

Mobility Services Transport: Problem Statement

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

There are ongoing activities in the networking community to develop solutions that aid in IP handover mechanisms between heterogeneous wired and wireless access systems including, but not limited to, IEEE 802.21. Intelligent access selection, taking into account link-layer attributes, requires the delivery of a variety of different information types to the terminal from different sources within the network and vice-versa. The protocol requirements for this signalling have both transport and security issues that must be considered. The signalling must not be constrained to specific link types, so there is at least a common component to the signalling problem, which is within the scope of the IETF. This document presents a problem statement for this core problem.

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. Requirements Language	3
3. Definition of Mobility Services	4
4. Deployment Scenarios for MoS	4
4.1. End-to-End Signalling and Transport over IP	5
4.2. End-to-End Signalling and Partial Transport over IP	5
4.3. End-to-End Network-to-Network Signalling	6
5. MoS Transport Protocol Splitting	7
5.1. Payload Formats and Extensibility Considerations	8
5.2. Requirements on the Mobility Service Transport Layer	8
6. Security Considerations	11
7. Conclusions	12
8. Acknowledgements	13
9. References	13
9.1. Normative References	13
9.2. Informative References	13
Contributors	14

1. Introduction

This document provides a problem statement for the exchange of information to support handover in heterogeneous link environments [1]. This mobility support service allows more sophisticated handover operations by making available information about network characteristics, neighboring networks and associated characteristics, indications that a handover should take place, and suggestions for suitable target networks to which to handover. The mobility support services are complementary to IP mobility mechanisms [4], [5], [6], [7], [8], [9] to enhance the overall performance and usability perception.

There are two key attributes to the handover support service problem for inter-technology handovers:

1. The Information: the information elements being exchanged. The messages could be of a different nature, such as information, commands to perform an action, or events informing of a change, potentially being defined following a common structure.

2. The Underlying Transport: the transport mechanism to support exchange of the information elements mentioned above. This transport mechanism includes information transport, discovery of peers, and the securing of this information over the network.

The initial requirement for this protocol comes from the need to provide a transport for the Media Independent Handover (MIH) protocol being defined by IEEE 802.21 [1], which is not bound to any specific link layer and can operate over more than one network-layer hop. The solution should be flexible to accommodate evolution in the MIH standard, and should also be applicable for other new mobility signalling protocols that have similar message patterns and discovery and transport requirements.

The structure of this document is as follows. Section 3 defines Mobility Services. Section 4 provides a simple model for the protocol entities involved in the signalling and their possible relationships. Section 5 describes a decomposition of the signalling problem into service-specific parts and a generic transport part. Section 5.2 describes more detailed requirements for the transport component. Section 6 provides security considerations. Section 7 summarizes the conclusions and open issues.

2. Terminology

The following abbreviations are used in the document:

MIH: Media Independent Handover

MN: Mobile Node

NN: Network Node, intended to represent some device in the network (the location of the node, e.g., in the access network, the home network is not specified, and for the moment it is assumed that they can reside anywhere).

EP: Endpoint, intended to represent the terminating endpoints of the transport protocol used to support the signalling exchanges between nodes.

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [2].

3. Definition of Mobility Services

As mentioned in the Introduction, mobility (handover) support in heterogeneous wireless environments requires functional components located either in the mobile terminal or in the network to exchange information and eventually to make decisions upon this information exchange. For instance, traditional host-based handover solutions could be complemented with more sophisticated network-centric solutions. Also, neighborhood discovery, potentially a complex operation in heterogeneous wireless scenarios, can result in a simpler step if implemented with a unified interface towards the access network.

In this document, the different supporting functions for Media Independent Handover (MIH) management are generally referred to as Mobility Services (MoS) that have different requirements for the transport protocol. These requirements and associated functionalities are the focus of this document. Speaking 802.21 terminology, MoS can be regarded as Information Services (IS), Event Services (ES), and Command Service (CS).

4. Deployment Scenarios for MoS

The deployment scenarios are outlined in the following sections.

Note: while MN-to-MN signalling exchanges are theoretically possible, these are not currently being considered.

The following scenarios are discussed for understanding the overall problem of transporting MIH protocol. Although these are all possible scenarios and MIH services can be delivered through link-layer specific solutions and/or through a "layer 3 or above" protocol, this problem statement focuses on the delivery of information for Mobility Services only for the latter case.

4.1. End-to-End Signalling and Transport over IP

In this case, the end-to-end signalling used to exchange the handover information elements (the Information Exchange) runs end-to-end between MN and NN. The underlying transport is also end-to-end.

