

Internet Engineering Task Force (IETF)
Request for Comments: 7416
Category: Informational
ISSN: 2070-1721

T. Tsao
R. Alexander
Eaton's Cooper Power Systems Business
M. Dohler
CTTC
V. Daza
A. Lozano
Universitat Pompeu Fabra
M. Richardson, Ed.
Sandelman Software Works
January 2015

A Security Threat Analysis for
the Routing Protocol for Low-Power and Lossy Networks (RPLs)

Abstract

This document presents a security threat analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs). The development builds upon previous work on routing security and adapts the assessments to the issues and constraints specific to low-power and lossy networks. A systematic approach is used in defining and evaluating the security threats. Applicable countermeasures are application specific and are addressed in relevant applicability statements.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7416>.

Copyright Notice

Copyright (c) 2015 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	3
2. Relationship to Other Documents	4
3. Terminology	5
4. Considerations on RPL Security	5
4.1. Routing Assets and Points of Access	6
4.2. The ISO 7498-2 Security Reference Model	8
4.3. Issues Specific to or Amplified in LLNs	10
4.4. RPL Security Objectives	12
5. Threat Sources	13
6. Threats and Attacks	13
6.1. Threats Due to Failures to Authenticate	14
6.1.1. Node Impersonation	14
6.1.2. Dummy Node	14
6.1.3. Node Resource Spam	15
6.2. Threats Due to Failure to Keep Routing Information Confidential	15
6.2.1. Routing Exchange Exposure	15
6.2.2. Routing Information (Routes and Network Topology) Exposure	15
6.3. Threats and Attacks on Integrity	16
6.3.1. Routing Information Manipulation	16
6.3.2. Node Identity Misappropriation	17
6.4. Threats and Attacks on Availability	18
6.4.1. Routing Exchange Interference or Disruption	18
6.4.2. Network Traffic Forwarding Disruption	18
6.4.3. Communications Resource Disruption	20
6.4.4. Node Resource Exhaustion	20

7.	Countermeasures	21
7.1.	Confidentiality Attack Countermeasures	21
7.1.1.	Countering Deliberate Exposure Attacks	21
7.1.2.	Countering Passive Wiretapping Attacks	22
7.1.3.	Countering Traffic Analysis	22
7.1.4.	Countering Remote Device Access Attacks	23
7.2.	Integrity Attack Countermeasures	24
7.2.1.	Countering Unauthorized Modification Attacks	24
7.2.2.	Countering Overclaiming and Misclaiming Attacks	24
7.2.3.	Countering Identity (including Sybil) Attacks	25
7.2.4.	Countering Routing Information Replay Attacks	25
7.2.5.	Countering Byzantine Routing Information Attacks	26
7.3.	Availability Attack Countermeasures	26
7.3.1.	Countering HELLO Flood Attacks and ACK Spoofing Attacks	27
7.3.2.	Countering Overload Attacks	27
7.3.3.	Countering Selective Forwarding Attacks	29
7.3.4.	Countering Sinkhole Attacks	29
7.3.5.	Countering Wormhole Attacks	30
8.	RPL Security Features	31
8.1.	Confidentiality Features	32
8.2.	Integrity Features	32
8.3.	Availability Features	33
8.4.	Key Management	34
9.	Security Considerations	34
10.	References	34
10.1.	Normative References	34
10.2.	Informative References	35
	Acknowledgments	39
	Authors' Addresses	40

1. Introduction

In recent times, networked electronic devices have found an increasing number of applications in various fields. Yet, for reasons ranging from operational application to economics, these wired and wireless devices are often supplied with minimum physical resources; the constraints include those on computational resources (RAM, clock speed, and storage) and communication resources (duty cycle, packet size, etc.) but also form factors that may rule out user-access interfaces (e.g., the housing of a small stick-on switch) or simply safety considerations (e.g., with gas meters). As a consequence, the resulting networks are more prone to loss of traffic and other vulnerabilities. The proliferation of these Low-Power and Lossy Networks (LLNs), however, are drawing efforts to examine and address their potential networking challenges. Securing the establishment and maintenance of network connectivity among these deployed devices becomes one of these key challenges.

