

TICTOC
Internet-Draft
Intended status: Informational
Expires: March 27, 2008

S. Bryant
Cisco Systems
Y(J). Stein
RAD Data Communications
September 24, 2007

TICTOC Problem Statement
draft-bryant-tictoc-probstat-01.txt

Status of this Memo

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 becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

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

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

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

This Internet-Draft will expire on March 27, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This Internet draft describes a number of applications that require accurate time and/or frequency, and elucidates difficulties related to the transfer of high quality time and frequency across an IP or MPLS Packet Switched Network. This issue is not addressed by any currently chartered IETF working group, and we therefore propose the formation of a new working group to be called Transmitting Timing over IP Connections and Transfer of Clock (TICTOC).

Internet-Draft

TICTOC

September 2007

Table of Contents

1.	Introduction	3
2.	Applications	3
2.1.	Time Service Applications	3
2.2.	Frequency Service Applications	4
3.	Existing Time and Frequency Transfer Mechanisms	7
3.1.	Radio-based Timing Transfer Methods	7
3.2.	Dedicated Wire-based Timing Transfer Methods	8
3.3.	Transfer Using Packet Networks	9
3.3.1.	The Packet Network Environment	11
4.	Problems with Existing Solutions	11
5.	Other Forums Working in this Problem Space	12
6.	Security Considerations	13
7.	Security Considerations	13
8.	IANA Considerations	14
9.	Acknowledgements	14
10.	Informative References	14
	Authors' Addresses	14
	Intellectual Property and Copyright Statements	16

Internet-Draft

TICTOC

September 2007

1. Introduction

There is an emerging need to distribute highly accurate time and frequency information over IP and over MPLS packet switched networks (PSNs). In this problem statement we give several examples of applications that require time and/or frequency information, and explain why their needs can not be satisfied by time/frequency transfer over PSNs using existing protocols. We review these existing protocols and the work being carried out in the IETF and in other forums. Finally, we list the objectives of a proposed Working Group.

2. Applications

2.1. Time Service Applications

There are many applications that need to know the time with greater precision than provided by available mechanisms, such as the current version of NTP [[RFC1305](#)]. These applications span a range of industries: telecommunications, financial, test and measurement, government, industrial etc. Preliminary studies indicate that the availability of high accuracy time as a commodity enables use of techniques that were previously considered impossible. We can, therefore, expect that the provision of high quality time through the network infrastructure will generate a spiral of new innovative applications that will in turn make greater demands on the quality of time delivered to the end-user.

The best-known example of an application that requires high quality time in the telecommunications sector is the need to measure one-way packet delay. Current implementations of NTP have accuracy of the order of 10 milliseconds. When NTP is used to characterize packet delay and packet delay variation, such a time-base cannot resolve any two party event with a resolution of better than 20 milliseconds. Contrast this with the characteristics of a 10 Gb/s link, 1 kilometer

long. On such a link, a minimum sized packet takes 50 nanoseconds to send, and it takes 6 microseconds to traverse the link. The performance of current NTP implementations is orders of magnitude worse than the duration of network forwarding events, and clearly insufficient to characterize them.

Although the measurement of the characteristics of a packet network is the best-known telecommunications example, there are other operational needs, notably synchronization at the MAC layer. The cable industry has recently defined a new intra-PoP time transmission mechanism for just this purpose (DOCSYS Timing Interface), and WiMAX is looking for relative time delivery to its transmitter sites.

In the test and measurement and industrial sector there is a desire to move from special purpose communications infrastructure with calibrated wiring run back to a centralized controller, to a distributed system, in which instructions are distributed in advance to be executed at a predetermined time, and in which measurements are taken remotely and communicated back to a common point for later correlation and analysis. Two examples of this tendency are described below.

In the printing industry there is a need to control operations in multi-stand printing machines. The paper travels through these machines at a speed of nearly 100 km/h. At these speeds, coordination error of 1 microsecond between operations taking place at different positions in the machine produces a 0.03mm color offset, which is visible to the naked eye and results in an unacceptable degradation in quality.

