the JM signal (in two generations).
2225.0	The calling gateway sends out the second packet of V.34 signalling as voice-band data (PCMU). Timestamp is $ts_0 + 17560$ and M bit is not set. The packet contains 240 8-bit samples.
2230.0	The called gateway sends out a packet containing no new events, but retransmissions of the final 9 bits of the JM signal.
2245.0	The called gateway begins to receive V.34 signalling from the called modem.
2255.0	The calling gateway sends out a third packet of V.34 signalling as voice-band data (PCMU). Timestamp is $ts_0 + 17800$ and M bit is not set. The packet contains 240 8-bit samples.
2260.0	The called gateway sends out a first packet of V.34 signalling as voice-band data (PCMU). Timestamp is $ts_0 + 17960$ and M bit is set to indicate the beginning of content after silence. The packet contains 120 samples. Packetization interval is shown here as continuing to be 30 ms. It could be less, but MUST NOT be more because that would make the silent period too long.
...	...

Table 9: Events for Example V.8 Scenario

4.1. Simultaneous Transmission of Events and Retransmitted Events Using RFC 2198 Redundancy

Negotiation of the transmission mode being described in this section would use SDP similar to the following:

```
m=audio 12343 RTP/AVP 99
a=rtpmap:99 pcmu/8000
m=audio 12345 RTP/AVP 100 101
a=rtpmap:100 red/8000/1
a=fmtp:100 101/101/101
a=rtpmap:101 telephone-event/8000
a=fmtp:101 0-15,32-41,43,46,48-49,52-68
```

This indicates two media streams, the first for G.711 (i.e., voice or voice-band data), the second for triply-redundant telephone events. As RFC 2198 notes, it is also possible for the sender to send telephone-event payloads without redundancy in the second stream, although the redundant form is the primary transmission mode. (It would be reasonable to send the interim ANSam reports without redundancy.) The set of telephone events supported includes the DTMF events (not relevant in this example), and all of the data events defined in this document. In fact, only event codes 34-35 and 37-40 are used in the example.

For the purpose of illustrating the use of RFC 2198 redundancy as well as showing the basic composition of the event reports, the second packet reporting JM signal bits (sent by the called gateway at 1690.0 ms) seems to be a good choice. This packet will also carry the second retransmission of the final /ANSam event report and the first retransmission of the initial 7 bits of the JM signal. The detailed content of the packet is shown in Figure 1. To see the contents of the successive generations more clearly, they are presented as if they were aligned on successive 32-bit boundaries. In fact, they are all offset by one octet, following on consecutively from the RFC 2198 header.

The M bit is set in the RTP header for the packet, as required for the coding of multiple events in the primary block of data. In fact, RFC 2198 implies that this is the correct behavior, but does not say so explicitly. The E bit is set for every event. It is possible that it would not be set for the final event in the primary block.

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3  4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|V=2|P|X| CC=0  |1|  PT=100      |   sequence number = seq0 + 48 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|
|           timestamp = ts0 + 13280
|-----+-----+-----+-----+-----+-----+-----+-----+
|           synchronization source (SSRC) identifier
+-----+-----+-----+-----+-----+-----+-----+-----+
|1| block PT=101| timestamp offset = 720 | block length = 4 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|1| block PT=101| timestamp offset = 267 | block length = 28 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|0| block PT=101|           (begin block for /ANSam ...)
+-----+-----+-----+-----+-----+-----+-----+-----+

           /ANSam block (second retransmission)

+-----+-----+-----+-----+-----+-----+-----+-----+
|   event = 35  |1|R| volume      |           duration = 533      |
+-----+-----+-----+-----+-----+-----+-----+-----+

           First 7 bits of JM ("1111111" in V.21 high channel)
           (first retransmission)

+-----+-----+-----+-----+-----+-----+-----+-----+
|   event = 40  |1|R| volume      |           duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+
/   (5 similar events, durations 27,26,27,27,26 respectively)   /
+-----+-----+-----+-----+-----+-----+-----+-----+
|   event = 40  |1|R| volume      |           duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+

           Next 9 bits of JM ("111000000" in V.21 high channel)
           (new content)

+-----+-----+-----+-----+-----+-----+-----+-----+
|   event = 40  |1|R| volume      |           duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+
/   (7 similar events, codes 40,40,39,39,39,39,39 and           /
/   durations 26,27,27,26,27,27,26 respectively)               /
+-----+-----+-----+-----+-----+-----+-----+-----+
|   event = 39  |1|R| volume      |           duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 1: Packet Contents, Redundant Events Only

Since all of the events in the above packet are consecutive and adjacent, it would have been permissible according to the telephone-event payload specification to carry them as a simple event payload without the RFC 2198 header. The advantage of the latter is that the receiving gateway can skip over the retransmitted events when processing the packet, unless it needs them.

