

ROLL
Internet-Draft
Intended status: Standards Track
Expires: December 30, 2010

T. Winter, Ed.

P. Thubert, Ed.
Cisco Systems
RPL Author Team
IETF ROLL WG
Jun 28, 2010

RPL: IPv6 Routing Protocol for Low power and Lossy Networks
draft-ietf-roll-rpl-10

Abstract

Low power and Lossy Networks (LLNs) are a class of network in which both the routers and their interconnect are constrained: LLN routers typically operate with constraints on (any subset of) processing power, memory and energy (battery), and their interconnects are characterized by (any subset of) high loss rates, low data rates and instability. LLNs are comprised of anything from a few dozen and up to thousands of routers, and support point-to-point traffic (between devices inside the LLN), point-to-multipoint traffic (from a central control point to a subset of devices inside the LLN) and multipoint-to-point traffic (from devices inside the LLN towards a central control point). This document specifies the IPv6 Routing Protocol for LLNs (RPL), which provides a mechanism whereby multipoint-to-point traffic from devices inside the LLN towards a central control point, as well as point-to-multipoint traffic from the central control point to the devices inside the LLN, is supported. Support for point-to-point traffic is also available.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 30, 2010.

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	6
1.1.	Design Principles	6
1.2.	Expectations of Link Layer Type	7
2.	Terminology	8
3.	Protocol Overview	11
3.1.	Topology	11
3.1.1.	Topology Identifiers	11
3.2.	Instances, DODAGs, and DODAG Versions	11
3.3.	Upward Routes and DODAG Construction	13
3.3.1.	Objective Function (OF)	14
3.3.2.	DODAG Repair	14
3.3.3.	Security	14
3.3.4.	Grounded and Floating DODAGs	14
3.3.5.	Local DODAGs	14
3.3.6.	Administrative Preference	15
3.3.7.	Datapath Validation and Loop Detection	15
3.3.8.	Distributed Algorithm Operation	15
3.4.	Downward Routes and Destination Advertisement	15
3.5.	Local DODAGs Route Discovery	16
3.6.	Routing Metrics and Constraints Used By RPL	16
3.6.1.	Loop Avoidance	17
3.6.2.	Rank Properties	18
3.7.	Traffic Flows Supported by RPL	20
3.7.1.	Multipoint-to-Point Traffic	21
3.7.2.	Point-to-Multipoint Traffic	21

3.7.3.	Point-to-Point Traffic	21
4.	RPL Instance	22
4.1.	RPL Instance ID	22
5.	ICMPv6 RPL Control Message	24
5.1.	RPL Security Fields	25

5.2.	DODAG Information Solicitation (DIS)	30
5.2.1.	Format of the DIS Base Object	30
5.2.2.	Secure DIS	31
5.2.3.	DIS Options	31
5.3.	DODAG Information Object (DIO)	31
5.3.1.	Format of the DIO Base Object	31
5.3.2.	Secure DIO	33
5.3.3.	DIO Options	33
5.4.	Destination Advertisement Object (DAO)	33
5.4.1.	Format of the DAO Base Object	34
5.4.2.	Secure DAO	34
5.4.3.	DAO Options	35
5.5.	Destination Advertisement Object Acknowledgement (DAO-ACK)	35
5.5.1.	Format of the DAO-ACK Base Object	35
5.5.2.	Secure DAO-ACK	36
5.5.3.	DAO-ACK Options	36
5.6.	Consistency Check (CC)	36
5.6.1.	Format of the CC Base Object	36
5.6.2.	CC Options	38
5.7.	RPL Control Message Options	38
5.7.1.	RPL ControlMessage Option Generic Format	38
5.7.2.	Pad1	39
5.7.3.	PadN	39
5.7.4.	Metric Container	40
5.7.5.	Route Information	40
5.7.6.	DODAG Configuration	42
5.7.7.	RPL Target	43
5.7.8.	Transit Information	45
5.7.9.	Solicited Information	46
5.7.10.	Prefix Information	48
6.	Sequence Counters	51
7.	Upward Routes	53
7.1.	DIO Base Rules	53
7.2.	Upward Route Discovery and Maintenance	53
7.2.1.	Neighbors and Parents within a DODAG Version	53

7.2.2.	Neighbors and Parents across DODAG Versions	54
7.2.3.	DIO Message Communication	58
7.3.	DIO Transmission	59
7.3.1.	Trickle Parameters	60
7.4.	DODAG Selection	60
7.5.	Operation as a Leaf Node	60
7.6.	Administrative Rank	61
8.	Downward Routes	62
8.1.	Destination Advertisement Parents	62
8.2.	Downward Route Discovery and Maintenance	62
8.3.	DAO Base Rules	63
8.4.	DAO Transmission Scheduling	64

8.5.	Triggering DAO Messages	64
8.6.	Structure of DAO Messages	65
8.7.	Non-storing Mode	65
8.8.	Storing Mode	66
8.9.	Path Control	67
8.10.	Multicast Destination Advertisement Messages	68
9.	Security Mechanisms	69
9.1.	Security Overview	69
9.2.	Installing Keys	70
9.3.	Joining a Secure Network	70
9.4.	Counter and Counter Compression	71
9.4.1.	Timestamp Counters	72
9.5.	Functional Description of Packet Protection	72
9.5.1.	Transmission of Outgoing Packets	72
9.5.2.	Reception of Incoming Packets	74
9.5.3.	Cryptographic Mode of Operation	76
9.6.	Coverage of Integrity and Confidentiality	77
10.	Packet Forwarding and Loop Avoidance/Detection	78
10.1.	Suggestions for Packet Forwarding	78
10.2.	Loop Avoidance and Detection	79
10.2.1.	Source Node Operation	80
10.2.2.	Router Operation	80
11.	Multicast Operation	83
12.	Maintenance of Routing Adjacency	85
13.	Guidelines for Objective Functions	86
13.1.	Objective Function Behavior	86
14.	Suggestions for Interoperation with Neighbor Discovery	88
15.	RPL Constants and Variables	89
16.	Manageability Considerations	91

16.1.	Introduction	91
16.2.	Configuration Management	92
16.2.1.	Initialization Mode	92
16.2.2.	DIO and DAO Base Message and Options Configuration .	92
16.2.3.	Protocol Parameters to be configured on every router in the LLN	93
16.2.4.	Protocol Parameters to be configured on every non-root router in the LLN	93
16.2.5.	Parameters to be configured on the DODAG root	94
16.2.6.	Configuration of RPL Parameters related to DAO-based mechanisms	95
16.2.7.	Default Values	96
16.3.	Monitoring of RPL Operation	96
16.3.1.	Monitoring a DODAG parameters	96
16.3.2.	Monitoring a DODAG inconsistencies and loop detection	97
16.4.	Monitoring of the RPL data structures	98
16.4.1.	Candidate Neighbor Data Structure	98
16.4.2.	Destination Oriented Directed Acyclic Graph (DAG)	

	Table	98
16.4.3.	Routing Table and DAO Routing Entries	99
16.5.	Fault Management	100
16.6.	Policy	100
16.7.	Liveness Detection and Monitoring	101
16.8.	Fault Isolation	102
16.9.	Impact on Other Protocols	102
16.10.	Performance Management	102
17.	Security Considerations	104
17.1.	Overview	104
18.	IANA Considerations	106
18.1.	RPL Control Message	106
18.2.	New Registry for RPL Control Codes	106
18.3.	New Registry for the Mode of Operation (MOP) DIO Control Field	107
18.4.	RPL Control Message Option	107
18.5.	Objective Code Point (OCP) Registry	108
18.6.	ICMPv6: Error in Source Routing Header	108
18.7.	Link-Local Scope multicast address	108
19.	Acknowledgements	110
20.	Contributors	111
21.	References	113

21.1.	Normative References	113
21.2.	Informative References	113
Authors' Addresses	117

[1.](#) Introduction

Low power and Lossy Networks (LLNs) consist of largely of constrained nodes (with limited processing power, memory, and sometimes energy when they are battery operated). These routers are interconnected by lossy links, typically supporting only low data rates, that are usually unstable with relatively low packet delivery rates. Another characteristic of such networks is that the traffic patterns are not simply point-to-point, but in many cases point-to-multipoint or multipoint-to-point. Furthermore such networks may potentially comprise up to thousands of nodes. These characteristics offer unique challenges to a routing solution: the IETF ROLL Working Group has defined application-specific routing requirements for a Low power and Lossy Network (LLN) routing protocol, specified in [\[RFC5867\]](#), [\[RFC5826\]](#), [\[RFC5673\]](#), and [\[RFC5548\]](#).

This document specifies the IPv6 Routing Protocol for Low power and lossy networks (RPL). Note that although RPL was specified according to the requirements set forth in the aforementioned requirement documents, its use is in no way limited to these applications.

1.1. Design Principles

RPL was designed with the objective to meet the requirements spelled out in [[RFC5867](#)], [[RFC5826](#)], [[RFC5673](#)], and [[RFC5548](#)].

A network may run multiple instances of RPL concurrently. Each such instance may serve different and potentially antagonistic constraints or performance criteria. This document defines how a single instance operates.

In order to be useful in a wide range of LLN application domains, RPL separates packet processing and forwarding from the routing optimization objective. Examples of such objectives include minimizing energy, minimizing latency, or satisfying constraints. This document describes the mode of operation of RPL. Other companion documents specify routing objective functions. A RPL implementation, in support of a particular LLN application, will include the necessary objective function(s) as required by the application.

A set of companion documents to this specification will provide further guidance in the form of applicability statements specifying a set of operating points appropriate to the Building Automation, Home Automation, Industrial, and Urban application scenarios.

1.2. Expectations of Link Layer Type

In compliance with the layered architecture of IP, RPL does not rely on any particular features of a specific link layer technology. RPL is designed to be able to operate over a variety of different link layers, including but not limited to, low power wireless or PLC (Power Line Communication) technologies.

Implementers may find [[RFC3819](#)] a useful reference when designing a link layer interface between RPL and a particular link layer technology.

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].

Additionally, this document uses terminology from [\[I-D.ietf-roll-terminology\]](#), and introduces the following terminology:

DAG: Directed Acyclic Graph. A directed graph having the property that all edges are oriented in such a way that no cycles exist. All edges are contained in paths oriented toward and terminating at one or more root nodes.

DAG root: A DAG root is a node within the DAG that has no outgoing edge. Because the graph is acyclic, by definition all DAGs must have at least one DAG root and all paths terminate at a DAG root.

Destination Oriented DAG (DODAG): A DAG rooted at a single destination, i.e. at a single DAG root (the DODAG root) with no outgoing edges.

DODAG root: A DODAG root is the DAG root of a DODAG.

Up: Up refers to the direction from leaf nodes towards DODAG roots, following DODAG edges. This follows the common terminology used in graphs and depth-first-search, where vertices further from the root are "deeper," or "down," and vertices closer to the root are "shallower," or "up."

Down: Down refers to the direction from DODAG roots towards leaf nodes, in the reverse direction of DODAG edges. This follows the common terminology used in graphs and depth-first-search, where vertices further from the root are "deeper," or "down," and vertices closer to the root are "shallower," or "up."

Rank: A node's Rank defines the node's individual position relative to other nodes with respect to a DODAG root. Rank strictly increases in the down direction and strictly decreases in the up direction. The exact way Rank is computed depends on the DAG's Objective Function (OF). The Rank may analogously track a simple topological distance, may be calculated as a function of link metrics, and may consider other properties such as constraints.

Objective Function (OF): Defines which routing metrics, optimization objectives, and related functions a DAG uses to compute Rank.

Objective Code Point (OCP): An identifier that indicates which Objective Function the DODAG uses.

RPLInstanceID: A unique identifier within a network. Two DODAGs with the same RPLInstanceID share the same Objective Function.

RPL Instance: A set of one or more DODAGs that share a RPLInstanceID. A RPL node can belong to at most one DODAG in a RPL Instance. Each RPL Instance operates independently of other RPL Instances. This document describes operation within a single RPL Instance.

DODAGID: The identifier of a DODAG root. The DODAGID must be unique within the scope of a RPL Instance in the LLN. The tuple (RPLInstanceID, DODAGID) uniquely identifies a DODAG.

DODAG Version: A specific sequence number iteration ("version") of a DODAG with a given DODAGID.

DODAGVersionNumber: A sequential counter that is incremented by the root to form a new Version of a DODAG. A DODAG Version is identified uniquely by the (RPLInstanceID, DODAGID, DODAGVersionNumber) tuple.

Goal: The Goal is a application specific goal that is defined outside the scope of RPL. Any node that roots a DODAG will need to know about this Goal to decide if the Goal can be satisfied or not. A typical Goal is to construct the DODAG according to a specific objective function and to keep connectivity to a set of hosts (e.g. to use an objective function that minimizes ETX and to be connected to a specific database host to store the collected data).

Grounded: A DODAG is grounded when the DODAG root can satisfy the Goal.

Floating: A DODAG is floating if is not Grounded. A floating DODAG is not expected to have the properties required to satisfy the goal. It may, however, provide connectivity to other nodes within the DODAG.

DODAG parent: A parent of a node within a DODAG is one of the immediate successors of the node on a path towards the DODAG

root. A DODAG parent's Rank is lower than the node's. (See [Section 3.6.2.1](#)).

Sub-DODAG The sub-DODAG of a node is the set of other nodes whose paths to the DODAG root pass through that node. Nodes in the sub-DODAG of a node have a greater Rank than that node itself. (See [Section 3.6.2.1](#))

As they form networks, LLN devices often mix the roles of 'host' and 'router' when compared to traditional IP networks. In this document, 'host' refers to an LLN device that can generate but does not forward RPL traffic, 'router' refers to an LLN device that can forward as well as generate RPL traffic, and 'node' refers to any RPL device, either a host or a router.

[3.](#) Protocol Overview

The aim of this section is to describe RPL in the spirit of [\[RFC4101\]](#). Protocol details can be found in further sections.

[3.1.](#) Topology

This section describes how the basic RPL topologies, and the rules by which these are constructed, i.e. the rules governing DODAG formation.

[3.1.1.](#) Topology Identifiers

RPL uses four identifiers to maintain the topology:

- o The first is a RPLInstanceID. A RPLInstanceID identifies a set of one or more DODAGs. All DODAGs in the same RPL Instance use the same Objective Function. A network may have multiple RPLInstanceIDs, each of which defines an independent set of DODAGs, which may be optimized for different OFs and/or applications. The set of DODAGs identified by a RPLInstanceID is called a RPL Instance.
- o The second is a DODAGID. The scope of a DODAGID is a RPL Instance. The combination of RPLInstanceID and DODAGID uniquely identifies a single DODAG in the network. A RPL Instance may have multiple DODAGs, each of which has a unique DODAGID.
- o The third is a DODAGVersionNumber. The scope of a DODAGVersionNumber is a DODAG. A DODAG is sometimes reconstructed from the DODAG root, by incrementing the DODAGVersionNumber. The combination of RPLInstanceID, DODAGID, and DODAGVersionNumber

uniquely identifies a DODAG Version.

- o The fourth is Rank. The scope of Rank is a DODAG Version. Rank establishes a partial order over a DODAG Version, defining individual node positions with respect to the DODAG root.

[3.2.](#) Instances, DODAGs, and DODAG Versions

A RPL Instance contains one or more Destination Oriented DAG (DODAG) roots. A RPL Instance may provide routes to certain destination prefixes, reachable via the DODAG roots or alternate paths within the DODAG. These roots may operate independently, or may coordinate over a non-LLN backchannel.

A RPL Instance may comprise:

- o a single DODAG with a single root
 - * For example, a DODAG optimized to minimize latency rooted at a single centralized lighting controller in a home automation application.
- o multiple uncoordinated DODAGs with independent roots (differing DODAGIDs)
 - * For example, multiple data collection points in an urban data collection application that do not have an always-on backbone suitable to coordinate to form a single DODAG, and further use the formation of multiple DODAGs as a means to dynamically and autonomously partition the network.
- o a single DODAG with a single virtual root coordinating LLN sinks (with the same DODAGID) over some non-LLN backbone
 - * For example, multiple border routers operating with a reliable backbone, e.g. in support of a 6LowPAN application, that are capable to act as logically equivalent sinks to the same DODAG.
- o a combination of the above as suited to some application scenario.

Each RPL packet has meta-data that associates it with a particular

RPLInstanceID and therefore RPL Instance.([Section 4](#)). The provisioning or automated discovery of a mapping between a RPLInstanceID and a type or service of application traffic is beyond the scope of this specification.

Figure 1 depicts an example of a RPL Instance comprising three DODAGs with DODAG Roots R1, R2, and R3. Figure 2 depicts how a DODAG version number increment leads to a new DODAG Version.

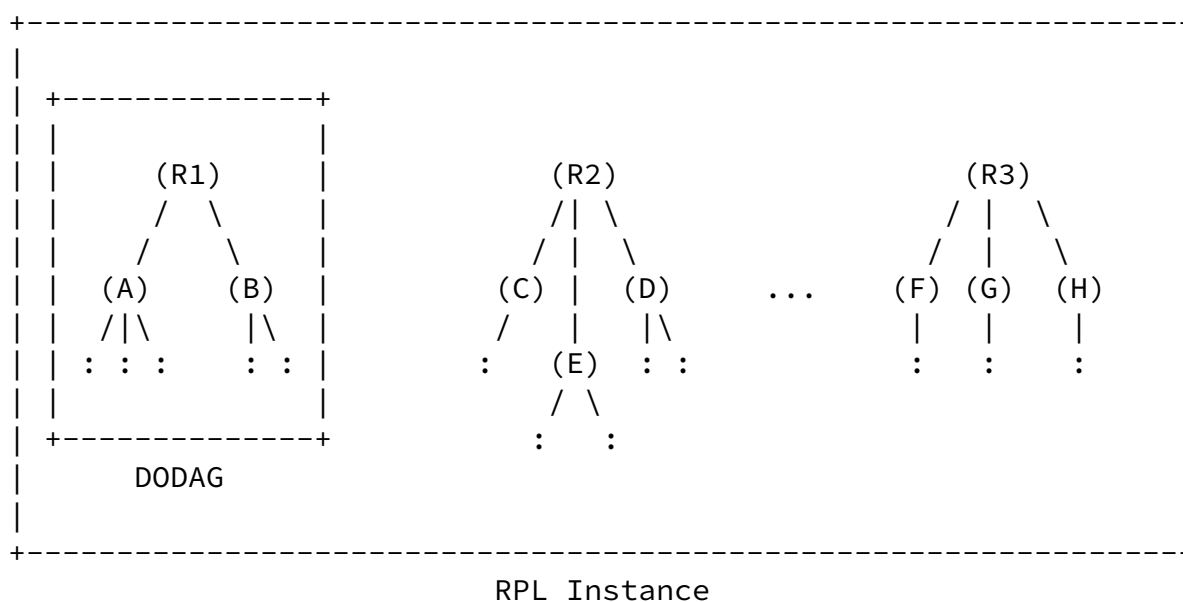


Figure 1: RPL Instance




Figure 2: DODAG Version

3.3. Upward Routes and DODAG Construction

RPL provisions routes up towards DODAG roots, forming a DODAG optimized according to an Objective Function (OF). RPL nodes construct and maintain these DODAGs through DODAG Information Object (DIO) messages.

3.3.1. Objective Function (OF)

The Objective Function (OF) defines how RPL nodes select and optimize routes within a RPL Instance. The OF is identified by an Objective Code Point (OCP) within the DIO Configuration option. An OF defines how nodes translate one or more metrics and constraints, which are themselves defined in [\[I-D.ietf-roll-routing-metrics\]](#), into a value called Rank, which approximates the node's distance from a DODAG root. An OF also defines how nodes select parents. Further details may be found in [Section 13](#), [\[I-D.ietf-roll-routing-metrics\]](#), [\[I-D.ietf-roll-of0\]](#), and related companion specifications.

[3.3.2.](#) DODAG Repair

A DODAG Root institutes a global repair operation by incrementing the DODAG Version Number. This initiates a new DODAG version. Nodes in the new DODAG version can choose a new position whose Rank is not constrained by their Rank within the old DODAG Version.

RPL also supports mechanisms which may be used for local repair within the DODAG version. The DIO message specifies the necessary parameters as configured from the DODAG root, as controlled by policy at the root.

[3.3.3.](#) Security

RPL supports message confidentiality and integrity. It is designed such that link-layer mechanisms can be used when available and appropriate, yet in their absence RPL can use its own mechanisms.

[3.3.4.](#) Grounded and Floating DODAGs

DODAGs can be grounded or floating: the DODAG root advertises which is the case. A grounded DODAG offers connectivity to hosts that are required for satisfying the application-defined goal. A floating DODAG is not expected to satisfy the goal and in most cases only provides routes to nodes within the DODAG. Floating DODAGs may be used, for example, to preserve inner connectivity during repair.

[3.3.5.](#) Local DODAGs

RPL nodes can optimize routes to a destination within an LLN by forming a local DODAG whose DODAG Root is the desired destination. Unlike global DAGs, which can consist of multiple DODAGs, local DAGs have one and only one DODAG and therefore one DODAG Root. Local DODAGs can be constructed on-demand.

[3.3.6.](#) Administrative Preference

An implementation/deployment may specify that some DODAG roots should be used over others through an administrative preference.

Administrative preference offers a way to control traffic and engineer DODAG formation in order to better support application requirements or needs.

[3.3.7.](#) Datapath Validation and Loop Detection

RPL uses a hop-by-hop IPv6 header to detect possible loops within a DODAG. Each data packet includes the Rank of the transmitter. An inconsistency between the routing decision for a packet (upward or downward) and the Rank relationship between the two nodes indicates a possible loop. On receiving such a packet, a node institutes a local repair operation.

[3.3.8.](#) Distributed Algorithm Operation

A high level overview of the distributed algorithm, which constructs the DODAG, is as follows:

- o Some nodes are configured to be DODAG roots, with associated DODAG configurations.
- o Nodes advertise their presence, affiliation with a DODAG, routing cost, and related metrics by sending link-local multicast DIO messages.
- o Nodes listen for DIOs and use their information to join a new DODAG, or to maintain an existing DODAG, as according to the specified Objective Function and Rank of their neighbors.
- o Nodes provision routing table entries, for the destinations specified by the DIO, via their DODAG parents in the DODAG version. Nodes MUST provision a DODAG parent as a default route for the associated instance. It is up to the end-to-end application to select the RPL instance to be associated to its traffic (should there be more than one instance) and thus the default route upwards when no longer-match exists.

[3.4.](#) Downward Routes and Destination Advertisement

RPL uses Destination Advertisement Object (DAO) messages to establish downward routes from DODAG roots. DAO messages are an optional feature for applications that require P2MP or P2P traffic. RPL supports two modes of downward traffic: storing (fully stateful) or non-storing (fully source routed). Any given RPL Instance is either

storing or non-storing. In both cases, P2P packets travel up to a DODAG Root then down to the final destination (unless the destination is on the upward route).

[3.5.](#) Local DODAGs Route Discovery

A RPL network can optionally support on-demand discovery of DODAGs to specific destinations within an LLN. Such local DODAGs behave slightly differently than global DODAGs.

