

Network Working Group
Request for Comments: 4541
Category: Informational

M. Christensen
Thrane & Thrane
K. Kimball
Hewlett-Packard
F. Solensky
Calix
May 2006

Considerations for Internet Group Management Protocol (IGMP)
and Multicast Listener Discovery (MLD) Snooping Switches

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.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo describes the recommendations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) snooping switches. These are based on best current practices for IGMPv2, with further considerations for IGMPv3- and MLDv2-snooping. Additional areas of relevance, such as link layer topology changes and Ethernet-specific encapsulation issues, are also considered.

1. Introduction

The IEEE bridge standard [BRIDGE] specifies how LAN packets are 'bridged', or as is more commonly used today, switched between LAN segments. The operation of a switch with respect to multicast packets can be summarized as follows. When processing a packet whose destination MAC address is a multicast address, the switch will forward a copy of the packet into each of the remaining network interfaces that are in the forwarding state in accordance with [BRIDGE]. The spanning tree algorithm ensures that the application of this rule at every switch in the network will make the packet accessible to all nodes connected to the network.

This behaviour works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is

intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. While nodes will rarely incur any processing overhead to filter packets addressed to unrequested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. In general, significant bandwidth can be wasted by flooding.

In recent years, a number of commercial vendors have introduced products described as "IGMP snooping switches" to the market. These devices do not adhere to the conceptual model that provides the strict separation of functionality between different communications layers in the ISO model, and instead utilize information in the upper level protocol headers as factors to be considered in processing at the lower levels. This is analogous to the manner in which a router can act as a firewall by looking into the transport protocol's header before allowing a packet to be forwarded to its destination address.

In the case of IP multicast traffic, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving packets addressed to the group address. This is in contrast to normal switch behavior where multicast traffic is typically forwarded on all interfaces.

Many switch datasheets state support for IGMP snooping, but no recommendations for this exist today. It is the authors' hope that the information presented in this document will supply this foundation.

The recommendations presented here are based on the following information sources: The IGMP specifications [RFC1112], [RFC2236] and [IGMPv3], vendor-supplied technical documents [CISCO], bug reports [MSOFT], discussions with people involved in the design of IGMP snooping switches, MAGMA mailing list discussions, and on replies by switch vendors to an implementation questionnaire.

Interoperability issues that arise between different versions of IGMP are not the focus of this document. Interested readers are directed to [IGMPv3] for a thorough description of problem areas.

The suggestions in this document are based on IGMP, which applies only to IPv4. For IPv6, Multicast Listener Discovery [MLD] must be used instead. Because MLD is based on IGMP, we do not repeat the entire description and recommendations for MLD snooping switches. Instead, we point out the few cases where there are differences from IGMP.

Note that the IGMP snooping function should apply only to IPv4 multicasts. Other multicast packets, such as IPv6, might be suppressed by IGMP snooping if additional care is not taken in the implementation as mentioned in the recommendations section. It is desired not to restrict the flow of non-IPv4 multicasts other than to the degree which would happen as a result of regular bridging functions. Likewise, MLD snooping switches are discouraged from using topological information learned from IPv6 traffic to alter the forwarding of IPv4 multicast packets.

2. IGMP Snooping Recommendations

The following sections list the recommendations for an IGMP snooping switch. The recommendation is stated and is supplemented by a description of a possible implementation approach. All implementation discussions are examples only and there may well be other ways to achieve the same functionality.

2.1. Forwarding rules

The IGMP snooping functionality is separated into a control section (IGMP forwarding) and a data section (Data forwarding).

2.1.1. IGMP Forwarding Rules

- 1) A snooping switch should forward IGMP Membership Reports only to those ports where multicast routers are attached.

Alternatively stated: a snooping switch should not forward IGMP Membership Reports to ports on which only hosts are attached. An administrative control may be provided to override this restriction, allowing the report messages to be flooded to other ports.

This is the main IGMP snooping functionality for the control path.

Sending membership reports to other hosts can result, for IGMPv1 and IGMPv2, in unintentionally preventing a host from joining a specific multicast group.

