Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 16, 2014

J. Ko
J. Jeong
J. Park
J. Jun
N. Kim
Electronics and Telecommunications Research Institute
O. Gnawali
University of Houston
February 12, 2014

RPL Routing Pathology In a Network With a Mix of Nodes Operating in Storing and Non-Storing Modes
draft-ko-roll-mix-network-pathology-04

Abstract

The RPL specification allows nodes running with storing or non-storing modes to operate in the same network. We describe how such a mix can result in network partitioning even when there are plenty of physical links available in the network. The partitioning affects both upwards (nodes to root) and downwards (root to leaf) traffic. This routing pathology stems from a recommendation made in the RPL specification forcing nodes with different modes of operation to join the RPL network as leaf nodes only. We propose a solution that modifies RPL by mandating that all the nodes parse and interpret source routing headers and storing mode nodes to sometimes act like a non-storing mode root by attaching source routing headers.

Table of Contents

1. Introduction

RPL [RFC6550] can operate in storing and non-storing modes. These modes introduce two different ways to perform downward routing. Downward routing is used when a node needs to send a packet to an arbitrary node (e.g., non-DODAG root node) in the network: the packet can go from a node "upward" towards the root and "downwards" to the final destination.

The RPL specification allows operating a network with a mix of storing and non-storing modes. RFC 6550 describes special rules to operate such a network: a node that operates with a different Mode of Operation (MOP) than the DODAG root will act as a leaf node in the network. The consensus was that it is unknown if the network would work properly because no one had designed such a network and was left to be explored in the future.

In this draft, we document a case in which we allow a mix of nodes operating in storing and non-storing modes to form a single network (e.g, despite having different MOPs) and show that RPL's two downwards routing modes, as it is, can cause a routing pathology. This pathology can partition the network, i.e., it can result in scenarios where nodes cannot send packets to the root and the root cannot send packets to the nodes even though these nodes have plenty of multi-hop physical connectivity in the network.

We propose one approach of modifying RPL to prevent this routing pathology. The methodology, introducing a new mode of operation (MOP), has been implemented and tested on an LLN testbed and in process of publication. It is possible there are more elegant approaches to prevent the pathology described.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

The terminologies used in this document are consistent with the terminologies described in [I-D.ietf-roll-terminology], [RFC6551], and [RFC6550].

3. Storing and Non-storing modes

Before we describe the routing pathology that arises due to the existence of a mix of nodes operating in storing and non-storing modes, we review the storing and non-storing downwards routing modes that RPL introduces.

In Storing mode, a node keeps a (not necessarily) complete list of (destination, nexthop) for nodes in its subtree. When a node receives a packet, it forwards the packet to the nexthop if the node finds the destination in the list. If it does not find the destination in the list, it forwards the packet to the preferred parent.

In Non-storing mode, if a packet does not have routing path in the header, it forwards the packet to the preferred parent. The root in this mode collects and maintains topology information of the network. If the packet makes it to the root, the root computes the path to the destination based on this topology information. The root puts this path in the header and sends it to the next hop. The nodes, upon receiving a packet with a path in the header, forward the packet to the next hop as indicated in the path in the header.

4. Routing Pathology

We first examine the effect of this routing pathology for routing collection traffic packets. Lets consider the following network topology.

A -> B -> N -> S -> Root (Storing)

Note that RPL indicates that, storing mode nodes and non-storing mode nodes use a different mode of operation (MOP) field. Furthermore, if the MOP supported from a DODAG root is not supported at a RPL node, the node can only participate in the RPL network as a leaf node. Say that the Root of the topology is a storing mode node. In this case, S (e.g., a storing mode node) can connect to the Root (a storing mode root) properly as a RPL router node. On the other hand, using the DIOs initiated at the Root, N (e.g., a non-storing mode node) will notice that the Root's MOP and its MOP is different; therefore, will only connect join the RPL network as a leaf node. As a result, A and B, which physically have connections to the root will not be able to join the RPL network. On a functional perspective, this means that both upwards (e.g., collection traffic) and downwards traffic to or from nodes A and B cannot reach the root or any other node in the RPL network. Thus, there is needless network partitioning. While this is an extreme case, other cases where using routes that non-storing mode nodes provide can help optimize the collection and downwards routes that RPL nodes form.

On a practical perspective, two RPL networks, each deployed for different purposes and use different MOPs, can meet at the geographically same area. Allowing these overlapping networks cooperate as a single dense network will improve the routing reliability with better connectivities. Since nodes of two networks can interoperate with each other, a network can migrate into the another network in case of root failures or power outages. It will also improve the network system robustness.

On a system design perspective, as the size of LLN deployments increase, using a homogeneous non-storing mode RPL network will introduce a high level of communication overhead, and a homogeneous network of strong mode nodes will require too much memory resources. By allowing a single mixed RPL network with computationally powerful nodes with route storing capabilities and low-cost nodes that do not need to maintain a routing table, it will construct a more efficient DODAG and resolve such system design issues.