[3.6.](#) Routing Metrics and Constraints Used By RPL

Routing metrics are used by routing protocols to compute shortest paths. Interior Gateway Protocols (IGPs) such as IS-IS ([\[RFC5120\]](#)) and OSPF ([\[RFC4915\]](#)) use static link metrics. Such link metrics can simply reflect the bandwidth or can also be computed according to a polynomial function of several metrics defining different link characteristics. Some routing protocols support more than one metric: in the vast majority of the cases, one metric is used per (sub)topology. Less often, a second metric may be used as a tie-breaker in the presence of Equal Cost Multiple Paths (ECMP). The optimization of multiple metrics is known as an NP complete problem and is sometimes supported by some centralized path computation engine.

In contrast, LLNs do require the support of both static and dynamic metrics. Furthermore, both link and node metrics are required. In the case of RPL, it is virtually impossible to define one metric, or even a composite metric, that will satisfy all use cases.

In addition, RPL supports constrained-based routing where constraints may be applied to both link and nodes. If a link or a node does not satisfy a required constraint, it is 'pruned' from the candidate list, thus leading to a constrained shortest path.

An Objective Function specifies the objectives used to compute the (constrained) path. Upstream and Downstream metrics may be merged or advertised separately depending on the OF and the metrics. When they are advertised separately, it may happen that the set of DIO parents is different from the set of DAO parents (a DAO parent is a node to which unicast DAO messages are sent). Yet, all are DODAG parents with regards to the rules for Rank computation.

The Objective Function itself is decoupled from the routing metrics and constraints used by RPL. Indeed, whereas the OF dictates rules such as DODAG parents selection, load balancing and so on, the set of metrics and/or constraints used to select a DODAG parent and thus

determine the preferred path are based on the information carried

within the DAG container option in DIO messages.

The set of supported link/node constraints and metrics is specified in [[I-D.ietf-roll-routing-metrics](#)].

Example 1: Shortest path: path offering the shortest end-to-end delay

Example 2: Constrained shortest path: the path that does not traverse any battery-operated node and that optimizes the path reliability

[3.6.1.](#) Loop Avoidance

RPL guarantees neither loop free path selection nor tight delay convergence times. In order to reduce control overhead, however, such as the cost of the count-to-infinity problem, RPL avoids creating loops when undergoing topology changes. Furthermore, RPL includes rank-based datapath validation mechanisms for detecting loops when they do occur. RPL uses this loop detection to ensure that packets make forward progress within the DODAG version and trigger repairs when necessary.

[3.6.1.1.](#) Greediness and Rank-based Instabilities

A node is greedy if it attempts to move deeper in the DODAG version, in order to increase the size of the parent set or improve some other metric. Moving deeper in within a DODAG version in this manner could result in instability and be detrimental to other nodes.

Once a node has joined a DODAG version, RPL disallows certain behaviors, including greediness, in order to prevent resulting instabilities in the DODAG version.

Suppose a node is willing to receive and process a DIO messages from a node in its own sub-DODAG, and in general a node deeper than itself. In this case, a possibility exists that a feedback loop is created, wherein two or more nodes continue to try and move in the DODAG version while attempting to optimize against each other. In some cases, this will result in instability. It is for this reason that RPL limits the cases where a node may process DIO messages from

deeper nodes to some forms of local repair. This approach creates an 'event horizon', whereby a node cannot be influenced beyond some limit into an instability by the action of nodes that may be in its own sub-DODAG.

[3.6.1.2.](#) DODAG Loops

A DODAG loop may occur when a node detaches from the DODAG and reattaches to a device in its prior sub-DODAG. This may happen in particular when DIO messages are missed. Strict use of the DODAG Version Number can eliminate this type of loop, but this type of loop may possibly be encountered when using some local repair mechanisms.

[3.6.1.3.](#) DAO Loops

A DAO loop may occur when the parent has a route installed upon receiving and processing a DAO message from a child, but the child has subsequently cleaned up the related DAO state. This loop happens when a No-Path (a DAO message that invalidates a previously announced prefix) was missed and persists until all state has been cleaned up. RPL includes an optional mechanism to acknowledge DAO messages, which may mitigate the impact of a single DAO message being missed. RPL includes loop detection mechanisms that may mitigate the impact of DAO loops and trigger their repair.

[3.6.2.](#) Rank Properties

The rank of a node is a scalar representation of the location of that node within a DODAG version. The rank is used to avoid and detect loops, and as such must demonstrate certain properties. The exact calculation of the rank is left to the Objective Function, and may depend on parents, link metrics, and the node configuration and policies.

The rank is not a cost metric, although its value can be derived from and influenced by metrics. The rank has properties of its own that are not necessarily those of all metrics:

Type: The rank is an abstract numeric value.

Function: The rank is the expression of a relative position within a DODAG version with regard to neighbors and is not necessarily a good indication or a proper expression of a distance or a cost to the root.

Stability: The stability of the rank determines the stability of the routing topology. Some dampening or filtering might be applied to keep the topology stable, and thus the rank does not necessarily change as fast as some physical metrics would. A new DODAG version would be a good opportunity to reconcile the discrepancies that might form over time between metrics and ranks within a DODAG version.

Winter, et al.

Expires December 30, 2010

[Page 18]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Properties: The rank is strictly monotonic, and can be used to validate a progression from or towards the root. A metric, like bandwidth or jitter, does not necessarily exhibit this property.

Abstract: The rank does not have a physical unit, but rather a range of increment per hop, where the assignment of each increment is to be determined by the Objective Function.

The rank value feeds into DODAG parent selection, according to the RPL loop-avoidance strategy. Once a parent has been added, and a rank value for the node within the DODAG has been advertised, the nodes further options with regard to DODAG parent selection and movement within the DODAG are restricted in favor of loop avoidance.

[3.6.2.1](#). Rank Comparison (DAGRank())

Rank may be thought of as a fixed point number, where the position of the radix point between the integer part and the fractional part is determined by MinHopRankIncrease. MinHopRankIncrease is the minimum increase in rank between a node and any of its DODAG parents. When an objective function computes rank, the objective function operates on the entire (i.e. 16-bit) rank quantity. When rank is compared, e.g. for determination of parent relationships or loop detection, the integer portion of the rank is to be used. The integer portion of the Rank is computed by the DAGRank() macro as follows, where

`floor(x)` is the function that evaluates to the greatest integer less than or equal to `x`:

$$\text{DAGRank}(\text{rank}) = \text{floor}(\text{rank}/\text{MinHopRankIncrease})$$

`MinHopRankIncrease` is provisioned at the DODAG Root and propagated in the DIO message. The default value of `MinHopRankIncrease` is `DEFAULT_MIN_HOP_RANK_INCREASE`. For efficient implementation the `MinHopRankIncrease` MUST be a power of 2. An implementation may configure a value `MinHopRankIncrease` as appropriate to balance between the loop avoidance logic of RPL (i.e. selection of eligible parents) and the metrics in use. A further effect of `MinHopRankIncrease` is to impact the number increments that are allowed before `INFINITE_RANK` is reached, i.e. to control how long it may take to count-to-infinity.

By convention in this document, using the macro `DAGRank(node)` may be interpreted as `DAGRank(node.rank)`, where `node.rank` is the rank value as maintained by the node.

A node A has a rank less than the rank of a node B if `DAGRank(A)` is less than `DAGRank(B)`.

A node A has a rank equal to the rank of a node B if `DAGRank(A)` is equal to `DAGRank(B)`.

A node A has a rank greater than the rank of a node B if `DAGRank(A)` is greater than `DAGRank(B)`.

[3.6.2.2](#). Rank Relationships

The computation of the rank MUST be done in such a way so as to maintain the following properties for any nodes M and N that are neighbors in the LLN:

`DAGRank(M)` is less than `DAGRank(N)`: In this case, the position of M is closer to the DODAG root than the position of N. Node M may safely be a DODAG parent for Node N without risk of creating a loop. Further, for a node N, all parents in the

DODAG parent set must be of rank less than DAGRank(N). In other words, the rank presented by a node N MUST be greater than that presented by any of its parents.

DAGRank(M) equals DAGRank(N): In this case the positions of M and N within the DODAG and with respect to the DODAG root are similar (identical). In some cases, Node M may be used as a successor by Node N, which however entails the chance of creating a loop (which must be detected and resolved by some other means).

DAGRank(M) is greater than DAGRank(N): In this case, the position of M is farther from the DODAG root than the position of N. Further, Node M may in fact be in the sub-DODAG of Node N. If node N selects node M as DODAG parent there is a risk to create a loop.

As an example, the rank could be computed in such a way so as to closely track ETX (Expected Transmission Count, a fairly common routing metric used in LLN and defined in [[I-D.ietf-roll-routing-metrics](#)]) when the objective function is to minimize ETX, or latency when the objective function is to minimize latency, or in a more complicated way as appropriate to the objective function being used within the DODAG.

[3.7.](#) Traffic Flows Supported by RPL

RPL supports three basic traffic flows: Multipoint-to-Point (MP2P), Point-to-Multipoint (P2MP), and Point-to-Point (P2P).

[3.7.1.](#) Multipoint-to-Point Traffic

Multipoint-to-Point (MP2P) is a dominant traffic flow in many LLN applications ([[RFC5867](#)], [[RFC5826](#)], [[RFC5673](#)], [[RFC5548](#)]). The destinations of MP2P flows are designated nodes that have some application significance, such as providing connectivity to the larger Internet or core private IP network. RPL supports MP2P traffic by allowing MP2P destinations to be reached via DODAG roots.

[3.7.2.](#) Point-to-Multipoint Traffic

Point-to-multipoint (P2MP) is a traffic pattern required by several

LLN applications ([[RFC5867](#)], [[RFC5826](#)], [[RFC5673](#)], [[RFC5548](#)]). RPL supports P2MP traffic by using a destination advertisement mechanism that provisions routes toward destinations (prefixes, addresses, or multicast groups), and away from roots. Destination advertisements can update routing tables as the underlying DODAG topology changes.

[3.7.3.](#) Point-to-Point Traffic

RPL DODAGs provide a basic structure for point-to-point (P2P) traffic. For a RPL network to support P2P traffic, a root must be able to route packets to a destination. Nodes within the network may also have routing tables to destinations. A packet flows towards a root until it reaches an ancestor that has a known route to the destination. As pointed out later in this document, in the most constrained case (when nodes cannot store routes), that common ancestor may be the DODAG root. In other cases it may be a node closer to both the source and destination.

RPL also supports the case where a P2P destination is a 'one-hop' neighbor.

RPL neither specifies nor precludes additional mechanisms for computing and installing potentially more optimal routes to support arbitrary P2P traffic.

[4.](#) RPL Instance

Within a given LLN, there may be multiple, logically independent RPL instances. A RPL node may belong to multiple RPL instances, and may act as a router in some and as a leaf in others. This document

describes how a single instance behaves.

There are two types of RPL Instances: local and global. Local RPL Instances are always a single DODAG whose singular root owns the corresponding DODAGID. Local RPL Instances can be used for constructing DODAGs that may be used by future on-demand routing solutions that are outside of the scope of this document. Global RPL Instances have one or more DODAGs and are typically long-lived. RPL divides the RPLInstanceID space between global and local instances to allow for both coordinated and unilateral allocation of RPLInstanceIDs.

The definition and provisioning of RPL instances are beyond the scope of this specification. Those operations are expected to be such that data packets coming from the outside of the RPL network can unambiguously be associated to at least one RPL instance, and be safely routed over any instance that would match the packet. Information used to match a packet to a RPL instance can typically be taken from fields in the IPv6 header, like the flow label, TOS bits, or destination address.

Control and data packets within RPL network are tagged to unambiguously identify what RPL Instance they are part of.

Every RPL control message has a RPLInstanceID field. Some RPL control messages, when referring to a local RPLInstanceID as defined below, may also include a DODAGID.

For data packets, the RPLInstanceID may be indicated in the flow label by the source of the packet. If it is not, then it is inferred and added by the RPL network ingress router in the RPL Hop-by-hop option ([\[I-D.hui-6man-rpl-option\]](#)) as further described in [Section 10.2](#)

[4.1.](#) RPL Instance ID

A global RPLInstanceID MUST be unique to the whole LLN. Mechanisms for allocating and provisioning global RPLInstanceID are out of scope for this document. There can be up to 128 global instance in the whole network, and up to 64 local instances per DODAGID.

A global RPLInstanceID is encoded in a RPLInstanceID field as follows:




Figure 3: RPL Instance ID field format for global instances

A local RPLInstanceID is autoconfigured by the node that owns the DODAGID and it MUST be unique for that DODAGID. In that case, the DODAGID MUST be a valid address of the root that is used as an endpoint of all communications within that instance.

A local RPLInstanceID is encoded in a RPLInstanceID field as follows:




Figure 4: RPL Instance ID field format for local instances

The D flag in a Local RPLInstanceID is always set to 0 in RPL control messages. It is used in data packets to indicate whether the DODAGID is the source or the destination of the packet. If the D flag is set to 1 then the destination address of the IPv6 packet MUST be the DODAGID. If the D flag is clear then the source address of the IPv6 packet MUST be the DODAGID.

5. ICMPv6 RPL Control Message

This document defines the RPL Control Message, a new ICMPv6 message. A RPL Control Message is identified by a code, and composed of a base that depends on the code, and a series of options.

A RPL Control Message has the scope of a link. The source address is a link local address. The destination address is either the RPL routers multicast address or a link local address. The RPL routers multicast address is a new address with a requested value of FF02::1:A (to be confirmed by IANA).

In accordance with [\[RFC4443\]](#), the RPL Control Message consists of an ICMPv6 header followed by a message body. The message body is comprised of a message base and possibly a number of options as illustrated in Figure 5.

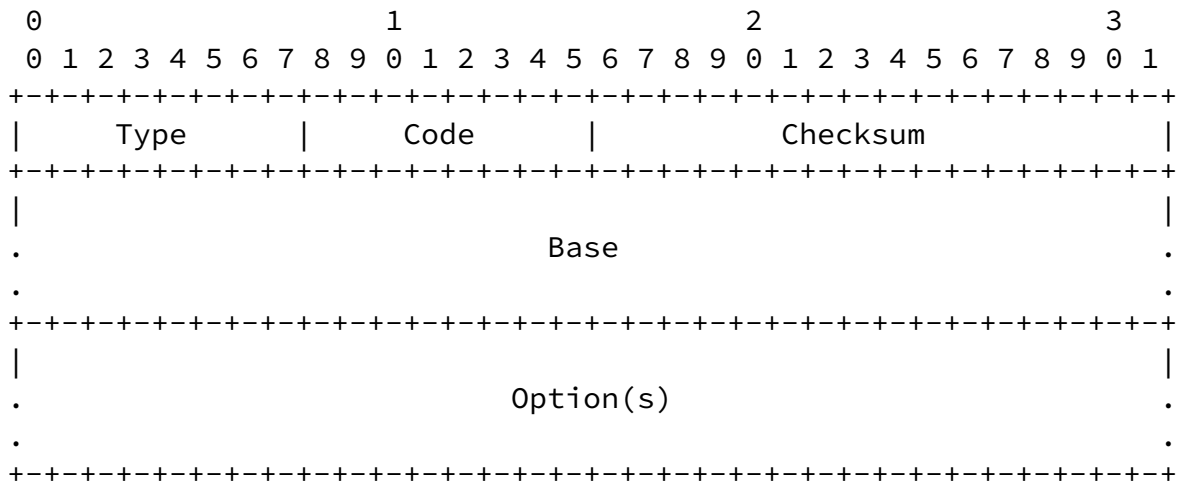


Figure 5: RPL Control Message

The RPL Control message is an ICMPv6 information message with a requested Type of 155 (to be confirmed by IANA).

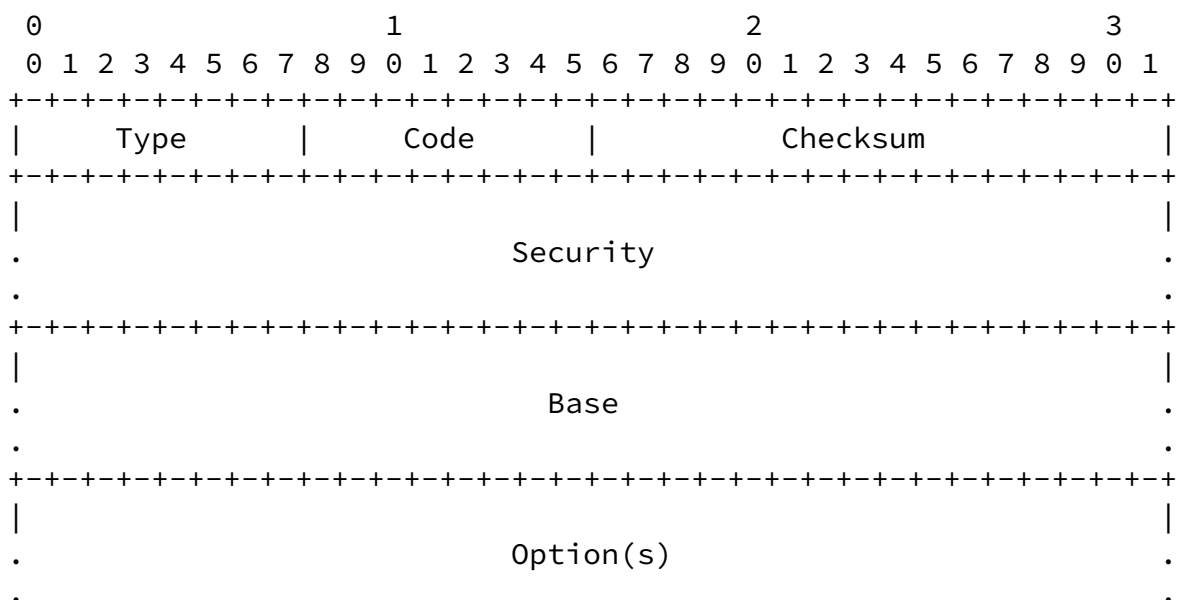
The Code field identifies the type of RPL Control Message. This document defines codes for the following RPL Control Message types (all codes are to be confirmed by the IANA [Section 18.2](#)):

- o 0x00: DODAG Information Solicitation ([Section 5.2](#))

- o 0x01: DODAG Information Object ([Section 5.3](#))
- o 0x02: Destination Advertisement Object ([Section 5.4](#))

- o 0x03: Destination Advertisement Object Acknowledgment ([Section 5.5](#))
- o 0x80: Secure DODAG Information Solicitation ([Section 5.2.2](#))
- o 0x81: Secure DODAG Information Object ([Section 5.3.2](#))
- o 0x82: Secure Destination Advertisement Object ([Section 5.4.2](#))
- o 0x83: Secure Destination Advertisement Object Acknowledgment ([Section 5.5.2](#))
- o 0x8A: Consistency Check ([Section 5.6](#))

The high order bit (0x80) of the code denotes whether the RPL message has security enabled. Secure RPL messages have a format to support confidentiality and integrity, illustrated in Figure 6.



+++++

Figure 6: Secure RPL Control Message

The remainder of this section describes the currently defined RPL Control Message Base formats followed by the currently defined RPL Control Message Options.

5.1. RPL Security Fields

Each RPL message has a secure version. The secure versions provide integrity and replay protection as well as optional confidentiality and delay protection. Because security covers the base message as

well as options, in secured messages the security information lies between the checksum and base, as shown in Figure Figure 6.

The format of the security section is as follows:

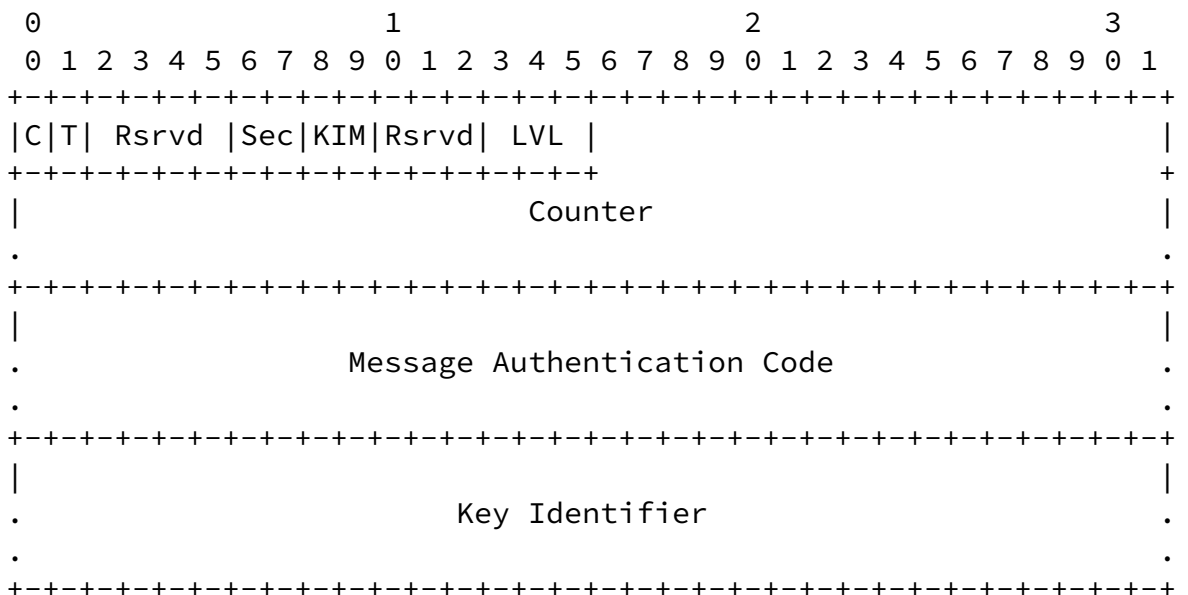


Figure 7: Security Section

All fields are considered as packet payload from a security processing perspective. The exact placement and format of message

integrity/authentication codes has not yet been determined.

Use of the Security section is further detailed in [Section 17](#).

Security Control Field: The Security Control Field has one flag and three fields:

Counter Compression (C): If the Counter Compression flag is set then the Counter field is compressed from 4 bytes into 1 byte. If the Counter Compression flag is clear then the Counter field is 4 bytes and uncompressed.

Counter is Time (T): If the Counter is Time flag is set then the Counter field is a timestamp. If the flag is cleared then the Counter is an incrementing counter. [Section 9.4](#) describes the details of the 'T' flag and Counter field.

Security Mode (Sec): The security algorithm field specifies what security mode and algorithms the network uses. Supported values of this field are as follows:

ID	Sec	Algorithm
0	00	CCM* with AES-128
1	01	Reserved
2	10	Reserved
3	11	Reserved

Security Mode (Sec) Encoding

Key Identifier Mode (KIM): The Key Identifier Mode field indicates whether the key used for packet protection is determined implicitly or explicitly and indicates the particular representation of the Key Identifier field.

The Key Identifier Mode is set one of the non-reserved values from the table below:

Mode	KIM	Meaning	Key Identifier Length (octets)
0	00	Group key used. Key determined by Key Index field. Key Source is not present. Key Index is present.	1
1	01	Per-pair key used.	0

		Key determined by source and destination of packet.	
		Key Source is not present.	
		Key Index is not present.	
2	10	Group key used.	9
		Key determined by Key Index and Key Source Identifier.	
		Key Source is present.	
		Key Index is present.	
3	11	Node's signature key used.	0/9
		If packet is encrypted, group key used. Group key determined by Key Index and Key Source Identifier.	
		Key Source may be present.	
		Key Index may be present.	