This document presents a threat analysis for securing the Routing Protocol for LLNs (RPL). The process requires two steps. First, the analysis will be used to identify pertinent security issues. The second step is to identify necessary countermeasures to secure RPL. As there are multiple ways to solve the problem and the specific trade-offs are deployment specific, the specific countermeasure to be used is detailed in applicability statements.

This document uses a model based on [ISO.7498-2.1989], which describes authentication, access control, data confidentiality, data integrity, and non-repudiation security services. This document expands the model to include the concept of availability. As explained below, non-repudiation does not apply to routing protocols.

Many of the issues in this document were also covered in the IAB Smart Object Workshop [RFC6574] and the IAB Smart Object Security Workshop [RFC7397].

This document concerns itself with securing the control-plane traffic. As such, it does not address authorization or authentication of application traffic. RPL uses multicast as part of its protocol; therefore, mechanisms that RPL uses to secure this traffic might also be applicable to the Multicast Protocol for Low-Power and Lossy Networks (MPL) control traffic as well: the important part is that the threats are similar.

2. Relationship to Other Documents

Routing Over Low-Power and Lossy (ROLL) networks has specified a set of routing protocols for LLNs [RFC6550]. A number of applicability texts describe a subset of these protocols and the conditions that make the subset the correct choice. The text recommends and motivates the accompanying parameter value ranges. Multiple applicability domains are recognized, including Building and Home and Advanced Metering Infrastructure. The applicability domains distinguish themselves in the way they are operated, by their performance requirements, and by the most probable network structures. Each applicability statement identifies the distinguishing properties according to a common set of subjects described in as many sections.

The common set of security threats herein are referred to by the applicability statements, and that series of documents describes the preferred security settings and solutions within the applicability statement conditions. This applicability statement may recommend more lightweight security solutions and specify the conditions under which these solutions are appropriate.

3. Terminology

This document adopts the terminology defined in [RFC6550], [RFC4949], and [RFC7102].

The terms "control plane" and "forwarding plane" are used in a manner consistent with Section 1 of [RFC6192].

The term "Destination-Oriented DAG (DODAG)" is from [RFC6550].

Extensible Authentication Protocol - Transport Layer Security (EAP-TLS) is defined in [RFC5216].

The Protocol for Carrying Authentication for Network Access (PANA) is defined in [RFC5191].

Counter with CBC-MAC (CCM) mode is defined in [RFC3610].

The term "sleepy node", introduced in [RFC7102], refers to a node that may sometimes go into a low-power state, suspending protocol communications.

The terms Service Set Identifier (SSID), Extended Service Set Identifier (ESSID), and Personal Area Network (PAN) refer to network identifiers, defined in [IEEE.802.11] and [IEEE.802.15.4].

Although this is not a protocol specification, 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 [RFC2119] in order to clarify and emphasize the guidance and directions to implementers and deployers of LLN nodes that utilize RPL.

4. Considerations on RPL Security

Routing security, in essence, ensures that the routing protocol operates correctly. It entails implementing measures to ensure controlled state changes on devices and network elements, both based on external inputs (received via communications) or internal inputs (physical security of the device itself and parameters maintained by the device, including, e.g., clock). State changes would thereby involve not only authorization of the injector's actions, authentication of injectors, and potentially confidentiality of routing data, but also proper order of state changes through timeliness, since seriously delayed state changes, such as commands or updates of routing tables, may negatively impact system operation. A security assessment can, therefore, begin with a focus on the assets [RFC4949] that may be the target of the state changes and the

access points in terms of interfaces and protocol exchanges through which such changes may occur. In the case of routing security, the focus is directed towards the elements associated with the establishment and maintenance of network connectivity.

