

Cisco Systems
Router-port Group Management Protocol (RGMP)

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 (2003). All Rights Reserved.

Abstract

This document describes the Router-port Group Management Protocol (RGMP). This protocol was developed by Cisco Systems and is used between multicast routers and switches to restrict multicast packet forwarding in switches to those routers where the packets may be needed.

RGMP is designed for backbone switched networks where multiple, high speed routers are interconnected.

1. Conventions used in this document

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 BCP 14, RFC 2119 [2].

2. Introduction

IGMP Snooping is a popular, but not well documented mechanism to restrict multicast traffic, in switched networks, to those ports that want to receive the multicast traffic. It dynamically establishes and terminates multicast group specific forwarding in switches that support this feature.

The main limitation of IGMP Snooping is that it can only restrict multicast traffic onto switch ports where receiving hosts are connected directly or indirectly via other switches. IGMP Snooping can not restrict multicast traffic to ports where at least one multicast router is connected. It must instead flood multicast traffic to these ports. Snooping on IGMP messages alone is an intrinsic limitation. Through it, a switch can only learn which multicast flows are being requested by hosts. A switch can not learn through IGMP which traffic flows need to be received by router ports to be routed because routers do not report these flows via IGMP.

In situations where multiple multicast routers are connected to a switched backbone, IGMP Snooping will not reduce multicast traffic load. Nor will it assist in increasing internal bandwidth through the use of switches in the network.

In switched backbone networks or exchange points, where predominantly routers are connected with each other, a large amount of multicast traffic may lead to unexpected congestion. It also leads to more resource consumption in the routers because they must discard the unwanted multicast traffic.

The RGMP protocol described in this document restricts multicast traffic to router ports. To effectively restrict traffic, it must be supported by both the switches and the routers in the network.

The RGMP message format resembles the IGMPv2 message format so that existing switches, capable of IGMP Snooping, can easily be enhanced with this feature. Its messages are encapsulated in IPv4 datagrams, with a protocol number of 2, the same as that of IGMP. All RGMP messages are sent with TTL 1, to destination address 224.0.0.25. This address has been assigned by IANA to carry messages from routers to switches [3].

RGMP is designed to work in conjunction with multicast routing protocols where explicit join/prune to the distribution tree is performed. PIM-SM [4] is an example of such a protocol.

The RGMP protocol specifies operations only for IP version 4 multicast routing. IP version 6 is not considered.

To keep RGMP simple, efficient and easy to implement, it is designed for switches to expect RGMP messages from only one source per port. For this reason, RGMP only supports a single RGMP enabled router to be connected directly to a port of an RGMP enabled switch. Such a topology should be customary when connecting routers to backbone switches and thus not pose a limitation on the deployment of RGMP.

All RGMP messages have the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type           |   Reserved   |           Checksum           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                   Group Address                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The reserved field in the message MUST be transmitted as zeros and ignored on receipt.

2.1 Type

There are four types of RGMP messages of concern to the router-switch interaction. The type codes are defined to be the highest values in an octet to avoid the re-use of already assigned IGMP type codes.

```

0xFF = Hello
0xFE = Bye
0xFD = Join a group
0xFC = Leave a group

```

Unrecognized message types should be silently ignored.

Note:

RGMP and the IANA assignment of address 224.0.0.25 for it predates RFC 3228 [9]. RGMP defines Type values which in RFC 3228 are assigned to protocol testing and experimentation. This is not an operational issue for RGMP itself because only RGMP packets use the IPv4 destination address 224.0.0.25. The Type values defined above are thus ONLY valid in conjunction with the RGMP destination address.

2.2. Checksum

Checksum covers the RGMP message (the entire IPv4 payload). The algorithm and handling of checksum are the same as those for IGMP messages as described in RFC 3376 [5].

2.3. Group Address

In an RGMP Hello or Bye message, the group address field is set to zero.

In an RGMP Join or Leave message, the group address field holds the IPv4 multicast group address of the group being joined or left.

2.4 IPv4 header

RGMP messages are sent by routers to switches. The source IPv4 address of an RGMP packet is the sending interface IPv4 address of the originating router. The destination IPv4 address of an RGMP packet is 224.0.0.25. Switches supporting RGMP need to listen to packets to this group.