Key Identifier Mode (KIM) Encoding

Security Level (LVL): The Security Level field indicates the provided packet protection. This value can be adapted on a per-packet basis and allows for varying levels of data authenticity and, optionally, for data confidentiality. The KIM field indicates whether signatures are used. The Security Level is set to one of the non-reserved values in the table below:

		Without Signatures		With Signatures	
ID	LVL	Attributes	Auth Len	Attributes	Sig Len
0	000	Reserved	N/A	Reserved	N/A

	1		001		MAC-32		4		Sign-32		40	
	2		010		MAC-64		8		Sign-64		44	
	3		011		Reserved		N/A		Sign-128		52	
	4		100		Reserved		N/A		Reserved		N/A	
	5		101		ENC-MAC-32		4		ENC-Sign-32		40	
	6		110		ENC-MAC-64		8		ENC-Sign-64		44	
	7		111		Reserved		N/A		ENC-Sign-128		52	
+-----+-----+-----+-----+-----+-----+-----+												

Security Level (LVL) Encoding

Counter: The Counter field indicates the non-repeating value (nonce) used with the cryptographic mechanism that implements packet protection and allows for the provision of semantic security. This value is compressed from 4 octets to 1 octet if the Counter Compression field of the Security Control Field is set to one.

Message Authentication Code: The Message Authentication Code field contains a cryptographic MAC. The length of the MAC is defined by a combination of the LVL and Sec fields: it can be 0, 4, or 8 octets long. In the case of Security Modes where the MAC is computed as part of the ciphertext (as in Security Mode 0, CCM*), the MAC field is zero bytes long.

Key Identifier: The Key Identifier field indicates which key was used to protect the packet. This field provides various levels of granularity of packet protection, including peer-to-peer keys, group keys, and signature keys. This field is represented as indicated by the Key Identifier Mode field and is formatted as follows:




Figure 8: Key Identifier

Key Source: The Key Source field, when present, indicates the logical identifier of the originator of a group key. When present this field is 8 bytes in length.


Key Index: The Key Index field, when present, allows unique identification of different keys with the same originator. It is the responsibility of each key originator to make sure that actively used keys that it issues have distinct key indices and that all key indices have a value unequal to 0x00. Value 0x00 is reserved for a pre-installed, shared key. When present this field is 1 byte in length.

Unassigned bits of the Security section are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

5.2. DODAG Information Solicitation (DIS)

The DODAG Information Solicitation (DIS) message may be used to solicit a DODAG Information Object from a RPL node. Its use is analogous to that of a Router Solicitation as specified in IPv6 Neighbor Discovery; a node may use DIS to probe its neighborhood for nearby DODAGs. [Section 7.3](#) describes how nodes respond to a DIS.

5.2.1. Format of the DIS Base Object



Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Figure 9: The DIS Base Object

Unassigned bits of the DIS Base are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[5.2.2.](#) Secure DIS

A Secure DIS message follows the format in Figure Figure 6, where the base format is the DIS message shown in Figure Figure 9.

[5.2.3.](#) DIS Options

The DIS message MAY carry valid options.

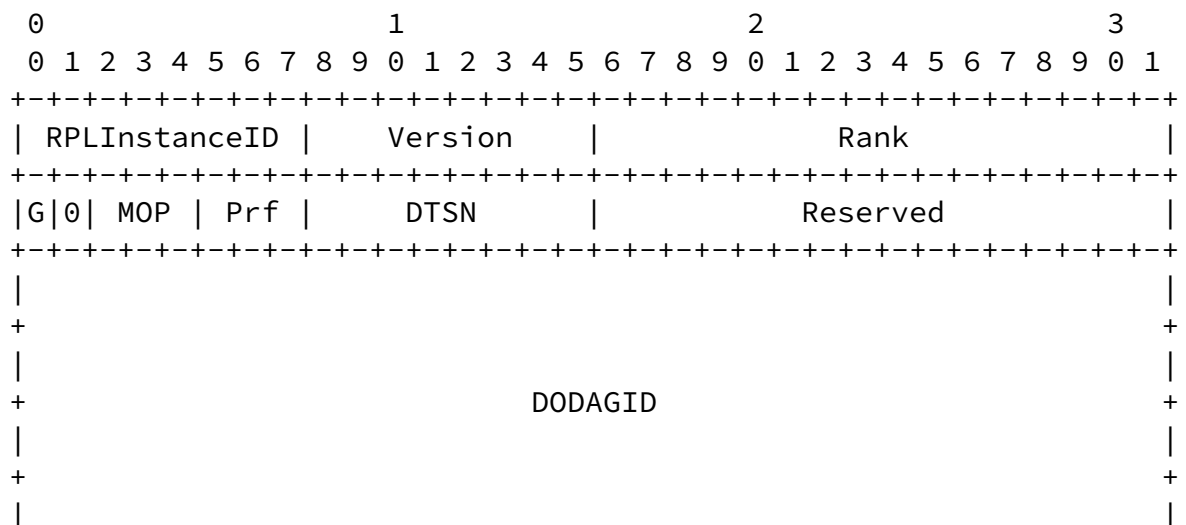
This specification allows for the DIS message to carry the following options:

```
0x00 Pad1
0x01 PadN
0x07 Solicited Information
```

[5.3.](#) DODAG Information Object (DIO)

The DODAG Information Object carries information that allows a node to discover a RPL Instance, learn its configuration parameters, select a DODAG parent set, and maintain the upward routing topology.

[5.3.1.](#) Format of the DIO Base Object



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Option(s)...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 10: The DIO Base Object

Control Field: The DAG Control Field has three flags and two fields:

Grounded (G): The Grounded (G) flag indicates whether the DODAG advertised can satisfy the application-defined goal. If the flag is set, the DODAG is grounded. If the flag is cleared, the DODAG is floating.

Mode of Operation (MOP): The Mode of Operation (MOP) field identifies the mode of operation of the RPL Instance as administratively provisioned at and distributed by the DODAG Root. All nodes who join the DODAG must be able to honor the MOP in order to fully participate as a router, or else they must only join as a leaf. MOP is encoded as in the table below:

MOP		Meaning
000		No downward routes maintained by RPL
001		Non storing mode
010		Storing without multicast support
011		Storing with multicast support
		All other values are reserved

A value of 000 indicates that destination advertisement messages are disabled and the DODAG maintains only upward routes

Mode of Operation (MOP) Encoding

DODAGPreference (Prf): A 3-bit unsigned integer that defines how preferable the root of this DODAG is compared to other DODAG roots within the instance. DAGPreference

ranges from 0x00 (least preferred) to 0x07 (most preferred). The default is 0 (least preferred). [Section 7.2](#) describes how DAGPreference affects DIO processing.

Version Number: 8-bit unsigned integer set by the DODAG root. [Section 7.2](#) describes the rules for version numbers and how they affect DIO processing.

Rank: 16-bit unsigned integer indicating the DODAG rank of the node sending the DIO message. [Section 7.2](#) describes how Rank is set and how it affects DIO processing.

RPLInstanceID: 8-bit field set by the DODAG root that indicates which RPL Instance the DODAG is part of.

Destination Advertisement Trigger Sequence Number (DTSN): 8-bit unsigned integer set by the node issuing the DIO message. The Destination Advertisement Trigger Sequence Number (DTSN) flag is used as part of the procedure to maintain downward routes. The details of this process are described in [Section 8](#).

DODAGID: 128-bit unsigned integer set by a DODAG root which uniquely identifies a DODAG. Possibly derived from the IPv6 address of the DODAG root.

Unassigned bits of the DIO Base are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[5.3.2](#). Secure DIO

A Secure DIO message follows the format in Figure Figure 6, where the base format is the DIS message shown in Figure Figure 10.

[5.3.3](#). DIO Options

The DIO message MAY carry valid options.

This specification allows for the DIO message to carry the following options:

- 0x00 Pad1
- 0x01 PadN
- 0x02 Metric Container
- 0x03 Routing Information
- 0x04 DODAG Configuration
- 0x08 Prefix Information

[5.4.](#) Destination Advertisement Object (DAO)

The Destination Advertisement Object (DAO) is used to propagate destination information upwards along the DODAG. The DAO message is unicast by the child to the selected parent(s). The DAO message may optionally, upon explicit request or error, be acknowledged by the parent with a Destination Advertisement Acknowledgement (DAO-ACK) message back to the child.

[5.4.1.](#) Format of the DAO Base Object




Figure 11: The DAO Base Object

RPLInstanceID: 8-bit field indicating the topology instance

associated with the DODAG, as learned from the DIO.

K: The 'K' flag indicates that the parent is expected to send a DAO-ACK back.

D: The 'D' flag indicates that the DODAGID field is present. This flag MUST be set when a local RPLInstanceID is used.

DAOSequence: Incremented at each unique DAO message, echoed in the DAO-ACK message.

DODAGID (optional): 128-bit unsigned integer set by a DODAG root which uniquely identifies a DODAG. This field is only present when the 'D' flag is set. This field is typically only present when a local RPLInstanceID is in use, in order to identify the DODAGID that is associated with the RPLInstanceID. When a global RPLInstanceID is in use this field need not be present.

Unassigned bits of the DAO Base are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[5.4.2.](#) Secure DAO

A Secure DAO message follows the format in Figure Figure 6, where the base format is the DAO message shown in Figure Figure 11.

[5.4.3.](#) DAO Options

The DAO message MAY carry valid options.

This specification allows for the DAO message to carry the following options:

- 0x00 Pad1
- 0x01 PadN
- 0x05 RPL Target
- 0x06 Transit Information

A special case of the DAO message, termed a No-Path, is used to clear downward routing state that has been provisioned through DAO operation. The No-Path carries a RPL Transit Information option, which identifies the destination to which the DAO is associated, with

a lifetime of 0x00000000 to indicate a loss of reachability.

5.5. Destination Advertisement Object Acknowledgement (DAO-ACK)

The DAO-ACK message is sent as a unicast packet by a DAO parent in response to a unicast DAO message from a child.

5.5.1. Format of the DAO-ACK Base Object

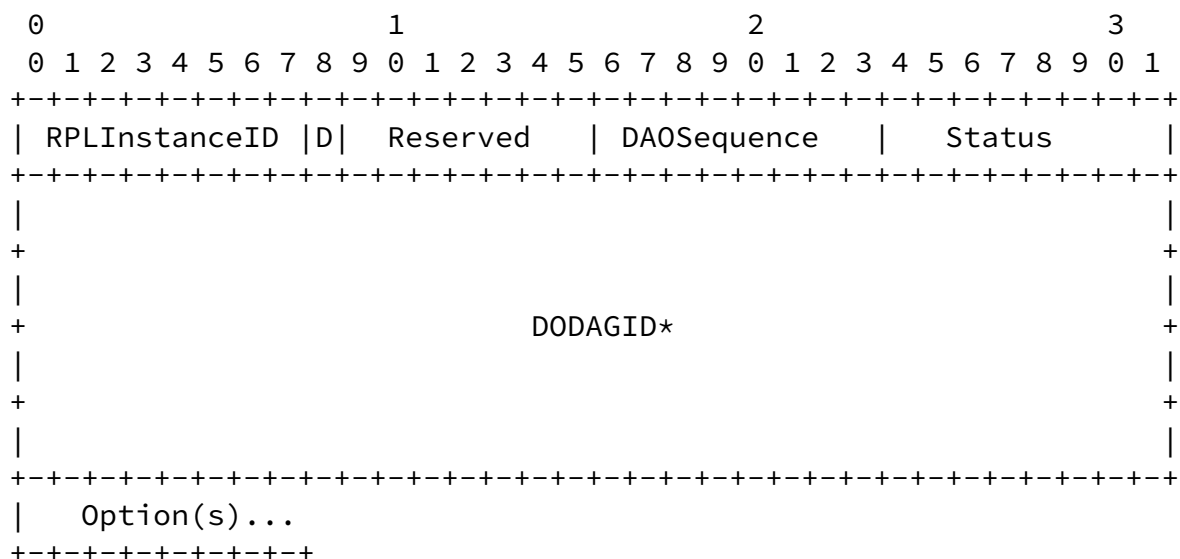


Figure 12: The DAO ACK Base Object

RPLInstanceID: 8-bit field indicating the topology instance associated with the DODAG, as learned from the DIO.

D: The 'D' flag indicates that the DODAGID field is present. This would typically only be set when a local RPLInstanceID is used.

DAOSequence: Incremented at each DAO message from a given child, echoed in the DAO-ACK by the parent. The DAOSequence serves in the parent-child communication and is not to be confused with the Transit Information option Sequence that is associated to a given target down the DODAG.

Status: Indicates the completion. 0 is unqualified acceptance, above 128 are rejection code indicating that the node should select an alternate parent.

DODAGID (optional): 128-bit unsigned integer set by a DODAG root which uniquely identifies a DODAG. This field is only present when the 'D' flag is set. This field is typically only present when a local RPLInstanceID is in use, in order to identify the DODAGID that is associated with the RPLInstanceID. When a global RPLInstanceID is in use this field need not be present.

Unassigned bits of the DAO-ACK Base are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[5.5.2.](#) Secure DAO-ACK

A Secure DAO-ACK message follows the format in Figure Figure 6, where the base format is the DAO-ACK message shown in Figure Figure 12.

[5.5.3.](#) DAO-ACK Options

This specification does not define any options to be carried by the DAO-ACK message.

[5.6.](#) Consistency Check (CC)

The CC message is used to check secure message counters and issue challenge/responses. A CC message MUST be sent as a secured RPL message.

A CC message (request or response) MUST NOT set the 'C' bit of the security section: CC messages always have full counters.

[5.6.1.](#) Format of the CC Base Object

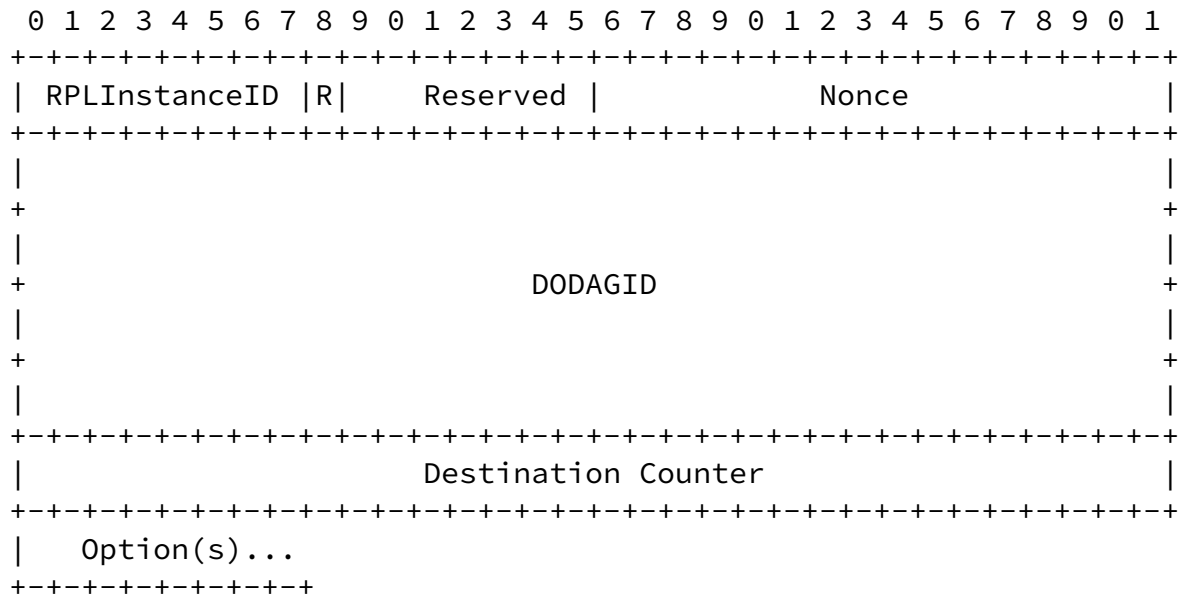


Figure 13: The CC Base Object

RPLInstanceID: 8-bit field indicating the topology instance associated with the DODAG, as learned from the DIO.

R: The 'R' flag indicates whether the CC message is a response. A message with the 'R' flag cleared is a request; a message with the 'R' flag set is a response. A CC message with the R bit set MUST NOT compress the security Counter field: the C bit of the security section MUST be 0.

Nonce: 16-bit unsigned integer set by a CC request. The corresponding CC response includes the same nonce value as the request.

Destination Counter: 32-bit unsigned integer value indicating the sender's estimate of the destination's current security Counter value. If the sender does not have an estimate, it SHOULD set the Destination Counter field to zero.

Unassigned bits of the CC Base are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

The Destination Counter value allows new or recovered nodes to resynchronize through CC message exchanges. This is important to ensure that a Counter value is not repeated for a given security key even in the event of devices recovering from a failure that created a loss of Counter state. For example, where a CC request or other RPL message is received with an initialized Counter within the message security section, the provision of the Incoming Counter within the CC

response message allows the requesting node to reset its Outgoing Counter to a value greater than the last value received by the responding node; the Incoming Counter will also be updated from the received CC response.

[5.6.2.](#) CC Options

The CC message MAY carry valid options. In the scope of this specification, there are no valid options for a CC message.

This specification allows for the CC message to carry the following options:

0x00 Pad1
0x01 PadN

[5.7.](#) RPL Control Message Options

[5.7.1.](#) RPL Control Message Option Generic Format

RPL Control Message Options all follow this format:




Figure 14: RPL Option Generic Format

Option Type: 8-bit identifier of the type of option. The Option Type values are to be confirmed by the IANA [Section 18.4](#).

Option Length: 8-bit unsigned integer, representing the length in octets of the option, not including the Option Type and Length fields.

Option Data: A variable length field that contains data specific to the option.

When processing a RPL message containing an option for which the Option Type value is not recognized by the receiver, the receiver MUST silently ignore the unrecognized option and continue to process the following option, correctly handling any remaining options in the message.

RPL message options may have alignment requirements. Following the

convention in IPv6, options with alignment requirements are aligned in a packet such that multi-octet values within the Option Data field

of each option fall on natural boundaries (i.e., fields of width n octets are placed at an integer multiple of n octets from the start of the header, for $n = 1, 2, 4$, or 8).

[5.7.2.](#) Pad1

The Pad1 option may be present in DIS, DIO, DAO, and DAO-ACK messages, and its format is as follows:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+
|   Type = 0   |
+---+---+---+---+
```

Figure 15: Format of the Pad 1 Option

The Pad1 option is used to insert one or two octets of padding into the message to enable options alignment. If more than one octet of padding is required, the PadN option should be used rather than multiple Pad1 options.

NOTE! the format of the Pad1 option is a special case - it has neither Option Length nor Option Data fields.

[5.7.3.](#) PadN

The PadN option may be present in DIS, DIO, DAO, and DAO-ACK messages, and its format is as follows:

```

0                               1                               2
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type = 1   | Option Length | 0x00 Padding...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Figure 16: Format of the Pad N Option

The PadN option is used to insert two or more octets of padding into the message to enable options alignment. PadN Option data MUST be ignored by the receiver.

Option Type: 0x01 (to be confirmed by IANA)

Option Length: For N (N > 1) octets of padding, the Option Length field contains the value N-2.

Option Data: For N (N > 1) octets of padding, the Option Data consists of N-2 zero-valued octets.

5.7.4. Metric Container

The Metric Container option may be present in DIO messages, and its format is as follows:




Figure 17: Format of the Metric Container Option

The Metric Container is used to report metrics along the DODAG. The Metric Container may contain a number of discrete node, link, and aggregate path metrics and constraints specified in [\[I-D.ietf-roll-routing-metrics\]](#) as chosen by the implementer.

The Metric Container MAY appear more than once in the same RPL control message, for example to accommodate a use case where the Metric Data is longer than 256 bytes. More information is in [\[I-D.ietf-roll-routing-metrics\]](#)

The processing and propagation of the Metric Container is governed by implementation specific policy functions.

Option Type: 0x02 (to be confirmed by IANA)

Option Length: The Option Length field contains the length in octets of the Metric Data.

Metric Data: The order, content, and coding of the Metric Container data is as specified in [I-D.ietf-roll-routing-metrics].

5.7.5. Route Information

The Route Information option may be present in DIO messages, and is equivalent in function to the IPv6 ND Route Information option as defined in [\[RFC4191\]](#). The format of the option is modified slightly

Winter, et al.

Expires December 30, 2010

[Page 40]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

(Type, Length, Prefix) in order to be carried as a RPL option as follows:

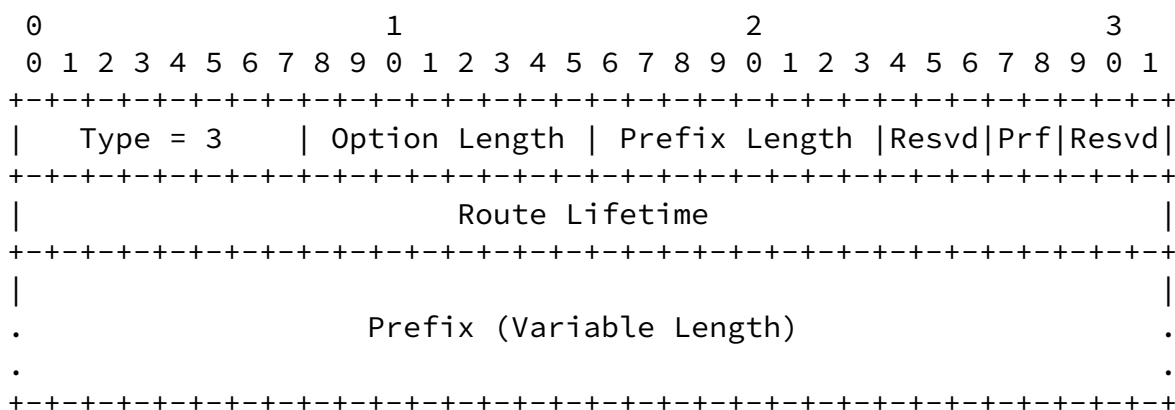


Figure 18: Format of the Route Information Option

The Route Information option is used to indicate that connectivity to the specified destination prefix is available from the DODAG root.

In the event that a RPL Control Message may need to specify connectivity to more than one destination, the Route Information option may be repeated.

[RFC4191] should be consulted as the authoritative reference with respect to the Route Information option. The field descriptions are transcribed here for convenience:

Option Type: 0x03 (to be confirmed by IANA)

Option Length: Variable, length of the option in octets excluding the Type and Length fields. Note that this length is expressed in units of single-octets, unlike in IPv6 ND.

Prefix Length 8-bit unsigned integer. The number of leading bits in the Prefix that are valid. The value ranges from 0 to 128. The Prefix field has the number of bytes inferred from the Option Length field, that must be at least the Prefix Length. Note that in RPL this means that the Prefix field may have lengths other than 0, 8, or 16.

Prf: 2-bit signed integer. The Route Preference indicates whether to prefer the router associated with this prefix over others, when multiple identical prefixes (for different routers) have been received. If the Reserved (10) value is received, the Route Information Option MUST be ignored.

Resvd: Two 3-bit unused fields. They MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Route Lifetime 32-bit unsigned integer. The length of time in seconds (relative to the time the packet is sent) that the prefix is valid for route determination. A value of all one bits (0xffffffff) represents infinity.