In the electrical power industry there is a need to improve the measurement of power flows in order to monitor and predict usage patterns. One proposal is to extensively deploy synchro-phasors in the power network and to correlate their output to determine demand. These devices need to be able to determine the time of measurement with an accuracy of 1 microsecond.

More generally, there is growing interest in clock synchronization in massively parallel sensor networks. Advances in wireless communications have enabled the development of low power miniature sensors that collect and disseminate data from their immediate environment. Although each sensor has limited processing power,

through distributed processing the network becomes capable of performing various tasks of data fusion, but only assuming a common time base can be established.

The examples cited above are a small illustration of a trend that will continue to grow as designers realize that better scaling can be achieved with action-in-the-future, measure-and-correlate-later approaches to systems design.

Closer to the core interests and expertise of the IETF there is an emerging opinion that the availability of time as a commodity may simplify the protocols that we use in distributed systems.

[2.2.](#) Frequency Service Applications

There are applications that require time with a greater precision than can easily be provided using available mechanisms. Cellular base-stations require a highly accurate frequency reference from which they derive transmission frequencies and operational timing.

Conventionally GSM and WCDMA base stations obtain this reference frequency by locking on to the E1/T1 that links them to the base station controller. With the replacement of TDM links with Packet Switched Networks (PSNs) such as Ethernet, IP or MPLS, this simple method of providing a frequency reference is lost, and frequency information must be made available in some other way.

Why must the frequency reference be so accurate? First and foremost the requirement is derived from the need for the radio frequencies to be accurate. Radio spectrum is a limited and valuable commodity that needs to be used as efficiently as possible. In GSM, transmission frequencies are allocated to a given cellular base station and its neighbors in such fashion as to ensure that they do not interfere with each other. If the radio network designer cannot rely on the accuracy of these frequencies, the spacing between the frequencies used by neighboring sites must be increased, with significant economic consequences.

There is an additional requirement derived from the need for smooth handover when a mobile station crosses from one cell to another. If the radio system designer can not guarantee that the preparations required for handover occur in a few milliseconds, then they must

allow the mobile station to consume frequency resources simultaneously in both cells in order to avoid service disruption. The preparations required involve agreement between the mobile and base stations on the new frequencies and time offsets; these agreements can be accomplished quickly when the two base stations' frequency references are the same to a high degree of accuracy.

Another application requiring highly accurate frequency distribution is TDM pseudowires. The PWE3 WG has produced three techniques for emulating traditional low-rate (E1, T1, E3, T3) TDM services over PSNs, namely SAToP [[RFC4553](#)], CESoPSN, and TDMoIP. The major technical barrier to universal acceptance of TDM pseudowires is the accuracy of its clock recovery.

TDM network standards for timing accuracy and stability are extremely demanding. These requirements are not capriciously dictated by standards bodies, rather they are critical to the proper functioning of a high-speed TDM network. Consider a TDM receiver utilizing its own clock when converting the physical signal back into a bit-stream. If the receive clock runs at precisely the same rate as the source clock, then the receiver need only determine the optimal sampling phase. However, with any mismatch of clock rates, no matter how small, bit slips will eventually occur. For example, if the receive clock is slower than the source clock by one part per million (ppm), then the receiver will output 999,999 bits for every 1,000,000 bits sent, thus deleting one bit. Similarly, if the receive clock is

faster than the source clock by one part per billion (ppb), the receiver will insert a spurious bit every billion bits. One bit slip every million bits may seem acceptable at first glance, but translates to a catastrophic two errors per second for a 2 Mb/s E1 signal. ITU-T recommendations permit a few bit slips per day for a low-rate 64 kb/s channel, but strive to prohibit bit slips entirely for higher-rate TDM signals.

In certain cases, such as "toll-bypass" or "carrier's carrier" links, the endpoints of the TDM PW are full TDM networks, and timing may (indeed must) be derived from the respective network clocks. Since each of these clocks is highly accurate, they necessarily agree to high order. However, TDM PWs are expected to increasingly replace native TDM links delivering services from core networks towards users, and here there is no alternative to provision of accurate

frequency information.