4.2. Simultaneous Transmission of Events and Voice-Band Data Using RFC 2198 Redundancy

Negotiation of the transmission mode being described in this section would use SDP similar to the following:

```
m=audio 12343 RTP/AVP 99 100 101
a=rtpmap:99 red/8000/1
a=fmtp:99 100/101/101/101
a=rtpmap:100 pcmu/8000
a=rtpmap:101 telephone-event/8000
a=fmtp:101 0-15,32-41,43,46,48-49,52-68
```

This indicates one media stream, with G.711 (i.e., voice or voice-band data) as the primary content, along with three blocks of telephone events. RFC 2198 requires that the more voluminous representation (i.e., the G.711) be the primary one. The most recent block of events covers the same time period as the voice-band data. The other two streams provide the first and second retransmissions of the events as in the previous example. Because G.711 is the primary content, the M bit for the packets will in general not be set, except after periods of silence.

Figure 2 shows the detailed packet content for the same sample point as in the previous figure, but including the G.711 content.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3  4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|V=2|P|X| CC=0  |0|  PT=99      |      sequence number = seq0 + 48 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|
|          timestamp = ts0 + 13280
|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|          synchronization source (SSRC) identifier
|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|1| block PT=101| timestamp offset = 720 | block length = 4 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|1| block PT=101| timestamp offset = 267 | block length = 28 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|1| block PT=101| timestamp offset = 0   | block length = 36 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|0| block PT=100|      (begin block for /ANSam ...)
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

/ANSam block (second retransmission)

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      event = 35 |1|R| volume      |      duration = 533      |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

First 7 bits of JM ("1111111" in V.21 high channel) (first retransmission)

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      event = 40 |1|R| volume      |      duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/      (5 similar events, durations 27,26,27,27,26 respectively) /
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      event = 40 |1|R| volume      |      duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Next 9 bits of JM ("111000000" in V.21 high channel) (new content)

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      event = 40 |1|R| volume      |      duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/      (7 similar events, codes 40,40,39,39,39,39,39 and /
/      durations 26,27,27,26,27,27,26 respectively) /
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      event = 39 |1|R| volume      |      duration = 27      |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

30 ms of G.711-encoded voice-band data (240 samples)

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Sample 1   | Sample 2   | Sample 3   | Sample 4   |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/                                     . . .                                     /
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Sample 237 | Sample 238 | Sample 239 | Sample 240 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: Packet Contents with Voice-Band Data Combined with Events

5. Security Considerations

The V.21 bit events defined in this document may be used to transmit user-sensitive data. This could include initial log-on sequences and application-level protocol exchanges as well as user content. As a result, such a usage of V.21 bit events entails, in the terminology of [16], threats to both communications and system security. The attacks of concern are:

- o confidentiality violations and password sniffing;
- o hijacking of data sessions through message insertion;
- o modification of the transmitted content through man-in-the-middle attacks;
- o denial of service by means of message insertion, deletion, and modification aimed at interference with the application protocol.

To prevent these attacks, the transmission of V.21 bit events MUST be given confidentiality protection. Message authentication and the protection of message integrity MUST also be provided. These address the threats posed by message insertion and modification. With these measures in place, RTP sequence numbers and the redundancy provided by the RFC 4733 procedures for transmission of events add protection against and some resiliency in the face of message deletion.

The other events defined in this document (and V.21 bit events within control sequences) are used only for the setup and control of sessions between data terminals or fax devices. While disclosure of these events would not expose user-sensitive data, it can potentially expose capabilities of the user equipment that could be exploited by attacks in the PSTN domain. Thus, confidentiality protection SHOULD be provided. The primary threat is denial of service, through injection of inappropriate signals at vulnerable points in the control sequence or through alteration or blocking of enough event

packets to disrupt that sequence. To meet the injection threat, message authentication and integrity protection MUST be provided.

The Secure Real-time Transport Protocol (SRTP) [3] meets the requirements for protection of confidentiality, message integrity, and message authentication described above. It SHOULD therefore be used to protect media streams containing the events described in this document.

Note that the appropriate method of key distribution for SRTP may vary with the specific application.

In some deployments, it may be preferable to use other means to provide protection equivalent to that provided by SRTP.