3. RGMP Protocol Description

3.1 RGMP Router side Protocol Description

Backbone switches use RGMP to learn which groups are desired at each of their ports. Multicast routers use RGMP to pass such information to the switches. Only routers send RGMP messages. They ignore received RGMP messages.

A Router enabled for RGMP on an interface periodically [Hello Interval] sends an RGMP Hello message to the attached network to indicate that it is RGMP enabled. When RGMP is disabled on a routers interface, it will send out an RGMP Bye message on that interface, indicating that it again wishes to receive IPv4 multicast traffic promiscuously from that interface.

When an interface is RGMP enabled, a router sends an RGMP Join message out through this interface to each group that it wants to receive traffic for from the interface. The router needs to periodically [Join Interval] re-send an RGMP Join for a group to indicate its continued desire to receive multicast traffic.

Routers supporting RGMP MUST NOT send RGMP Join or Leave messages for groups 224.0.0.x (x=0...255), 224.0.1.39 and 224.0.1.40. The latter two are known as cisco-rp-announce and cisco-rp-discovery [3].

When a router no longer needs to receive traffic for a particular group, it sends an RGMP Leave message for the group. For robustness, the router MAY send more than one such message.

If IPv4 multicast packets for an undesired group are received at a router from a switch, the router MAY send a RGMP Leave message for that group to the switch. These messages are called data-triggered RGMP Leave messages and the router SHOULD rate-limit them. The router MAY suppress sending a data triggered RGMP Leave message if it has a desired group that has the same MAC destination address as the undesired group. (See RFC 1112 [6] for MAC ambiguity.) Such suppression of data triggered RGMP Leave messages SHOULD be configurable if supported.

3.2 RGMP Switch side Protocol Description

A switch enabled for RGMP on a network consumes RGMP messages received from ports of the network and processes them as described below. If enabled for RGMP, the switch must NOT forward/flood received RGMP messages out to other ports of the network.

RGMP on a switch operates on a per port basis, establishing per-group forwarding state on RGMP enabled ports. A port reverts into an RGMP enabled port upon receipt of an RGMP Hello message on the port, and a timer [5 * Hello Interval] is started. This timer is restarted by each RGMP Hello message arriving on the port. If this timer expires or if it is removed by the arrival of an RGMP Bye message, then the port reverts to its prior state of multicast traffic forwarding.

Correct deployment of RGMP is one RGMP enabled router directly connected to a port on a switch that supports RGMP. The port on the switch MAY want to keep track of the IPv4 originator address of the RGMP Hello and Bye messages it receives on that port. In the event it receives multiple IPv4 originating addresses in RGMP messages on one port, the switch MAY generate an alert to notify the administrator. The switch MAY also have a configuration option that will allow for the operator to disable RGMP and have the switch fall back to flooding IPv4 multicast on that port, although this is a potentially dangerous option.

By default, connecting two or more RGMP enabled routers to a switch port will cause intermittent black holing of IPv4 multicast traffic towards these routers. Black holing occurs when a RGMP Leave is received from one router while the other router is still joined.

This malfunction is not only easily recognized by the actual users connected through the routers, but it also adheres to the principle that a failure situation causes less traffic than more. Reverting to flooding easily maintains the illusion that everything is working perfectly. The exception is that the traffic constraining benefits

of RGMP are not realized. This suggests that congestion happens at a much later time than the misconfiguration and can then not easily be correlated with the cause anymore.

Because routers supporting RGMP are not required to send RGMP Join or Leave messages for groups 224.0.0.x (x=0...255), 224.0.1.39 and 224.0.1.40, RGMP enabled ports always need to receive traffic for these groups. Traffic for other groups is initially not forwarded to an RGMP enabled port.