In this context there are two types of frequency distribution being studied. In the first type the frequency reference used by the TDM source is distributed downstream, as in the native TDM service. In the "common clock" scenario highly accurate frequency information is distributed from a central server to both ends of the emulated TDM link. By placing in the protocol overhead timestamps based on the common clock, the receiver can accurately recover the TDM source clock.

While it is true that services designed for PSN (e.g. VoIP) transport are less dependent on frequency accuracy, there are still cases where such services need accurate frequency distribution. For example, when interconnecting tradition telephones via VoIP links, users expect these links to support legacy services, such as facsimile and dial-up data modems. The optimal technique for supporting these services is by provision of relay functions, e.g. T.38 fax-relay and V.150 modem-relay, that terminate the analog transmissions on both sides and transfer data content over the PSN. However, provision of relay services is computationally expensive, often requires expensive DSP-capable media gateways, and can only support known modem types. In many deployments old fax machines or proprietary data modems or secure voice telephones are used, and the modulations and handshake protocols are not recognized by the relays provided. In such cases the solution is to carry these transmissions over "clear channel" or Voice Band Data (VBD), i.e. to send raw samples of the audio in packets over the PSN.

The problem with clear channel transfer of data over PSNs is that the end points expect a non-intrusive analog channel between them, over which they implicitly transfer timing information. The receiver can thus continually lock onto the transmitter's frequency, and the

transmission can continue for an unlimited period without interruption. When employing clear channel, the frequency signal seen by the receiver is influenced by the destination gateway's clock used to convert the data samples back to analog form. If the source and destination gateways' clocks do not agree to a high degree of accuracy, the receiver does not properly lock onto the transmitter's clock, leading to disruption of the data reception. In typical cases a modem conversation transferred over clear channel may drop after

only several minutes, and fax reception may be interrupted after several pages have been received.

3. Existing Time and Frequency Transfer Mechanisms

In this section we will discuss existing mechanisms for transfer of time information, frequency information, or both. It should be noted that a sufficiently accurate time transfer service may be used to derive an accurate frequency transfer. Indeed, this is exactly what happens in a GPS disciplined frequency standard. On the other hand, an accurate frequency transfer service, while itself unable to transfer absolute time, is usually used to support and improve the performance of the time transfer service. Indeed, implementations of NTP or IEEE 1588 clients can be considered to consist of two phases. First, a local oscillator is locked to the server's frequency using incoming information from the incoming packets, and then the local time set based on the server's time and the propagation latency. By maintaining a local frequency source, the client requires relatively infrequent updates, and can continue functioning during short periods of network outage. Moreover, it can be shown that this method results in significantly better time transfer accuracy than methods that do not discipline a local clock.

Time transfer mechanisms can be divided into three classes. The first class consists of mechanisms that use radio frequency transport, while the second mechanism uses dedicated "wires" (which for our purposes include optical fibers). The third, which will be our main focus, exploits a Packet Switched Network for transfer of timing information. Advantages and disadvantages of these three methods are discussed in the following subsections.

3.1. Radio-based Timing Transfer Methods

The transfer of time by radio transmission is one of the oldest methods available, and is still the most accurate wide area method. In particular, there are two navigation in wide use that can be used for time transfer, The LONg RANge Navigation (LORAN) terrestrial radio system, and the Global Navigation Satellite System (GNSS). In both cases the user needs to be able to receive the transmitted

signal, requiring access to a suitable antenna. In certain

situations, e.g. basement communications rooms and urban canyons, the required signal may not be receivable.

Radio systems have high accuracy, far better than what we will later see can be achieved by existing PSN technologies. However coverage is limited; eLORAN for example only covers North America, and GPS does not have good coverage near the poles.

Although civilian use is sanctioned, the GPS was developed and is operated by the U.S. Department of Defense as a military system. For this reason there are political concerns that rules out its use in certain countries. The European Union is working on an alternative system called Galileo, which will be run as a commercial enterprise. In addition, GPS has some well-documented multi-hour outages, and is considered vulnerable to jamming. One major PTT also reports that they see a 2% per year failure rate for the antenna/receiver/clock-out chain.