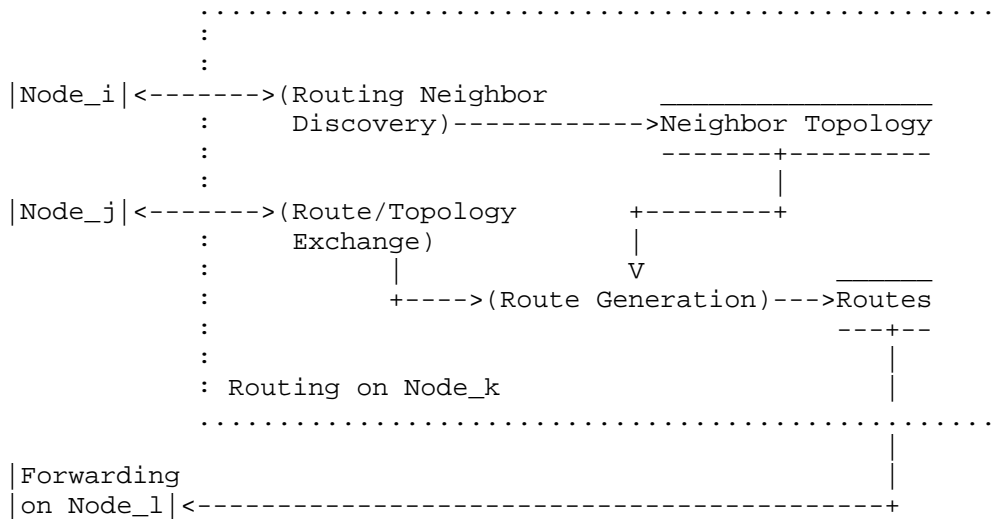
This section sets the stage for the development of the analysis by applying the systematic approach proposed in [Myagmar2005] to the routing security, while also drawing references from other reviews and assessments found in the literature, particularly [RFC4593] and [Karlof2003] (i.e., selective forwarding, wormhole, and sinkhole attacks). The subsequent subsections begin with a focus on the elements of a generic routing process that is used to establish routing assets and points of access to the routing functionality. Next, the security model based on [ISO.7498-2.1989] is briefly described. Then, consideration is given to issues specific to or amplified in LLNs. This section concludes with the formulation of a set of security objectives for RPL.

4.1. Routing Assets and Points of Access

An asset is an important system resource (including information, process, or physical resource); the access to and corruption or loss of an asset adversely affects the system. In the control-plane context, an asset is information about the network, processes used to manage and manipulate this data, and the physical devices on which this data is stored and manipulated. The corruption or loss of these assets may adversely impact the control plane of the network. Within the same context, a point of access is an interface or protocol that facilitates interaction between control-plane assets. Identifying these assets and points of access will provide a basis for enumerating the attack surface of the control plane.

A level-0 data flow diagram [Yourdon1979] is used here to identify the assets and points of access within a generic routing process. The use of a data flow diagram allows for a clear and concise model of the way in which routing nodes interact and process information; hence, it provides a context for threats and attacks. The goal of the model is to be as detailed as possible so that corresponding assets, points of access, and processes in an individual routing protocol can be readily identified.

Figure 1 shows that nodes participating in the routing process transmit messages to discover neighbors and to exchange routing information; routes are then generated and stored, which may be maintained in the form of the protocol forwarding table. The nodes use the derived routes for making forwarding decisions.



Notation:

(Proc) A process Proc

topology A structure storing neighbor adjacency (parent/child)

routes A structure storing the forwarding information base (FIB)

|Node_n| An external entity Node_n

-----> Data flow

Figure 1: Data Flow Diagram of a Generic Routing Process

Figure 1 shows the following:

- o Assets include
 - * routing and/or topology information;
 - * route generation process;
 - * communication channel resources (bandwidth);

- * node resources (computing capacity, memory, and remaining energy); and
 - * node identifiers (including node identity and ascribed attributes such as relative or absolute node location).
- o Points of access include
- * neighbor discovery;
 - * route/topology exchange; and
 - * node physical interfaces (including access to data storage).

A focus on the above list of assets and points of access enables a more directed assessment of routing security; for example, it is readily understood that some routing attacks are in the form of attempts to misrepresent routing topology. Indeed, the intention of the security threat analysis is to be comprehensive. Hence, some of the discussion that follows is associated with assets and points of access that are not directly related to routing protocol design but are nonetheless provided for reference since they do have direct consequences on the security of routing.

