Workgroup: ROLL
Updates: 6550, 8138 (if approved)
Published: 12 December 2019
Intended Status: Standards Track
Expires: 14 June 2020
Authors: P. Thubert, Ed.   L. Zhao
         Cisco Systems   Cisco Systems

Configuration option for RFC 8138

Abstract

This document complements RFC 8138 and dedicates a bit in the RPL configuration option defined in RFC 6550 to indicate whether RFC 8138 compression is used within the RPL instance.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The transition to [RFC8138] in a network can only be done when all nodes support the specification. In a mixed case with both RFC8138-capable and non-capable nodes, the compression should be turned off.

This document complements RFC 8138 and dedicates a bit in the RPL configuration option to indicate whether RFC 8138 compression should be used within the RPL instance. When the bit is not set, source nodes that support RFC 8138 should refrain from using the compression unless the information is superseded by configuration.

2. BCP 14

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 BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Updating RFC 6550

RPL defines a configuration option that is registered to IANA in section 20.14. of [RFC6550]. This specification defines a new flag "Enable RFC8138 Compression" (T) that is encoded in one of the reserved control bits in the option. The new flag is set to turn on the use of the compression of RPL artifacts with RFC 8138. The bit position of the "T" flag is indicated in Section 6.

Section 6.3.1. of [RFC6550] defines a 3-bit Mode of Operation (MOP) in the DIO Base Object. The new "T" flag is defined only for MOP value between 0 to 6. For a MOP value of 7 or above, the flag MAY indicate something different and MUST NOT be interpreted as "Enable RFC8138 Compression" unless the specification of the MOP indicates to do so.

4. Updating RFC 8138

This document specifies controls that enable and disable the use of the [RFC8138] compression in a RPL Instance. Arguably, this could have been done in [RFC8138] itself.

A node that supports this specification SHOULD source packets in the compressed form using [RFC8138] if the new "T" flag is set in the RPL configuration option from its parents. Failure to do so will result in larger packets, yields higher risks of loss and may cause a fragmentation.

A node that supports this specification SHOULD refrain from sourcing packets in the compressed form using [RFC8138] if the "T" flag is reset. This behaviour can be overridden by a configuration of the node in order to cope with intermediate implementations of the root that support [RFC8138] but not this specification and cannot set the "T" flag.

The decision of using RFC 8138 to compress a packet is made at the source depending on its capabilities and its knowledge of the state of the "T" flag. A router MUST forward the packet in the form that the source used, either compressed or uncompressed. A router that encapsulates a packet is the source of the resulting packet and the rules above apply to it in that case.

5. Transition Scenarios

A node that supports [RFC8138] but not this specification can only be used in an homogeneous network and an upgrade requires a "flag day" where all nodes are updated and then the network is rebooted with implicitely RFC 8138 compression turned on with the "T" flag set on.

A node that supports this specification can work in a network with RFC 8138 compression turned on or off with the "T" flag set accordingly and in a network in transition from off to on or on to off (see Section 5.1).

A node that does not support [RFC8138] can interoperate with a node that supports this specification in a network with RFC 8138 compression turned off. But it cannot forward compressed packets and therefore it cannot act as a router in a network with RFC 8138 compression turned on. It may remain connected to that network as a leaf and generate uncompressed packets as long as imcoming packets are decapsulated by the parent and delivered in uncompressed form.

[RFC6550] states that "Nodes other than the DODAG root MUST NOT modify this information when propagating the DODAG Configuration option". In other words, the configuration option is a way for the root to configure the LLN nodes but it cannot be used by a parent to advertise its capabilities down the DODAG. It results whether a parent supports RFC 8138 is not known by the child with the current level of specifications, and a child cannot favor a parent based on a particular support.

Sections 8.5 and 9.2 of [RFC6550] also suggests that a RPL-aware node may attach to a DODAG as a leaf node only, e.g., when a node does not support the Mode of Operation of a RPL Instance, the Objective Function (OF) as indicated by the Objective Code Point (OCP) or some other parameters in the configuration option. But the node is also free to refrain from joining an Instance when a parameter is not suitable. This means that changing the OCP in a DODAG can be used to force nodes that do not support a particular feature to join as leaf only. This specification reiterates that a node that is configured to operate in an Instance but does not support a value for a known parameter that is mandatory for routing MUST NOT operate as a router but MAY still joins as a leaf. Note that a legacy node will not recognize when a reserved field is now used and will not turn to a leaf when that happens.

The intent for this specification is to perform a migration once and for all without the need for a flag day. In particular it is not the intention to undo the setting of the "T" flag, and though it is possible to roll back (see Section 5.4), adding nodes that do not support [RFC8138] after a roll back may be problematic if the roll back is not fully complete (see caveats in Section 5.2).