While a radio-based timing service may be acceptable for some sites, it is frequently impractical to use on a per equipment basis. Hence, some form of local timing distribution is usually also required.

3.2. Dedicated Wire-based Timing Transfer Methods

The use of dedicated networks in the wide area does not scale well. Such services were available in the past, but for reasons of cost and accuracy have been superseded by GPS based solutions.

In the local area, one new technique is emerging as a mechanism for time transport, namely DOCSIS Timing Interface / Telecommunications Timing Interface (DTI/TTI). DTI was designed by DOCSIS for the distribution of time in a cable head-end in support of media access control. Time transfer is packet-based over a multi-stage hub and spoke dedicated network. It uses a single twisted-pair in half-duplex to eliminate inaccuracies due to the length differences between the pairs in a multi-pair cable. TTI is a development of DTI designed to provide synchronization in a telephone local office. Accuracy for DTI is better than 5 nanoseconds and range is 100 feet for DTI. This increases to 600 feet for TTI at some reduction in packet rate and hence time quality.

The DTI/TTI approach is applicable for special applications, but the need for a dedicated network imposes significant drawbacks for the general time transfer case.

Synchronous Ethernet is a technique that has recently been proposed for providing frequency distribution over Ethernet links. Modern

dedicated-media full-duplex Ethernet, in both copper and optical physical layer variants, transmits continuously. One can thus elect to derive the physical layer transmitter clock from a high quality frequency reference, instead of the conventional 100 ppm crystal-derived transmitter rate. The receiver at the other end of the link automatically locks onto the physical layer clock of the received signal, and thus itself gain access to a highly accurate and stable frequency reference. Then, in TDM fashion, this receiver could lock the transmission clock of its other ports to this frequency reference.

The ITU-T is presently working on a specification for synchronous Ethernet. Apart from some necessary higher layer packet based configuration and OAM operations, the solution is entirely physical layer, and has no impact on higher layers.

At first sight it would seem that the only application of synchronous Ethernet was in frequency transfer (it has no intrinsic time transfer mechanism). However, the quality of packet-based time transfer mechanism can be considerably enhanced if used in conjunction with synchronous Ethernet as a frequency reference.

3.3. Transfer Using Packet Networks

When using a PSN to transfer timing, a server sends timing information in the form of packets to one or multiple clients. When there are multiple clients, the timing packets may be multicast. Software in the client recovers the frequency and/or time of the server based on the packet arrival time and the packet contents.

There are two well-known protocols capable of running over a general-purpose packet network, NTP [[RFC1305](#)], and IEEE 1588 [[1588](#)]. NTP is the product of the IETF, and is currently undergoing revision to version 4. IEEE 1588 (a product of the IEEE Test and Measurement community) is specified in a limited first version, and the second version (1588v2) is in the detailed design stage.

NTP is widely deployed, but existing implementations deliver accuracy on the order of 10 milliseconds. This accuracy is not adequate for the applications described above. NTP suffers from the fact that it was designed to operate over the Internet, and the routers and switches used in the best effort Internet make no special concessions to NTP for preservation of time transfer accuracy. Furthermore, typical update rates are low and can not be significantly increased due to scalability issues in the server. In addition most NTP time servers and time receivers have a relatively unsophisticated

implementation that further degrades the final time quality.

IEEE 1588v1 was a time transfer protocol that exclusively used a fairly crude multicast technique. It is widely anticipated that wide scale deployment of IEEE1588 will be based on 1588v2. The information exchange component of IEEE 1588 is a protocol known as Precision Time Protocol (PTP).

IEEE 1588v2 can be considered to consist of several components:

1. A configuration and control protocol
2. A time transfer protocol
3. A time correction protocol
4. Physical mapping

The configuration and control protocol is based on the multicast approach of IEEE 1588v1 (multicast IP with recommended TTL=1, UDP, IEEE1588 payload with equipment identifier in the payload). The rationale for this approach was that the equipment needed to be "plug and play" (no configuration), was required to map to physical media other than Ethernet, and had to have a very low memory and processor footprint.