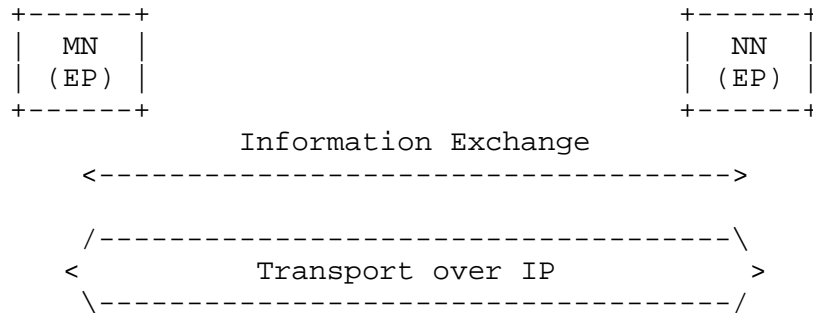


Figure 1: End-to-End Signalling and Transport

4.2. End-to-End Signalling and Partial Transport over IP

As before, the Information Exchange runs end-to-end between the MN and the second NN. However, in this scenario, some transport means other than IP are used from the MN to the first NN, and the transport over IP is used only between NNs. This is analogous to the use of EAP end-to-end between Supplicant and Authentication Server, with an upper-layer multihop protocol, such as Remote Authentication Dial-In User Service (RADIUS), used as a backhaul transport protocol between an Access Point and the Authentication Server.

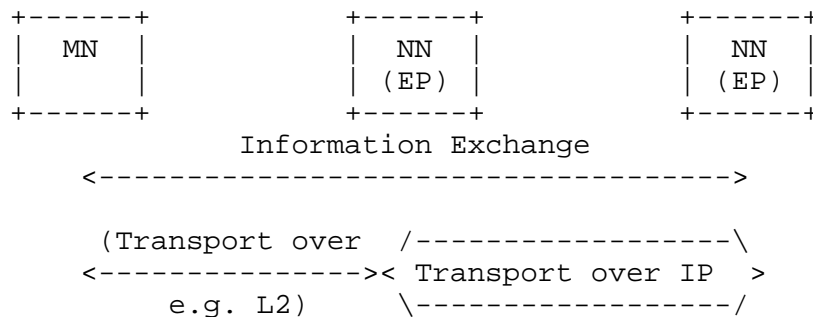


Figure 2: Partial Transport

4.3. End-to-End Network-to-Network Signalling

In this case, NN to NN signalling is envisioned. Such a model should allow different network components to gather information from each other. This is useful for instance in conditions where network components need to make decisions and instruct mobile terminals of operations to be executed.

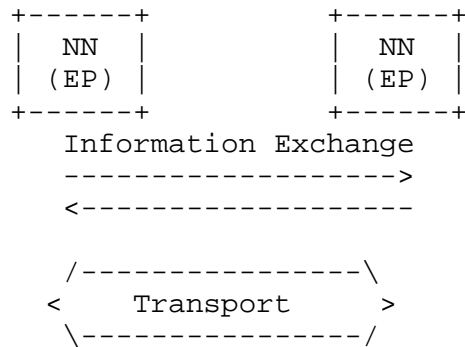


Figure 3: Information Exchange between Different NNs

Network nodes exchange information about the status of connected terminals.

5. MoS Transport Protocol Splitting

Figure 4 shows a model where the Information Exchanges are implemented by a signalling protocol specific to a particular mobility service, and these are relayed over a generic transport layer (the Mobility Service Transport Layer).