4.2. The ISO 7498-2 Security Reference Model

At the conceptual level, security within an information system, in general, and applied to RPL in particular is concerned with the primary issues of authentication, access control, data confidentiality, data integrity, and non-repudiation. In the context of RPL:

Authentication

Authentication involves the mutual authentication of the routing peers prior to exchanging route information (i.e., peer authentication) as well as ensuring that the source of the route data is from the peer (i.e., data origin authentication). LLNs can be drained by unauthenticated peers before configuration per [RFC5548]. Availability of open and untrusted side channels for new joiners is required by [RFC5673], and strong and automated authentication is required so that networks can automatically accept or reject new joiners.

Access Control

Access Control provides protection against unauthorized use of the asset and deals with the authorization of a node.

Confidentiality

Confidentiality involves the protection of routing information as well as routing neighbor maintenance exchanges so that only authorized and intended network entities may view or access it. Because LLNs are most commonly found on a publicly accessible shared medium, e.g., air or wiring in a building, and are sometimes formed ad hoc, confidentiality also extends to the neighbor state and database information within the routing device since the deployment of the network creates the potential for unauthorized access to the physical devices themselves.

Integrity

Integrity entails the protection of routing information and routing neighbor maintenance exchanges, as well as derived information maintained in the database, from unauthorized modifications, insertions, deletions, or replays to be addressed beyond the routing protocol.

Non-repudiation

Non-repudiation is the assurance that the transmission and/or reception of a message cannot later be denied. The service of non-repudiation applies after the fact; thus, it relies on the logging or other capture of ongoing message exchanges and signatures. Routing protocols typically do not have a notion of repudiation, so non-repudiation services are not required. Further, with the LLN application domains as described in [RFC5867] and [RFC5548], proactive measures are much more critical than retrospective protections. Finally, given the significant practical limits to ongoing routing transaction logging and storage and individual device digital signature verification for each exchange, non-repudiation in the context of routing is an unsupportable burden that bears no further consideration as an RPL security issue.

It is recognized that, besides those security issues captured in the ISO 7498-2 model, availability is a security requirement:

Availability

Availability ensures that routing information exchanges and forwarding services are available when they are required for the functioning of the serving network. Availability will apply to maintaining efficient and correct operation of routing and neighbor discovery exchanges (including needed information) and forwarding services so as not to impair or limit the network's central traffic flow function.

It should be emphasized here that for RPL security, the above requirements must be complemented by the proper security policies and enforcement mechanisms to ensure that security objectives are met by a given RPL implementation.

4.3. Issues Specific to or Amplified in LLNs

The requirements work detailed in Urban Requirements [RFC5548], Industrial Requirements [RFC5673], Home Automation [RFC5826], and Building Automation [RFC5867] have identified specific issues and constraints of routing in LLNs. The following is a list of observations from those requirements and evaluations of their impact on routing security considerations.

Limited energy, memory, and processing node resources

As a consequence of these constraints, the need to evaluate the kinds of security that can be provided needs careful study. For instance, security provided at one level could be very memory efficient yet might also be very energy costly for the network (as a whole) if it requires significant effort to synchronize the security state. Synchronization of security states with sleepy nodes [RFC7102] is a complex issue. A non-rechargeable battery-powered node may well be limited in energy for its lifetime: once exhausted, it may well never function again.

Large scale of rolled out network

The possibly numerous nodes to be deployed make manual on-site configuration unlikely. For example, an urban deployment can see several hundreds of thousands of nodes being installed by many installers with a low level of expertise. Nodes may be installed and not activated for many years, and additional nodes may be added later on, which may be from old inventory. The lifetime of the network is measured in decades, and this complicates the operation of key management.

Autonomous operations

Self-forming and self-organizing are commonly prescribed requirements of LLNs. In other words, a routing protocol designed for LLNs needs to contain elements of ad hoc networking and, in most cases, cannot rely on manual configuration for initialization or local filtering rules. Network topology/ownership changes, partitioning or merging, and node replacement can all contribute to complicating the operations of key management.

Highly directional traffic

Some types of LLNs see a high percentage of their total traffic traverse between the nodes and the LLN Border Routers (LBRs) where the LLNs connect to non-LLNs. The special routing status of and the greater volume of traffic near the LBRs have routing security consequences as a higher-valued attack target. In fact, when Point-to-MultiPoint (P2MP) and MultiPoint-to-Point (MP2P) traffic represents a majority of the traffic, routing attacks consisting of advertising incorrect preferred routes can cause serious damage.