The time transfer protocol is a standard two-way time transfer approach used in other packet-based approaches. Like all such approaches it is subject to inaccuracies due to variable store and forward delays in the packet switches, and due to the assumption of symmetric propagation delays. The time transfer packets (in both directions) may be operated in a multicast or unicast mode.

The time correction protocol is used to correct for propagation, store and forward delays in the packet switches. This again may be operated multicast or unicast. This mechanism requires some level of hop-by-hop hardware support. This mechanism may also be considered a concept in its own right and may be adapted to enhance other packet time transfer protocols such as NTP.

The base 1588 specification describes how the PTP operates over the

Ethernet/IP/UDP protocol stack. Annexes are planned that describe PTP operation over pure layer 2 Ethernet, over IP without UDP, over MPLS, and over a number of specialist media.

The mappings of interest for telecommunications are PTP over UDP/IP, PTP over MPLS, and perhaps PTP over Ethernet, all in unicast mode only. Issues of a suitable control management and OAM environment for these applications are largely in abeyance, as are considerations about the exact nature of the network environment.

It is also worth noting the existence of a second IEEE effort, IEEE 802.1AS. This group is specifying the protocol and procedures to ensure synchronization across Bridged and Virtual Bridged Local Area Networks for time sensitive applications such as audio and video. For these LAN media the transmission delays are assumed to be fixed and symmetrical. IEEE 802.1AS specifies the use of IEEE 1588 specifications where applicable in the context of IEEE Standards 802.1D and 802.1Q. Synchronization to an externally provided timing signal (e.g., a recognized timing standard such as UTC or TAI) is not part of this standard but is not precluded. IEEE 802.1AS will specify how stations attached to bridged LANs to meet the respective jitter, wander, and time synchronization requirements for time-sensitive applications.

[3.3.1.](#) The Packet Network Environment

Packet delay variation, propagation asymmetry, and maximum permissible packet rate all have a significant bearing on the accuracy with which the client is able to determine absolute time. Thus the network environment has a large bearing on the quality of time that can be delivered.

Packet delay variation can to some extent be addressed by traffic engineering, thus time transfer with a service provider network in which suitable traffic engineering techniques had been applied might reasonably be expected to deliver a higher quality time service than can be achieved between two arbitrary hosts connected to the Internet. Greater gains can probably be obtained by deploying equipment that incorporates IEEE 1588 style on-the-fly packet timestamp correction, or follow-up message mechanisms that report the packet storage and forward delays to the client. However one can only be sure that such techniques are available along the entire path

in a well-controlled environment.

The packet rate between the time-server and its client also has a bearing on the quality of the time transfer, because at a higher rate the smart filter has a better chance of extracting the "good" packets. In a controlled environment it is possible to ensure that there is adequate bandwidth, and that the server is not overloaded. In such an environment the onus moves from protecting the server from overload, to ensuring that the server can satisfy the needs of all of the clients.

[4.](#) Problems with Existing Solutions

An obvious candidate for clock distribution is NTP or some upgrade thereof. While the time resolution provided by NTP is extremely

good, the accuracy attainable by existing NTP implementations does not satisfy the needs of the most demanding of the applications, mainly due to update rate and the particular client/server method employed.

The new IEEE 1588v2 protocol also addresses these needs, but has been largely designed around a well-controlled LAN environment. A 1588 server in unicast mode needs to save state information for each client, a solution that does not scale well to deployment sizes envisioned. In addition, 1588 specifies hardware upgrades in order to perform well in an IP network.

Synchronous Ethernet only satisfies the need for frequency distribution, and even then only over one physical Ethernet link at a time. In order to use synchronous Ethernet in a network, all network elements must be upgraded to support synchronous operation at the physical layer. Even when hardware can be upgraded, only frequency is delivered, and there is still a need to develop a time transfer protocol.