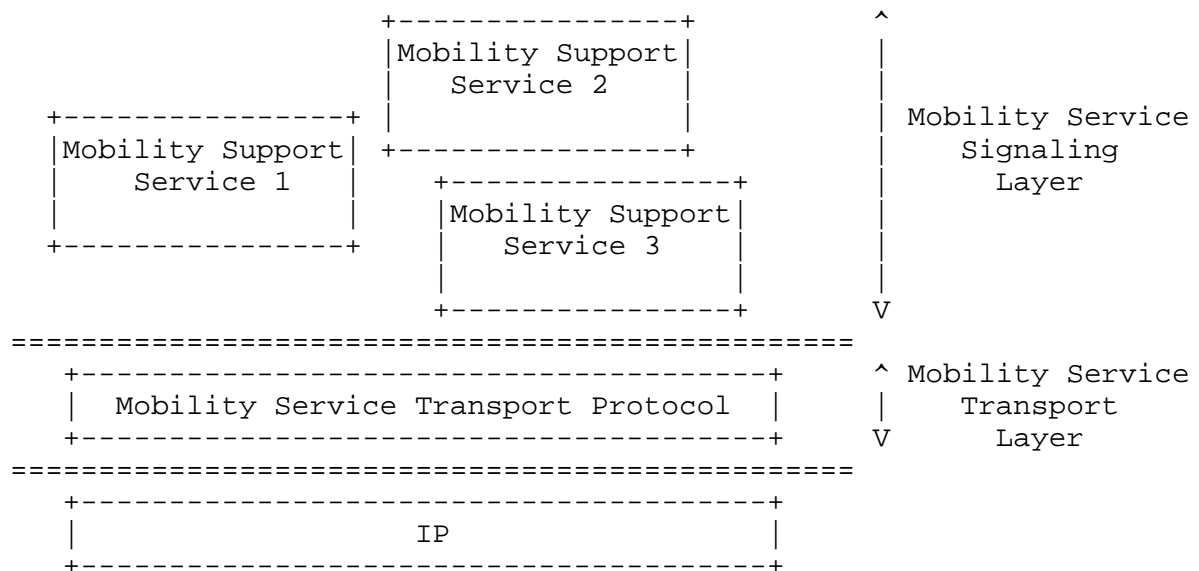


Figure 4: Handover Services over IP

The Mobility Service Transport Layer provides certain functionality (outlined in Section 5.2) to the higher-layer mobility support services in order to support the exchange of information between communicating Mobility Service functions. The transport layer effectively provides a container capability to mobility support services, as well as any required transport and security operations required to provide communication, without regard to the protocol semantics and data carried in the specific Mobility Services.

The Mobility Support Services themselves may also define certain protocol exchanges to support the exchange of service-specific information elements. It is likely that the responsibility for defining the contents and significance of the information elements is the responsibility of standards bodies other than the IETF. Example Mobility Services include the Information Services, Event Services, and Command Services.

5.1. Payload Formats and Extensibility Considerations

The format of the Mobility Service Transport Protocol (MSTP) is as follows:

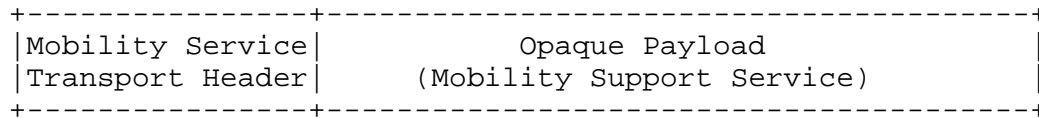


Figure 5: Protocol Structure

This figure shows the case for an MIH message that is smaller than the MTU of the path to the destination. A larger payload may require the transport protocol to transparently fragment and reassemble the MIH message.

The opaque payload encompasses the Mobility Support Service (MSTP) information that is to be transported. The definition of the Mobility Service Transport Header is something that is best addressed within the IETF. MSTP does not inspect the payload, and any required information will be provided by the MSTP users.

5.2. Requirements on the Mobility Service Transport Layer

The following section outlines some of the general transport requirements that should be supported by the Mobility Service Transport Protocol. Analysis has suggested that at least the following need to be taken into account:

Discovery: MNs need the ability to either discover nodes that support certain services or discover services provided by a certain node. The service discovery can be dealt with using messages as defined in [1]. This section refers to node-discovery in either scenario. There are no assumptions about the location of these Mobility Service nodes within the network. Therefore, the discovery mechanism needs to operate across administrative boundaries. Issues such as speed of discovery, protection against spoofing, when discovery needs to take place, and the length of time over which the discovery information may remain valid; all need to be considered. Approaches include:

- * Hard coding information into the MN, indicating either the IP address of the NN, or information about the NN that can be resolved onto an IP address. The configuration information could be managed dynamically, but assumes that the NN is independent of the access network to which the MN is currently attached.

- * Pushing information to the MN, where the information is delivered to the MN as part of other configuration operations, for example, via DHCP or Router Discovery exchange. The benefit of this approach is that no additional exchanges with the network would be required, but the limitations associated with modifying these protocols may limit applicability of the solution.
- * MN dynamically requesting information about a node, which may require both MN and NN support for a particular service discovery mechanism. This may require additional support by the access network (e.g., multicast or anycast) even when it may not be supporting the service directly itself.

Numerous directory and configuration services already exist, and reuse of these mechanisms may be appropriate. There is an open question about whether multiple methods of discovery would be needed, and whether NNs would also need to discover other NNs. The definition of a service also needs to be determined, including the granularity of the description. For example, IEEE 802.21 specifies three different types of Mobility Services (Information Services, Command Services, and Event Services) that can be located in different portions of the network. An MN could therefore run a discovery procedure of any service running in the (home or visited) network or could run a discovery procedure for a specific service.