When an IGMPv1 or IGMPv2 host receives a membership report for a group address that it intends to join, the host will suppress its own membership report for the same group. This join or message suppression is a requirement for IGMPv1 and IGMPv2 hosts. However, if a switch does not receive a membership report from the host it will not forward multicast data to it.

This is not a problem in an IGMPv3-only network because there is no suppression of IGMP Membership reports.

The administrative control allows IGMP Membership Report messages to be processed by network monitoring equipment such as packet analyzers or port replicators.

The switch supporting IGMP snooping must maintain a list of multicast routers and the ports on which they are attached. This list can be constructed in any combination of the following ways:

- a) This list should be built by the snooping switch sending Multicast Router Solicitation messages as described in IGMP Multicast Router Discovery [MRDISC]. It may also snoop Multicast Router Advertisement messages sent by and to other nodes.
- b) The arrival port for IGMP Queries (sent by multicast routers) where the source address is not 0.0.0.0.

The 0.0.0.0 address represents a special case where the switch is proxying IGMP Queries for faster network convergence, but is not itself the Querier. The switch does not use its own IP address (even if it has one), because this would cause the Queries to be seen as coming from a newly elected Querier. The 0.0.0.0 address is used to indicate that the Query packets are NOT from a multicast router.

- c) Ports explicitly configured by management to be IGMP-forwarding ports, in addition to or instead of any of the above methods to detect router ports.
- 2) IGMP networks may also include devices that implement "proxy-reporting", in which reports received from downstream hosts are summarized and used to build internal membership states. Such proxy-reporting devices may use the all-zeros IP Source-Address when forwarding any summarized reports upstream. For this reason, IGMP membership reports received by the snooping switch must not be rejected because the source IP address is set to 0.0.0.0.
 - 3) The switch that supports IGMP snooping must flood all unrecognized IGMP messages to all other ports and must not attempt to make use of any information beyond the end of the network layer header.

In addition, earlier versions of IGMP should interpret IGMP fields as defined for their versions and must not alter these fields when forwarding the message. When generating new messages, a given IGMP version should set fields to the appropriate values for its

own version. If any fields are reserved or otherwise undefined for a given IGMP version, the fields should be ignored when parsing the message and must be set to zeroes when new messages are generated by implementations of that IGMP version. An exception may occur if the switch is performing a spoofing function, and is aware of the settings for new or reserved fields that would be required to correctly spoof for a different IGMP version.

The reason to worry about these trivialities is that IGMPv3 overloads the old IGMP query message using the same type number (0x11) but with an extended header. Therefore there is a risk that IGMPv3 queries may be interpreted as older version queries by, for example, IGMPv2 snooping switches. This has already been reported [IETF56] and is discussed in section 2.2.

- 4) An IGMP snooping switch should be aware of link layer topology changes caused by Spanning Tree operation. When a port is enabled or disabled by Spanning Tree, a General Query may be sent on all active non-router ports in order to reduce network convergence time. Non-Querier switches should be aware of whether the Querier is in IGMPv3 mode. If so, the switch should not spoof any General Queries unless it is able to send an IGMPv3 Query that adheres to the most recent information sent by the true Querier. In no case should a switch introduce a spoofed IGMPv2 Query into an IGMPv3 network, as this may create excessive network disruption.

If the switch is not the Querier, it should use the 'all-zeros' IP Source Address in these proxy queries (even though some hosts may elect to not process queries with a 0.0.0.0 IP Source Address). When such proxy queries are received, they must not be included in the Querier election process.

- 5) An IGMP snooping switch must not make use of information in IGMP packets where the IP or IGMP headers have checksum or integrity errors. The switch should not flood such packets but if it does, it should also take some note of the event (i.e., increment a counter). These errors and their processing are further discussed in [IGMPv3], [MLD] and [MLDv2].
- 6) The snooping switch must not rely exclusively on the appearance of IGMP Group Leave announcements to determine when entries should be removed from the forwarding table. It should implement a membership timeout mechanism such as the router-side functionality of the IGMP protocol as described in the IGMP and MLD specifications (See Normative Reference section for IGMPv1-3 and MLDv1-2) on all its non-router ports. This timeout value should be configurable.