[5.](#) Other Forums Working in this Problem Space

The NTP WG is the IETF group working on time distribution, but is presently only documenting NTPv4 and is not working on new algorithms

or protocols. It is expected that many participants of the NTP WG will participate in the TICTOC effort.

The PWE3 WG has discussed frequency distribution for the TDM PW application, however it is not chartered to develop protocols for this purpose. It is expected that participants of the PWE3 WG who were active in the TDM PW discussions will participate in the TICTOC effort.

The work that is underway outside the IETF is either complementary to this proposal, or less general than the proposal proposed by the TICTOC work proposal.

The IEEE 1588 task force is working on a new version of their protocol that will run over more types of PSNs, and is planning to conclude its development work in the near future. The protocol to be specified contains elements that will be of use in an IETF environment, but is unlikely to be regarded as being a complete, robust solution in such an environment. If the IEEE 1588 structure is deemed to be a suitable platform, then the IETF could contribute an Internet profile, including a complete distributed systems environment suitable for our purposes. Alternatively, the IETF could perhaps borrow some of the delay correction mechanisms and

incorporate them into a development of a new version of NTP.

In addition, IEEE 802.1AS is working on Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks, basing itself on the IEEE 1588 standard.

ITU-T SG15 Question 13 has produced Recommendation G.8261 "Timing and synchronization aspects in packet networks" [[G8261](#)]. This Recommendation defines requirements for various scenarios, outlines the functionality of frequency distribution elements, and provides measurement guidelines. It does not specify algorithms to be used for attaining the performance needed. It does define requirements for status synchronization messages, but does not otherwise define a protocol (although work is in progress). To date the ITU-T has focused on Ethernet infrastructure, but this is likely to extend to an MPLS environment. Two new work items, G.paclock and G.pacmod extend the work, and in particular, G.pacmod intends to introduce time transfer. This is an area where the IETF, with its expertise in

IP and MPLS networks, may co-operate with the ITU.

6. Security Considerations

Time and frequency services are a significant element of network infrastructure, and are critical for certain emerging applications. Hence time and frequency transfer services **MUST** be protected from being compromised. The most significant threat is a false time or frequency server being accepted instead of a true one, thus enabling a hacker to bring down critical services.

Any protection mechanism must be designed in such a way that it does not degrade the quality of the time transfer. Such a mechanism **SHOULD** also be relatively lightweight, as client restrictions often dictate a low processing and memory footprint, and because the server may have extensive fan-out.

7. Security Considerations

Timing distribution is highly sensitive to packet delay, and can thus can deteriorate under congestion conditions. Furthermore the disciplining of the client's oscillator (the sole component of frequency transfer, and a critical component of time transfer) is a function that should not be disrupted. When the service is disrupted the client needs to go into "holdover" mode, and its accuracy will consequently be degraded. Depending on the relative quality of the client's clock and the required quality after disciplining, a relatively high packet rate may be required.

Timing transfer packets should always be sent using the highest class of service, and when possible should be sent over a traffic engineered path.

When the network goes into congestion it should try to avoid discarding time transfer packets until the situation is critical. Work performed by the IETF PWE3 WG on congestion would seem to be applicable to this problem area.

8. IANA Considerations

No IANA actions are required as a result of the publication of this document.

9. Acknowledgements

The authors wish to thank Laurent Montini for valuable comments.

10. Informative References

- [1588] IEEE, "1588-2002 Standard for A Precision Clock Synchronization Protocol for Networked Measurement and Control Systems".
- [G8261] ITU-T, "Recommendation G.8261/Y.1361 - Timing and synchronization aspects in packet networks", May 2006.
- [RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", [RFC 1305](#), March 1992.
- [RFC4553] Vainshtein, A. and YJ. Stein, "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", [RFC 4553](#), June 2006.

Authors' Addresses

Stewart Bryant
Cisco Systems
250 Longwater Ave., Green Park
Reading RG2 6GB
United Kingdom

Email: stbryant@cisco.com

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

ISRAEL

Phone: +972 3 645-5389

Email: yaakov_s@rad.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

Acknowledgment

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