Prefix Variable-length field containing an IP address or a prefix of an IP address. The Prefix Length field contains the number of valid leading bits in the prefix. The bits in the prefix after the prefix length (if any) are reserved and MUST be initialized to zero by the sender and ignored by the receiver. Note that in RPL this field may have lengths other than 0, 8, or 16.

Unassigned bits of the Route Information option are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

5.7.6. DODAG Configuration

The DODAG Configuration option may be present in DIO messages, and its format is as follows:




Figure 19: Format of the DODAG Configuration Option

The DODAG Configuration option is used to distribute configuration information for DODAG Operation through the DODAG.

The information communicated in this option is generally static and unchanging within the DODAG, therefore it is not necessary to include in every DIO. This information is configured at the DODAG Root and distributed throughout the DODAG with the DODAG Configuration Option. Nodes other than the DODAG Root MUST NOT modify this information when propagating the DODAG Configuration option. This option MAY be included occasionally by the DODAG Root (as determined by the DODAG Root), and MUST be included in response to a unicast request, e.g. a unicast DODAG Information Solicitation (DIS) message.

Option Type: 0x04 (to be confirmed by IANA)

Option Length: 8 bytes

Authentication Enabled (A): One bit describing the security mode of the network. The bit describe whether a node must authenticate with a key authority before joining the network as a router. If the DIO is not a secure DIO, the 'A' bit MUST be zero.

Path Control Size (PCS): 3-bit unsigned integer used to configure the number of bits that may be allocated to the Path Control

field (see [Section 8.9](#)). Note that as used a value of 1 is added to this field, i.e. a PCS value of 0 results in 1 active bit in the Path Control field. The default value of PCS is DEFAULT_PATH_CONTROL_SIZE.

DIOIntervalDoublings: 8-bit unsigned integer used to configure I_{max} of the DIO trickle timer (see [Section 7.3.1](#)).

DIOIntervalMin: 8-bit unsigned integer used to configure I_{min} of the DIO trickle timer (see [Section 7.3.1](#)).

DIORedundancyConstant: 8-bit unsigned integer used to configure k of the DIO trickle timer (see [Section 7.3.1](#)).

MaxRankIncrease: 16-bit unsigned integer used to configure DAGMaxRankIncrease, the allowable increase in rank in support of local repair. If DAGMaxRankIncrease is 0 then this mechanism is disabled.

MinHopRankInc 16-bit unsigned integer used to configure MinHopRankIncrease as described in [Section 3.6.2.1](#).

Objective Code Point (OCP) 16-bit unsigned integer. The OCP field identifies the OF and is managed by the IANA.

[5.7.7](#). RPL Target

The RPL Target option format is as follows:

```

      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 = 5      | Option Length |   Reserved   | Prefix Length |

```




Figure 20: Format of the RPL Target Option

The RPL Target Option is used to indicate a target IPv6 address, prefix, or multicast group that is reachable or queried along the DODAG. In a DIO, the RPL Target Option identifies a resource that the root is trying to reach. In a DAO, the RPL Target option indicates reachability.

A set of one or more Transit Information options MAY directly follow the Target option in a DAO message in support of constructing source routes in a non-storing mode of operation [[I-D.hui-6man-rpl-routing-header](#)]. When the same set of Transit Information options apply equally to a set of DODAG Target options, the group of Target options MUST appear first, followed by the Transit Information options which apply to those Targets.

The RPL Target option may be repeated as necessary to indicate multiple targets.

Option Type: 0x05 (to be confirmed by IANA)

Option Length: Variable, length of the option in octets excluding the Type and Length fields.

Prefix Length: 8-bit unsigned integer. Number of valid leading bits in the IPv6 Prefix.

Target Prefix: Variable-length field identifying an IPv6 destination address, prefix, or multicast group. The Prefix Length field contains the number of valid leading bits in the prefix. The bits in the prefix after the prefix length (if any) are reserved and MUST be set to zero on transmission and MUST be ignored on receipt.

5.7.8. Transit Information

The Transit Information option may be present in DAO messages, and its format is as follows:

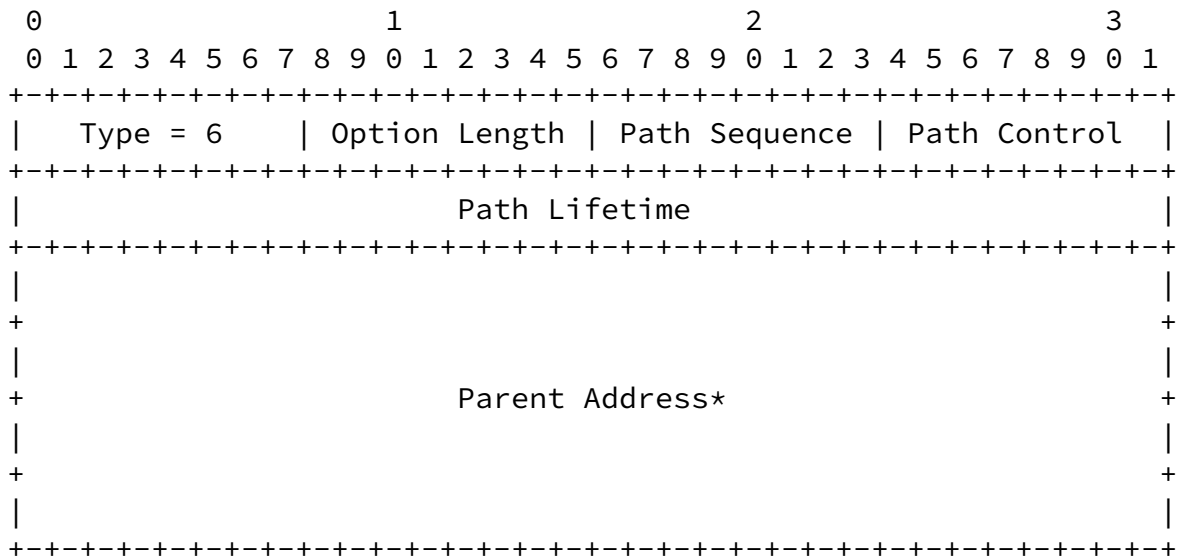


Figure 21: Format of the Transit Information option

The Transit Information option is used for a node to indicate attributes for a path to one or more destinations. The destinations are indicated as by one or more Target options that immediately precede the Transit Information option(s).

The Transit Information option can be used for a node to indicate its DODAG parents to an ancestor that is collecting DODAG routing information, typically for the purpose of constructing source routes. In the non-storing mode of operation this ancestor will be the DODAG Root, and this option is carried by the DAO message. The option length is used to determine whether the Parent Address is present or not.

A non-storing node that has more than one DAO parent MAY include a Transit Information option for each DAO parent as part of the non-storing Destination Advertisement operation. The node may code the Path Control field in order to signal a preference among parents.

One or more Transit Information options MUST be preceded by one or more RPL Target options. In this manner the RPL Target option indicates the child node, and the Transit Information option(s) enumerate the DODAG parents.

A typical non-storing node will use multiple Transit Information options, and it will send the DAO thus formed to only one parent that will forward it to the root. A typical storing node will use one Transit Information option with no parent field, and will send the DAO thus formed to multiple parents.

Option Type: 0x06 (to be confirmed by IANA)

Option Length: Variable, depending on whether or not Parent Address is present.

Path-Sequence: 8-bit unsigned integer. When a RPL Target option is issued by the node that owns the Target Prefix (i.e. in a DAO message), that node sets the Path-Sequence and increments the Path-Sequence each time it issues a RPL Target option.

Path Control: 8-bit bitfield. The Path Control field limits the number of DAO-Parents to which a DAO message advertising connectivity to a specific destination may be sent, as well as providing some indication of relative preference. The limit provides some bound on overall DAO fan-out in the LLN. The leftmost bit is associated with a path that contains a most-preferred link, and the subsequent bits are ordered down to the rightmost bit which is least preferred.

Path Lifetime: 32-bit unsigned integer. The length of time in seconds (relative to the time the packet is sent) that the prefix is valid for route determination. A value of all one bits (0xFFFFFFFF) represents infinity. A value of all zero bits (0x00000000) indicates a loss of reachability. This is referred to as a No-Path in this document.

Parent Address (optional): IPv6 Address of the DODAG Parent of the node originally issuing the Transit Information Option. This field may not be present, as according to the DODAG Mode of Operation and indicated by the Transit Information option length.

Unassigned bits of the Transit Information option are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

5.7.9. Solicited Information

The Solicited Information option may be present in DIS messages, and its format is as follows:

Winter, et al.

Expires December 30, 2010

[Page 46]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

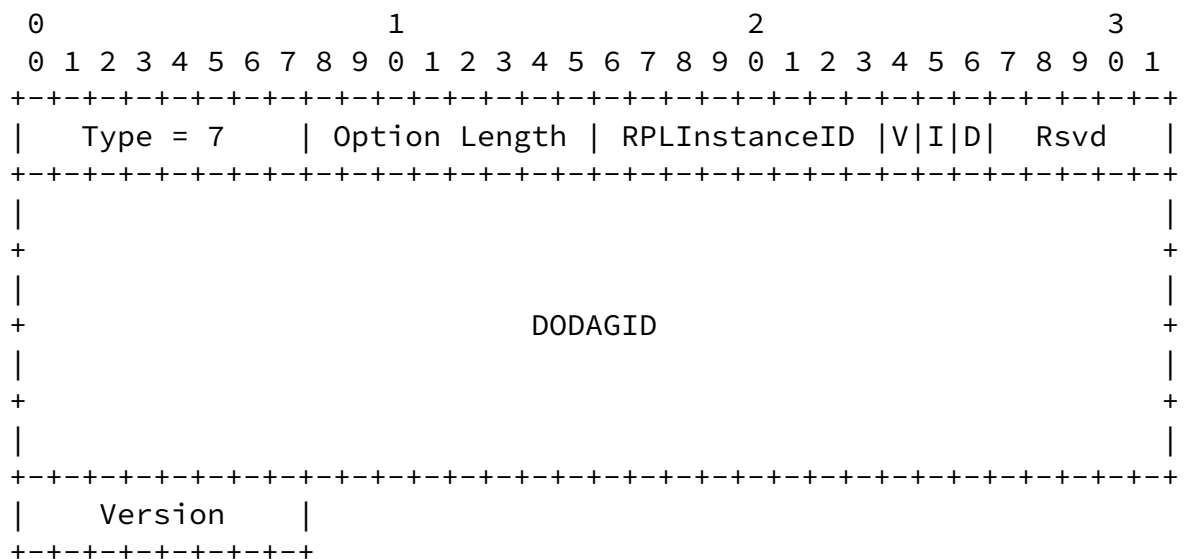


Figure 22: Format of the Solicited Information Option

The Solicited Information option is used for a node to request DIO messages from a subset of neighboring nodes. The Solicited Information option may specify a number of predicate criteria to be matched by a receiving node. These predicates affect whether a node resets its DIO trickle timer, as described in [Section 7.3](#)

Option Type: 0x07 (to be confirmed by IANA)

Option Length: 19 bytes

Control Field: The Solicited Information option Control Field has three flags:

V: If the V flag is set then the Version field is valid and

a node matches the predicate if its DODAGVersionNumber matches the requested version. If the V flag is clear then the Version field is not valid and the Version field MUST be set to zero on transmission and ignored upon receipt.

- I: If the I flag is set then the RPLInstanceID field is valid and a node matches the predicate if it matches the requested RPLInstanceID. If the I flag is clear then the RPLInstanceID field is not valid and the RPLInstanceID field MUST be set to zero on transmission and ignored upon receipt.

- D: If the D flag is set then the DODAGID field is valid and a node matches the predicate if it matches the requested DODAGID. If the D flag is clear then the DODAGID field is not valid and the DODAGID field MUST be set to zero on transmission and ignored upon receipt.

Version: 8-bit unsigned integer containing the DODAG Version number that is being solicited when valid.

RPLInstanceID: 8-bit unsigned integer containing the RPLInstanceID that is being solicited when valid.

DODAGID: 128-bit unsigned integer containing the DODAGID that is being solicited when valid.

Unassigned bits of the Solicited Information option are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[5.7.10](#). Prefix Information

The Prefix Information option may be present in DIO messages, and is equivalent in function to the IPv6 ND Prefix Information option as defined in [\[RFC4861\]](#). The format of the option is modified slightly (Type, Length, Prefix) in order to be carried as a RPL option as follows:

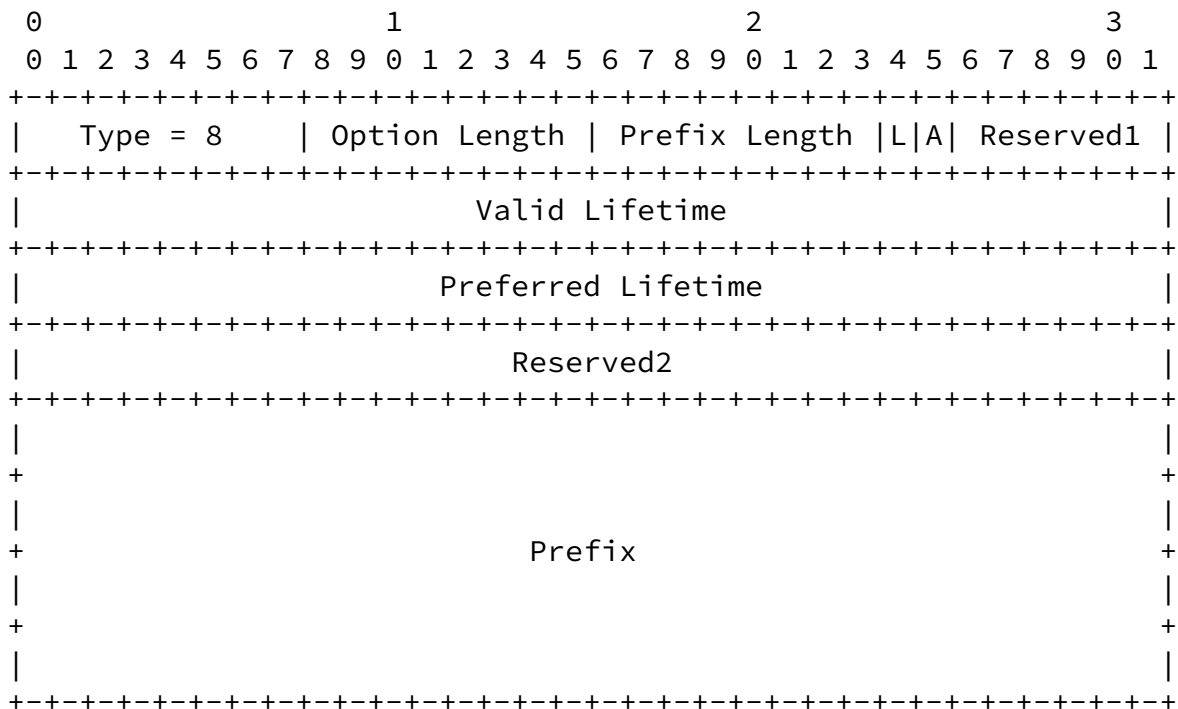


Figure 23: Format of the Prefix Information Option

The Prefix Information option may be used to distribute the prefix in use inside the DODAG, e.g. for address autoconfiguration.

[RFC4861] should be consulted as the authoritative reference with respect to the Prefix Information option. The field descriptions are transcribed here for convenience:

Option Type: 0x08 (to be confirmed by IANA)

Option Length: 30. Note that this length is expressed in units of single-octets, unlike in IPv6 ND.

Prefix Length 8-bit unsigned integer. The number of leading bits in the Prefix that are valid. The value ranges from 0 to 128. The prefix length field provides necessary information for on-link determination (when combined with the L flag in the prefix information option). It also assists with address

autoconfiguration as specified in [[RFC4862](#)], for which there may be more restrictions on the prefix length.

- L 1-bit on-link flag. When set, indicates that this prefix can be used for on-link determination. When not set the advertisement makes no statement about on-link or off-link properties of the prefix. In other words, if the L flag is not set a host MUST NOT conclude that an address derived from the prefix is off-link. That is, it MUST NOT update a previous indication that the address is on-link.
- A 1-bit autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address configuration as specified in [[RFC4862](#)].

Reserved1 6-bit unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Valid Lifetime 32-bit unsigned integer. The length of time in seconds (relative to the time the packet is sent) that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. The Valid Lifetime is also used by [[RFC4862](#)].

Preferred Lifetime 32-bit unsigned integer. The length of time in seconds (relative to the time the packet is sent) that addresses generated from the prefix via stateless address autoconfiguration remain preferred [[RFC4862](#)]. A value of all one bits (0xffffffff) represents infinity. See [[RFC4862](#)].

Note that the value of this field MUST NOT exceed the Valid Lifetime field to avoid preferring addresses that are no longer valid.

Reserved2 This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Prefix An IP address or a prefix of an IP address. The Prefix Length field contains the number of valid leading bits in the prefix. The bits in the prefix after the prefix length are reserved and MUST be initialized to zero by the sender and ignored by the receiver. A router SHOULD NOT send a prefix

option for the link-local prefix and a host SHOULD ignore such a prefix option. A non-storing node SHOULD refrain from advertising a prefix till it owns an address of that prefix, and then it SHOULD advertise its full address in this field, to be used by its children in the Parent Address field of the Transit Information Option

Unassigned bits of the Prefix Information option are reserved. They MUST be set to zero on transmission and MUST be ignored on reception.

[6.](#) Sequence Counters

This section describes the general scheme for bootstrap and operation of sequence counters in RPL, such as the DODAGVersionNumber in the DIO message, the DAOSequence in the DAO message, and the Path-

Sequence in the Transit Information option.

RPL sequence counters are subdivided in a 'lollipop' fashion ([[Perlman83](#)]), where the values from 128 and greater are used as a linear sequence to indicate a restart and bootstrap the counter, and the values less than or equal to 127 used as a circular sequence number space of size 128 as in [[RFC1982](#)]. Consideration is given to the mode of operation when transitioning from the linear region to the circular region. Finally, when operating in the circular region, if sequence numbers are detected to be too far apart then they are not comparable, as detailed below.

A window of comparison, `SEQUENCE_WINDOW` = 16, is configured based on a value of 2^N , where $N=4$.

For a given sequence counter,

1. The sequence counter SHOULD be initialized to an implementation defined value which is 128 or greater prior to use. A recommended value is 240 (256 - `SEQUENCE_WINDOW`).
2. When a sequence counter increment would cause the sequence counter to increment beyond its maximum value, the sequence counter MUST wrap back to zero. When incrementing a sequence counter greater than or equal to 128, the maximum value is 255. When incrementing a sequence counter less than 128, the maximum value is 127.
3. When comparing two sequence counters, the following rules MUST be applied:
 1. When a first sequence counter A is in the interval [0..127] and a second sequence counter B is in [128..255]:
 1. If $B-A$ is less than or equal to `SEQUENCE_WINDOW`, then B is greater than A, A is less than B, and the two are not equal.
 2. If $B-A$ is greater than `SEQUENCE_WINDOW`, then A is greater than B, B is less than A, and the two are not equal.

2. In the case where both sequence counters to be compared are less than or equal to 127, and in the case where both sequence counters to be compared are greater than or equal to 128:
 1. If the absolute magnitude of difference between the two sequence counters is less than or equal to SEQUENCE_WINDOW, then a comparison as described in [\[RFC1982\]](#) is used to determine the relationships greater than, less than, and equal
 2. If the absolute magnitude of difference of the two sequence counters is greater than SEQUENCE_WINDOW, then a desynchronization has occurred and the two sequence numbers are not comparable.
4. If two sequence numbers are determined to be not comparable, i.e. the results of the comparison are not defined, then a node should consider the comparison as if it has evaluated in such a way so as to give precedence to the sequence number that has most recently been observed to increment. Failing this, the node should consider the comparison as if it has evaluated in such a way so as to minimize the resulting changes to its own state.

[7.](#) Upward Routes

This section describes how RPL discovers and maintains upward routes. It describes the use of DODAG Information Objects (DIOs), the messages used to discover and maintain these routes. It specifies how RPL generates and responds to DIOs. It also describes DODAG Information Solicitation (DIS) messages, which are used to trigger DIO transmissions.

[7.1.](#) DIO Base Rules

1. For the following DIO Base fields, a node that is not a DODAG root MUST advertise the same values as its preferred DODAG parent (defined in [Section 7.2.1](#)). Therefore, if a DODAG root does not change these values, every node in a route to that DODAG root eventually advertises the same values for these fields. These fields are:
 1. Grounded (G)
 2. Mode of Operation (MOP)
 3. DAGPreference (Prf)
 4. Version
 5. RPLInstanceID
 6. DODAGID
2. A node MAY update the following fields at each hop:
 1. Rank
 2. DTSN
3. The DODAGID field each root sets MUST be unique within the RPL Instance.

[7.2.](#) Upward Route Discovery and Maintenance

Upward route discovery allows a node to join a DODAG by discovering neighbors that are members of the DODAG of interest and identifying a set of parents. The exact policies for selecting neighbors and parents is implementation-dependent and driven by the OF. This section specifies the set of rules those policies must follow for interoperability.

[7.2.1.](#) Neighbors and Parents within a DODAG Version

RPL's upward route discovery algorithms and processing are in terms of three logical sets of link-local nodes. First, the candidate neighbor set is a subset of the nodes that can be reached via link-local multicast. The selection of this set is implementation-dependent and OF-dependent. Second, the parent set is a restricted subset of the candidate neighbor set. Finally, the preferred parent,

a set of size one, is an element of the parent set that is the preferred next hop in upward routes.

More precisely:

1. The DODAG parent set MUST be a subset of the candidate neighbor set.
2. A DODAG root MUST have a DODAG parent set of size zero.
3. A node that is not a DODAG root MAY maintain a DODAG parent set of size greater than or equal to one.
4. A node's preferred DODAG parent MUST be a member of its DODAG parent set.
5. A node's rank MUST be greater than all elements of its DODAG parent set.
6. When Neighbor Unreachability Detection (NUD), or an equivalent mechanism, determines that a neighbor is no longer reachable, a RPL node MUST NOT consider this node in the candidate neighbor set when calculating and advertising routes until it determines that it is again reachable. Routes through an unreachable neighbor MUST be removed from the routing table.

These rules ensure that there is a consistent partial order on nodes within the DODAG. As long as node ranks do not change, following the above rules ensures that every node's route to a DODAG root is loop-free, as rank decreases on each hop to the root.

The OF can guide candidate neighbor set and parent set selection, as discussed in [[I-D.ietf-roll-routing-metrics](#)] and [[I-D.ietf-roll-of0](#)].

[7.2.2.](#) Neighbors and Parents across DODAG Versions

The above rules govern a single DODAG version. The rules in this section define how RPL operates when there are multiple DODAG versions:

[7.2.2.1](#). DODAG Version

1. The tuple (RPLInstanceID, DODAGID, DODAGVersionNumber) uniquely defines a DODAG Version. Every element of a node's DODAG parent set, as conveyed by the last heard DIO message from each DODAG parent, MUST belong to the same DODAG version. Elements of a node's candidate neighbor set MAY belong to different DODAG Versions.

