

AVT
Internet-Draft
Expires: June 20, 2005

R. Even
Polycom
December 20, 2004

RTP Payload Format for H.261 Video Streams
draft-ietf-avt-rfc2032-bis-04.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on June 20, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

This memo describes a scheme to packetize an H.261 video stream for transport using the Real-time Transport Protocol, RTP, with any of the underlying protocols that carry RTP.

This specification is a product of the Audio/Video Transport working group within the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list and/or the authors.

Even

Expires June 20, 2005

[Page 1]

Table of Contents

1.	Introduction	3
2.	Terminology	4
3.	Structure of the packet stream	5
3.1	Overview of the ITU-T recommendation H.261	5
3.2	Considerations for packetization	5
4.	Specification of the packetization scheme	7
4.1	Usage of RTP	7
4.2	Recommendations for operation with hardware codecs	9
5.	Packet loss issues	11
6.	Payload Format Parameters	13
6.1	IANA Considerations	13
6.1.1	Registration of MIME media type video/H261	13
6.2	SDP Parameters	14
6.2.1	Usage with the SDP Offer Answer Model	14
7.	Backward Compatibility to RFC2032	16
7.1	Optional H.261-specific control packets	16
7.2	New SDP optional parameters	16
8.	Security Considerations	17
9.	Acknowledgements	18
10.	Changes from previous versions	19
10.1	changes from RFC 2032 >	19
10.2	Changes from bis-00	19
11.	References	20
11.1	Normative References	20
11.2	Informative References	20
	Author's Address	21
	Intellectual Property and Copyright Statements	22

Even

Expires June 20, 2005

[Page 2]

1. Introduction

The ITU-T recommendation H.261[H261] specifies the encoding used by ITU-T compliant video-conference codecs. Although these encoding were originally specified for fixed data rate ISDN circuits, experiments [[INRIA](#)], [[MICE](#)] have shown that they can also be used over packet-switched networks such as the Internet.

The purpose of this memo is to specify the RTP payload format for encapsulating H.261 video streams in RTP[RFC3550].

This document updates [RFC2032](#). The intention is to move the new RFC to draft standard.

Even

Expires June 20, 2005

[Page 3]

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119](#)[\[RFC2119\]](#) and indicate requirement levels for compliant RTP implementations.

3. Structure of the packet stream

3.1 Overview of the ITU-T recommendation H.261

The H.261 coding is organized as a hierarchy of groupings. The video stream is composed of a sequence of images, or frames, which are themselves organized as a set of Groups of Blocks (GOB). Note that H.261 "pictures" are referred as "frames" in this document. Each GOB holds a set of 3 lines of 11 macro blocks (MB). each MB of 16x16 pixels the luminance information is encoded for each pixel, sub divided into four 8x8 groups. While the chrominance information is subsampled, for 2x2 pixels, resulting in color resolution of 8x8 within the MB. These components and the codes representing their sampled values are as defined in the ITU-R Recommendation 601 [BT601].

This grouping is used to specify information at each level of the hierarchy:

- At the frame level, one specifies information such as the delay from the previous frame, the image format, and various indicators.
- At the GOB level, one specifies the GOB number and the default quantifier that will be used for the MBs.
- At the MB level, one specifies which blocks are present and which did not change, and optionally a quantifier and motion vectors.

Blocks which have changed are encoded by computing the discrete cosine transform (DCT) of their coefficients, which are then quantized and Huffman encoded (Variable Length Codes).

The H.261 Huffman encoding includes a special "GOB start" pattern, which is a word of 16 bits, 0000 0000 0000 0001. This pattern is included at the beginning of each GOB header (and also at the beginning of each frame header) to mark the separation between two GOBs, and is in fact used as an indicator that the current GOB is terminated. The encoding also includes a stuffing pattern, composed of seven zero bits followed by four bits with a value of one; that stuffing pattern can only be entered between the encoding of MBs, or just before the GOB separator.

3.2 Considerations for packetization

H.261 codecs designed for operation over ISDN circuits produce a bit stream composed of several levels of encoding specified by H.261 and companion recommendations. The bits resulting from the Huffman encoding are arranged in 512-bit frames, containing 2 bits of

Even

Expires June 20, 2005

[Page 5]

synchronization, 492 bits of data and 18 bits of error correcting code. The 512-bit frames are then interlaced with an audio stream and transmitted over px64 kbps circuits according to specification H.221 [[H221](#)].