2.1.2. Data Forwarding Rules

- 1) Packets with a destination IP address outside 224.0.0.X which are not IGMP should be forwarded according to group-based port membership tables and must also be forwarded on router ports.

This is the main IGMP snooping functionality for the data path. One approach that an implementation could take would be to maintain separate membership and multicast router tables in software and then "merge" these tables into a forwarding cache.

- 2) Packets with a destination IP (DIP) address in the 224.0.0.X range which are not IGMP must be forwarded on all ports.

This recommendation is based on the fact that many host systems do not send Join IP multicast addresses in this range before sending or listening to IP multicast packets. Furthermore, since the 224.0.0.X address range is defined as link-local (not to be routed), it seems unnecessary to keep the state for each address in this range. Additionally, some routers operate in the 224.0.0.X address range without issuing IGMP Joins, and these applications would break if the switch were to prune them due to not having seen a Join Group message from the router.

- 3) An unregistered packet is defined as an IPv4 multicast packet with a destination address which does not match any of the groups announced in earlier IGMP Membership Reports.

If a switch receives an unregistered packet, it must forward that packet on all ports to which an IGMP router is attached. A switch may default to forwarding unregistered packets on all ports. Switches that do not forward unregistered packets to all ports must include a configuration option to force the flooding of unregistered packets on specified ports.

In an environment where IGMPv3 hosts are mixed with snooping switches that do not yet support IGMPv3, the switch's failure to flood unregistered streams could prevent v3 hosts from receiving their traffic. Alternatively, in environments where the snooping switch supports all of the IGMP versions that are present, flooding unregistered streams may cause IGMP hosts to be overwhelmed by multicast traffic, even to the point of not receiving Queries and failing to issue new membership reports for their own groups.

It is encouraged that snooping switches at least recognize and process IGMPv3 Join Reports, even if this processing is limited to the behavior for IGMPv2 Joins, i.e., is done without considering

any additional "include source" or "exclude source" filtering. When IGMPv3 Joins are not recognized, a snooping switch may incorrectly prune off the unregistered data streams for the groups (as noted above); alternatively, it may fail to add in forwarding to any new IGMPv3 hosts if the group has previously been joined as IGMPv2 (because the data stream is seen as already having been registered).

- 4) All non-IPv4 multicast packets should continue to be flooded out to all remaining ports in the forwarding state as per normal IEEE bridging operations.

This recommendation is a result of the fact that groups made up of IPv4 hosts and IPv6 hosts are completely separate and distinct groups. As a result, information gleaned from the topology between members of an IPv4 group would not be applicable when forming the topology between members of an IPv6 group.

- 5) IGMP snooping switches may maintain forwarding tables based on either MAC addresses or IP addresses. If a switch supports both types of forwarding tables then the default behavior should be to use IP addresses. IP address based forwarding is preferred because the mapping between IP multicast addresses and link-layer multicast addresses is ambiguous. In the case of Ethernet, there is a multiplicity of 1 Ethernet address to 32 IP addresses [RFC1112].
- 6) Switches which rely on information in the IP header should verify that the IP header checksum is correct. If the checksum fails, the information in the packet must not be incorporated into the forwarding table. Further, the packet should be discarded.
- 7) When IGMPv3 "include source" and "exclude source" membership reports are received on shared segments, the switch needs to forward the superset of all received membership reports on to the shared segment. Forwarding of traffic from a particular source *S* to a group *G* must happen if at least one host on the shared segment reports an IGMPv3 membership of the type INCLUDE(*G*, *Slist1*) or EXCLUDE(*G*, *Slist2*), where *S* is an element of *Slist1* and not an element of *Slist2*.

The practical implementation of the (*G*,*S1*,*S2*,...) based data forwarding tables are not within the scope of this document. However, one possibility is to maintain two (*G*,*S*) forwarding lists: one for the INCLUDE filter where a match of a specific (*G*,*S*) is required before forwarding will happen, and one for the EXCLUDE filter where a match of a specific (*G*,*S*) will result in no forwarding.

