August 2022

MULTICAST OPERATIONS, ADMINISTRATION, AND MANAGEMENT (OAM) TECHNIQUES UTILIZING PROTOCOL INDEPENDENT MULTICAST (PIM) FLOODING MECHANISMS

Ramakrishnan Chokkanathapuram Sundaram

Stig Venaas

Abhiram Srinivasan

MULTICAST OPERATIONS, ADMINISTRATION, AND MANAGEMENT (OAM) TECHNIQUES UTILIZING PROTOCOL INDEPENDENT MULTICAST (PIM) FLOODING MECHANISMS

AUTHORS:
Ramakrishnan Chokkanathapuram Sundaram
Stig Venaas
Abhiram Srinivasan

ABSTRACT

Multicast networks are often complex and to provide a visualization of traffic flows within a multicast network often involves the full knowledge of a distribution tree for the network. Further, isolating problems within a multicast network can involve tracing of multiple nodes across the distribution tree. Techniques presented provide efficient multicast tree discovery through Protocol Independent Multicast (PIM) flooding mechanisms, which can be further used to facilitate network visualizations and fault isolation within a network.

DETAILED DESCRIPTION

Providing flow visualizations and fault isolation within multicast networks can be challenging. However, network operators deploying complex Internet Protocol (IP) fabric networks often desire the ability to monitor flow path visualizations and easily troubleshoot network failures/issues.

Current in-situ Operations, Administration, and Management (iOAM) techniques provide data path telemetry and path tracing methods from a last-hop router (LHR) using multicast tracing functionality, such as 'mtracev2'. Other telemetry techniques exist to visualize a multicast tree by exporting route data that can be stored in a data management entity, object store, or the like. These techniques can be applied to different problem scenarios, however, not all problem scenarios can be addressed, or it may not be intuitive how some problem scenarios may be addressed using such currently existing techniques.

There is a need to provide lightweight visualization from a network switch or any remote dashboard agent that can quickly provide a tree visualization and drive real time observability. However, providing an on-demand path trace using current tracing functions requires network administrators to know the LHRs in order to trigger a path trace, which can be difficult. Further, with current tracing functions only one branch can be traced at a time. Thus, building a topology using current tracing functions can be cumbersome. Additionally, current iOAM techniques utilized on a data path can be used only along the path that data traffic travels, which needs a data flow. Thus, iOAM cannot be used to track any intermediate broken branches due to data path failures and cannot provide for the ability to pre-visualize a tree when there is no existing multicast tree based on intended receiver points.

This submission proposes techniques that can be used for dynamic tree visualization, role discovery, troubleshooting, and/or the like through the use of PIM flooding mechanisms that are both simple and can be utilized in existing PIM architectures.

Current Internet Engineering Task Force (IETF) Request For Comments (RFC) documents, such as RFC 8364, provide various source discovery techniques that have been well-received by IP fabric network operators as such techniques can simplify IP fabric networks and operations thereof. PIM flooding mechanisms described herein are generic to such techniques.

Consider an example scenario in which an IP fabric network operator desires to discover all end points and all source nodes for a given group or a given Source group combination and understand router roles utilizing a technique that is easy to use and that facilitates fast response times to easily perform minimal triage for network issues/failures.

Provided herein is a technique that can address such desires of an IP fabric network operator through a PIM flooding mechanism that can provide for the ability to discover a multicast tree from any node in a given network topology. This technique does not involve knowledge of a source node or a receiver node, which often may not be known.

The PIM flooding mechanism according to this submission involves a tree trace message type that can be used to discover multicast trees. A tree discover message can be sent across a domain in a manner that is similar to Rendezvous Point (RP) mapping messages or PIM flooding-based source discovery messages are propagated. Figure 1, below, illustrates an example network topology in which such a message can be initiated.

