T. Pusateri
Juniper Networks

INTERNET DRAFT
draft-ietf-idmr-dvmrp-v3-as-01
July 2000
Expires: January 26, 2001

Distance Vector Multicast Routing Protocol Applicability Statement

Status of this Memo

Abstract

This document provides a framework for the use of Distance Vector Multicast Routing Protocol (DVMRP) Version 3 within a multicast routing domain. It is an interior gateway protocol designed to be used within an autonomous system.

1. Intended use

DVMRP [Pusa99] can be characterized as a broadcast and prune multicast routing protocol. This means that IP multicast datagrams are forwarded to all possible receivers and then unwanted data traffic is pruned back to the minimum tree necessary to reach all of the current receivers.

This provides a data driven mechanism for all last hop routers to learn about all active multicast sources. If the last hop routers have local group members for the multicast group associated with the source, they simply do nothing and the multicast data will continue to arrive.

If the last hop routers do not have corresponding group members for the active source, then they send control messages back up the multicast delivery tree toward the source to prune their branch from the tree.

The prunes sent upstream toward the source time out and are periodically refreshed to reduce the amount of state that is stored. The periodic broadcast and prune nature of the protocol makes it well suited for use within a multicast domain or autonomous system where bandwidth is plentiful. It is NOT intended for use as an EGP across multicast domains or for interconnecting multicast domains.

DVMRP's advantages compared with other multicast routing protocols include:

a. Pure source specific multicast distribution trees provide a simple model to deploy and troubleshoot.

b. Join latency is minimized since new sources are automatically sent to all receivers.

c. DVMRP builds a broadcast tree per source network with the routing updates, so it doesn't have to flood links that are not part of the broadcast tree for a source.

d. DVMRP uses its own topology discovery mechanism allowing both faster adaptation to change and a more stable steady state.

e. Discovery mechanisms relying on expanding ring searches are fully supported.

Of course, these benefits do not come without a cost. The broadcast nature of DVMRP makes it unsuitable for connecting leaf domains via low speed links. Additionally, the cost of keeping prune state for unwanted data can be high depending on the mix of receivers.

Additionally, DVMRP operates on the premise that all data will be broadcast throughout the domain and unwanted data will be pruned back toward the source. If a DVMRP domain is connected downstream from an explicit join multicast domain, then without additional mechanism, no data would ever be broadcast into the DVMRP domain from the explicit join domain. For this purpose, Domain Wide Reports [Fenn00] were invented to inform border routers between explicit join domains and broadcast and prune domains about the active group membership within the broadcast and prune domain.

2. Scalability/Applicability

A routing protocol need only scale within the limits of its intended use. In this case, DVMRP need only scale within the context of multicast routing domain or autonomous system.

There are several factors that limit the scalability of DVMRP. Fortunately, these factors are well understood and by making these explicit within this applicability statement, it should be possible to avoid negative side-effects.

Convergence time

Since DVMRP uses a distance vector (also called Bellman-Ford after its authors) routing algorithm to distribute source information, its scope is limited by the convergence time required to propagate updated source information across the routing domain. In order to limit the convergence time, DVMRP has imposed a limit on the number of hops across a domain by setting the infinity metric to 32 hops.

Count to Infinity

All distance vector protocols are susceptible to the well known "count to infinity" problem [Perl92]. However, by requiring the use of split horizon with poison reverse, the chances of encountering this are significantly lowered.

3. State Analysis

As stated earlier, DVMRP broadcasts new multicast data to all receivers and then prunes back the unwanted data based on group membership information. To further reduce the amount of broadcasted data, DVMRP builds the broadcast tree by determining the single reverse path from a receiver back to the source and pruning duplicate paths.

This truncated broadcast tree is precalculated using poison reverse route updates as each router participates in a designated forwarder election per source network to determine the upstream router.

This mechanism prevents duplicate datagrams from arriving on multi-access networks at the expense of more state in the router. DVMRP requires designated forwarder election state to be kept per prefix per interface and prune state to be kept per neighbor per (group, source prefix) pair.

4. Control Traffic Analysis

Independent of the amount of multicast data being forwarded, DVMRP will always send some control traffic. First, there is a DVMRP probe message sent every 10 seconds per interface for neighbor discovery and keep-alive functions. Additionally, there are periodic route exchanges that advertise source networks and assist in building the multicast delivery trees. These route reports are resent every 60 seconds but are spread out across the report interval to reduce the processing load. The number of route report packets sent is dependent on the number of prefixes being reported.

The amount of prune and graft messages sent is a function of the number of active senders and receivers for the data being forwarded. Grafts are only sent to undo the effects of previous prunes. Prunes are sent in response to unwanted multicast data. The lifetime of a prune which is specific to a group and source network is approximately 2 hours. This long lifetime greatly limits the amount of prune control traffic that is sent.

5. References

[Fenn00] Fenner, W., "Domain Wide Multicast Group Membership Reports", Work in Progress, July 2000.

[Perl92] Perlman, R., "Interconnections: Bridges and Routers", Addison-Wesley, 1992, pp. 210-211.

[Pusa99] Pusateri, T., "Distance-Vector Multicast Routing Protocol Version 3", Work In Progress, September 1999.

6. Author's Address

Thomas Pusateri
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: (408) 734-7690
EMail: pusateri@juniper.net

7. Acknowledgments

The author would like to acknowledge Dave Thaler, Bill Fenner, Thomas Maufer, and Ravi Shekhar for their review and comments.

8. Full Copyright Statement

Table of Contents

1. Intended Use 2
2. Scalability/Applicability 2
3. State Analysis 3
4. Control Traffic Analysis 3
5. References 4
6. Author's Address 4
7. Acknowledgments 4
8. Full Copyright Statement 5