2. A node is a member of a DODAG version if every element of its DODAG parent set belongs to that DODAG version, or if that node is the root of the corresponding DODAG.
3. A node MUST NOT send DIOs for DODAG versions of which it is not a member.
4. DODAG roots MAY increment the DODAGVersionNumber that they advertise and thus move to a new DODAG version. When a DODAG root increments its DODAGVersionNumber, it MUST follow the conventions of Serial Number Arithmetic as described in [Section 6](#).
5. Within a given DODAG, a node that is not a root MUST NOT advertise a DODAGVersionNumber higher than the highest DODAGVersionNumber it has heard. Higher is defined as the greater-than operator in [Section 6](#).
6. Once a node has advertised a DODAG version by sending a DIO, it MUST NOT be member of a previous DODAG version of the same DODAG (i.e. with the same RPLInstanceID, the same DODAGID, and a lower DODAGVersionNumber). Lower is defined as the less-than operator in [Section 6](#).

When the DODAG parent set becomes empty on a node that is not a root, (i.e. the last parent has been removed, causing the node to no longer be associated with that DODAG), then the DODAG information should not be suppressed until after the expiration of an implementation-

specific local timer in order to observe if the DODAGVersionNumber has been incremented, should any new parents appear for the DODAG. This will help protect against the possibility of loops that may occur of that node were to inadvertently rejoin the old DODAG version in its own prior sub-DODAG.

As the DODAGVersionNumber is incremented, a new DODAG Version spreads outward from the DODAG root. A parent that advertises the new DODAGVersionNumber cannot belong to the sub-DODAG of a node advertising an older DODAGVersionNumber. Therefore a node can safely add a parent of any Rank with a newer DODAGVersionNumber without forming a loop.

Exactly when a DODAG Root increments the DODAGVersionNumber is implementation and application-dependent and outside the scope of this document. Examples include incrementing the DODAGVersionNumber periodically, upon administrative intervention, or on application-level detection of lost connectivity or DODAG inefficiency.

After a node transitions to and advertises a new DODAG Version, the

rules above make it unable to advertise the previous DODAG Version (prior DODAGVersionNumber) once it has committed to advertising the new DODAG Version.

[7.2.2.2](#). DODAG Roots

1. A DODAG root without possibility to satisfy the application-defined goal MUST NOT set the Grounded bit.
2. A DODAG root MUST advertise a rank of ROOT_RANK.
3. A node whose DODAG parent set is empty MAY become the DODAG Root of a floating DODAG. It MAY also set its DAGPreference such that it is less preferred.

In a deployment that uses a backbone link to federate a number of LLN roots, it is possible to run RPL over that backbone and use one router as a "backbone root". The backbone root is the virtual root of the DODAG, and exposes a rank of BASE_RANK over the backbone. All the LLN roots that are parented to that backbone root, including the backbone root if it also serves as LLN root itself, expose a rank of

ROOT_RANK to the LLN. These virtual roots are part of the same DODAG and advertise the same DODAGID. They coordinate DODAGVersionNumbers and other DODAG parameters with the virtual root over the backbone.

[7.2.2.3.](#) DODAG Selection

The objective function of a DAG determines how a node selects its neighbor set, parent set, and preferred parents. This selection implicitly also decides the DODAG within a DAG. Such selection can include administrative preference (Prf) as well as metrics or other considerations.

If a node has the option to join a more preferred DODAG while still meeting other optimization objectives, then the node will generally seek to join the more preferred DODAG as determined by the OF. All else being equal, it is left to the implementation to determine which DODAG is most preferred.

[7.2.2.4.](#) Rank and Movement within a DODAG Version

1. A node MUST NOT advertise a Rank less than or equal to any member of its parent set within the DODAG Version.
2. A node MAY advertise a Rank lower than its prior advertisement within the DODAG Version.

3. Let L be the lowest rank within a DODAG version that a given node has advertised. Within the same DODAG Version, that node MUST NOT advertise an effective rank higher than $L + \text{DAGMaxRankIncrease}$. INFINITE_RANK is an exception to this rule: a node MAY advertise an INFINITE_RANK within a DODAG version without restriction. If a node's Rank would be higher than allowed by $L + \text{DAGMaxRankIncrease}$, when it advertises Rank it MUST advertise its Rank as INFINITE_RANK.
4. A node MAY, at any time, choose to join a different DODAG within a RPL Instance. Such a join has no rank restrictions, unless that different DODAG is a DODAG Version of which this node has previously been a member, in which case the rule of the previous bullet (3) must be observed. Until a node transmits a DIO

indicating its new DODAG membership, it MUST forward packets along the previous DODAG.

5. A node MAY, at any time after hearing the next DODAGVersionNumber advertised from suitable DODAG parents, choose to migrate to the next DODAG Version within the DODAG.

Conceptually, an implementation is maintaining a DODAG parent set within the DODAG Version. Movement entails changes to the DODAG parent set. Moving up does not present the risk to create a loop but moving down might, so that operation is subject to additional constraints.

When a node migrates to the next DODAG Version, the DODAG parent set needs to be rebuilt for the new version. An implementation could defer to migrate for some reasonable amount of time, to see if some other neighbors with potentially better metrics but higher rank announce themselves. Similarly, when a node jumps into a new DODAG it needs to construct new a DODAG parent set for this new DODAG.

If a node needs to move down a DODAG that it is attached to, increasing its Rank, then it MAY poison its routes and delay before moving as described in [Section 7.2.2.5](#).

[7.2.2.5](#). Poisoning

1. A node poisons routes by advertising a Rank of INFINITE_RANK.
2. A node MUST NOT have any nodes with a Rank of INFINITE_RANK in its parent set.

Although an implementation may advertise INFINITE_RANK for the purposes of poisoning, doing so is not the same as setting Rank to INFINITE_RANK. For example, a node may continue to send data packets

whose meta-data include a Rank that is not INFINITE_RANK yet still advertise INFINITE_RANK in its DIOs.

[7.2.2.6](#). Detaching

1. A node unable to stay connected to a DODAG within a given DODAG version MAY detach from this DODAG version. A node that detaches

becomes root of its own floating DODAG and SHOULD immediately advertise this new situation in a DIO as an alternate to poisoning.

[7.2.2.7.](#) Following a Parent

1. If a node receives a DIO from one of its DODAG parents, indicating that the parent has left the DODAG, that node SHOULD stay in its current DODAG through an alternative DODAG parent, if possible. It MAY follow the leaving parent.

A DODAG parent may have moved, migrated to the next DODAG Version, or jumped to a different DODAG. A node should give some preference to remaining in the current DODAG, if possible via an alternate parent, but ought to follow the parent if there are no other options.

[7.2.3.](#) DIO Message Communication

When an DIO message is received, the receiving node must first determine whether or not the DIO message should be accepted for further processing, and subsequently present the DIO message for further processing if eligible.

1. If the DIO message is malformed, then the DIO message is not eligible for further processing and a node MUST silently discard it.
2. If the sender of the DIO message is a member of the candidate neighbor set and the DIO message is not malformed, the node MUST process the DIO.

[7.2.3.1.](#) DIO Message Processing

As DIO messages are received from candidate neighbors, the neighbors may be promoted to DODAG parents by following the rules of DODAG discovery as described in [Section 7.2](#). When a node places a neighbor into the DODAG parent set, the node becomes attached to the DODAG through the new DODAG parent node.

The most preferred parent should be used to restrict which other nodes may become DODAG parents. Some nodes in the DODAG parent set

may be of a rank less than or equal to the most preferred DODAG parent. (This case may occur, for example, if an energy constrained device is at a lesser rank but should be avoided as per an optimization objective, resulting in a more preferred parent at a greater rank).

[7.3.](#) DIO Transmission

RPL nodes transmit DIOs using a Trickle timer ([\[I-D.ietf-roll-trickle\]](#)). A DIO from a sender with a lower DAGRank that causes no changes to the recipient's parent set, preferred parent, or Rank SHOULD be considered consistent with respect to the Trickle timer.

The following packets and events MUST be considered inconsistencies with respect to the Trickle timer, and cause the Trickle timer to reset:

- o When a node detects an inconsistency when forwarding a packet, as detailed in [Section 10.2](#).
- o When a node receives a multicast DIS message without a Solicited Information option.
- o When a node receives a multicast DIS with a Solicited Information option and the node matches all of the predicates in the Solicited Information option.
- o When a node joins a new DODAG Version (e.g. by updating its DODAGVersionNumber, joining a new RPL Instance, etc.)

Note that this list is not exhaustive, and an implementation MAY consider other messages or events to be inconsistencies.

A node SHOULD NOT reset its DIO trickle timer in response to unicast DIS messages. When a node receives a unicast DIS without a Solicited Information option, it MUST unicast a DIO to the sender in response. This DIO MUST include a DODAG Configuration option. When a node receives a unicast DIS message with a Solicited Information option, if it satisfies the predicates of the Solicited Information option it MUST unicast a DIO to the sender in response. This unicast DIO MUST include a DODAG Configuration Option. Thus a node may transmit a unicast DIS message to a potential DODAG parent in order to probe for DODAG Configuration and other parameters.

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

[7.3.1.](#) Trickle Parameters

The configuration parameters of the trickle timer are specified as follows:

Imin: learned from the DIO message as $(2^{\text{DIOIntervalMin}})\text{ms}$. The default value of DIOIntervalMin is DEFAULT_DIO_INTERVAL_MIN.

Imax: learned from the DIO message as DIOIntervalDoublings. The default value of DIOIntervalDoublings is DEFAULT_DIO_INTERVAL_DOUBLINGS.

k: learned from the DIO message as DIORedundancyConstant. The default value of DIORedundancyConstant is DEFAULT_DIO_REDUNDANCY_CONSTANT. In RPL, when k has the value of 0x00 this is to be treated as a redundancy constant of infinity in RPL, i.e. Trickle never suppresses messages.

[7.4.](#) DODAG Selection

The DODAG selection is implementation and OF dependent. Nodes SHOULD prefer to join DODAGs for RPLInstanceIDs advertising OCPs and destinations compatible with their implementation specific objectives. In order to limit erratic movements, and all metrics being equal, nodes SHOULD keep their previous selection. Also, nodes SHOULD provide a means to filter out a parent whose availability is detected as fluctuating, at least when more stable choices are available.

When connection to a grounded DODAG is not possible or preferable for security or other reasons, scattered DODAGs MAY aggregate as much as possible into larger DODAGs in order to allow connectivity within the LLN.

A node SHOULD verify that bidirectional connectivity and adequate link quality is available with a candidate neighbor before it considers that candidate as a DODAG parent.

[7.5.](#) Operation as a Leaf Node

In some cases a RPL node may attach to a DODAG as a leaf node only. One example of such a case is when a node does not understand the RPL Instance's OF or advertised metric/constraint. As specified in

[Section 16.6](#) related to policy function, the node may either join the DODAG as a leaf node or may not join the DODAG. As mentioned in [Section 16.5](#), it is then recommended to log a fault.

A leaf node does not extend DODAG connectivity but in some cases the

leaf node may still need to transmit DIOs on occasion, in particular when the leaf node may not have always been acting as a leaf node and an inconsistency is detected.

A node operating as a leaf node must obey the following rules:

1. It MUST NOT transmit DIOs containing the DAG Metric Container.
2. Its DIOs MUST advertise a DAGRank of INFINITE_RANK.
3. It MAY suppress DIO transmission, except DIO transmission MUST NOT be suppressed when DIO transmission has been triggered due to detection of inconsistency when a packet is being forwarded or in response to a unicast DIS message.
4. It MAY transmit unicast DAOs as described in [Section 8.2](#).
5. It MAY transmit multicast DAOs to the '1 hop' neighborhood as described in [Section 8.10](#).

A particular case that requires a leaf node to send a DIO is if that leaf node was a prior member of another DODAG and another node forwards a message assuming the old topology, triggering an inconsistency. The leaf node needs to transmit a DIO in order to repair the inconsistency. Note that due to the lossy nature of LLNs, even though the leaf node may have optimistically poisoned its routes by advertising a rank of INFINITE_RANK in the old DODAG prior to becoming a leaf node, that advertisement may have become lost and a leaf node must be capable to send a DIO later in order to repair the inconsistency.

In general it is not expected that such a leaf node would advertise itself as a router.

[7.6](#). Administrative Rank

In some cases it might be beneficial to adjust the rank advertised by a node beyond that computed by the OF based on some implementation specific policy and properties of the node. For example, a node that has limited battery should be a leaf unless there is no other choice, and may then augment the rank computation specified by the OF in order to expose an exaggerated rank.

[8.](#) Downward Routes

This section describes how RPL discovers and maintains downward routes. RPL constructs and maintains downward routes with Destination Advertisement Object (DAO) messages. Downward routes support of P2MP flows, from the DODAG roots toward the leaves. Downward routes also support P2P flows: P2P messages can flow to a DODAG Root through an upward route, then away from the DODAG Root to a destination through a downward route.

This specification describes the two modes a RPL Instance may choose from for maintaining downward routes. In the first mode, call "storing," nodes store downward routing tables for their sub-DODAG. Each hop on a downward route in a storing network examines its routing table to decide on the next hop. In the second mode, called "non-storing," nodes do not store downward routing tables. Downward packets are routed with source routes populated by a DODAG Root.

RPL allows a simple one-hop P2P optimization for both storing and non-storing networks. A node may send a P2P packet destined to a one-hop neighbor directly to that node.

[8.1.](#) Destination Advertisement Parents

To establish downward routes, RPL nodes send DAO messages upwards. The next hop destinations of these DAO messages are called DAO parents. The collection of a node's DAO parents is called the DAO parent set.

- o A node's DAO parent set MUST be a subset of its DODAG parent set.
- o A node MUST NOT unicast DAOs to nodes that are not DAO parents.
- o A node MAY link-local multicast DAO messages.
- o The IPv6 Source Address of a DAO message MUST be the link local address of the sending node.
- o If a node sends a DAO to one DAO parent, it MUST send a DAO with the same DAOSequence to all other DAO parents.

The selection of DAO parents is implementation and objective function specific.

[8.2.](#) Downward Route Discovery and Maintenance

Destination Advertisement may be configured to be entirely disabled, or operate in either a storing or non-storing mode, as reported in

the MOP in the DIO message.

1. All nodes who join a DODAG MUST abide by the MOP setting from the root. Nodes that do not have the capability to fully participate as a router MAY join the DODAG as a leaf.
2. If the MOP is 000, indicating no downward routing, nodes MUST NOT transmit DAO messages, and MAY ignore DAO messages.
3. In non-storing mode, the DODAG Root MUST store source routing table entries for all destinations learned from DAOs.
4. In storing mode, all non-root, non-leaf nodes MUST store routing table entries for all destinations learned from DAOs.

A DODAG can have one of several possible modes of operation, as defined by the MOP field. Either it does not support downward routes, it supports downward routes through source routing from DODAG Roots, or it supports downward routes through in-network routing tables. When downward routes are supported through in-network routing tables, the multicast operation defined in this specification may or may not be supported, also as indicated by the MOP field. As

of this specification RPL does not support mixed-mode operation, where some nodes source route and other store routing tables: future extensions to RPL may support this mode of operation.

[8.3.](#) DAO Base Rules

1. Each time a node generates a new DAO, the DAOSequence field MUST increment by at least one since the last generated DAO.
2. Each time a node link-local multicasts a DAO, the DAOSequence field MUST increment by one since the last link local multicast DAO.
3. The RPLInstanceID and DODAGID fields of a DAO MUST be the same value as the members of the node's parent set and the DIOs it transmits.
4. A node MAY set the K flag in a unicast DAO message to solicit a unicast DAO-ACK in response in order to confirm the attempt. A node receiving a unicast DAO message with the K flag set SHOULD respond with a DAO-ACK. A node receiving a DAO message without the K flag set MAY respond with a DAO-ACK, especially to report an error condition.
5. Nodes SHOULD ignore DAOs without newer sequence numbers and MUST NOT process them further.

Unlike the Version field of a DIO, which is incremented only by a DODAG Root and repeated unchanged by other nodes, DAOSequence values are unique to each node. The sequence number space for unicast and multicast DAO messages can be either the same or distinct.

[8.4.](#) DAO Transmission Scheduling

Because DAOs flow upwards, receiving a unicast DAO can trigger sending a unicast DAO.

1. On receiving a unicast DAO with a new DAOSequence, a node SHOULD send a DAO. It SHOULD NOT send this DAO immediately. It SHOULD delay sending the DAO in order to aggregate DAO information from other nodes for which it is a DAO parent.

2. A node SHOULD delay sending a DAO with a timer (DelayDAO). Receiving a DAO starts the DelayDAO timer. DAOs received while the DelayDAO timer is active do not reset the timer. When the DelayDAO timer expires, the node sends a DAO.
3. When a node adds a node to its DAO parent set, it SHOULD schedule a DAO transmission.

DelayDAO's value and calculation is implementation-dependent.

8.5. Triggering DAO Messages

Nodes can trigger their sub-DODAG to send DAO messages. Each node maintains a DAO Trigger Sequence Number (DTSN), which it communicates through DIO messages.

1. If a node hears one of its DAO parents increment its DTSN, the node MUST schedule a DAO transmission using rules in [Section 8.3](#) and [Section 8.4](#).
2. In non-storing mode, if a node hears one of its DAO parents increment its DTSN, the node MUST increment its own DTSN.

In a storing mode of operation, a storing node MAY increment DTSN in order to reliably trigger a set of DAO updates from its immediate children, as part of routine routing table updates and maintenance. In a storing mode of operation it is not necessary to trigger DAO updates from the entire sub-DODAG, since that state information will percolate hop-by-hop up the DODAG in the storing mode of operation.

In a non-storing mode of operation, a DTSN increment will also cause the immediate children of a node to increment their DTSN in turn, triggering a set of DAO updates from the entire sub-DODAG. In a non-

storing mode of operation typically only the root would independently increment the DTSN when a DAO refresh is needed but a global repair (such as by incrementing DODAGVersionNumber) is not desired. In a non-storing mode of operation typically all non-root nodes would only increment their DTSN when their parent(s) are observed to do so.

In the case of triggered DAOs, selecting a proper DAODelay can greatly reduce the number of DAOs transmitted. The trigger flows

down the DODAG; in the best case the DAOs flow up the DODAG such that leaves send DAOs first, with each node sending a DAO only once. Such a scheduling could be approximated by setting DAODelay inversely proportional to Rank. Note that this suggestion is intended as an optimization to allow efficient aggregation -- it is not required for correct operation in the general case.

[8.6.](#) Structure of DAO Messages

DAOs follow a common structure in both storing and non-storing networks. Later sections describe further details for each mode of operation.

1. RPL nodes MUST include one or more RPL Target Options in each DAO they transmit. One RPL Target Option MUST have a prefix that includes the node's IPv6 address if that node needs the DODAG to provision downward routes to that node.
2. A RPL Target Option in a unicast DAO MUST be followed by a Transit Information Option.
3. Multicast DAOs MUST NOT include Transit Information options.
4. If a node receives a DAO that does not follow the above three rules, it MUST discard the DAO without further processing.

[8.7.](#) Non-storing Mode

In non-storing mode, RPL routes messages downward using source routing. The following rule applies to nodes that are in non-storing mode. Storing mode has a separate set of rules, described in [Section 8.8](#).

1. The Parent Address field of a Transit Information Option MUST contain one or more addresses. All of these addresses MUST be addresses of DAO parents of the sender.
2. On receiving a unicast DAO, a node MUST forward the DAO upwards. This forwarding MAY use any parent in the parent set. Note that this forwarding may be delayed in support of aggregation as

described below, but that such a delay is not required if a

node's resources do not support it.

3. When a node removes a node from its DAO parent set, it MAY generate a new DAO with an updated Transit Information option.

In non-storing mode, a node uses DAOs to report its DAO parents to the DODAG Root. The DODAG Root can piece together a downward route to a node by using DAO parent sets from each node in the route. The purpose of this per-hop route calculation is to minimize traffic when DAO parents change. If nodes reported complete source routes, then on a DAO parent change the entire sub-DODAG would have to send new DAOs to the DODAG Root. Therefore, in non-storing mode, a node can send a single DAO, although it might choose to send more than one DAO to each of multiple DAO parents.

Nodes aggregate DAOs by sending a single DAO with multiple RPL Target Options. Each RPL Target Option has its own, immediately following, Transit Information options.

[8.8](#). Storing Mode

In storing mode, RPL routes messages downward by the IPv6 destination address. The following rule apply to nodes that are in storing mode:

1. The Parent Address field of a Transmit Information option MUST be empty.
2. On receiving a unicast DAO, a node MUST compute if the DAO would change the set of prefixes that the node itself advertises. If so, the node MUST generate a new DAO and transmit it, following the rules in [Section 8.4](#). Such a change includes receiving a No-Path DAO.
3. When a node generates a new DAO, it SHOULD unicast it to each of its DAO parents. It MUST NOT unicast the DAO to nodes that are not DAO parents.
4. When a node removes a node from its DAO parent set, it SHOULD send a No-Path DAO ([Section 5.4.3](#)) to that removed DAO parent to invalidate the existing route.
5. If messages to an advertised downwards address suffer from a forwarding error, neighbor unreachable detected (NUD), or similar failure, a node MAY mark the address as unreachable and generate an appropriate No-Path DAO.

DAOs advertise what destination addresses and prefixes a node has

routes to. Unlike in non-storing mode, these DAOs do not communicate information about the routes themselves: that information is stored within the network and is implicit from the IPv6 source address. When a storing node generates a DAO, it uses the stored state of DAOs it has received to produce a set of RPL Target options and their associated Transmit Information options.

Because this information is stored within a network, in storing mode DAOs are communicated directly to DAO parents, who store this information.

[8.9.](#) Path Control

A DAO message from a node contains one or more Target Options. Each Target Option specifies either the node's prefix, a prefix of addresses reachable outside the LLN, or a destination in the node's sub-DODAG. The Path Control field of the Transit Information option allows nodes to request multiple downward routes. A node constructs the Path Control field of a Transit Information option as follows:

1. The bit width of the path control field **MUST** be equal to the value $(PCS + 1)$, where PCS is specified in the control field of the DODAG Configuration Option. Bits greater than or equal to the value $(PCS + 1)$ **MUST** be cleared on transmission and **MUST** be ignored on reception. Bits below that value are considered "active" bits.
2. For a RPL Target option describing a node's own address or a prefix outside the LLN, at least one active bit of the Path Control field **MUST** be set. More active bits of the Path Control field **MAY** be set.
3. If a node receives multiple DAOs with the same RPL Target option, it **MUST** bitwise-OR the Path Control fields it receives. This aggregated bitwise-OR represents the number of downward routes the prefix requests.
4. When a node sends a DAO to one of its DAO parents, it **MUST** select one or more of the set, active bits in the aggregated Path Control field. The DAO it transmits to its parent **MUST** have these active bits set and all other active bits cleared.
5. For the RPL Target option and DAOSequence number, the DAOs a node sends to different DAO parents **MUST** have disjoint sets of active Path Control bits. A node **MUST NOT** set the same active bit on DAOs to two different DAO parents.