When transmitting over the Internet, we will directly consider the output of the Huffman encoding. All the bits produced by the Huffman encoding stage will be included in the packet. We will not carry the 512-bit frames, as protection against bit errors can be obtained by other means. Similarly, we will not attempt to multiplex audio and video signals in the same packets, as UDP and RTP provide a much more efficient way to achieve multiplexing.

Directly transmitting the result of the Huffman encoding over an unreliable stream of UDP datagrams would, however, have poor error resistance characteristics. The result of the hierarchical structure of H.261 bit stream is that one needs to receive the information present in the frame header to decode the GOBs, as well as the information present in the GOB header to decode the MBs. Without precautions, this would mean that one has to receive all the packets that carry an image in order to properly decode its components.

If each image could be carried in a single packet, this requirement would not create a problem. However, a video image or even one GOB by itself can sometimes be too large to fit in a single packet. Therefore, the MB is taken as the unit of fragmentation. Packets must start and end on a MB boundary, i.e. a MB cannot be split across multiple packets. Multiple MBs may be carried in a single packet when they will fit within the maximal packet size allowed. This practice is recommended to reduce the packet send rate and packet overhead.


To allow each packet to be processed independently for efficient resynchronization in the presence of packet losses, some state information from the frame header and GOB header is carried with each packet to allow the MBs in that packet to be decoded. This state information includes the GOB number in effect at the start of the packet, the macroblock address predictor (i.e. the last MBA encoded in the previous packet), the quantizer value in effect prior to the start of this packet (GQUANT, MQANT or zero in case of a beginning of GOB) and the reference motion vector data (MVD) for computing the true MVDs contained within this packet. The bit stream cannot be fragmented between a GOB header and MB 1 of that GOB.

Moreover, since the compressed MB may not fill an integer number of octets, the data header contains two three-bit integers, SBIT and EBIT, to indicate the number of unused bits in the first and last octets of the H.261 data, respectively.

Even

Expires June 20, 2005

[Page 6]