Information from a trusted source: The MN uses the Mobility Service information to make decisions about what steps to take next. It is essential that there is some way to ensure that the information received is from a trustworthy source. This requirement should reuse trust relationships that have already been established in the network, for example, on the relationships established by the Authentication, Authorization, and Accounting (AAA) infrastructure after a mutual authentication, or on the certificate infrastructure required to support SEND [10]. Section 6 provides a more complete analysis.

Security association management: A common security association negotiation method, independent of any specific MSTP user, should be implemented between the endpoints of the MSTP. The solution must also work in the case of MN mobility.

Secure delivery: The Mobility Service information must be delivered securely (integrity and confidentiality) between trusted peers, where the transport may pass through untrusted intermediate nodes and networks. The Mobility Service information should also be protected against replay attacks and denial-of-service attacks.

Low latency: Some of the Mobility Services generate time-sensitive information. Therefore, there is a need to deliver the information over quite short timescales, and the required lifetime of a connection might be quite short-lived. As an example, the frequency of messages defined in [1] varies according to the MIH service type. It is expected that Events and Commands messages arrive at an interval of hundreds of milliseconds in order to capture quick changes in the environment and/or process handover commands. On the other hand, Information Service messages are mainly exchanged each time a new network is visited that may be in the order of hours or days. For reliable delivery, short-lived connections could be set up as needed, although there is a connection setup latency associated with this approach. Alternatively, a long-lived connection could be used, but this requires advanced warning of being needed and some way to maintain the state associated with the connection. It also assumes that the relationships between devices supporting the mobility service are fairly stable. Another alternative is connectionless operation, but this has interactions with other requirements, such as reliable delivery.

Reliability: Reliable delivery for some of the Mobility Services may be essential, but it is difficult to trade this off against the low latency requirement. It is also quite difficult to design a robust, high-performance mechanism that can operate in heterogeneous environments, especially one where the link characteristics can vary quite dramatically. There are two main approaches that could be adopted:

1. Assume the transport cannot be guaranteed to support reliable delivery. In this case, the Mobility Support Service itself will have to provide a reliability mechanism (at the MIH level) to allow communicating endpoints to acknowledge receipt of information.
2. Assume the underlying transport will provide reliable delivery. There is no need in this case to provide reliability at the MIH level.

Guidelines provided in [3] are being considered while writing this document.

Congestion Control: A Mobility Service may wish to transfer small or large amounts of data, placing different requirements for congestion control in the transport. As an example, the MIH message [1] size varies widely from about 30 bytes (for a broadcast capability discovery request) to be normally less than 64 KB, but may be greater than 64KB (for an IS MIH_Get_Information

response primitive). A typical MIH message size for the Events and Commands Services service ranges between 50 to 100 bytes. The solution should consider different congestion control mechanisms depending on the amount of data generated by the application (MIH) as suggested in [3].

Fragmentation and reassembly: ES and CS messages are small in nature, are sent frequently, and may wish trade reliability in order to satisfy the tight latency requirements. On the other hand, IS messages are more resilient in terms of latency constraints, and some long IS messages could exceed the MTU of the path to the destination. Depending on the choice of the transport protocol, different fragmentation and reassembly strategies are required.

Multihoming: For some Information Services exchanged with the MN, there is a possibility that the request and response messages could be carried over two different links. For example, a handover command request is on the current link while the response could be delivered on the new link. It is expected that the transport protocol is capable of receiving information via multiple links. It is also expected that the MSTP user combines information belonging to the same session/transaction. When mobility is applied, the underlying IP mobility mechanism should provide session continuity when required.

IPv4 and IPv6 support: The MSTP must support both IPv4 and IPv6 including NAT traversal for IPv4 networks and firewall pass-through for IPv4 and IPv6 networks.

6. Security Considerations

Network-supported Mobility Services aim at improving decision making and management of dynamically connected hosts.

Information Services may not require authorization of the client, but both Event and Command Services may authenticate message sources, particularly if they are mobile. Network-side service entities will typically need to provide proof of authority to serve visiting devices. Where signalling or radio operations can result from received messages, significant disruption may result from processing bogus or modified messages. The effect of processing bogus messages depends largely upon the content of the message payload, which is handled by the handover services application. Regardless of the variation in effect, message delivery mechanisms need to provide protection against tampering, spoofing, and replay attacks.