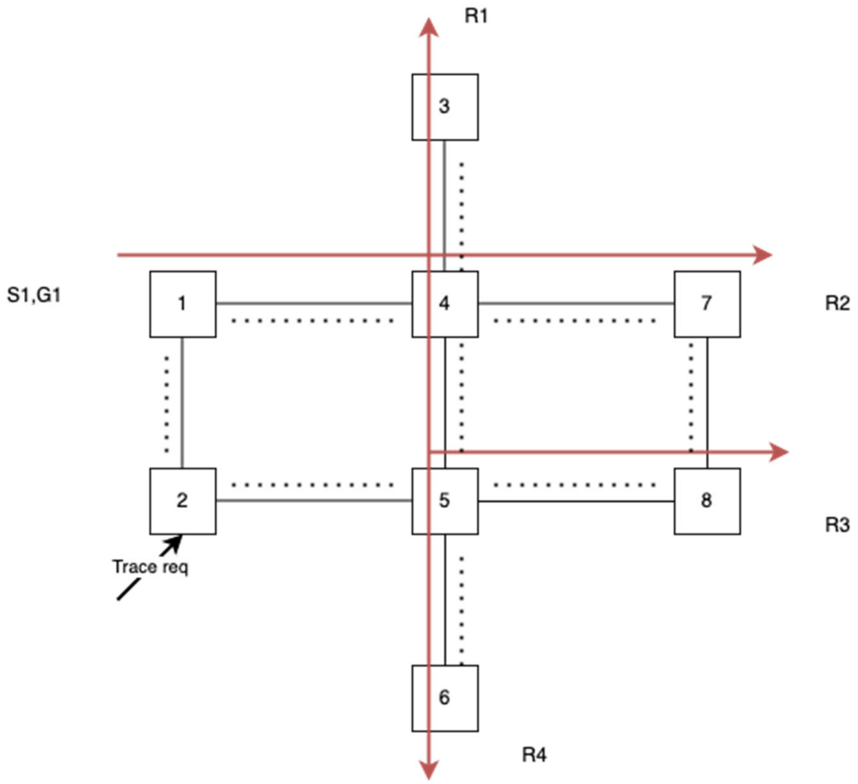

Figure 1: Example Network Topology

In one example, consider that a request can be formatted as:

> *{MSG_type: "request", Trace originator: ip_r, trace collector: ip_c, Trace_id: id, flow identification : (S,G)| (*,G), Aux cmds: [Stats, BW ..]}*

Nodes that participate in the multicast tree can respond to the requestor node identified in the request, as illustrated in Figure 2, below, in which information included in the responses can be used to reconstruct the tree. In one example, a response from each node that participates in a tree can be formatted as:

> *{MSG_type: " response", response_node : router_id, trace_id: id, flow identification: (S,G)| (*,G), Tree_id: (S,G), incoming : (router_id, interface_id), outgoing: [(router_id: interface_id1,..], role : (FHR| RP| LHR .), Aux cmd data: (..)}*