Even

Expires June 20, 2005

[Page 7]

```

|                                     H.261 stream ...
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The H.261 header is defined as following:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|SBIT |EBIT |I|V| GOBN |   MBAP   |  QUANT  |  HMVD   |  VMVD   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The fields in the H.261 header have the following meanings:

Start bit position (SBIT): 3 bits

Number of most significant bits that should be ignored in the first data octet.

End bit position (EBIT): 3 bits

Number of least significant bits that should be ignored in the last data octet.

INTRA-frame encoded data (I): 1 bit

Set to 1 if this stream contains only INTRA-frame coded blocks. Set to 0 if this stream may or may not contain INTRA-frame coded blocks. The meaning of this bit should not be changed during the course of the RTP session.

Motion Vector flag (V): 1 bit

Set to 0 if motion vectors are not used in this stream. Set to 1 if motion vectors may or may not be used in this stream. The meaning of this bit should not be changed during the course of the session.

GOB number (GOBN): 4 bits

Encodes the GOB number in effect at the start of the packet. Set to 0 if the packet begins with a GOB header.

Macroblock address predictor (MBAP): 5 bits

Encodes the macroblock address predictor (i.e. the last MBA encoded in the previous packet). This predictor ranges from 0-32 (to predict the valid MBAs 1-33), but because the bit stream cannot be fragmented between a GOB header and MB 1, the predictor at the start of the

Even

Expires June 20, 2005

[Page 8]

packet shall not be 0. Therefore, the range is 1-32, which is biased by -1 to fit in 5 bits. For example, if MBAP is 0, the value of the MBA predictor is 1. Set to 0 if the packet begins with a GOB header.

Quantizer (QUANT): 5 bits

Quantizer value (MQANT or GQUANT) in effect prior to the start of this packet. Set to 0 if the packet begins with a GOB header.

Horizontal motion vector data (HMVD): 5 bits

Reference horizontal motion vector data (MVD). Set to 0 if V flag is 0 or if the packet begins with a GOB header, or when the MTYPE of the last MB encoded in the previous packet was not MC. HMVD is encoded as a 2's complement number, and '10000' corresponding to the value -16 is forbidden (motion vector fields range from +/-15).

Vertical motion vector data (VMVD): 5 bits

Reference vertical motion vector data (MVD). Set to 0 if V flag is 0 or if the packet begins with a GOB header, or when the MTYPE of the last MB encoded in the previous packet was not MC. VMVD is encoded as a 2's complement number, and '10000' corresponding to the value -16 shall not be used (motion vector fields range from +/-15).

Note that the I and V flags are hint flags, i.e. they can be inferred from the bit stream. They are included to allow decoders to make optimizations that would not be possible if these hints were not provided before bit stream was decoded. Therefore, these bits cannot change for the duration of the stream. A conformant implementation can always set V=1 and I=0.

The H.261 stream SHALL be used without BCH error correction and without error correction framing.

4.2 Recommendations for operation with hardware codecs

Packetizers for hardware codecs can trivially figure out GOB boundaries using the GOB-start pattern included in the H.261 data. (Note that software encoders already know the boundaries.) The cheapest packetization implementation is to packetize at the GOB level all the GOBs that fit in a packet. But when a GOB is too large, the packetizer has to parse it to do MB fragmentation. (Note that only the Huffman encoding must be parsed and that it is not necessary to fully decompress the stream, so this requires relatively little processing; example implementations can be found in some public H.261 codecs such as IVS [[IVS](#)] and VIC [[VIC](#)].) It is recommended that MB level fragmentation be used when feasible in

Even

Expires June 20, 2005

[Page 9]

order to obtain more efficient packetization. Using this fragmentation scheme reduces the output packet rate and therefore reduces the overhead.

At the receiver, the data stream can be depacketized and directed to a hardware codec's input. If the hardware decoder operates at a fixed bit rate, synchronization may be maintained by inserting the stuffing pattern between MBs (i.e., between packets) when the packet arrival rate is slower than the bit rate.

5. Packet loss issues

On the Internet, most packet losses are due to network congestion rather than transmission errors. Using UDP, no mechanism is available at the sender to know if a packet has been successfully received. It is up to the application, i.e. coder and decoder, to handle the packet loss. Each RTP packet includes a sequence number field which can be used to detect packet loss.

H.261 uses the temporal redundancy of video to perform compression. This differential coding (or INTER-frame coding) is sensitive to packet loss. After a packet loss, parts of the image may remain corrupt until all corresponding MBs have been encoded in INTRA-frame mode (i.e. encoded independently of past frames). There are several ways to mitigate packet loss:

- (1) One way is to use only INTRA-frame encoding and MB level conditional replenishment. That is, only MBs that change (beyond some threshold) are transmitted.
- (2) Another way is to adjust the INTRA-frame encoding refreshment rate according to the packet loss observed by the receivers. The H.261 recommendation specifies that a MB is INTRA-frame encoded at least every 132 times it is transmitted. However, the INTRA-frame refreshment rate can be raised in order to speed the recovery when the measured loss rate is significant.