6. Path control bits SHOULD be allocated in order of preference, such that the most significant bits, or groupings of bits, are allocated to the most preferred DAO parents as determined by the node.
7. In a non-storing mode of operation, a node MAY pass DAOs through without performing any further processing on the Path Control field.
8. A node MUST NOT unicast a DAO that has no active bits in the Path Control field set.

The Path Control field allows a node to bound how many downward routes will be generated to it. It sets a number of bits in the Path Control field equal to the maximum number of downward routes it prefers. Each bit is sent to at most one DAO parent; clusters of bits can be sent to a single DAO parent for it to divide among its own DAO parents.

[8.10](#). Multicast Destination Advertisement Messages

A special case of DAO operation, distinct from unicast DAO operation, is multicast DAO operation which may be used to populate '1-hop' routing table entries.

1. A node MAY multicast a DAO message to the link-local scope all-nodes multicast address FF02::1.
2. A multicast DAO message MUST be used only to advertise information about self, i.e. prefixes directly connected to or owned by this node, such as a multicast group that the node is subscribed to or a global address owned by the node.
3. A multicast DAO message MUST NOT be used to relay connectivity information learned (e.g. through unicast DAO) from another node.
4. Information obtained from a multicast DAO MAY be installed in the routing table and MAY be propagated by a node in unicast DAOs.

5. A node MUST NOT perform any other DAO related processing on a received multicast DAO, in particular a node MUST NOT perform the actions of a DAO parent upon receipt of a multicast DAO.
 - o The multicast DAO may be used to enable direct P2P communication, without needing the RPL routing structure to relay the packets.
 - o The multicast DAO does not presume any DODAG relationship between the emitter and the receiver.

[9.](#) Security Mechanisms

This section describes the generation and processing of secure RPL messages. The high order bit of the RPL message code identifies whether a RPL message is secure or not. In addition to secure versions of basic control messages (DIS, DIO, DAO, DAO-Ack), RPL has several messages which are relevant only in networks with security enabled.

[9.1.](#) Security Overview

RPL supports three security modes:

- o Insecure. In this security mode, RPL uses insecure DIS, DIO, DAO, and DAO-Ack messages.
- o Pre-installed. In this security mode, RPL uses secure messages. To join a RPL Instance, a node must have a pre-installed key. Nodes use this to provide message confidentiality, integrity, and authenticity. A node may, using this preinstalled key, join the RPL network as either a host or a router.
- o Authenticated. In this security mode, RPL uses secure messages. To join a RPL Instance, a node must have a pre-installed key. Node use this key to provide message confidentiality, integrity, and authenticity. Using this preinstalled key, a node may join the network as a host only. To join the network as a router, a node must obtain a second key from a key authority. This key authority can authenticate that the requester is allowed to be a router before providing it with the second key.

Whether or not the RPL Instance uses insecure mode is signaled by

whether it uses secure RPL messages. Whether a secured network uses the pre-installed or authenticated mode is signaled by the 'A' bit of the DAG Configuration option.

RPL uses CCM* -- Counter with CBC-MAC (Cipher Block Chaining Message Authentication Code) -- as the cryptographic basis for its security[RFC3610]. In this specification, CCM uses AES-128 as its underlying cryptographic algorithm. There are bits reserved in the security section to specify other algorithms in the future.

All secured RPL messages have a message authentication code (MAC). Secured RPL messages optionally also have encryption protection for confidentiality. Secured RPL message formats support both integrated encryption/authentication schemes (e.g., CCM*) as well as schemes that separately encrypt and authenticate packets.

[9.2.](#) Installing Keys

Authenticated mode requires a would-be router to dynamically install new keys once they have joined a network as a host.

The exact message exchange to obtain such keys is TBD. It will involve communication with a key authority, possibly, using the pre-installed shared key. The key authority can apply a security policy to decide whether to grant the would-be-router a new key. These keys may have lifetimes (start and end times) associated with them, which nodes that support timestamps (described in [Section 9.4.1](#)) can use.

[9.3.](#) Joining a Secure Network

RPL security assumes that a node wishing to join a secured network has been preconfigured with a shared key for communicating with neighbors and the RPL root. To join a secure RPL network, a node either listens for secure DIOs or triggers secure DIOs by sending a secure DIS. In addition to the DIO/DIS rules in [Section 7](#), secure DIO and DIS messages have these rules:

1. If sent, this initial secure DIS MUST NOT set the C bit, MUST set the KIM field to 0 (00), and MUST set the LVL field to 1 (001). The key used MUST be the preconfigured group key (Key Index 0x00).

2. When a node resets its Trickle timer in response to a secure DIS ([Section 7.3](#)), the next DIO it transmits MUST be a secure DIO with the same security configuration as the secure DIS. If a node receives multiple secure DIS messages before it transmits a DIO, the secure DIO MUST have the same security configuration as the last DIS it is responding to.
3. When a node sends a DIO in response to a unicast secure DIS ([Section 7.3](#)), the DIO MUST be a secure DIO.

The above rules allow a node to join a secured RPL Instance using the preconfigured shared key. Once a node has joined the DODAG using the preconfigured shared key, the 'A' bit of the Configuration option determines its capabilities. If the 'A' bit of the Configuration is cleared, then nodes can use this preinstalled, shared key to exchange messages normally: it can issue DIOs, DAOs, etc.

If the 'A' bit of the Configuration option is set:

1. A node MUST NOT advertise a Rank besides INFINITE_RANK in secure DIOs secured with Key Index 0x00. If a node receives a secure DIO that advertises a Rank besides INFINITE_RANK and is secured

with Key Index 0x00, it MUST discard the message without further processing.

2. Secure DAOs using Key Index 0x00 MUST NOT have a RPL Target option with a prefix besides the node's address. If a node receives a secured DAO using the preinstalled, shared key where the RPL Target option does not match the IPv6 source address, it MUST discard the secured DAO without further processing.

The above rules mean that in RPL Instances where the 'A' bit is set, using Key Index 0x00 a node can join the RPL Instance as a host but not a router. A node must communicate with a key authority to obtain a key that will enable it to act as a router. Obtaining this key might require authentication on one or both ends. This message exchange is TBD.

[9.4.](#) Counter and Counter Compression

Every secured RPL packet has a Counter field. Depending on whether the 'C' bit is set, this Counter field can be 1 or 4 bits. RPL nodes send CC messages to force uncompressed Counter values, protect against replay attacks and synchronize counters.

1. If a node is sending a secured RPL packet, and the Counter value of the packet is more than 255 greater than the last secured packet to the destination address, the node MUST NOT set the 'C' bit of the security section of the packet.
2. If a node receives a secure RPL message with the C bit set and is uncertain of the 32-bit counter value, it MAY send a CC message with the R bit cleared to obtain an uncompressed counter value. The Nonce field of the CC message SHOULD be a random or pseudorandom number.
3. If a node receives a unicast CC message with the R bit cleared, and it is a member of or is in the process of joining the associated DODAG, it SHOULD respond with a unicast CC message to the sender. This response MUST have the C bit of the security section cleared, MUST have the R bit set, and MUST have the same Nonce, RPLInstanceID and DODAGID fields as the message it received.
4. If a node receives a multicast CC message, it MUST discard the message with no further processing.

These rules allow nodes to compress the Counter when destinations who received the prior packet can determine the full counter value. If a node cannot determine the full counter value, it can request the full

counter with a CC message.

[9.4.1](#). Timestamp Counters

In the simplest case, the Counter value is an unsigned integer that a node increments by one or more on each secured RPL transmission. The Counter MAY represent a timestamp that has the following properties:

1. The timestamp MUST be at least six octets long.
2. The timestamp MUST be in 1kHz (millisecond) granularity.

3. The timestamp start time MUST be January 1, 2010, 12:00:00AM UTC.
4. If the Counter represents such as timestamp, the Counter value MUST be a value computed as follows. Let T be the timestamp, S be the start time of the key in use, and E be the end time of the key in use. Both S and E are represented using the same 3 rules as the timestamp described above. If $E > T < S$, then the Counter is invalid and a node MUST NOT generate a packet. Otherwise, the Counter value is equal to $T-S$.
5. If the Counter represents such a timestamp, a node MAY set the 'T' flag of the security section of secured RPL packets.
6. If the Counter field does not present such a timestamp, then a node MUST NOT set the 'T' flag.
7. If a node does not have a local timestamp that satisfies the above requirements, it MUST ignore the 'T' flag.

If a node supports such timestamps and it receives a message with the 'T' flag set, it MAY apply the temporal check on the received message described in [Section 9.5.2.1](#). If a node receives a message without the 'T' flag set, it MUST NOT apply this temporal check. A node's security policy MAY, for application reasons, include rejecting all messages without the 'T' flag set.

[9.5](#). Functional Description of Packet Protection

[9.5.1](#). Transmission of Outgoing Packets

Given an outgoing RPL control packet and required security protection, this section describes how RPL generates the secured packet to transmit. It also describes the order of cryptographic operations to provide the required protection.

The requirement for security protection and the level of security to

be applied to an outgoing RPL packet shall be determined by the node's security policy database. The configuration of this security policy database for outgoing packet processing is TBD (it may, for example, be defined through DIO Configuration or through out-of-band

administrative router configuration).

Where secured RPL messages are to be transmitted, a RPL node MUST set the security section (C, T, Sec, KIM, and LVL) in the outgoing RPL packet to describe the protection level and security settings that are applied (see [Section 5.1](#)). The Security subfield bit of the RPL message Code field MUST be set to indicate the secure RPL message.

The Counter value used in constructing the Nonce to secure the outgoing packet MUST be an increment of the last Counter transmitted to the particular destination address. Where a Counter for the intended destination address has not been established, the Counter value MUST be initialized to zero and sent as a Full Counter for the initial RPL message transmission.

Where a Counter is currently maintained for outgoing messages to the intended destination address, the Compressed Counter (indicated with the 'C' bit set) MUST be transmitted within the secured RPL message, provided the message is not a RPL Consistency Check message. The current Full Counter (indicated with the 'C' bit cleared) for the given destination address SHALL always be used when the outgoing packet is a Consistency Check (challenge or response) message. Where a Counter for the intended destination address does not exist, the initialized (zero-value), Full Counter MUST be transmitted within the initial RPL control message. Where security policy specifies the application of delay protection, the Timestamp Counter used in constructing the Nonce to secure the outgoing packet MUST be incremented according to the rules in [Section 9.4.1](#). Where a Timestamp Counter is applied (indicated with the 'T' flag set) the locally maintained Time Counter MUST be included as part of the transmitted secured RPL message.

The cryptographic algorithm used in securing the outgoing packet shall be specified by the node's security policy database and MUST be indicated in the value of the Sec field set within the outgoing message.

The security policy for the outgoing packet shall determine the applicable Key Identifier Mode (KIM) and Key Identifier specifying the security key to be used for the cryptographic packet processing, including the optional use of signature keys (see [Section 5.1](#)). The security policy will also specify the level of protection (LVL) in the form of authentication or authentication and encryption, and potential use of signatures that shall apply to the outgoing packet.

Where encryption is applied, a node MUST replace the original packet payload with that payload encrypted using the security protection, key, and nonce specified in the security section of the packet.

All secured RPL messages include integrity protection. In conjunction with the security algorithm processing, a node derives a Message Authentication Code (MAC) that MUST be included as part of the outgoing secured RPL packet.

[9.5.2.](#) Reception of Incoming Packets

This section describes the reception and processing of a secured RPL packet. Given an incoming secured RPL packet, where the Security subfield bit of the RPL message Code field is set, this section describes how RPL generates an unencrypted version of the packet and validates its integrity.

The receiver uses the RPL security control fields to determine the necessary packet security processing. If the described level of security for the message type and originator does not meet locally maintained security policies, a node MAY discard the packet without further processing. These policies can include security levels, keys used, source identifiers, or the lack of timestamp-based counters (as indicated by the 'T' flag). The configuration of the security policy database for incoming packet processing is TBD (it may, for example, be defined through DIO Configuration or through out-of-band administrative router configuration).

Where the message security level (LVL) indicates an encrypted RPL message, the node uses the key information identified through the KIM field as well as the Nonce as input to the message payload decryption processing. The Nonce shall be derived from the message Counter field and other received and locally maintained information (see [Section 9.5.3.1](#)). The plaintext message contents shall be obtained by invoking the inverse cryptographic mode of operation specified by the Sec field of the received packet.

The receiver shall use the Nonce and identified key information to check the integrity of the incoming packet. If the integrity check fails against the received message authentication code (MAC), a node MUST discard the packet.

If a Compressed Counter is received and the node does not currently have an incoming Counter currently maintained for the originator of the message, the node MUST send a Consistency Check request to the message source to update the Counters.

If an initialized (zero value) Full Counter is received in a secured

RPL message and the receiving node currently has an incoming Counter currently maintained for the originator of the message, the node MUST initiate a Counter resynchronization by sending a Consistency Check response message (see [Section 5.6.1](#)) to the message source. The Consistency Check response message shall be protected with the current full outgoing Counter maintained for the particular node address. That outgoing Counter will be included within the security section of the message while the incoming Counter will be included within the Consistency Check message payload.

Based on the specified security policy a node MAY apply replay protection for a received RPL message. The replay check MUST be performed following the authentication of the received packet. The full Counter, as obtained from the incoming packet or as derived from the received Compressed Counter shall be compared against the watermark of the incoming Counter maintained for the given origination node address. If the received message Counter value is non-zero and less than the maintained incoming Counter watermark a potential packet replay is indicated and the node MUST discard the incoming packet.

If delay protection is specified as part of the incoming packet security policy checks, the Timestamp Counter is used to validate the timeliness of the received RPL message. If the incoming message Timestamp Counter value indicates a message transmission time prior to the locally maintained transmission time Counter for the originator address, a replay violation is indicated and the node MUST discard the incoming packet. If the received Timestamp Counter value indicates a message transmission time that is earlier than the Current time less the acceptable packet delay, a delay violation is indicated and the node MUST discard the incoming packet.

Once a message has been decrypted, where applicable, and has successfully passed its integrity check, replay, and optionally delay protection checks, the node can update its local security information, such as the source's expected Counter value for counter compression and replay comparison.

A node MUST NOT update its security information on receipt of a message that fails security policy checks or other applied integrity,

replay, or delay checks.

[9.5.2.1. Timestamp Key Checks](#)

If the 'T' flag of a message is set and a node has a local timestamp that follows the requirements in [Section 9.4.1](#), then a node MAY check the temporal consistency of the message. The node computes the transmit time of the message by adding the Counter value to the start

time of the associated key. If this transmit time is past the end time of the key, the node MAY discard the message without further processing. If the transmit time is too far in the past or future compared to the local time on the receiver, it MAY discard the message without further processing.

[9.5.3. Cryptographic Mode of Operation](#)

The cryptographic mode of operation used is based on the CCM mode of operation and the block-cipher AES-128[RFC3610]. This mode of operation is widely supported by existing implementations and coincides with the CCM* mode of operation[CCMStar]. CCM mode requires a nonce.

[9.5.3.1. Nonce](#)

A RPL node constructs a CCM nonce as follows:

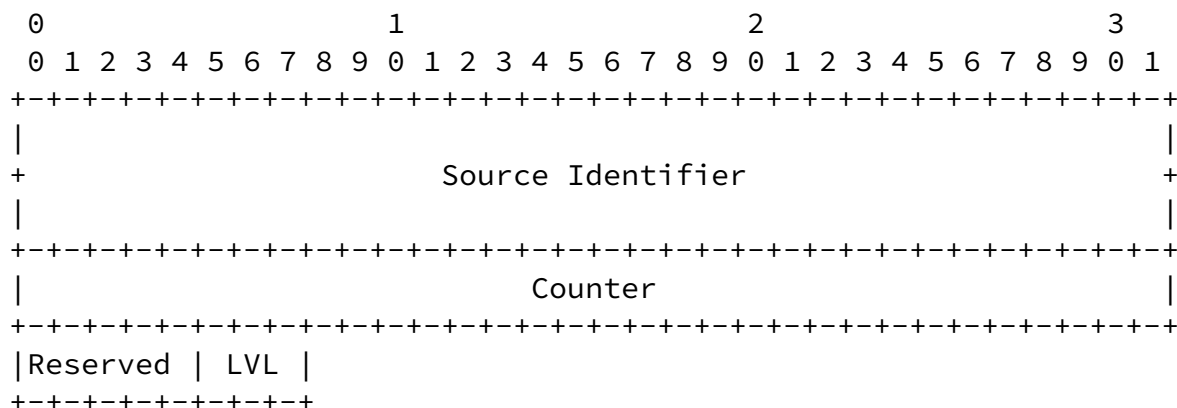


Figure 24: CCM* Nonce

Source Identifier: 8 bytes. Source Identifier is set to the logical identifier of the originator of the protected packet.

Counter: 4 bytes. Counter is set to the (uncompressed) value of the corresponding field in the Security option of the RPL control message.

Security Level (LVL): 3 bits. Security Level is set to the value of the corresponding field in the Security option of the RPL control message.

Unassigned bits of the nonce are reserved. They MUST be set to zero when constructing the nonce.

All fields of the nonce shall be represented in most-significant-octet and most-significant-bit first order.

[9.5.3.2](#). Signatures

If the Key Identification Mode (KIM) mode indicates the use of signatures (a value of 3), then a node appends a signature to the data payload of the packet. The Security Level (LVL) field describes the length of this signature.

The signature scheme in RPL for Security Mode 00 is an instantiation of the ECPVS signature scheme[X9.92]. It uses as an elliptic curve the named curve K-283[X9.92]. It uses CCM* mode[CCMStar] as the encryption scheme with M=0 (as a stream-cipher). It uses the Matyas-Meyer-Oseas unkeyed hash function[AppliedCryptography]. It uses the key derivation function based on this unkeyed hash function specified in Section 5.6.3 of [X9.63-2001], and the message encoding rule of [Section 7.8](#) or ANSI X9.92 [X9.92]. PadLen is a non-negative integer set to M-OctCurve, where OctCurve is the byte-length of the curve in question (with K-283, one has OctCurve=36).

Let 'a' be a concatenation of a six-byte representation of Counter and the message header. The packet payload is a concatenation of packet data 'c' and the signature 's'. This signature scheme is invoked with visible and recoverable message parts a and c, whereas the signature verification is invoked with as received visible and

message representative a, c, and with signature s.

[9.6.](#) Coverage of Integrity and Confidentiality

For a RPL ICMPv6 message, the entire packet is within the scope of RPL security. The message authentication code is calculated over the entire IPv6 packet. This calculation is done before any compression that lower layers may apply. The IPv6 and ICMPv6 headers are never encrypted. The body of the RPL ICMPv6 message MAY be encrypted, starting from the first byte after the security section and continuing to the end of the packet.

[10.](#) Packet Forwarding and Loop Avoidance/Detection

[10.1.](#) Suggestions for Packet Forwarding

When forwarding a packet to a destination, precedence is given to selection of a next-hop successor as follows:

1. This specification only covers how a successor is selected from the DODAG version that matches the RPLInstanceID marked in the IPv6 header of the packet being forwarded. Routing outside the instance can be done as long as additional rules are put in place such as strict ordering of instances and routing protocols to protect against loops.
2. If a local administrative preference favors a route that has been learned from a different routing protocol than RPL, then use that successor.
3. If the packet header specifies a source route, then use that

route [[I-D.hui-6man-rpl-routing-header](#)]. If the node fails to forward the packet with that specified source route, then that packet SHOULD be dropped. The node MAY log an error. The node MAY send an ICMPv6 Error in Source Routing Header message to the source of the packet [Section 18.6](#).

4. If there is an entry in the routing table matching the destination that has been learned from a multicast destination advertisement (e.g. the destination is a one-hop neighbor), then use that successor.
5. If there is an entry in the routing table matching the destination that has been learned from a unicast destination advertisement (e.g. the destination is located down the sub-DODAG), then use that successor. If there are DAO Path Control bits associated with multiple successors, then consult the Path Control bits to order the successors by preference when choosing.
6. If there is a DODAG version offering a route to a prefix matching the destination, then select one of those DODAG parents as a successor according to the OF and routing metrics.
7. Any other as-yet-unattempted DODAG parent may be chosen for the next attempt to forward a unicast packet when no better match exists.
8. Finally the packet is dropped. ICMP Destination Unreachable may be invoked (an inconsistency is detected).

TTL must be decremented when forwarding.

Note that the chosen successor MUST NOT be the neighbor that was the predecessor of the packet (split horizon), except in the case where it is intended for the packet to change from an up to a down flow, such as switching from DIO routes to DAO routes as the destination is neared.

[10.2](#). Loop Avoidance and Detection

RPL loop avoidance mechanisms are kept simple and designed to minimize churn and states. Loops may form for a number of reasons,

e.g. control packet loss. RPL includes a reactive loop detection technique that protects from meltdown and triggers repair of broken paths.

RPL loop detection uses information that is placed into the packet. A future version of this specification will detail how this information is carried with the packet (e.g. a hop-by-hop option ([[I-D.hui-6man-rpl-option](#)]) or summarized somehow into the flow label). For the purpose of RPL operations, the information carried with a packet is constructed follows:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|R|F|0|0|0|0|0| RPLInstanceID |           SenderRank           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

RPL Packet Information

Down '0' bit: 1-bit flag indicating whether the packet is expected to progress up or down. A router sets the '0' bit when the packet is expect to progress down (using DAO routes), and resets it when forwarding towards the root of the DODAG version. A host or RPL leaf node MUST set the bit to 0.

Rank-Error 'R' bit: 1-bit flag indicating whether a rank error was detected. A rank error is detected when there is a mismatch in the relative ranks and the direction as indicated in the '0' bit. A host or RPL leaf node MUST set the bit to 0.

Forwarding-Error 'F' bit: 1-bit flag indicating that this node can not forward the packet further towards the destination. The 'F' bit might be set by a child node that does not have a route to destination for a packet with the down '0' bit set. A host

or RPL leaf node MUST set the bit to 0.

RPLInstanceID: 8-bit field indicating the DODAG instance along which the packet is sent.

SenderRank: 16-bit field set to zero by the source and to DAGRank(rank) by a router that forwards inside the RPL network.

[10.2.1.](#) Source Node Operation

If the source is aware of the RPLInstanceID that is preferred for the packet, then it MUST set the RPLInstanceID field associated with the packet accordingly, otherwise it MUST set it to the RPL_DEFAULT_INSTANCE.

[10.2.2.](#) Router Operation

[10.2.2.1.](#) Instance Forwarding

Instance IDs are used to avoid loops between DODAGs from different origins. DODAGs that constructed for antagonistic constraints might contain paths that, if mixed together, would yield loops. Those loops are avoided by forwarding a packet along the DODAG that is associated to a given instance.

The RPLInstanceID is associated by the source with the packet. This RPLInstanceID MUST match the RPL Instance onto which the packet is placed by any node, be it a host or router. For traffic originating outside of the RPL domain there may be a mapping occurring at the gateway into the RPL domain, possibly based on an encoding within the flow label. This aspect of RPL operation is to be clarified in a future version of this specification.

The source of the packet might be aware of the RPL network, of the constraints imposed on OFs, and of associated Instance IDs. In that case, the source of the packet MAY tag the flow label with the RPLInstanceID, in which case it is used in that form within the RPL network.