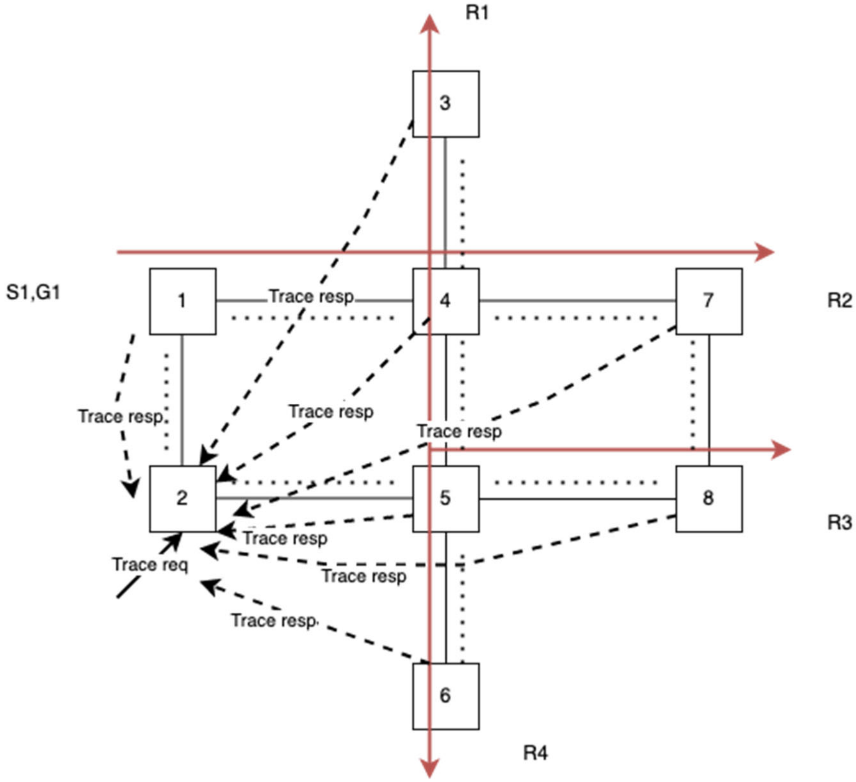

Figure 2: Example Node Responses

The request can also be used to determine a point or condition for a response, such as aggregate information, responding from a last hop router node, responding on specific conditions, and/or the like. Information indicating a node's role, such as first hop router, last hop router, intermediate hop, RP, etc. can also be mapped into a node's response. Further, in some instances, responses can include additional data associated with data path details, such as flow statistics, interface statistics, and/or the like, which can be used in the preliminary diagnosis of network problems.

A collector, which can be the requestor or other agent will collect the trace data response from participating nodes and will organize the responses to form the tree. The representation from the collector can be visualized in graphic or any readable format.

Figure 3, below, illustrates an example multicast tree that can be determined utilizing the technique provided herein for the network topology of Figure 1 based on node responses that can be obtained by a collector, as shown in Figure 2.

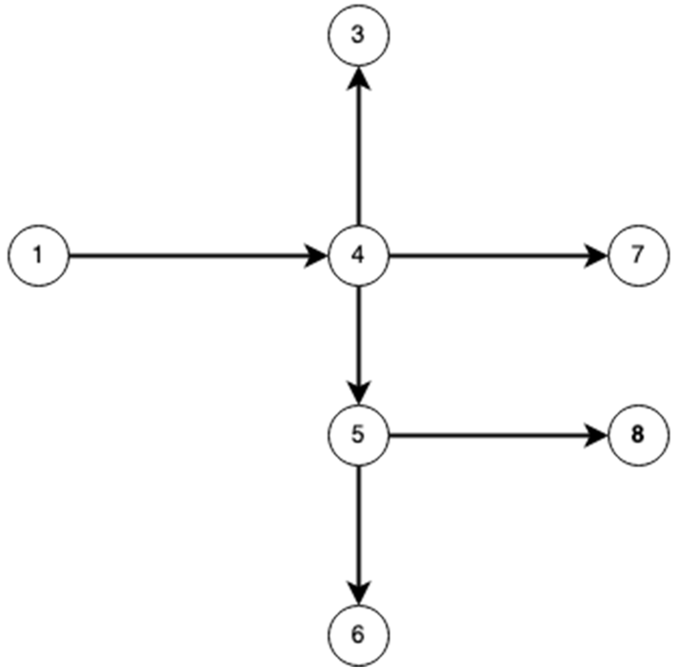

Figure 3: Example Multicast Tree Representation Determined from Node Responses

The multicast tree representation, such as the representation shown in Figure 3, can help to visualize any non-attached branches, which can be used for troubleshooting issues, such as reverse path null (rpf-null) issues, join not sent issues, forwarding plane drops, or the like.

Since the request/response messages can be formatted in a Type-Length-Value (TLV) format, any additional requests/responses can be easily added to gather additional network data. For example, some messages can be provided to perform a consistency check on each node, collect specific statistics, determine flow liveliness, etc. In some instances, certain TLVs can also be used to drive any specific instruction to the router.

Accordingly, techniques proposed herein can be used for dynamic tree visualization, role discovery, troubleshooting, and/or the like through the use of PIM flooding mechanisms that are both simple and can be utilized in existing PIM architectures.