2.2. IGMP Snooping-Related Problems

A special problem arises in networks consisting of IGMPv3 routers as well as IGMPv2 and IGMPv3 hosts interconnected by an IGMPv2 snooping switch as recently reported [IETF56]. The router will continue to maintain IGMPv3 even in the presence of IGMPv2 hosts, and thus the network will not converge on IGMPv2. But it is likely that the IGMPv2 snooping switch will not recognize or process the IGMPv3 membership reports. Groups for these unrecognized reports will then either be flooded (with all of the problems that may create for hosts in a network with a heavy multicast load) or pruned by the snooping switch.

Therefore, it is recommended that in such a network, the multicast router be configured to use IGMPv2. If this is not possible, and if the snooping switch cannot recognize and process the IGMPv3 membership reports, it is instead recommended that the switch's IGMP snooping functionality be disabled, as there is no clear solution to this problem.

3. IPv6 Considerations

In order to avoid confusion, the previous discussions have been based on the IGMP protocol which only applies to IPv4 multicast. In the case of IPv6, most of the above discussions are still valid with a few exceptions that we will describe here.

The control and data forwarding rules in the IGMP section can, with a few considerations, also be applied to MLD. This means that the basic functionality of intercepting MLD packets, and building membership lists and multicast router lists, is the same as for IGMP.

In IPv6, the data forwarding rules are more straight forward because MLD is mandated for addresses with scope 2 (link-scope) or greater. The only exception is the address FF02::1 which is the all hosts link-scope address for which MLD messages are never sent. Packets with the all hosts link-scope address should be forwarded on all ports.

MLD messages are also not sent regarding groups with addresses in the range FF00::/15 (which encompasses both the reserved FF00::/16 and node-local FF01::/16 IPv6 address spaces). These addresses should never appear in packets on the link.

Equivalent to the IPv4 behaviors regarding the null IP Source address, MLD membership reports must not be rejected by an MLD snooping switch because of an unspecified IP source address (::). Additionally, if a non-Querier switch spoofs any General Queries (as

addressed in Section 2.1 above, for Spanning Tree topology changes), the switch should use the null IP source address (::) when sending said queries. When such proxy queries are received, they must not be included in the Querier election process.

The three major differences between IPv4 and IPv6 in relation to multicast are:

- The IPv6 protocol for multicast group maintenance is called Multicast Listener Discovery [MLDv2]. MLDv2 uses ICMPv6 message types instead of IGMP message types.
- The RFCs [IPV6-ETHER] and [IPV6-FDDI] describe how 32 of the 128 bit DIP addresses are used to form the 48 bit DMAC addresses for multicast groups, while [IPV6-TOKEN] describes the mapping for token ring DMAC addresses by using three low-order bits. The specification [IPV6-1394] makes use of a 6 bit channel number.
- Multicast router discovery is accomplished using the Multicast Router Discovery Protocol (MRDISC) defined in [MRDISC].

The IPv6 packet header does not include a checksum field. Nevertheless, the switch should detect other packet integrity issues such as address version and payload length consistencies if possible. When the snooping switch detects such an error, it must not include information from the corresponding packet in the MLD forwarding table. The forwarding code should instead drop the packet and take further reasonable actions as advocated above.


The fact that MLDv2 is using ICMPv6 adds new requirements to a snooping switch because ICMPv6 has multiple uses aside from MLD. This means that it is no longer sufficient to detect that the next-header field of the IP header is ICMPv6 in order to identify packets relevant for MLD snooping. A software-based implementation which treats all ICMPv6 packets as candidates for MLD snooping could easily fill its receive queue and bog down the CPU with irrelevant packets. This would prevent the snooping functionality from performing its intended purpose and the non-MLD packets destined for other hosts could be lost.