A router that injects a data packet into the RPL network MUST tag the packet by inserting a RPL Hop-by-hop option as specified in [\[I-D.hui-6man-rpl-option\]](#). If the RPLInstanceID is not present in flow label of the data packet, the ingress router that injects the packet into the RPL network MUST add a RPLInstanceID field to the RPL Hop-by-hop option.

A router that forwards a packet to outside the RPL network MUST remove the RPL Hop-by-hop option.

When a router receives a packet that specifies a given RPLInstanceID and the node can forward the packet along the DODAG associated to that instance, then the router MUST do so and leave the RPLInstanceID value unchanged.

If any node can not forward a packet along the DODAG associated to the RPLInstanceID, then the node SHOULD discard the packet and send an ICMP error message.

10.2.2.2. DAG Inconsistency Loop Detection

The DODAG is inconsistent if the direction of a packet does not match the rank relationship. A receiver detects an inconsistency if it receives a packet with either:

- the '0' bit set (to down) from a node of a higher rank.

- the '0' bit reset (for up) from a node of a lesser rank.

When the DODAG root increments the DODAGVersionNumber a temporary rank discontinuity may form between the next version and the prior version, in particular if nodes are adjusting their rank in the next version and deferring their migration into the next version. A router that is still a member of the prior version may choose to forward a packet to a (future) parent that is in the next version. In some cases this could cause the parent to detect an inconsistency because the rank-ordering in the prior version is not necessarily the same as in the next version and the packet may be judged to not be making forward progress. If the sending router is aware that the chosen successor has already joined the next version, then the sending router MUST update the SenderRank to INFINITE_RANK as it forwards the packets across the discontinuity into the next DODAG version in order to avoid a false detection of rank inconsistency.

One inconsistency along the path is not considered as a critical error and the packet may continue. But a second detection along the path of a same packet should not occur and the packet is dropped.

This process is controlled by the Rank-Error bit associated with the packet. When an inconsistency is detected on a packet, if the Rank-Error bit was not set then the Rank-Error bit is set. If it was set the packet is discarded and the trickle timer is reset.

10.2.2.3. DAO Inconsistency Loop Detection and Recovery

A DAO inconsistency happens when router that has an down DAO route via a child that is a remnant from an obsolete state that is not matched in the child. With DAO inconsistency loop recovery, a packet

can be used to recursively explore and cleanup the obsolete DAO states along a sub-DODAG.

In a general manner, a packet that goes down should never go up again. If DAO inconsistency loop recovery is applied, then the router SHOULD send the packet back to the parent that passed it with the Forwarding-Error 'F' bit set and the 'O' bit left untouched. Otherwise the router MUST silently discard the packet.

[10.2.2.4.](#) Forward Path Recovery

Upon receiving a packet with a Forwarding-Error bit set, the node MUST remove the routing states that caused forwarding to that neighbor, clear the Forwarding-Error bit and attempt to send the packet again. The packet may be sent to an alternate neighbor. If that alternate neighbor still has an inconsistent DAO state via this node, the process will recurse, this node will set the Forwarding-Error 'F' bit and the routing state in the alternate neighbor will be cleaned up as well.

[11](#). Multicast Operation

This section describes further the multicast routing operations over an IPv6 RPL network, and specifically how unicast DAOs can be used to relay group registrations up. Wherever the following text mentions Multicast Listener Discovery (MLD), one can read MLDv1 ([\[RFC2710\]](#)) or MLDv2 ([\[RFC3810\]](#)).

Nodes that support the RPL storing mode of operation SHOULD also support multicast DAO operations as described below. Nodes that only support the non-storing mode of operation are not expected to support this section.

The multicast operation is controlled by the MOP field in the DIO.

If the MOP field requires multicast support, then a node that joins the RPL network as a router must operate as described in this section for multicast signaling and forwarding within the RPL network. A node that does not support the multicast operation required by the MOP field can only join as a leaf.

If the MOP field does not require multicast support, then multicast is handled by some other way that is out of scope for this specification. (Examples may include as a series of unicast copies or limited-scope flooding)

As is traditional, a listener uses a protocol such as MLD with a router to register to a multicast group.

Along the path between the router and the DODAG root, MLD requests are mapped and transported as DAO messages within the RPL protocol; each hop coalesces the multiple requests for a same group as a single DAO message to the parent(s), in a fashion similar to proxy IGMP, but recursively between child router and parent up to the root.

A router might select to pass a listener registration DAO message to its preferred parent only, in which case multicast packets coming back might be lost for all of its sub-DODAG if the transmission fails over that link. Alternatively the router might select to copy additional parents as it would do for DAO messages advertising unicast destinations, in which case there might be duplicates that the router will need to prune.

As a result, multicast routing states are installed in each router on the way from the listeners to the root, enabling the root to copy a multicast packet to all its children routers that had issued a DAO message including a DAO for that multicast group, as well as all the attached nodes that registered over MLD.

For unicast traffic, it is expected that the grounded root of an DODAG terminates RPL and MAY redistribute the RPL routes over the external infrastructure using whatever routing protocol is used in the other routing domain. For multicast traffic, the root MAY proxy MLD for all the nodes attached to the RPL domain (this would be needed if the multicast source is located in the external infrastructure). For such a source, the packet will be replicated as it flows down the DODAG based on the multicast routing table entries installed from the DAO message.

For a source inside the DODAG, the packet is passed to the preferred parents, and if that fails then to the alternates in the DODAG. The packet is also copied to all the registered children, except for the one that passed the packet. Finally, if there is a listener in the external infrastructure then the DODAG root has to further propagate the packet into the external infrastructure.

As a result, the DODAG Root acts as an automatic proxy Rendezvous Point for the RPL network, and as source towards the Internet for all multicast flows started in the RPL LLN. So regardless of whether the root is actually attached to the Internet, and regardless of whether the DODAG is grounded or floating, the root can serve inner multicast streams at all times.

12. Maintenance of Routing Adjacency

The selection of successors, along the default paths up along the DODAG, or along the paths learned from destination advertisements down along the DODAG, leads to the formation of routing adjacencies that require maintenance.

In IGPs such as OSPF [[RFC4915](#)] or IS-IS [[RFC5120](#)], the maintenance of a routing adjacency involves the use of Keepalive mechanisms (Hellos) or other protocols such as BFD ([[RFC5880](#)]) and MANET Neighborhood Discovery Protocol (NHDP [[I-D.ietf-manet-nhdp](#)]). Unfortunately, such an approach is not desirable in constrained environments such as LLN and would lead to excessive control traffic in light of the data traffic with a negative impact on both link loads and nodes resources. Overhead to maintain the routing adjacency should be minimized. Furthermore, it is not always possible to rely on the link or transport layer to provide information of the associated link state. The network layer needs to fall back on its own mechanism.

Thus RPL makes use of a different approach consisting of probing the neighbor using a Neighbor Solicitation message (see [[RFC4861](#)]). The

reception of a Neighbor Advertisement (NA) message with the "Solicited Flag" set is used to verify the validity of the routing adjacency. Such mechanism MAY be used prior to sending a data packet. This allows for detecting whether or not the routing adjacency is still valid, and should it not be the case, select another feasible successor to forward the packet.

[13.](#) Guidelines for Objective Functions

An Objective Function (OF) allows for the selection of a DODAG to join, and a number of peers in that DODAG as parents. The OF is used to compute an ordered list of parents. The OF is also responsible to compute the rank of the device within the DODAG version.

The Objective Function is indicated in the DIO message using an Objective Code Point (OCP), and indicates the method that must be used to construct the DODAG. The Objective Code Points are specified in [[I-D.ietf-roll-of0](#)], and related companion specifications.

[13.1.](#) Objective Function Behavior

Most Objective Functions are expected to follow the same abstract behavior:

- o The parent selection is triggered each time an event indicates that a potential next hop information is updated. This might happen upon the reception of a DIO message, a timer elapse, all DODAG parents are unavailable, or a trigger indicating that the state of a candidate neighbor has changed.
 - o An OF scans all the interfaces on the device. Although there may typically be only one interface in most application scenarios, there might be multiple of them and an interface might be configured to be usable or not for RPL operation. An interface can also be configured with a preference or dynamically learned to be better than another by some heuristics that might be link-layer dependent and are out of scope. Finally an interface might or not match a required criterion for an Objective Function, for instance a degree of security. As a result some interfaces might be completely excluded from the computation, while others might be more or less preferred.
 - o An OF scans all the candidate neighbors on the possible interfaces to check whether they can act as a router for a DODAG. There might be multiple of them and a candidate neighbor might need to pass some validation tests before it can be used. In particular, some link layers require experience on the activity with a router to enable the router as a next hop.
 - o An OF computes self's rank by adding to the rank of the candidate a value representing the relative locations of self and the candidate in the DODAG version.
- * The increase in rank must be at least MinHopRankIncrease.

- * To keep loop avoidance and metric optimization in alignment, the increase in rank should reflect any increase in the metric value. For example, with a purely additive metric such as ETX, the increase in rank can be made proportional to the increase in the metric.
- * Candidate neighbors that would cause self's rank to increase

are not considered for parent selection

- o Candidate neighbors that advertise an OF incompatible with the set of OF specified by the policy functions are ignored.
- o As it scans all the candidate neighbors, the OF keeps the current best parent and compares its capabilities with the current candidate neighbor. The OF defines a number of tests that are critical to reach the objective. A test between the routers determines an order relation.
 - * If the routers are equal for that relation then the next test is attempted between the routers,
 - * Else the best of the two routers becomes the current best parent and the scan continues with the next candidate neighbor
 - * Some OFs may include a test to compare the ranks that would result if the node joined either router
- o When the scan is complete, the preferred parent is elected and self's rank is computed as the preferred parent rank plus the step in rank with that parent.
- o Other rounds of scans might be necessary to elect alternate parents. In the next rounds:
 - * Candidate neighbors that are not in the same DODAG are ignored
 - * Candidate neighbors that are of greater rank than self are ignored
 - * Candidate neighbors of an equal rank to self are ignored for parent selection
 - * Candidate neighbors of a lesser rank than self are preferred

14. Suggestions for Interoperation with Neighbor Discovery

This specification directly borrows the Prefix Information Option (PIO) and the Routing Information Option (RIO) from IPv6 ND. It is envisioned that as future specifications build on this base that there may be additional cause to leverage parts of IPv6 ND. This section provides some suggestions for future specifications.

First and foremost RPL is a routing protocol. One should take great care to preserve architecture when mapping functionalities between RPL and ND. RPL is for routing only. That said, there may be persuading technical reasons to allow for sharing options between RPL and IPv6 ND in a particular implementation/deployment.

In general the following guidelines apply:

- o RPL Type codes must be allocated from the RPL Control Message Options registry.
- o RPL Length fields must be expressed in units of single octets, as opposed to ND Length fields which are expressed in units of 8 octets.
- o RPL Options are generally not required to be aligned to 8 octet boundaries.
- o When mapping/transposing an IPv6 ND option for redistribution as a RPL option, any padding octets should be removed when possible. For example, the Prefix Length field in the PIO is sufficient to describe the length of the Prefix field. When mapping/transposing a RPL option for redistribution as an IPv6 ND option, any such padding octets should be restored. This procedure must be unambiguous.

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

15. RPL Constants and Variables

Following is a summary of RPL constants and variables:

BASE_RANK This is the rank for a virtual root that might be used to coordinate multiple roots. **BASE_RANK** has a value of 0.

ROOT_RANK This is the rank for a DODAG root. **ROOT_RANK** has a value of `MinHopRankIncrease` (as advertised by the DODAG root), such that `DAGRank(ROOT_RANK)` is 1.

INFINITE_RANK This is the constant maximum for the rank. **INFINITE_RANK** has a value of 0xFFFF.

RPL_DEFAULT_INSTANCE This is the `RPLInstanceID` that is used by this protocol by a node without any overriding policy. **RPL_DEFAULT_INSTANCE** has a value of 0.

DEFAULT_PATH_CONTROL_SIZE This is the default value used to configure PCS in the DODAG Configuration Option, which dictates the number of significant bits in the Path Control field of the Transit Information option. **DEFAULT_PATH_CONTROL_SIZE** has a value of 0. This configures the simplest case-- limiting the fan-out to 1 and limiting a node to send a DAO message to only one parent.

DEFAULT_DIO_INTERVAL_MIN This is the default value used to configure `Imin` for the DIO trickle timer. **DEFAULT_DIO_INTERVAL_MIN** has a value of 3. This configuration results in `Imin` of 8ms.

DEFAULT_DIO_INTERVAL_DOUBLINGS This is the default value used to configure `Imax` for the DIO trickle timer. **DEFAULT_DIO_INTERVAL_DOUBLINGS** has a value of 20. This configuration results in a maximum interval of 2.3 hours.

DEFAULT_DIO_REDUNDANCY_CONSTANT This is the default value used to configure `k` for the DIO trickle timer. **DEFAULT_DIO_REDUNDANCY_CONSTANT** has a value of 10. This configuration is a conservative value for trickle suppression mechanism.

DEFAULT_MIN_HOP_RANK_INCREASE This is the default value of `MinHopRankIncrease`. **DEFAULT_MIN_HOP_RANK_INCREASE** has a value

of 256. This configuration results in an 8-bit wide integer part of Rank.

DIO Timer One instance per DODAG that a node is a member of. Expiry triggers DIO message transmission. Trickle timer with variable interval in $[0, \text{DIOIntervalMin} \cdot 2^{\text{DIOIntervalDoublings}}]$. See [Section 7.3.1](#)

DAG Version Increment Timer Up to one instance per DODAG that the node is acting as DODAG root of. May not be supported in all implementations. Expiry triggers increment of `DODAGVersionNumber`, causing a new series of updated DIO message to be sent. Interval should be chosen appropriate to propagation time of DODAG and as appropriate to application requirements (e.g. response time vs. overhead).

DelayDAO Timer Up to one instance per DAO parent (the subset of DODAG parents chosen to receive destination advertisements) per DODAG. Expiry triggers sending of DAO message to the DAO parent. See [Section 8.4](#)

RemoveTimer Up to one instance per DAO entry per neighbor (i.e. those neighbors that have given DAO messages to this node as a DODAG parent) Expiry triggers a change in state for the DAO entry, setting up to do unreachable (No-Path) advertisements or immediately deallocating the DAO entry if there are no DAO parents.

[16.](#) Manageability Considerations

The aim of this section is to give consideration to the manageability of RPL, and how RPL will be operated in a LLN. The scope of this section is to consider the following aspects of manageability: configuration, monitoring, fault management, accounting, and performance of the protocol in light of the recommendations set forth in [[RFC5706](#)].

[16.1.](#) Introduction

Most of the existing IETF management standards are Structure of Management Information (SMI) based data models (MIB modules) to monitor and manage networking devices.

For a number of protocols, the IETF community has used the IETF Standard Management Framework, including the Simple Network Management Protocol [[RFC3410](#)], the Structure of Management Information [[RFC2578](#)], and MIB data models for managing new protocols.

As pointed out in [[RFC5706](#)], the common policy in terms of operation and management has been expanded to a policy that is more open to a set of tools and management protocols rather than strictly relying on a single protocol such as SNMP.

In 2003, the Internet Architecture Board (IAB) held a workshop on Network Management [[RFC3535](#)] that discussed the strengths and weaknesses of some IETF network management protocols and compared

them to operational needs, especially configuration.

One issue discussed was the user-unfriendliness of the binary format of SNMP [[RFC3410](#)]. In the case of LLNs, it must be noted that at the time of writing, the CoRE Working Group is actively working on resource management of devices in LLNs. Still, it is felt that this section provides important guidance on how RPL should be deployed, operated, and managed.

As stated in [[RFC5706](#)], "A management information model should include a discussion of what is manageable, which aspects of the protocol need to be configured, what types of operations are allowed, what protocol-specific events might occur, which events can be counted, and for which events an operator should be notified". These aspects are discussed in detail in the following sections.

RPL will be used on a variety of devices that may have resources such as memory varying from a very few Kbytes to several hundreds of Kbytes and even Mbytes. When memory is highly constrained, it may

not be possible to satisfy all the requirements listed in this section. Still it is worth listing all of these in an exhaustive fashion, and implementers will then determine which of these requirements could be satisfied according to the available resources on the device.

[16.2.](#) Configuration Management

[16.2.1.](#) Initialization Mode

"Architectural Principles of the Internet" [[RFC1958](#)], [Section 3.8](#), states: "Avoid options and parameters whenever possible. Any options and parameters should be configured or negotiated dynamically rather than manually. This especially true in LLNs where the number of devices may be large and manual configuration is infeasible. This has been taken into account in the design of RPL whereby the DODAG root provides a number of parameters to the devices joining the DODAG, thus avoiding cumbersome configuration on the routers and potential sources of misconfiguration (e.g. values of trickle timers, ...). Still there are additional RPL parameters that a RPL implementation should allow to be configured, which are discussed in this section."

[16.2.1.1.](#) DIS mode of operation upon boot-up

When a node is first powered up:

1. The node may decide to stay silent, waiting to receive DIO messages from DODAG of interest (advertising a supported OF and metrics/constraints) and not send any multicast DIO messages until it has joined a DODAG.
2. The node may decide to send one or more DIS messages (optionally requesting DIO for a specific DODAG) message as an initial probe for nearby DODAGs, and in the absence of DIO messages in reply after some configurable period of time, the node may decide to root a floating DODAG and start sending multicast DIO messages.

A RPL implementation SHOULD allow configuring the preferred mode of operation listed above along with the required parameters (in the second mode: the number of DIS messages and related timer).

[16.2.2.](#) DIO and DAO Base Message and Options Configuration

RPL specifies a number of protocol parameters considering the large spectrum of applications where it will be used. That said, particular attention has been given to limiting the number of these parameters that must be configured on each RPL router. Instead, a

number of the default values can be used, and when required these parameters can be provided by the DODAG root thus allowing for dynamic parameter setting.

A RPL implementation SHOULD allow configuring the following routing protocol parameters. As pointed out above, note that a large set of parameters is configured on the DODAG root.

[16.2.3.](#) Protocol Parameters to be configured on every router in the LLN

- o RPLInstanceID [DIO message, in DIO base message]. Although the RPLInstanceID must be configured on the DODAG root, it must also be configured as a policy on every node in order to determine whether or not the node should join a particular DODAG. Note that a second RPLInstance can be configured on the node, should it

become root of a floating DODAG.

- o Objective Code Point (OCP)
- o List of supported metrics: [[I-D.ietf-roll-routing-metrics](#)] specifies a number of metrics and constraints used for the DODAG formation. Thus a RPL implementation should allow configuring the list of metrics that a node can accept and understand. If a DIO is received with a metric and/or constraint that is not understood or supported, as specified in [Section 7.5](#), the node would join as a leaf node.
- o DODAGID [DIO, DIO base option] and [DAO message when the D flag of the DAO message is set).
- o Route Information (and preference) [DIO message, in Route Information option]
- o Solicited Information [DIS message, in Solicited Information option]. Note that an RPL implementation SHOULD allow configuring when such messages should be sent and under which circumstances, along with the value of the RPLInstance ID, V/I/D flags.
- o K flag [DAO message, in DAO base message].
- o MOP (Mode of Operation) [DIO message, in DIO base message]

[16.2.4](#). Protocol Parameters to be configured on every non-root router in the LLN

- o Target prefix [DAO, in RPL Target option and DIO messages]

- o Transit information [DAO, Transit information option]: A RPL implementation SHOULD allow configuring whether a non-storing node provides the transit information in DAO messages.

A node whose DODAG parent set is empty may become the DODAG root of a floating DODAG. It may also set its DAGPreference such that it is less preferred. Thus a RPL implementation MUST allow configuring the set of actions that the node should initiate in this case:

- o Start its own (floating) DODAG: the new DODAGID must be configured in addition to its DAGPreference
- o Poison the broken path (see procedure in [Section 7.2.2.5](#))
- o Trigger a local repair

[16.2.5](#). Parameters to be configured on the DODAG root

In addition, several other parameters are configured only on the DODAG root and advertised in options carried in DIO messages.

As specified in [Section 7.3](#), a RPL implementation makes use of trickle timers to govern the sending of DIO messages. The operation of the trickle algorithm is determined by a set of configurable parameters, which MUST be configurable and that are then advertised by the DODAG root along the DODAG in DIO messages.

- o DIOIntervalDoublings [DIO, in DODAG configuration option]
- o DIOIntervalMin [DIO, in DODAG configuration option]
- o DIORedundancyConstant [DIO, in DODAG configuration option]

In addition, a RPL implementation SHOULD allow for configuring the following set of RPL parameters:

- o Path Control Size [DIO, in DODAG configuration option]
- o MinHopRankIncrease [DIO, in DODAG configuration option]
- o The following fields: MOP (Mode of Operation), DODAGPreference field [DIO message, DIO Base object]
- o Route information (list of prefixes with preference) [DIO message, in Route Information option]
- o The T flag allows for triggering a refresh of the downward routes. A RPL implementation SHOULD support manual setting of the T flag

or upon the occurrence of a set of event such as the expiration of

a configurable periodic timer.

- o List of metrics and constraints used for the DODAG.
- o Prefix information along with valid and preferred lifetime and the L and A flags. [DIO message, Prefix Information option]. A RPL implementation SHOULD allow configuring if the Prefix Information Option must be carried with the DIO message to distribute the prefix information for auto-configuration. In that case, the RPL implementation MUST allow the list of prefixes to be advertised in the Prefix Information Option along with the corresponding flags.

DAG Root behavior: in some cases, a node may not want to permanently act as a floating DODAG root if it cannot join a grounded DODAG. For example a battery-operated node may not want to act as a floating DODAG root for a long period of time. Thus a RPL implementation MAY support the ability to configure whether or not a node could act as a floating DODAG root for a configured period of time.

DAG Version Number Increment: a RPL implementation may allow by configuration at the DODAG root to refresh the DODAG states by updating the DODAGVersionNumber. A RPL implementation SHOULD allow configuring whether or not periodic or event triggered mechanisms are used by the DODAG root to control DODAGVersionNumber change (which triggers a global repair as specified in [Section 3.3.2](#)).

[16.2.6](#). Configuration of RPL Parameters related to DAO-based mechanisms

DAO messages are optional and used in DODAGs that require downward routing operation. This section deals with the set of parameters related to DAO message and provides recommendations on their configuration.

An implementation SHOULD bound the time that the entry is allocated in the UNREACHABLE state. Upon the equivalent expiry of the related timer (RemoveTimer), the entry SHOULD be suppressed. Thus a RPL implementation MAY allow for the configuration of the RemoveTimer.

While the entry is in the UNREACHABLE state a node SHOULD make a reasonable attempt to report a No-Path to each of the DAO parents. That number of attempts MAY be configurable.

When the associated Retry Counter for a REACHABLE(Pending) entry reaches a maximum threshold, the entry is placed into the UNREACHABLE state and No-Path should be scheduled to send to the node's DAO Parents. The maximum threshold MAY be configurable.

An implementation should support rate-limiting the sending of DAO messages. The related parameters MAY be configurable.

When scheduling to send a DAO, an implementation should equivalently start a timer (DelayDAO) to delay sending the DAO, thus helping to potentially aggregate DAOs. The DelayDAO timer MAY be configurable.