5.1. Inconsistent State While Migrating

When the "T" flag is turned on in the configuration option by the root, the information slowly percolates through the DODAG as the DIO gets propagated. Some nodes will see the flag and start sourcing packets in the compressed form while other nodes in the same instance are still not aware of it. Conversely, in non-storing mode, the root will start using RFC 8138 with a SRH-6LoRH that routes all the way to the last router or possibly to the leaf, if the leaf supports RFC 8138.

This is why it is required that all the routers in the Instance support [RFC8138] at the time of the switch, and all nodes that do not support [RFC8138] only operate as leaves.

Setting the "T" flag is ultimately the responsibility of the network administrator. In a case of upgrading a network to turn the compression on, the network SHOULD be operated with the "T" flag reset until all targeted nodes are upgraded to support this specification. Section 5.2 and Section 5.3 provide possible transition scenarios where this can be enforced.

5.2. Single Instance Scenario

In a single instance scenario, nodes that support RFC 8138 are configured with a new OCP, that may use the same OF operation or a variation of it. when it finally sets the "T" flag, the root also migrates to the new OCP. As a result, nodes that do not support RFC 8138 join as leaves and do not forward packets anymore. The leaves generate packets without compression. The parents - which supports RFC 8138 - may encapsulate the packets using RFC 8138 if needed. The other way around, the root encapsulates packets to the leaves all the way to the parent, which decapsulates and distribute the uncompresses inner packet to the leaf.

This scenario presents a number of caveats:

*The method consumes an extra OCP. It also requires a means to signal the capabilities of the leaf, e.g., using "RPL Mode of Operation extension" [MOP-EXT].

*If an implementation does not move to a leaf mode when the OCP is changed to an unknown one, then the node may be stalled.

*If the only possible parents of a node are nodes that do not support RFC 8138, then that node will loose all its parent at the time of the migration and it will be stalled until a parent is deployed with the new capability.

*Nodes that only support RFC8138 for forwarding may not parse the RPI in native form. If such nodes are present, the parent needs to encapsulate with RFC8138.

5.3. Double Instance Scenario

An alternate to the Single Instance Scenario is to deploy an additional Instance for the nodes that support [RFC8138]. The two instances operate as ships-in-the-night as specified in [RFC6550]. The preexisting Instance that does not use [RFC8138], whereas the new Instance does. This is signaled by the "T" flag which is only set in the configuration option in DIO messages in the new Instance.

Nodes that support RFC 8138 participate to both Instances but favor the new Instance for the traffic that they source. On the other hand, nodes that only support the uncompressed format would either not be configured for the new instance, or would be configured to join it as leaves only.

This method eliminates the risks of nodes being stalled that are described in Section 5.2 but requires implementations to support at least two RPL Instances and demands management capabilities to introduce new Instances and deprecate old ones.

5.4. Rolling Back

After downgrading a network to turn the [RFC8138] compression off, the administrator SHOULD make sure that all nodes have converged to the "T" flag reset before allowing nodes that do not support the compression in the network (see caveats in Section 5.2).

It is RECOMMENDED to only deploy nodes that support [RFC8138] in a network where the compression is turned on. A node that does not support [RFC8138] MUST only be used as a leaf.

6. IANA Considerations

This specification updates the Registry for the "DODAG Configuration Option Flags" that was created for [RFC6550] as follows:

Bit Number	Capability Description	Reference
2	Turn on RFC8138 Compression (T)	THIS RFC

Table 1: New DODAG Configuration Option Flag

7. Security Considerations

Turning the "T" flag on before some routers are upgraded may cause a loss of packets. The new bit is protected as the rest of the configuration so this is just one of the many attacks that can happen if an attacker manages to inject a corrupted configuration.

Turning the "T" flag on and off may create inconsistencies in the network but as long as all nodes are upgraded to RFC 8138 support they will be able to forward both forms. The draft insists that the source is responsible for selecting whether the packet is compressed or not, and all routers must use the format that the source selected. So the result of an inconsistency is merely that both forms will be present in the network, at an additional cost of bandwidth for packets in the uncompressed form.

8. Acknowledgments

9. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., 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, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

10. Informative References

[RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <https://www.rfc-editor.org/info/rfc8138>.

[MOP-EXT] Jadhav, R., Thubert, P., and M. Richardson, "Mode of Operation extension and Capabilities", Work in Progress, Internet-Draft, draft-ietf-roll-mopex-cap-01, 2 November 2019, <https://tools.ietf.org/html/draft-ietf-roll-mopex-cap-01>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D, 45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Li Zhao
Cisco Systems, Inc
Xinsi Building, No. 926 Yi Shan Rd
SHANGHAI
200233
China

Email: liz3@cisco.com