A solution is either to require that the snooping switch looks further into the packets, or to be able to detect a multicast DMAC address in conjunction with ICMPv6. The first solution is desirable when a configuration option allows the administrator to specify which ICMPv6 message types should trigger a CPU redirect and which should not. The reason is that a hardcoding of message types is inflexible for the introduction of new message types. The second solution introduces the risk that new protocols that use ICMPv6 and multicast

DMAC addresses could be incorrectly identified as MLD. It is suggested that solution one is preferred when the configuration option is provided. If this is not the case, then the implementor should seriously consider making it available since Neighbor Discovery messages would be among those that fall into this false positive case and are vital for the operational integrity of IPv6 networks.

The mapping from IP multicast addresses to multicast DMAC addresses introduces a potentially enormous overlap. The structure of an IPv6 multicast address is shown in the figure below. As a result, there are 2 ** (112 - 32), or more than 1.2e24 unique DIP addresses which map into a single DMAC address in Ethernet and FDDI. This should be compared to 2**5 in the case of IPv4.

Initial allocation of IPv6 multicast addresses, as described in [RFC3307], however, cover only the lower 32 bits of group ID. While this reduces the problem of address ambiguity to group IDs with different flag and scope values for now, it should be noted that the allocation policy may change in the future. Because of the potential overlap it is recommended that IPv6 address based forwarding is preferred to MAC address based forwarding.



4. IGMP Questionnaire

As part of this work, the following questions were asked on the MAGMA discussion list and were sent to known switch vendors implementing IGMP snooping. The individual contributions have been anonymized upon request and do not necessarily apply to all of the vendors' products.

The questions were:

- Q1 Do your switches perform IGMP Join aggregation? In other words, are IGMP joins intercepted, absorbed by the hardware/software so that only one Join is forwarded to the querier?
- Q2 Is multicast forwarding based on MAC addresses? Would datagrams addressed to multicast IP addresses 224.1.2.3 and 239.129.2.3 be forwarded on the same ports-groups?

- Q3 Is it possible to forward multicast datagrams based on IP addresses (not routed)? In other words, could 224.1.2.3 and 239.129.2.3 be forwarded on different port-groups with unaltered TTL?
- Q4 Are multicast datagrams within the range 224.0.0.1 to 224.0.0.255 forwarded on all ports whether or not IGMP Joins have been sent?
- Q5 Are multicast frames within the MAC address range 01:00:5E:00:00:01 to 01:00:5E:00:00:FF forwarded on all ports whether or not IGMP joins have been sent?
- Q6 Does your switch support forwarding to ports on which IP multicast routers are attached in addition to the ports where IGMP Joins have been received?
- Q7 Is your IGMP snooping functionality fully implemented in hardware?
- Q8 Is your IGMP snooping functionality partly software implemented?
- Q9 Can topology changes (for example spanning tree configuration changes) be detected by the IGMP snooping functionality so that for example new queries can be sent or tables can be updated to ensure robustness?

The answers were:

	Switch Vendor					
	1	2	3	4	5	6
Q1 Join aggregation	x	x	x		x	x
Q2 Layer-2 forwarding	x	x	x	x	(1)	
Q3 Layer-3 forwarding	(1)		(1)		(1)	x
Q4 224.0.0.X aware	(1)	x	(1)	(2)	x	x
Q5 01:00:5e:00:00:XX aware	x	x	x	(2)	x	x
Q6 Mcast router list	x	x	x	x	x	x
Q7 Hardware implemented						
Q8 Software assisted	x	x	x	x	x	x
Q9 Topology change aware	x	x	x	x		(2)

x Means that the answer was Yes.

- (1) In some products (typically high-end) Yes; in others No.
- (2) Not at the time that the questionnaire was received but expected in the near future.

5. References

5.1. Normative References

- [BRIDGE] IEEE Std. 802.1D-2004 IEEE Standard for Local and metropolitan area networks, Media Access Control (MAC) Bridges
- [IGMPv3] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [IPV6-1394] Fujisawa, K. and A. Onoe, "Transmission of IPv6 Packets over IEEE 1394 Networks", RFC 3146, October 2001.
- [IPV6-ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [IPV6-FDDI] Crawford, M., "Transmission of IPv6 Packets over FDDI Networks", RFC 2467, December 1998.
- [IPV6-TOKEN] Crawford, M., Narten, T., and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", RFC 2470, December 1998.
- [MLD] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [MLDv2] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [MRDISC] Haberman, B. and J. Martin, "Multicast Router Discovery", RFC 4286, December 2005.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, August 2002.