- (3) The fastest way to repair a corrupted image is to request an INTRA-frame coded image refreshment after a packet loss is detected. One means to accomplish this is for the decoder to send to the coder a list of packets lost. The coder can decide to encode every MB of every GOB of the following video frame in INTRA-frame mode (i.e. Full INTRA-frame encoded), or if the coder can deduce from the packet sequence numbers which MBs were affected by the loss, it can save bandwidth by sending only those MBs in INTRA-frame mode. This mode is particularly efficient in point-to-point connection or when the number of decoders is low.

The H.261 specific control packets FIR and NACK as described in [RFC2032](#) SHALL not be used to request image refreshment. Old implementations are encourage to use the methods describes in this section. Image refreshment may be needed due to packet loss or due to application requirements. An example of application requirement may be the change of the speaker in a voice-activated multipoint video switching conference. There are two methods that can be used for requesting image refreshment. The first method is by using the Extended RTP Profile for RTCP-based Feedback and sending RTCP generic control packets as described in

Even

Expires June 20, 2005

[Page 11]

[draft-ietf-avt-rtcp-feedback](#)[[rtcp-feedback](#)]. The second method is by using the application protocol specific commands like H.245[ITU.H245] FastUpdateRequest.

6. Payload Format Parameters

This section updates the H.261 media type described in [RFC3555](#) [[RFC3555](#)].

This section specifies optional parameters that MAY be used to select optional features of the payload format. The parameters are specified here as part of the MIME subtype registration for the ITU-T H.261 codec. A mapping of the parameters into the Session Description Protocol (SDP) [[RFC2327](#)] is also provided for those applications that use SDP. Multiple parameters SHOULD be expressed as a MIME media type string, in the form of a semi-colon separated list of parameters.

6.1 IANA Considerations

This section describes the MIME types and names associated with this payload format. The section registers the MIME types, as per [RFC2048](#) [[RFC2048](#)]

6.1.1 Registration of MIME media type video/H261

MIME media type name: video

MIME subtype name: H261

Required parameters: None

Optional parameters:

CIF: This parameter has the format of parameter=value. It describes the maximum supported frame rate for CIF resolution. permissible value are integer values 1 to 4 and it means that the maximum rate is 29.97/ specified value

QCIF: This parameter has the format of parameter=value. It describes the maximum supported frame rate for QCIF resolution. permissible value are integer values 1 to 4 and it means that the maximum rate is 29.97/ specified value

D: specifies support for still image graphics according to H.261 annex D.

Encoding considerations:

This type is only defined for transfer via RTP [[RFC3550](#)]

Security considerations: See [Section 8](#)

Even

Expires June 20, 2005

[Page 13]

Interoperability considerations: These are receiver options, current implementations will not send any optional parameters in their SDP. They will ignore the optional parameters and will encode the H.261 stream without annex D. Most decoders support at least QCIF resolutions and they are expected to be available almost in every H.261 based video application.

Published specification: RFC yyy

Applications which use this media type:

Audio and video streaming and conferencing tools.

Additional information: none

Person and email address to contact for further information :

Roni Even: roni.even@polycom.co.il

Intended usage: COMMON

Author/Change controller:

Roni Even

6.2 SDP Parameters

The MIME media type video/H261 string is mapped to fields in the Session Description Protocol (SDP) as follows:

- o The media name in the "m=" line of SDP MUST be video.
- o The encoding name in the "a=rtpmap" line of SDP MUST be H261 (the MIME subtype).
- o The clock rate in the "a=rtpmap" line MUST be 90000.
- o The optional parameters "CIF", "QCIF" and "D" if any, SHALL be included in the "a=fmtp" line of SDP. These parameters are expressed as a MIME media type string, in the form of as a semi-colon separated list of parameters

6.2.1 Usage with the SDP Offer Answer Model

When offering H.261 over RTP using SDP in an Offer/Answer model[RFC3264] the following considerations are necessary.

Codec options: (D) This option MUST NOT appear unless the sender of

Even

Expires June 20, 2005

[Page 14]

this SDP message is able to decode this option. This option shall be considered as a receiver's capability even when send in a "sendonly" offer.

Picture sizes and MPI:

Supported picture sizes and their corresponding minimum picture interval (MPI) information for H.261 can be combined. All picture sizes may be advertised to the other party, or only a subset of it. A terminal shall announce only those picture sizes (with their MPIs) which it is willing to receive. For example, MPI=2 means that maximum (decodeable) picture rate per second is about 15.

If the receiver does not specify the picture size /MPI optional parameter then it SHOULD be ready to receive QCIF resolution with MPI=1.

An example of media representation in SDP is as follows: (CIF at 15 frames per second, QCIF at 30 frames per second and annex D