Unfortunately, in such cases, the pathology that we discuss above can arise and cause downwards packets to be dropped and even more, restrict the formation of efficient collection routes.

5. Fixing the Pathology

We describe one way to fix RPL to prevent the pathology described above, while acknowledging that there might be more elegant solutions. In this approach, we acknowledge the fact that non-storing mode nodes are more likely to have strict resource limitations compared to nodes implementing the storing mode. Therefore, we make sure that the most of the required additional capabilities occur at the storing mode nodes rather than the non-storing mode nodes.

1. A new mode of operation (MOP) that allows a node to choose either to implement the storing or non-storing features, or both. The changes below are made compared to the original storing and non-storing modes.

2. Require storing and non-storing nodes to implement source routing header parsing capabilities.

3. RPL DODAG Root nodes supporting this MOP should have the capability to store routes (similar to the non-storing mode option) and also identify storing mode node children nodes.

4. Non-storing nodes send hop-by-hop DAO.

5. Storing nodes keep a table of all the DAO senders and a flag indicating if each of those sender is operating in storing or non-storing mode. This requires allocating one of the bits in the DAO message for a node to indicate if it is operating in storing or non-storing mode.

6. Change the forwarding mechanism in the storing mode node when it receives a downward bound packet:

7.

    1. If packet does not have source routing header and the next hop is a storing-mode node, forward as in [RFC6550]. If the next hop is a non-storing node, insert the source routing header [RFC6554] into the packet and forward, i.e., act like a non-storing root.

    2. Using the flag indicating the storing status of nodes in its sub-DODAG, a node constructing a source routing header MAY choose to construct a source routing header only up to the next storing mode node.

3. If the incoming packet has a source routing header, a storing mode node SHOULD obey the route specified in the source routing header to comply with the strict source routing requirements in [RFC6554].

If there is a mix of storing and non-storing nodes, we should also be more aggressive about loop detection. More aggressive loop detection will quickly remove the looping packets from the network. Even with the implementation of this suggestion, nodes beyond storing / non-storing nodes will still remain unreachable.

6. Considerations for Route Management

Given the lossy nature of the links and the variability in link quality on a temporal scale, the DAO parent(s) that were originally selected can change over time. In such cases, nodes take the following steps.

1. Upon realizing a change in the DAO parent set, a RPL node SHOULD send a DAO with an increased Path Sequence Number to its ``new" DAO parent set.

2. Next, the node experiencing the change in its DAO parents, SHOULD send a No-Path DAO message with an increased Path Sequence Number to the set of DAO parents that have been deleted.

3. When nodes receive a DAO that contains a modified Transit Information option for a specific Target option, they SHOULD propagate this information to its DAO parents while respecting the Path Control field in the Transit Information option.

4. When a node without the route storing capability fails in forwarding a packet using the information provided by a source routing header, the node initiates an ICMPv6 error message to the node that constructed the source routing header.

This process alerts nodes keeping the routing state of the DAO initiating nodes with the up-to-date DAO parent information. Nodes with route storing capabilities will process the information and nodes without this capability will pass the information up the DODAG. When using the flag indicating the storing status, the propagation of the updated DAO and ICMPv6 error message MAY stop at the first node with route storing capability since this node takes the role of managing the routing state for the target node.

7. Acknowledgements

8. IANA Considerations

9. Security Considerations

Future work.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.

[RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, March 2012.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, March 2012.

10.2. Informative References

[I-D.ietf-roll-terminology] Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-05 (work in progress), March 2011.

Authors' Addresses

JeongGil Ko
Electronics and Telecommunications Research Institute
218 Gajeong-Ro
Yuseong-Gu, Daejeon 305-700
Korea

Phone: +82-42-860-5824
Email: jeonggil.ko@etri.re.kr

Jongsoo Jeong
Electronics and Telecommunications Research Institute
218 Gajeong-Ro
Yuseong-Gu, Daejeon 305-700
Korea

Phone: +82-42-860-1806
Email: jsjeong@etri.re.kr

Jongjun Park
Electronics and Telecommunications Research Institute
218 Gajeong-Ro
Yuseong-Gu, Daejeon 305-700
Korea

Phone: +82-42-860-5413
Email: juny@etri.re.kr

Jong Arm Jun
Electronics and Telecommunications Research Institute
218 Gajeong-Ro
Yuseong-Gu, Daejeon 305-700
Korea

Phone: +82-42-860-4835
Email: jajun@etri.re.kr

Naesoo Kim
Electronics and Telecommunications Research Institute
218 Gajeong-Ro
Yuseong-Gu, Daejeon 305-700
Korea

Phone: +82-42-860-5214
Email: nskim@etri.re.kr

Omprakash Gnawali
University of Houston
PGH 577, University of Houston
Houston, TX 77204
USA

Phone: +1-713-743-3356
Email: gnawali@cs.uh.edu