6. IANA Considerations

This document adds the events in Table 10 to the registry established by RFC 4733 [5].

Event Code	Event Name	Reference
23	CRdSeg: second segment of V.8 bis CRd signal	RFC 4734
24	CReSeg: second segment of V.8 bis CRe signal	RFC 4734
25	MRdSeg: second segment of V.8 bis MRd signal	RFC 4734
26	MReSeg: second segment of V.8 bis MRe signal	RFC 4734
27	V32AC: A pattern of bits modulated at 4800 bits/s, emitted by a V.32/V.32bis answering terminal upon detection of the AA pattern.	RFC 4734
28	V8bISeg: first segment of initiating V.8 bis signal	RFC 4734
29	V8bRSeg: first segment of responding V.8 bis signal	RFC 4734

30	V21L300: 300 bits/s low channel V.21 indication	RFC 4734
31	V21H300: 300 bits/s high channel V.21 indication	RFC 4734
32	ANS (V.25 Answer tone). Also known as CED (T.30 Called tone).	RFC 4734
33	/ANS (V.25 Answer tone after phase shift). Also known as /CED (T.30 Called tone after phase shift)	RFC 4734
34	ANSam (V.8 amplitude modified Answer tone)	RFC 4734
35	/ANSam (V.8 amplitude modified Answer tone after phase shift)	RFC 4734
36	CNG (T.30 Calling tone)	RFC 4734
37	V.21 channel 1 (low channel), '0' bit	RFC 4734
38	V.21 channel 1, '1' bit. Also used for ESiSeg (second segment of V.8 bis ESi signal).	RFC 4734
39	V.21 channel 2, '0' bit	RFC 4734
40	V.21 channel 2, '1' bit. Also used for ESrSeg (second segment of V.8 bis ESr signal).	RFC 4734
49	CT (V.25 Calling Tone)	RFC 4734
52	ANS2225: 2225-Hz indication for text telephony	RFC 4734
53	CI (V.8 Call Indicator signal preamble)	RFC 4734
54	V.21 preamble flag (T.30)	RFC 4734
55	V21L110: 110 bits/s V.21 indication for text telephony	RFC 4734
56	B103L300: Bell 103 low channel indication for text telephony	RFC 4734

57	V23Main: V.23 main channel indication for text telephony	RFC 4734
58	V23Back: V.23 back channel indication for text telephony	RFC 4734
59	Baud4545: 45.45 bits/s Baudot indication for text telephony	RFC 4734
60	Baud50: 50 bits/s Baudot indication for text telephony	RFC 4734
61	VBDGen: Tone patterns indicative of use of an unidentified modem type	RFC 4734
62	XCIMark: A pattern of bits modulated in the V.23 main channel, emitted by a V.18 calling terminal.	RFC 4734
63	V32AA: A pattern of bits modulated at 4800 bits/s, emitted by a V.32/V.23bis calling terminal.	RFC 4734

Table 10: Data-Related Additions to RFC 4733 Telephony Event Registry

7. Acknowledgements

Scott Petrack was the original author of RFC 2833. Henning Schulzrinne later loaned his expertise to complete the document, but Scott must be credited with the energy behind the idea of a compact encoding of tones over IP.

Gunnar Hellstrom and Keith Chu provided particularly useful comments helping to shape the present document. Amiram Allouche and Ido Benda drew the authors' attention to the value of including events for V.32/V.32bis in the document, and Yaakov Stein confirmed details of operation of this modem.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Perkins, C., Kouvelas, I., Hodson, O., Hardman, V., Handley, M., Bolot, J., Vega-Garcia, A., and S. Fosse-Parisis, "RTP Payload for Redundant Audio Data", RFC 2198, September 1997.
- [3] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [4] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [5] Schulzrinne, H. and T. Taylor, "RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals", RFC 4733, December 2006.
- [6] International Telecommunication Union, "Echo suppressors", ITU-T Recommendation G.164, November 1988.
- [7] International Telecommunication Union, "Echo cancellers", ITU-T Recommendation G.165, March 1993.
- [8] International Telecommunication Union, "Procedures for document facsimile transmission in the general switched telephone network", ITU-T Recommendation T.30, July 2003.
- [9] International Telecommunication Union, "Procedures for starting sessions of data transmission over the public switched telephone network", ITU-T Recommendation V.8, November 2000.
- [10] International Telecommunication Union, "Procedures for the identification and selection of common modes of operation between data circuit-terminating equipments (DCEs) and between data terminal equipments (DTEs) over the public switched telephone network and on leased point-to-point telephone-type circuits", ITU-T Recommendation V.8 bis, November 2000.
- [11] International Telecommunication Union, "Operational and interworking requirements for {DCEs operating in the text telephone mode", ITU-T Recommendation V.18, November 2000.