RGMP Join and Leave messages are accepted if they arrive on an RGMP enabled port, otherwise they will be discarded. Upon acceptance of an RGMP Join message, the switch MUST start forwarding traffic for the group to the port. Upon acceptance of an RGMP Leave message, the switch SHOULD stop forwarding traffic for the group to that port. The switch's ability to stop forwarding traffic for a group may be limited, for example, because of destination MAC based forwarding in the switch. Therefore, it is necessary for the switch to always forward traffic for all MAC-ambiguous IPv4 multicast groups (see [6] for MAC-ambiguity).

To stop forwarding of traffic to a group in the event of lost RGMP Leave message(s), a switch MAY time out RGMP forwarding state on a port for a group [5 * Join Interval] after the last RGMP Join for that group has been received on the port.

Without any layer 2 IPv4 multicast filtering methods running, a switch needs to flood multicast traffic to all ports. If a switch does actually run one or more mechanisms beside RGMP to filter IPv4 multicast traffic, such as IGMP snooping [10], then by default it will not flood IPv4 multicast traffic to all ports anymore. Instead, the switch will try to determine which ports still needs to receive all IPv4 multicast traffic by default, and which ports do not.

Compliance with this specification requires that a switch MUST be able to elect a port for flooding through the presence of PIM Hello messages [4] arriving from the port and also through a manual configuration option. In addition, the switch SHOULD recognize a port connected to a router by other appropriate protocol packets or dedicated IPv4 multicast router discovery mechanisms such as MRDISC [11]. The manual configuration is required to support routers not supporting PIM or other methods recognized by the switch.

Further mechanisms for IPv4 multicast traffic restriction may also be used on RGMP enabled ports. In this case, forwarding for a group on the port must be established if either mechanism requires it, and it must only be removed if no mechanism requires it anymore.

4. Operational Notes

4.1. Support for networks with multiple switches

To be simple to implement on switches and resilient in face of potential layer 2 network topology changes, RGMP does not specify how to restrict multicast traffic on links connecting switches amongst each other. With just RGMP being used, multicast traffic will thus be flooded on inter-switch links within a network if at least one router is connected to each of the switches.

This happens implicitly because the switch will not flood/forward received RGMP messages out to the inter-switch link and thus the switch on the other end will only recognize the port as a router port via the PIM Hello messages flooded by the switches. Manual configuration for inter-switch links may be required if non-PIM routers are being used, depending on the other capabilities of the switch.

If appropriate, a switch can send out RGMP messages on ports to make it look like an RGMP enabled router to a potential switch at the other end of the link. This would constrain IPv4 multicast traffic between switches, but this type of "RGMP Spoofing" by the switch is outside the scope of this specification.

4.2. Interoperability with RGMP-incapable routers

Since RGMP messages received at a switch only affect the state of their ingress ports, the traffic restriction is applied there only. RGMP-incapable routers will receive multicast traffic for all multicast groups.

4.3. RGMP and multicast routing protocols

One result of the simplicity of RGMP are its restrictions in supporting specific routing protocols. The following paragraphs list a few known restrictions.

A router running RGMP on a switched network will not receive traffic for a multicast group unless it explicitly requests it via RGMP Join messages (besides those group ranges specified to be flooded in 3.1). For this reason, it is not possible to run a protocol like PIM Dense-Mode or DVMRP across an RGMP enabled network with RGMP enabled routers.

In Bidir-PIM, a router elected to be the DF must not be enabled for RGMP on the network, because it unconditionally needs to forward traffic received from the network towards the RP. If a router is not the DF for any group on the network, it can be enabled for RGMP on that network.

In PIM-SM, directly connected sources on the network can not be supported if the elected DR is running RGMP, because this DR needs to unconditionally receive traffic from directly connected sources to trigger the PIM-SM registering process on the DR. In PIM-SSM, directly connected sources can be supported with RGMP enabled routers.

Both in PIM-SM and PIM-SSM, upstream routers forwarding traffic into the switched network need to send RGMP Joins for the group in support of the PIM assert process.

5. List of timers and default values

5.1. Hello Interval

The Hello Interval is the interval between RGMP Hello messages sent by an RGMP-enabled router to an RGMP-enabled switch. Default: 60 seconds.

5.2. Join Interval

The Join Interval is the interval between periodic RGMP Join messages sent by an RGMP-enabled router to an RGMP-enabled switch for a given group address. Default: 60 seconds.