5.2. Informative References

- [CISCO] Cisco Tech Notes, "Multicast In a Campus Network: CGMP and IGMP snooping", <http://www.cisco.com/warp/public/473/22.html>
- [IETF56] Briefing by Dave Thaler, Microsoft, presented to the MAGMA WG at the 56'th IETF meeting in San Francisco, <http://www.ietf.org/proceedings/03mar/index.html>
- [MSOFT] Microsoft support article Q223136, "Some LAN Switches with IGMP Snooping Stop Forwarding Multicast Packets on RRAS Startup", <http://support.microsoft.com/support/articles/Q223/1/36.ASP>

6. Security Considerations

Under normal network operation, the snooping switch is expected to improve overall network performance by limiting the scope of multicast flooding to a smaller portion of the local network. In the event of forged IGMP messages, the benefits of using a snooping switch might be reduced or eliminated.

Security considerations for IGMPv3 at the network layer of the protocol stack are described in [IGMPv3]. The introduction of IGMP snooping functionality does not alter the handling of multicast packets by the router as it does not make use of link layer information.

There are, however, changes in the way that the IGMP snooping switch handles multicast packets within the local network. In particular:

- A Query message with a forged source address which is less than that of the current Querier could cause snooping switches to forward subsequent Membership reports to the wrong network interface. It is for this reason that IGMP Membership Reports should be sent to all multicast routers as well as the current Querier.
- It is possible for a host on the local network to generate Current-State Report Messages that would cause the switch to incorrectly believe that there is a multicast listener on the same network segment as the originator of the forged message. This will cause unrequested multicast packets to be forwarded into the network segments between the source and the router. If the router requires that all Multicast Report messages be authenticated as described in section 9.4 of [IGMPv3], it will discard the forged Report message from the host inside the network in the same way

that it would discard one which originates from a remote location. It is worth noting that if the router accepts unauthenticated Report messages by virtue of them having arrived over a network interface associated with the internal network, investigating the affected network segments will quickly narrow the search for the source of the forged messages.

- As noted in [IGMPv3], there is little motivation for an attacker to forge a Membership report message since joining a group is generally an unprivileged operation. The sender of the forged Membership report will be the only recipient of the multicast traffic to that group. This is in contrast to a shared LAN segment (HUB) or network without snooping switches, where all other hosts on the same segment would be unable to transmit when the network segment is flooding the unwanted traffic.

The worst case result for each attack would remove the performance improvements that the snooping functionality would otherwise provide. It would, however, be no worse than that experienced on a network with switches that do not perform multicast snooping.

7. Acknowledgements

We would like to thank Martin Bak, Les Bell, Yiqun Cai, Ben Carter, Paul Congdon, Toerless Eckert, Bill Fenner, Brian Haberman, Edward Hilquist, Hugh Holbrook, Kevin Humphries, Isidor Kouvelas, Pekka Savola, Suzuki Shinsuke, Jaff Thomas, Rolland Vida, and Margaret Wasserman for comments and suggestions on this document.

Furthermore, the following companies are acknowledged for their contributions: 3Com, Alcatel, Cisco Systems, Enterasys Networks, Hewlett-Packard, Vitesse Semiconductor Corporation, Thrane & Thrane. The ordering of these names do not necessarily correspond to the column numbers in the response table.

Authors' Addresses

Morten Jagd Christensen
Thrane & Thrane
Lundtoftegaardsvej 93 D
2800 Lyngby
DENMARK

EEmail: mjc@tt.dk

Karen Kimball
Hewlett-Packard
8000 Foothills Blvd.
Roseville, CA 95747
USA

EEmail: karen.kimball@hp.com

Frank Solensky
Calix
43 Nanog Park
Acton, MA 01720
USA

EEmail: frank.solensky@calix.com

Full Copyright Statement

Copyright (C) The Internet Society (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 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 provided by the IETF Administrative Support Activity (IASA).