See also Recommendation V.18 Amendment 1, Nov. 2002.

[12] International Telecommunication Union, "300 bits per second duplex modem standardized for use in the general switched telephone network", ITU-T Recommendation V.21, November 1988.

[13] International Telecommunication Union, "Automatic answering equipment and general procedures for automatic calling equipment on the general switched telephone network including procedures for disabling of echo control devices for both manually and automatically established calls", ITU-T Recommendation V.25, October 1996.

See also Corrigendum 1 to Recommendation V.25, Jul. 2001.

[14] International Telecommunication Union, "A family of 2-wire, duplex modems operating at data signalling rates of up to 9600 bit/s for use on the general switched telephone network and on leased telephone-type circuits", ITU-T Recommendation V.32, March 1993.

[15] International Telecommunication Union, "A duplex modem operating at data signalling rates of up to 14 400 bit/s for use on the general switched telephone network and on leased point-to-point 2-wire telephone-type circuits", ITU-T Recommendation V.32bis, February 1991.

8.2. Informative References

[16] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[17] Kreuter, R., "RTP Payload Format for a 64 kbit/s Transparent Call", RFC 4040, April 2005.

[18] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", RFC 4103, June 2005.

[19] International Telecommunication Union, "Pulse code modulation (PCM) of voice frequencies", ITU-T Recommendation G.711, November 1988.

[20] International Telecommunication Union, "Terminal for low bit-rate multimedia communication", ITU-T Recommendation H.324, March 2002.

[21] International Telecommunication Union, "Procedures for real-time Group 3 facsimile communication over IP networks", ITU-T Recommendation T.38, July 2003.

- [22] International Telecommunication Union, "International interworking for videotex services", ITU-T Recommendation T.101, November 1994.
- [23] International Telecommunication Union, "Data protocols for multimedia conferencing", ITU-T Recommendation T.120, July 1996.
- [24] International Telecommunication Union, "A 2-wire modem for facsimile applications with rates up to 14 400 bit/s", ITU-T Recommendation V.17, February 1991.
- [25] International Telecommunication Union, "600/1200-baud modem standardized for use in the general switched telephone network", ITU-T Recommendation V.23, November 1988.
- [26] International Telecommunication Union, "4800/2400 bits per second modem standardized for use in the general switched telephone network", ITU-T Recommendation V.27ter, November 1988.
- [27] International Telecommunication Union, "9600 bits per second modem standardized for use on point-to-point 4-wire leased telephone-type circuits", ITU-T Recommendation V.29, November 1988.
- [28] International Telecommunication Union, "A modem operating at data signalling rates of up to 33 600 bit/s for use on the general switched telephone network and on leased point-to-point 2-wire telephone-type circuits", ITU-T Recommendation V.34, February 1998.
- [29] International Telecommunication Union, "A digital modem and analogue modem pair for use on the Public Switched Telephone Network (PSTN) at data signalling rates of up to 56 000 bit/s downstream and up to 33 600 bit/s upstream", ITU-T Recommendation V.90, September 1998.
- [30] International Telecommunication Union, "A digital modem operating at data signalling rates of up to 64 000 bit/s for use on a 4-wire circuit switched connection and on leased point-to-point 4-wire digital circuits", ITU-T Recommendation V.91, May 1999.
- [31] International Telecommunication Union, "Enhancements to Recommendation V.90", ITU-T Recommendation V.92, November 2000.

- [32] International Telecommunication Union, "Modem-over-IP networks: Procedures for the end-to-end connection of V-series DCEs", ITU-T Recommendation V.150.1, January 2003.
- [33] International Telecommunication Union, "Procedures for supporting voice-band data over IP networks", ITU-T Recommendation V.152, January 2005.
- [34] Telecommunications Industry Association, "A Frequency Shift Keyed Modem for Use on the Public Switched Telephone Network", ANSI TIA- 825-A-2003, April 2003.

Authors' Addresses

Henning Schulzrinne
Columbia U.
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
US

EMail: schulzrinne@cs.columbia.edu

Tom Taylor
Nortel
1852 Lorraine Ave
Ottawa, Ontario K1H 6Z8
Canada

EMail: taylor@nortel.com

Full Copyright Statement

Copyright (C) The IETF Trust (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, THE IETF TRUST, 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 currently provided by the Internet Society.