6. Security Considerations

The RGMP protocol assumes that physical port security can be guaranteed for switch ports from which RGMP messages are accepted. Physical port security for RGMP means that physical measures will ensure that such ports are dedicatedly connected to one system which acts as an RGMP capable router. This is also the recommended configuration to best leverage the benefits of the RGMP protocol (e.g., avoiding unwanted third-party IPv4 multicast traffic arriving on said ports).

RGMP specific DoS attacks arise from forged RGMP messages. If more than one system is connected to a port of the RGMP switch, then one system may forge RGMP messages and affect the operations of the other system(s) on the same port. This is a potential security risk.

When physical security ensures that only one system is connected to a RGMP capable port on a switch, then forged messages from this system itself can take effect. Such forged messages can always be avoided by system local measures.

We consider the ramifications of a forged message of each type:

Hello Message:

A forged RGMP Hello message can restrict multicast data towards a non-RGMP enabled router on the same port. This effectively introduces a blackholing DoS attack.

Leave Message:

A forged RGMP Leave message can restrict IPv4 multicast traffic for individual groups toward the port. The effect is a possible blackholing DoS attack similar to an RGMP Hello Message except that it does not affect all IPv4 multicast traffic but only that of the groups indicated in the forged messages. It will also only affect a port if there officially is only one RGMP enabled router connected to it (i.e., if the port is RGMP enabled).

Bye Message:

A forged RGMP Bye message can turn the port into being RGMP-disabled. This could, indirectly, cause a DoS attack based on the port getting overloaded with IPv4 multicast traffic if the network bandwidth of the port was provisioned with the expectation that RGMP will suppress unwanted IPv4 multicast messages.

This type of DoS attack simply re-establishes a port behavior as if RGMP was not configured and invalidates the benefit of RGMP. This, however, does not introduce an issue that would not have been there without RGMP in the first place.

Join Message:

A forged RGMP Join message could attract undesired multicast packets to the port where it is received from. The effect is similar to an RGMP Bye Message except that it does not affect all IPv4 multicast traffic only the groups indicated in the forged messages. The message will affect a port only if there officially is only one RGMP enabled router connected to it (i.e., if the port is RGMP enabled).

7. Normative References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [4] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P. and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", RFC 2362, June 1998.
- [5] Cain, B., Deering, S., Kouvelas, I., Fenner, W. and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [6] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, August 1989.
- [7] ANSI/IEEE Std 802.1D 1998 Edition, "Media Access Control (MAC) Bridges", 1998.

8. Informative References

- [3] Internet Multicast Addresses,
<http://www.iana.org/assignments/multicast-addresses>
- [8] Farinacci D., Tweedly D., Speakman T., "Cisco Group Management Protocol (CGMP)", 1996/1997
<ftp://ftpeng.cisco.com/ipmulticast/specs/cgmp.txt>
- [9] Fenner, B., "IANA Considerations for IPv4 Internet Group Management Protocol (IGMP)", RFC 3228, February 2002.
- [10] Christensen, M. and F. Solensky, "IGMP and MLD snooping switches", Work In Progress.
- [11] Biswas, S., Cain, B. and B. Haberman, "IGMP Multicast Router Discovery", Work In Progress.

9. Acknowledgments

The authors would like to thank Gorrry Fairhurst, Bill Fenner, Giovanni Meo, Mike Norton, Pavlin Radoslavov and Alex Zinin for their review of the document and their suggestions.

Appendix A. Intellectual Property Rights

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

Appendix B. Comparison with GARP/GMRP

This appendix is not part of the RGMP specification but is provided for information only.

GARP/GMRP (defined in IEEE 802.1D [7]) is the ANSI/ISO/IEC/IEEE protocol suite to constrain ethernet multicast traffic in bridged ethernet networks. As such it is also a possible alternative to RGMP for the purpose of constraining multicast traffic towards router ports. This appendix will explain the motivation not to rely on GARP/GMRP and how GARP/GMRP and RGMP differ.