While it might seem that nodes higher up in the acyclic graph (i.e., those with lower rank) should be secured in a stronger fashion, it is not, in general, easy to predict which nodes will occupy those positions until after deployment. Issues of redundancy and inventory control suggest that any node might wind up in such a sensitive attack position, so all nodes are to be capable of being fully secured.

In addition, even if it were possible to predict which nodes will occupy positions of lower rank and provision them with stronger security mechanisms, in the absence of a strong authorization model, any node could advertise an incorrect preferred route.

Unattended locations and limited physical security

In many applications, the nodes are deployed in unattended or remote locations; furthermore, the nodes themselves are often built with minimal physical protection. These constraints lower the barrier of accessing the data or security material stored on the nodes through physical means.

Support for mobility

On the one hand, only a limited number of applications require the support of mobile nodes, e.g., a home LLN that includes nodes on wearable health care devices or an industry LLN that includes nodes on cranes and vehicles. On the other hand, if a routing protocol is indeed used in such applications, it will clearly need to have corresponding security mechanisms.

Additionally, nodes may appear to move from one side of a wall to another without any actual motion involved, which is the result of changes to electromagnetic properties, such as the opening and closing of a metal door.

Support for multicast and anycast

Support for multicast and anycast is called out chiefly for large-scale networks. Since application of these routing mechanisms in autonomous operations of many nodes is new, the consequence on security requires careful consideration.

The above list considers how an LLN's physical constraints, size, operations, and variety of application areas may impact security. However, it is the combinations of these factors that particularly stress the security concerns. For instance, securing routing for a large number of autonomous devices that are left in unattended locations with limited physical security presents challenges that are not found in the common circumstance of administered networked routers. The following subsection sets up the security objectives for the routing protocol designed by the ROLL WG.

4.4. RPL Security Objectives

This subsection applies the ISO 7498-2 model to routing assets and access points, taking into account the LLN issues, to develop a set of RPL security objectives.

Since the fundamental function of a routing protocol is to build routes for forwarding packets, it is essential to ensure that:

- o routing/topology information integrity remains intact during transfer and in storage;
- o routing/topology information is used by authorized entities; and
- o routing/topology information is available when needed.

In conjunction, it is necessary to be assured that:

- o Authorized peers authenticate themselves during the routing neighbor discovery process.
- o The routing/topology information received is generated according to the protocol design.

However, when trust cannot be fully vested through authentication of the principals alone, i.e., concerns of an insider attack, assurance of the truthfulness and timeliness of the received routing/topology information is necessary. With regard to confidentiality, protecting the routing/topology information from unauthorized exposure may be desirable in certain cases but is in itself less pertinent, in general, to the routing function.

One of the main problems of synchronizing security states of sleepy nodes, as listed in the last subsection, lies in difficulties in authentication; these nodes may not have received the most recent update of security material in time. Similarly, the issues of minimal manual configuration, prolonged rollout and delayed addition of nodes, and network topology changes also complicate key management. Hence, routing in LLNs needs to bootstrap the authentication process and allow for a flexible expiration scheme of authentication credentials.

The vulnerability brought forth by some special-function nodes, e.g., LBRs, requires the assurance, particularly in a security context, of the following:

- o The availability of communication channels and node resources.
- o The neighbor discovery process operates without undermining routing availability.

There are other factors that are not part of RPL but directly affect its function. These factors include a weaker barrier of accessing the data or security material stored on the nodes through physical means; therefore, the internal and external interfaces of a node need to be adequate for guarding the integrity, and possibly the confidentiality, of stored information, as well as the integrity of routing and route generation processes.

Each individual system's use and environment will dictate how the above objectives are applied, including the choices of security services as well as the strengths of the mechanisms that must be implemented. The next two sections take a closer look at how the RPL security objectives may be compromised and how those potential compromises can be countered.

5. Threat Sources

[RFC4593] provides a detailed review of the threat sources: outsiders and Byzantine. RPL has the same threat sources.