Sensitive and identifying information about a mobile device may be exchanged during handover-service message exchange. Since handover decisions are to be made based upon message exchanges, it may be possible to trace a user's movement between cells, or predict future movements, by inspecting handover service messages. In order to prevent such tracking, message confidentiality and message integrity should be available. This is particularly important because many mobile devices are associated with only one user, since divulging of such information may violate the user's privacy. Additionally, identifying information may be exchanged during security association construction. As this information may be used to trace users across cell boundaries, identity protection should be available, if possible, when establishing source addresses (SAs).

In addition, the user should not have to disclose its identity to the network (any more than it needed to during authentication) in order to access the Mobility Support Services. For example, if the local network is just aware that an anonymous user with a subscription to "example.com" is accessing the network, the user should not have to divulge their true identity in order to access the Mobility Support Services available locally.

Finally, the NNs themselves will potentially be subject to denial-of-service attacks from MNs, and these problems will be exacerbated if operation of the Mobility Service protocols imposes a heavy computational load on the NNs. The overall design has to consider at what stage (e.g., discovery, transport layer establishment, and service-specific protocol exchange) denial-of-service prevention or mitigation should be built in.

7. Conclusions

This document outlined a broad problem statement for the signalling of information elements across a network to support Mobility Services. In order to enable this type of signalling service, a need for a generic transport solution with certain transport and security properties was outlined. Whilst the motivation for considering this problem has come from work within IEEE 802.21, a desirable goal is to ensure that solutions to this problem are applicable to a wider range of Mobility Services.

It would be valuable to establish realistic performance goals for the solution to this common problem (i.e., transport and security aspects) using experience from previous IETF work in this area and knowledge about feasible deployment scenarios. This information could then be used as an input to other standards bodies in assisting them to design Mobility Services with feasible performance requirements.

Much of the functionality required for this problem is available from existing IETF protocols or combination thereof. This document takes no position on whether an existing protocol can be adapted for the solution or whether new protocol development is required. In either case, we believe that the appropriate skills for development of protocols in this area lie in the IETF.

8. Acknowledgements

Thanks to Subir Das, Juan Carlos Zuniga, Robert Hancock, and Yoshihiro Ohba for their input. Thanks to the IEEE 802.21 chair, Vivek Gupta, for coordinating the work and supporting the IETF liaison. Thanks to all IEEE 802.21 WG folks who contributed to this document indirectly.

9. References

9.1. Normative References

- [1] "Draft IEEE Standard for Local and Metropolitan Area Networks: Media Independent Handover Services", IEEE LAN/MAN Draft IEEE P802.21/D07.00, July 2007.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

- [3] Eggert, L. and G. Fairhurst, "UDP Usage Guidelines for Application Designers", Work in Progress.
- [4] 3GPP, "3GPP system architecture evolution (SAE): Report on technical options and conclusions", 3GPP TR 23.882 0.10.1, February 2006.
- [5] Perkins, C., Ed., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [6] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [7] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.
- [8] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, June 2006.

- [9] Koodli, R., Ed., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.
- [10] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

Contributors' Addresses

Eleanor Hepworth
Siemens Roke Manor Research
Roke Manor
Romsey, SO51 5RE
UK

EMail: eleanor.hepworth@roke.co.uk

Srinivas Sreemanthula
Nokia Research Center
6000 Connection Dr.
Irving, TX 75028
USA

EMail: srinivas.sreemanthula@nokia.com

Yoshihiro Ohba
Toshiba America Research, Inc.
1 Telcordia Drive
Piscataway NJ 08854
USA

EMail: yohba@tari.toshiba.com

Vivek Gupta
Intel Corporation
2111 NE 25th Avenue
Hillsboro, OR 97124
USA

Phone: +1 503 712 1754
EMail: vivek.g.gupta@intel.com

Jouni Korhonen
TeliaSonera Corporation.
P.O.Box 970
FIN-00051 Sonera
FINLAND

Phone: +358 40 534 4455
EMail: jouni.korhonen@teliasonera.com

Rui L.A. Aguiar
Instituto de Telecomunicacoes Universidade de Aveiro
Aveiro 3810
Portugal

Phone: +351 234 377900
EMail: ruilaa@det.ua.pt

Sam(Zhongqi) Xia
Huawei Technologies Co., Ltd
HuaWei Bld., No.3 Xinx Rd. Shang-Di Information Industry Base
100085
Hai-Dian District Beijing, P.R. China

Phone: +86-10-82836136
EMail: xiazhongqi@huawei.com

Authors' Addresses

Telemaco Melia, Editor
Cisco Systems International Sarl
Avenue des Uttins 5
1180 Rolle
Switzerland (FR)

Phone: +41 21 822718
EMail: tmelia@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