The key factor in rolling out GARP/GMRP would have been to completely replace IGMP Snooping. This was the design goal of GARP/GMRP. For efficient operations, IGMP Snooping requires hardware filtering support in the switch (to differentiate between hosts membership reports and actual IPv4 multicast traffic). Especially in many older switches this support does not exist. Vendors tried to find a way around this issue to provide the benefit of constraining IPv4 multicast traffic in a switched LAN without having to build more expensive switch hardware. GARP/GMRP is one protocol resulting from this. CGMP from Cisco is another one. While CGMP solves the problem without requiring changes to the host stack software, GARP/GMRP requires support for it by the host stack.

Up to date GARP/GMRP has so far not made significant inroads into deployed solutions. IGMP Snooping (and CGMP) are the norm for this environment. In result, GARP/GMRP can not necessarily be expected to be supported by layer 2 switches. In addition, GARP/GMRP does not address clearly the issues RGMP tries to solve. On one hand, GARP/GMRP provides much more functionality and as such complexity as immediately required. On the other hand, GARP/GMRP is limited by being a standard predominantly for the Ethernet scope.

Beyond the process and applicability reasons, the main differences between GARP/GMRP and RGMP are as follows:

- o GARP/GMRP switches/systems need to send and listen/react to GARP/GMRP messages. In RGMP, routers only need to send RGMP messages and switches only need to listen to them. This protocol approach is meant to simplify implementation, operations and troubleshooting.
- o The same switch running RGMP in a backbone network will likely see more states than running on the edge only doing IGMP Snooping, making it preferable to keep the amount of per group processing and memory requirements in RGMP more in bounds than possible in IGMP Snooping and GARP/GMRP: In GARP/GMRP, a (multiple) timer based state-machines needs to be maintained on a per ethernet group address, in RGMP timer maintenance is completely optional and there are only two states per group (joined or not joined).
- o GARP/GMRP is an ethernet level protocol from the IEEE. It supports to constrain traffic for ethernet addresses (groups). RGMP does constrain traffic for IPv4 multicast groups. Today this is even beyond the capabilities of typical switch platforms used as layer2 switches. Extensions to support further entities are likely easier to come by through extensions to RGMP than to GARP/GMRP.
- o RGMP shares the basic packet format with IGMP (version 2) and is as such easy to add to router and switch platforms that already support IGMP and IGMP Snooping respectively. This is especially true for switches that in hardware can differentiate between IGMP protocol type packets and other IPv4 multicast traffic sent to the same (or a MAC ambiguous) group. In addition, due to the state simplicity of RGMP it is easy to integrate IGMP Snooping and RGMP operations in the IPv4 multicast control and forwarding plane of a switch.
- o GARP/GMRP supports more than one system (host/router) on a switch port which is one reason for its complexity. In RGMP, this configuration is explicitly not supported: More than one router per switched port is not only not a common scenario in today's switches layer 2 networks, it is also an undesired configuration when unwanted IPv4 multicast traffic is to be kept away from routers.

- o GARP/GMRP defines how to constrain multicast traffic between switches, another reason for its complexity. RGMP does not explicitly support this as part of the protocol because of the following reasons:
 - o It is not necessary to include this function as part of the RGMP protocol description because switch implementations can transparently decide to support this function (see 4.1 about this "RGMP Spoofing").
 - o Important deployments through which large amounts of IPv4 multicast are moved today are typically single switch MIX - Multicast Internet eXchange points.
 - o Avoiding congestion on inter-switch links in general is more complex than simply constraining IPv4 multicast traffic to paths where it is needed. With or without IPv4 multicast, the aggregate bandwidth needed between switches can easily be the aggregate required bandwidth to routers on either sides. For this reason, inter-switch bandwidth is most often appropriately over provisioned. In addition, the likelihood for receiving routers to be only on the sources side of an inter-switch link is in general deployments rather low. The cases where traffic constraint on inter-switch links is required and helpful is thus limited and can in most cases be avoided or worked around. Moving the network to a layer 3 routed network is often the best solution, supporting RGMP-Spoofing (see section 4.1) is another one.

Appendix C. Possible future extensions / comparison to PIM Snooping

This appendix is not part of the RGMP specification but is provided for information only.