```
m=video 49170/2 RTP/AVP 31
```

```
a=rtpmap:31 H261/900000
```

```
a=fmtp:31 CIF=2;QCIF=3;D
```

Even

Expires June 20, 2005

[Page 15]

7. Backward Compatibility to [RFC2032](#)

The current draft updates [RFC2032](#). This section will address the major backward compatibility issues.

7.1 Optional H.261-specific control packets

[RFC 2032](#) defined two H.261-specific RTCP control packets, "Full INTRA-frame Request" and "Negative Acknowledgement". Support of these control packets was optional. The H.261-specific control packets differ from normal RTCP packets in that they are not transmitted to the normal RTCP destination transport address for the RTP session (which is often a multicast address). Instead, these control packets are sent directly via unicast from the decoder to the encoder. The destination port for these control packets is the same port that the encoder uses as a source port for transmitting RTP (data) packets. Therefore, these packets may be considered "reverse" control packets. This new RFC suggest generic methods to address the same requirement. The authors of the drafts are not aware of products that supports these control packets. Since these are optional features new implementations SHALL ignores them and they SHALL not be used by new implementations.

7.2 New SDP optional parameters

The draft adds new optional parameters to the H261 payload type. Since these are optional parameters we expect old implementation to ignore these parameters while new implementations that will receive the H261 payload type with no parameters will behave as if the other side can accept H.261 at QCIF resolution and 30 frames per second.

Even

Expires June 20, 2005

[Page 16]

8. Security Considerations

RTP packets using the payload format defined in this specification are subject to the security considerations discussed in the RTP specification [[RFC3550](#)], and any appropriate RTP profile (for example [[RFC3551](#)]). This implies that confidentiality of the media streams is achieved by encryption. Because the data compression used with this payload format is applied end-to-end, encryption may be performed after compression so there is no conflict between the two operations.

A potential denial-of-service threat exists for data encoding using compression techniques that have non-uniform receiver-end computational load. The attacker can inject pathological datagrams into the stream which are complex to decode and cause the receiver to be overloaded. The usage of authentication of at least the RTP packet is RECOMMENDED

As with any IP-based protocol, in some circumstances a receiver may be overloaded simply by the receipt of too many packets, either desired or undesired. Network-layer authentication may be used to discard packets from undesired sources, but the processing cost of the authentication itself may be too high.

Even

Expires June 20, 2005

[Page 17]

9. Acknowledgements

This is to acknowledge the authors of [RFC2032](#) Thierry Turletti and Christian Huitema. Special thanks for the work done by Petri Koskelainen from Nokia and Nermeen Ismail from Cisco who helped with drafting the text for the new MIME types.

10. Changes from previous versions

10.1 changes from [RFC 2032](#)>

The changes from the [RFC 2032](#) are:

1. The H.261 MIME type is now in the payload specification.

Added optional parameters to the H.261 MIME type

3. Editorial changes to be in line with RFC editing procedures

10.2 Changes from bis-00

The changes from bis-00 are:

1. Deprecated the H.261 specific control packets
2. Change the separator in the SDP parameters to be semi-colon.

11. References

11.1 Normative References

- [H261] International Telecommunications Union, "Video codec for audiovisual services at p x 64 kbit/s", ITU Recommendation H.261, March 1993.
- [RFC2048] Freed, N., Klensin, J. and J. Postel, "Multipurpose Internet Mail Extensions (MIME) Part Four: Registration Procedures", [BCP 13](#), [RFC 2048](#), November 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2327] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", [RFC 2327](#), April 1998.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, [RFC 3551](#), July 2003.
- [RFC3555] Casner, S. and P. Hoschka, "MIME Type Registration of RTP Payload Formats", [RFC 3555](#), July 2003.
- [rtcp-feedback] Ott, J. and S. Wenger, "Extended RTP Profile for RTCP-based Feedback(RTP/AVPF)", [draft-ietf-avt-rtcp-feedback-08](#) (work in progress), January 2004.

11.2 Informative References

- [BT601] International Telecommunications Union, "Studio encoding parameters of digital television for standard 4:3 and wide-screen 16:9 aspect ratios", ITU-R Recommendation BT.601-5, October 1995.
- [H221] International Telecommunications Union, "Frame structure for a 64 to 1920 kbit/s channel in audiovisual

Even

Expires June 20, 2005

[Page 20]

teleservices", ITU Recommendation H.221, May 1999.

- [IEEE] Pingali, S., Towsley, D. and JF. Kurose, "A comparison of sender-initiated and receiver-initiated reliable multicast protocols", IEEE GLOBECOM 94, 1994.
- [INRIA] Turletti, T., "H.261 software codec for videoconferencing over the Internet", INRIA Research Report 1834, January 1993.
- [ITU.H245] International Telecommunications Union, "CONTROL PROTOCOL FOR MULTIMEDIA COMMUNICATION", ITU Recommendation H.245, 2003.
- [IVS] Turletti, T., "INRIA Videoconferencing tool (IVS), available by anonymous ftp from zenon.inria.fr in the "rodeo/ivs/last_version" directory. See also URL <http://www.inria.fr/rodeo/ivs.html>".
- [MICE] Sasse, MA., Bilting, U., Schultz, CD. and T. Turletti, "Remote Seminars through MultiMedia Conferencing: Experiences from the MICE project", Proc. INET'94/JENC5, Prague pp. 251/1-251/8, June 1994.
- [VIC] MacCanne, S., "VIC Videoconferencing tool, available by anonymous ftp from ee.lbl.gov in the "conferencing/vic" directory".

Author's Address

Roni Even
Polycom
94 Derech Em Hamoshavot
Petach Tikva 49130
Israel

EMail: roni.even@polycom.co.il

Even

Expires June 20, 2005

[Page 21]

Intellectual Property Statement

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.

Disclaimer of Validity

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

Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Even

Expires June 20, 2005

[Page 22]