[16.2.7.](#) Default Values

This document specifies default values for the following set of RPL variables:

- DEFAULT_PATH_CONTROL_SIZE
- DEFAULT_DIO_INTERVAL_MIN
- DEFAULT_DIO_INTERVAL_DOUBLINGS
- DEFAULT_DIO_REDUNDANCY_CONSTANT
- DEFAULT_MIN_HOP_RANK_INCREASE

It is recommended to specify default values in protocols; that being said, as discussed in [[RFC5706](#)], default values may make less and less sense. RPL is a routing protocol that is expected to be used in a number of contexts where network characteristics such as the number of nodes, link and nodes types are expected to vary significantly. Thus, these default values are likely to change with the context and as the technology will evolve. Indeed, LLNs' related technology (e.g. hardware, link layers) have been evolving dramatically over the past few years and such technologies are expected to change and evolve considerably in the coming years.

The proposed values are not based on extensive best current practices and are considered to be conservative.

[16.3.](#) Monitoring of RPL Operation

Several RPL parameters should be monitored to verify the correct operation of the routing protocol and the network itself. This section lists the set of monitoring parameters of interest.

[16.3.1.](#) Monitoring a DODAG parameters

A RPL implementation SHOULD provide information about the following parameters:

- o DODAG Version number [DIO message, in DIO base message]
- o Status of the G flag [DIO message, in DIO base message]

- o Status of the MOP field [DIO message, in DIO base message]

- o Value of the DTSN [DIO message, in DIO base message]
- o Value of the rank [DIO message, in DIO base message]
- o DAOSequence: Incremented at each unique DAO message, echoed in the DAO-ACK message [DAO and DAO-ACK messages]
- o Route Information [DIO message, Route Information option] (list of IPv6 prefixes per parent along with lifetime and preference)
- o Trickle parameters:
 - * DIOIntervalDoublings [DIO, in DODAG configuration option]
 - * DIOIntervalMin [DIO, in DODAG configuration option]
 - * DIORedundancyConstant [DIO, in DODAG configuration option]
- o Path Control Size [DIO, in DODAG configuration option]
- o MinHopRankIncrease [DIO, in DODAG configuration option]

Values that may be monitored only on the DODAG root

- o Transit Information [DAO, Transit Information option]: A RPL implementation SHOULD allow configuring whether the set of received Transit Information options should be displayed on the DODAG root. In this case, the RPL database of received Transit Information should also contain: the path-sequence, path control, path lifetime and parent address.

[16.3.2.](#) Monitoring a DODAG inconsistencies and loop detection

Detection of DODAG inconsistencies is particularly critical in RPL networks. Thus it is recommended for a RPL implementation to provide appropriate monitoring tools. A RPL implementation SHOULD provide a counter reporting the number of a times the node has detected an inconsistency with respect to a DODAG parent, e.g. if the DODAGID has changed.

When possible more granular information about inconsistency detection should be provided. A RPL implementation MAY provide counters reporting the number of following inconsistencies:

- o Packets received with 0 bit set (to down) from a node with a higher rank

Winter, et al.

Expires December 30, 2010

[Page 97]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

- o Packets received with 0 bit reset (to up) from a node with a lower rank
- o Number of packets with the F bit set
- o Number of packets with the R bit set

[16.4.](#) Monitoring of the RPL data structures

[16.4.1.](#) Candidate Neighbor Data Structure

A node in the candidate neighbor list is a node discovered by the some means and qualified to potentially become a parent (with high enough local confidence). A RPL implementation SHOULD provide a way to monitor the candidate neighbor list with some metric reflecting local confidence (the degree of stability of the neighbors) as measured by some metrics.

A RPL implementation MAY provide a counter reporting the number of times a candidate neighbor has been ignored, should the number of candidate neighbors exceeds the maximum authorized value.

[16.4.2.](#) Destination Oriented Directed Acyclic Graph (DAG) Table

For each DODAG, a RPL implementation is expected to keep track of the following DODAG table values:

- o RPLInstanceID
- o DODAGID
- o DODAGVersionNumber

- o Rank
- o Objective Code Point
- o A set of DODAG Parents
- o A set of prefixes offered upwards along the DODAG
- o Trickle timers used to govern the sending of DIO messages for the DODAG
- o List of DAO parents
- o DTSN

- o Node status (router versus leaf)

A RPL implementation SHOULD allow for monitoring the set of parameters listed above.

[16.4.3.](#) Routing Table and DAO Routing Entries

A RPL implementation maintains several information elements related to the DODAG and the DAO entries (for storing nodes). In the case of a non storing node, a limited amount of information is maintained (the routing table is mostly reduced to a set of DODAG parents along with characteristics of the DODAG as mentioned above) whereas in the case of storing nodes, this information is augmented with routing entries.

A RPL implementation SHOULD provide the ability to monitor the following parameters:

- o Next Hop (DODAG parent)
- o Next Hop Interface
- o Path metrics value for each DODAG parent

A DAO Routing Table Entry conceptually contains the following

elements (for storing nodes only):

- o Advertising Neighbor Information
- o IPv6 Address
- o Interface ID to which DAO Parents has this entry been reported
- o Retry Counter
- o Logical equivalent of DAO Content:
 - * DAO Sequence
 - * DAO Lifetime
 - * DAO Path Control
- o Destination Prefix (or Address or Mcast Group)

A RPL implementation SHOULD provide information about the state of each DAO Routing Table entry states.

[16.5.](#) Fault Management

Fault management is a critical component used for troubleshooting, verification of the correct mode of operation of the protocol, network design, and is also a key component of network performance monitoring. A RPL implementation SHOULD allow providing the following information related to fault managements:

- o Memory overflow along with the cause (e.g. routing tables overflow, ...)
- o Number of times a packet could not be sent to a DODAG parent flagged as valid
- o Number of times a packet has been received for which the router did not have a corresponding RPLInstanceID
- o Number of times a local repair procedure was triggered

- o Number of times a global repair was triggered by the DODAG root
- o Number of received malformed messages
- o Number of seconds with packets to forward and no next hop (DODAG parent)
- o Number of seconds without next hop (DODAG parent)
- o Number of times a node has joined a DODAG as a leaf because it received a DIO with metric/constraint not understood and it was configured to join as a leaf node in this case (see [Section 16.6](#)).

It is RECOMMENDED to report faults via at least error log messages. Other protocols may be used to report such faults.

[16.6](#). Policy

Policy rules can be used by a RPL implementation to determine whether or not the node is allowed to join a particular DODAG advertised by a neighbor by means of DIO messages.

This document specifies operation within a single DODAG. A DODAG is characterized by the following tuple (RPLInstanceID, DODAGID). Furthermore, as pointed out above, DIO messages are used to advertise other DODAG characteristics such as the routing metrics and constraints used to build to the DODAG and the Objective Function in use (specified by OCP).

The first policy rules consists of specifying the following conditions that a RPL node must satisfy to join a DODAG:

- o RPLInstanceID
- o DODAGID
- o List of supported routing metrics and constraints
- o Objective Function (OCP values)

A RPL implementation MUST allow configuring these parameters and SHOULD specify whether the node must simply ignore the DIO if the advertised DODAG is not compliant with the local policy or whether the node should join as the leaf node if only the list of supported routing metrics and constraints, and the OF is not supported.

A RPL implementation SHOULD allow configuring the set of acceptable or preferred Objective Functions (OF) referenced by their Objective Codepoints (OCPs) for a node to join a DODAG, and what action should be taken if none of a node's candidate neighbors advertise one of the configured allowable Objective Functions, or if the advertised metrics/constraint is not understood/supported. Two actions can be taken in this case:

- o The node joins the DODAG as a leaf node as specified in [Section 7.5](#)
- o The node does not join the DODAG

A node in an LLN may learn routing information from different routing protocols including RPL. It is in this case desirable to control via administrative preference which route should be favored. An implementation SHOULD allow for specifying an administrative preference for the routing protocol from which the route was learned.

Internal Data Structures: some RPL implementations may limit the size of the candidate neighbor list in order to bound the memory usage, in which case some otherwise viable candidate neighbors may not be considered and simply dropped from the candidate neighbor list.

A RPL implementation MAY provide an indicator on the size of the candidate neighbor list.

[16.7.](#) Liveness Detection and Monitoring

By contrast with several other routing protocols, RPL does not define any 'keep-alive' mechanisms to detect routing adjacency failure: this

is in most cases, because such a mechanism may be too expensive in terms of bandwidth and even more importantly energy (a battery operated device could not afford to send periodic Keep alive). Still RPL requires mechanisms to detect that a neighbor is no longer

reachable: this can be performed by using mechanisms such as NUD (Neighbor Unreachability Detection) or even some form of Keep-alive that are outside of this document.

[16.8.](#) Fault Isolation

It is RECOMMENDED to quarantine neighbors that start emitting malformed messages at unacceptable rates.

[16.9.](#) Impact on Other Protocols

RPL has very limited impact on other protocols. Where more than one routing protocol is required on a router such as a LBR, it is expected for the device to support routing redistribution functions between the routing protocols to allow for reachability between the two routing domains. Such redistribution SHOULD be governed by the use of user configurable policy.

With regards to the impact in terms of traffic on the network, RPL has been designed to limit the control traffic thanks to mechanisms such as Trickle timers ([Section 7.3](#)). Thus the impact of RPL on other protocols should be extremely limited.

[16.10.](#) Performance Management

Performance management is always an important aspect of a protocol and RPL is not an exception. Several metrics of interest have been specified by the IP Performance Monitoring (IPPM) Working Group: that being said, they will be hardly applicable to LLN considering the cost of monitoring these metrics in terms of resources on the devices and required bandwidth. Still, RPL implementation MAY support some of these, and other parameters of interest are listed below:

- o Number of repairs and time to repair in seconds (average, variance)
- o Number of times and duration during which a devices could not forward a packet because of a lack of reachable neighbor in its routing table
- o Monitoring of resources consumption by RPL itself in terms of bandwidth and required memory

- o Number of RPL control messages sent and received

[17.](#) Security Considerations

[17.1.](#) Overview

From a security perspective, RPL networks are no different from any other network. They are vulnerable to passive eavesdropping attacks and potentially even active tampering when physical access to a wire is not required to participate in communications. The very nature of ad hoc networks and their cost objectives impose additional security constraints, which perhaps make these networks the most difficult environments to secure. Devices are low-cost and have limited capabilities in terms of computing power, available storage, and power drain; and it cannot always be assumed they have neither a trusted computing base nor a high-quality random number generator aboard. Communications cannot rely on the online availability of a fixed infrastructure and might involve short-term relationships between devices that may never have communicated before. These constraints might severely limit the choice of cryptographic algorithms and protocols and influence the design of the security architecture because the establishment and maintenance of trust relationships between devices need to be addressed with care. In addition, battery lifetime and cost constraints put severe limits on the security overhead these networks can tolerate, something that is of far less concern with higher bandwidth networks. Most of these security architectural elements can be implemented at higher layers and may, therefore, be considered to be outside the scope of this standard. Special care, however, needs to be exercised with respect to interfaces to these higher layers.

The security mechanisms in this standard are based on symmetric-key and public-key cryptography and use keys that are to be provided by higher layer processes. The establishment and maintenance of these keys are outside the scope of this standard. The mechanisms assume a secure implementation of cryptographic operations and secure and authentic storage of keying material.

The security mechanisms specified provide particular combinations of the following security services:

Data confidentiality: Assurance that transmitted information is only

disclosed to parties for which it is intended.

Data authenticity: Assurance of the source of transmitted information (and, hereby, that information was not modified in transit).

Winter, et al.

Expires December 30, 2010

[Page 104]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Replay protection: Assurance that a duplicate of transmitted information is detected.

Timeliness (delay protection): Assurance that transmitted information was received in a timely manner.

The actual protection provided can be adapted on a per-packet basis and allows for varying levels of data authenticity (to minimize security overhead in transmitted packets where required) and for optional data confidentiality. When nontrivial protection is required, replay protection is always provided.

Replay protection is provided via the use of a non-repeating value (nonce) in the packet protection process and storage of some status information for each originating device on the receiving device, which allows detection of whether this particular nonce value was used previously by the originating device. In addition, so-called delay protection is provided amongst those devices that have a loosely synchronized clock on board. The acceptable time delay can be adapted on a per-packet basis and allows for varying latencies (to facilitate longer latencies in packets transmitted over a multi-hop communication path).

Cryptographic protection may use a key shared between two peer devices (link key) or a key shared among a group of devices (group key), thus allowing some flexibility and application-specific tradeoffs between key storage and key maintenance costs versus the cryptographic protection provided. If a group key is used for peer-to-peer communication, protection is provided only against outsider devices and not against potential malicious devices in the key-sharing group.

Data authenticity may be provided using symmetric-key based or

public-key based techniques. With public-key based techniques (via signatures), one corroborates evidence as to the unique originator of transmitted information, whereas with symmetric-key based techniques data authenticity is only provided relative to devices in a key-sharing group. Thus, public-key based authentication may be useful in scenarios that require a more fine-grained authentication than can be provided with symmetric-key based authentication techniques alone, such as with group communications (broadcast, multicast), or in scenarios that require non-repudiation.

Winter, et al.

Expires December 30, 2010

[Page 105]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

[18.](#) IANA Considerations

[18.1.](#) RPL Control Message

The RPL Control Message is an ICMP information message type that is to be used carry DODAG Information Objects, DODAG Information Solicitations, and Destination Advertisement Objects in support of RPL operation.

IANA has defined an ICMPv6 Type Number Registry. The suggested type value for the RPL Control Message is 155, to be confirmed by IANA.

[18.2.](#) New Registry for RPL Control Codes

IANA is requested to create a registry, RPL Control Codes, for the Code field of the ICMPv6 RPL Control Message.

New codes may be allocated only by an IETF Consensus action. Each code should be tracked with the following qualities:

- o Code
- o Description
- o Defining RFC

Three codes are currently defined:

Code	Description	Reference
0x00	DODAG Information Solicitation	This document
0x01	DODAG Information Object	This document
0x02	Destination Advertisement Object	This document
0x03	Destination Advertisement Object Acknowledgment	This document
0x80	Secure DODAG Information Solicitation	This document
0x81	Secure DODAG Information Object	This document

0x82	Secure Destination Advertisement Object	This document
0x83	Secure Destination Advertisement Object Acknowledgment	This document

RPL Control Codes

[18.3.](#) New Registry for the Mode of Operation (MOP) DIO Control Field

IANA is requested to create a registry for the Mode of Operation (MOP) DIO Control Field, which is contained in the DIO Base.

New fields may be allocated only by an IETF Consensus action. Each field should be tracked with the following qualities:

- o Mode of Operation

- o Capability description
- o Defining RFC

Three values are currently defined:

MOP	Description	Reference
000	No downward routes maintained by RPL	This document
001	Non-Storing mode of operation	This document
010	Storing mode of operation with no multicast support	This document
011	Storing mode of operation with multicast support	This document

DIO Base Flags

[18.4.](#) RPL Control Message Option

IANA is requested to create a registry for the RPL Control Message Options

Value	Meaning	Reference
0	Pad1	This document
1	PadN	This document
2	DAG Metric Container	This Document
3	Routing Information	This Document
4	DODAG Configuration	This Document

5	RPL Target	This Document
6	Transit Information	This Document
7	Solicited Information	This Document
8	Prefix Information	This Document

RPL Control Message Options

[18.5.](#) Objective Code Point (OCP) Registry

IANA is requested to create a registry to manage the codespace of the Objective Code Point (OCP) field.

No OCP codepoints are defined in this specification.

[18.6.](#) ICMPv6: Error in Source Routing Header

In some cases RPL will return an ICMPv6 error message when a message cannot be delivered as specified by its source routing header. This ICMPv6 error message is "Error in Source Routing Header"

IANA has defined an ICMPv6 "Code" Fields Registry for ICMPv6 Message Types. ICMPv6 Message Type 1 describes "Destination Unreachable" codes. The "Error in Source Routing Header" code is suggested to be allocated from the ICMPv6 Code Fields Registry for ICMPv6 Message Type 1, with a suggested code value of 7, to be confirmed by IANA.

[18.7.](#) Link-Local Scope multicast address

The rules for assigning new IPv6 multicast addresses are defined in [\[RFC3307\]](#). This specification requires the allocation of a new permanent multicast address with a link local scope for RPL routers,

with a suggested value of FF02::1:A, to be confirmed by IANA.

[19.](#) Acknowledgements

The authors would like to acknowledge the review, feedback, and comments from Roger Alexander, Emmanuel Baccelli, Dominique Barthel, Yusuf Bashir, Yoav Ben-Yehezkel, Phoebus Chen, Mischa Dohler, Mathilde Durvy, Joakim Eriksson, Omprakash Gnawali, Manhar Goindi, Mukul Goyal, Ulrich Herberg, Anders Jagd, JeongGil (John) Ko, Quentin Lampin, Jerry Martocci, Matteo Paris, Alexandru Petrescu, Joseph Reddy, Don Sturek, Joydeep Tripathi, and Nicolas Tsiftes.

The authors would like to acknowledge the guidance and input provided by the ROLL Chairs, David Culler and JP Vasseur.

The authors would like to acknowledge prior contributions of Robert Assimiti, Mischa Dohler, Julien Abeille, Ryuji Wakikawa, Teco Boot, Patrick Wetterwald, Bryan McLaughlin, Carlos J. Bernardos, Thomas Watteyne, Zach Shelby, Caroline Bontoux, Marco Molteni, Billy Moon, Jim Bound, Yanick Pouffary, Henning Rogge and Arsalan Tavakoli, whom have provided useful design considerations to RPL.

RPL Security Design, found in [Section 9](#), [Section 17](#), and elsewhere throughout the document, is primarily the contribution of the Security Design Team: Tzeta Tsao, Roger Alexander, Dave Ward, Philip Levis, Kris Pister, and Rene Struik.

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

[20.](#) Contributors

RPL is the result of the contribution of the following members of the RPL Author Team, including the editors, and additional contributors as listed below:

JP Vasseur
Cisco Systems, Inc
11, Rue Camille Desmoulins
Issy Les Moulineaux, 92782
France

Email: jpv@cisco.com

Thomas Heide Clausen
LIX, Ecole Polytechnique, France

Phone: +33 6 6058 9349
Email: T.Clausen@computer.org
URI: <http://www.ThomasClausen.org/>

Philip Levis
Stanford University
358 Gates Hall, Stanford University
Stanford, CA 94305-9030
USA

Email: pal@cs.stanford.edu

Richard Kelsey
Ember Corporation
Boston, MA
USA

Phone: +1 617 951 1225
Email: kelsey@ember.com

Jonathan W. Hui

Arch Rock Corporation
501 2nd St. Ste. 410
San Francisco, CA 94107
USA

Email: jhui@archrock.com

Winter, et al.

Expires December 30, 2010

[Page 111]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Kris Pister
Dust Networks
30695 Huntwood Ave.
Hayward, 94544
USA

Email: kpister@dustnetworks.com

Anders Brandt
Sigma Designs
Emdrupvej 26A, 1.
Copenhagen, DK-2100
Denmark

Email: abr@sdesigns.dk

R. Struik

Email: rstruik.ext@gmail.com

Stephen Dawson-Haggerty
UC Berkeley
Soda Hall, UC Berkeley
Berkeley, CA 94720
USA

Email: stevedh@cs.berkeley.edu

[21.](#) References

[21.1.](#) Normative References

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

[21.2.](#) Informative References

[AppliedCryptography]

Menezes, AJ., van Oorschot, PC., and SA. Vanstone,
"Handbook of Applied Cryptography", CRC Press , 1997.

- [CCMStar] IEEE, "IEEE Std. 802.15.4-2006, IEEE Standard for Information Technology - Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks - Specific requirements Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)", IEEE Press Revision of IEEE Std 802.15.4-2003, 2006.

[I-D.hui-6man-rpl-option]

Hui, J. and J. Vasseur, "RPL Option for Carrying RPL Information in Data-Plane Datagrams",
[draft-hui-6man-rpl-option-01](#) (work in progress),
June 2010.

[I-D.hui-6man-rpl-routing-header]

Hui, J., Vasseur, J., and D. Culler, "An IPv6 Routing Header for Source Routes with RPL", [draft-hui-6man-rpl-routing-header-02](#) (work in progress), June 2010.

[I-D.ietf-manet-nhdp]

Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", [draft-ietf-manet-nhdp-12](#) (work in progress), March 2010.

[I-D.ietf-roll-of0]

Thubert, P., "RPL Objective Function 0", [draft-ietf-roll-of0-02](#) (work in progress), June 2010.

[I-D.ietf-roll-routing-metrics]

Vasseur, J., Kim, M., Networks, D., and H. Chong, "Routing Metrics used for Path Calculation in Low Power and Lossy Networks", [draft-ietf-roll-routing-metrics-07](#) (work in progress), June 2010.

Winter, et al.

Expires December 30, 2010

[Page 113]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

[I-D.ietf-roll-terminology]

Vasseur, J., "Terminology in Low power And Lossy Networks", [draft-ietf-roll-terminology-03](#) (work in progress), March 2010.

[I-D.ietf-roll-trickle]

Levis, P., Clausen, T., Hui, J., and J. Ko, "The Trickle Algorithm", [draft-ietf-roll-trickle-01](#) (work in progress), April 2010.

[Perlman83]

Perlman, R., "Fault-Tolerant Broadcast of Routing Information", North-Holland Computer Networks 7: 395-405, 1983, <<http://www.cs.illinois.edu/~pbg/courses/cs598fa09/readings/p83.pdf>>.

[RFC1958] Carpenter, B., "Architectural Principles of the Internet", [RFC 1958](#), June 1996.

[RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", [RFC 1982](#), August 1996.

- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, [RFC 2578](#), April 1999.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.
- [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", [RFC 3307](#), August 2002.
- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", [RFC 3410](#), December 2002.
- [RFC3535] Schoenwaelder, J., "Overview of the 2002 IAB Network Management Workshop", [RFC 3535](#), May 2003.
- [RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", [RFC 3610](#), September 2003.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC3819] Karn, P., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L.

Winter, et al.

Expires December 30, 2010

[Page 114]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

Wood, "Advice for Internet Subnetwork Designers", [BCP 89](#), [RFC 3819](#), July 2004.

- [RFC4101] Rescorla, E. and IAB, "Writing Protocol Models", [RFC 4101](#), June 2005.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", [RFC 4191](#), November 2005.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,

"Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.

- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), June 2007.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", [RFC 5548](#), May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", [RFC 5706](#), November 2009.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5826](#), April 2010.
- [RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", [RFC 5867](#), June 2010.

Winter, et al.

Expires December 30, 2010

[Page 115]

Internet-Draft

[draft-ietf-roll-rpl-10](#)

Jun 2010

- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), June 2010.

[X9.63-2001]

"ANSI X9.63-2001, Public Key Cryptography for the Financial Services Industry - Key Agreement and Key Transport Using Elliptic Curve Cryptography", 2001.

[X9.92] "ANSI X9.92, Public Key Cryptography for the Financial Services Industry - Digital Signature Algorithms Giving Partial Message Recovery - Part 1: Elliptic Curve Pintsov-Vanstone Signatures (ECPVS)", 2009.

Authors' Addresses

Tim Winter (editor)

Email: wintert@acm.org

Pascal Thubert (editor)
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 497 23 26 34

Email: pthubert@cisco.com

RPL Author Team

IETF ROLL WG

Email: rpl-authors@external.cisco.com