This appendix presents a discussion of possible extensions to RGMP. Included are points on why the extensions are not included and in addition a motivation for RGMP in comparison to (PIM) snooping.

- o Support for multiple switches

As discussed in "RGMP Spoofing", chapter 4.1 and GARP/GMRP comparison in Appendix B.

- o Support for SSM

While RGMP works with PIM-SSM, it does not have explicit messages for the router to selectively join to (S,G) channels individually. Instead the router must RGMP join to all (Si,G) channels by

joining to G. Extending RGMP to include (S,G) Join/Leaves is feasible. However, currently the majority of switches do not support actual traffic constraining on a per channel basis. In addition, the likelihood for actual channel collision (two SSM channels using the same group) will only become an issue when SSM is fully deployed.

- o Support for IPv6

RGMP could easily be extended to support IPv6 by mapping the RGMP packet format into the MLD/IPv6 packet format. This was not done for this specification because most switches today do not even support MLD snooping.

- o Support for multiple routers per port

As discussed in Appendix B. This is probably one extension that should be avoided. Multiple RGMP router per port are inappropriate for efficient multicast traffic constraintment.

- o Support for non-join based protocols / protocol elements

For protocols like PIM dense-mode, DVMRP or Bidir-PIM DF routers, additional RGMP messages may be added to allow routers to indicate that certain group (ranges) traffic need to be flooded from (dense-mode) or to (Bidir-PIM) them.

- o Support for multi-policy switching

In Multicast Exchange Points (MIXes) environments situations exist where different downstream routers for policy reasons need to receive the same traffic flow from different upstream routers.

This problem could be solved by actually providing an upstream neighbor field in RGMP Join/Leave messages. The RGMP switch would then forward traffic from one upstream router only to those downstream routers who want to have the traffic from exactly this upstream router. This extension would best go in hand with changes to the layer 3 routing protocol run between the routers.

As previously mentioned, RGMP was designed to be easy to implement and to support simple layer2 switches. Implementations could also be applied to switches beyond layer 2. If all the above possible future extensions were to be supported by an evolution of RGMP, it would be questionable whether such a protocol could be any less complex than actually snooping into the layer3 IPv4 routing protocol run between routers in a switched LAN.

From the perspective of protocol architecture it is certainly more appropriate to have a separate protocol like RGMP or GARP/GMRP for this purpose. Then again, the more complex the requirements are, the more duplication of effort is involved and snooping seems to become a more attractive option.

Even though there exists one predominant routing Protocol, PIM, in IPv4 multicast, routing with PIM in itself is extremely complex for a switch to snoop into. PIM has two main versions, different modes - sparse, dense, bidir, ssm, join / prune / graft messages (depending on the mode of the group), various PIM Hello options, different versions of asserts, two dynamic mode announcement protocols (BSR, AutoRP), and finally supports both IPv4 and IPv6.

A switch snooping into PIM is very likely to implement just a subset of this feature set, making it very hard for the user to determine what level of actual traffic constraint is achieved unless a clear specification exists for the implementation (or better the method per se.). In addition, there is always the danger that such a snooping implementation may break newer features of the routing protocol that it was not designed to handle (likely because they could not have been predicted). For example, this can happen with switches using IGMP (v2) snooping implementations that are being subjected to IGMP version 3 messages - they break IGMPv3.

In summary, with PIM still evolving, the approach taken by RGMP is the safest one for the immediate problems at hand, and extensions like those listed should be considered in time for actual demand. (PIM) snooping is a valid alternative once the total amount of features that need to be supported makes it an equally attractive solution (with respect to complexity) to a dedicated protocol and if its functions are well defined to allow predicting its effects - but always at the price of possible incompatibilities with upcoming PIM protocol extensions unless support for layer 2 switches is explicitly considered in moving PIM protocols forward.

Authors' Addresses

Ishan Wu
cisco Systems
170 West Tasman Drive
San Jose, CA 95134

Phone: (408) 526-5673
EMail: iwu@cisco.com

Toerless Eckert
cisco Systems
170 West Tasman Drive
San Jose, CA 95134

Phone: (408) 853-5856
Email: eckert@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