6. Threats and Attacks

This section outlines general categories of threats under the ISO 7498-2 model and highlights the specific attacks in each of these categories for RPL. As defined in [RFC4949], a threat is "a potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm."

Per [RFC3067], an attack is "an assault on system security that derives from an intelligent threat, i.e., an intelligent act that is a deliberate attempt (especially in the sense of a method or technique) to evade security services and violate the security policy of a system."

The subsequent subsections consider the threats and the attacks that can cause security breaches under the ISO 7498-2 model to the routing assets and via the routing points of access identified in Section 4.1. The assessment reviews the security concerns of each routing asset and looks at the attacks that can exploit routing points of access. The threats and attacks identified are based on the routing model analysis and associated review of the existing literature. The source of the attacks is assumed to be from either inside or outside attackers. While some attackers inside the network will be using compromised nodes and, therefore, are only able to do what an ordinary node can ("node-equivalent"), other attacks may not be limited in memory, CPU, power consumption, or long-term storage. Moore's law favors the attacker with access to the latest capabilities, while the defenders will remain in place for years to decades.

6.1. Threats Due to Failures to Authenticate

6.1.1. Node Impersonation

If an attacker can join a network using any identity, then it may be able to assume the role of a legitimate (and existing node). It may be able to report false readings (in metering applications) or provide inappropriate control messages (in control systems involving actuators) if the security of the application is implied by the security of the routing system.

Even in systems where there is application-layer security, the ability to impersonate a node would permit an attacker to direct traffic to itself. This may permit various on-path attacks that would otherwise be difficult, such as replaying, delaying, or duplicating (application) control messages.

6.1.2. Dummy Node

If an attacker can join a network using any identify, then it can pretend to be a legitimate node, receiving any service legitimate nodes receive. It may also be able to report false readings (in metering applications), provide inappropriate authorizations (in control systems involving actuators), or perform any other attacks that are facilitated by being able to direct traffic towards itself.

6.1.3. Node Resource Spam

If an attacker can join a network with any identity, then it can continuously do so with new (random) identities. This act may drain down the resources of the network (battery, RAM, bandwidth). This may cause legitimate nodes of the network to be unable to communicate.

6.2. Threats Due to Failure to Keep Routing Information Confidential

The assessment in Section 4.2 indicates that there are attacks against the confidentiality of routing information at all points of access. This threat may result in disclosure, as described in Section 3.1.2 of [RFC4593], and may involve a disclosure of routing information.

6.2.1. Routing Exchange Exposure

Routing exchanges include both routing information as well as information associated with the establishment and maintenance of neighbor state information. As indicated in Section 4.1, the associated routing information assets may also include device-specific resource information, such as available memory, remaining power, etc., that may be metrics of the routing protocol.

The routing exchanges will contain reachability information, which would identify the relative importance of different nodes in the network. Nodes higher up in the DODAG, to which more streams of information flow, would be more interesting targets for other attacks, and routing exchange exposures could identify them.

6.2.2. Routing Information (Routes and Network Topology) Exposure

Routes (which may be maintained in the form of the protocol forwarding table) and neighbor topology information are the products of the routing process that are stored within the node device databases.

The exposure of this information will allow attackers to gain direct access to the configuration and connectivity of the network, thereby exposing routing to targeted attacks on key nodes or links. Since routes and neighbor topology information are stored within the node device, attacks on the confidentiality of the information will apply to the physical device, including specified and unspecified internal and external interfaces.

The forms of attack that allow unauthorized access or disclosure of the routing information will include:

- o Physical device compromise.
- o Remote device access attacks (including those occurring through remote network management or software/field upgrade interfaces).

Both of these attack vectors are considered a device-specific issue and are out of scope for RPL to defend against. In some applications, physical device compromise may be a real threat, and it may be necessary to provide for other devices to securely detect a compromised device and react quickly to exclude it.

6.3. Threats and Attacks on Integrity

The assessment in Section 4.2 indicates that information and identity assets are exposed to integrity threats from all points of access. In other words, the integrity threat space is defined by the potential for exploitation introduced by access to assets available through routing exchanges and the on-device storage.

6.3.1. Routing Information Manipulation

Manipulation of routing information that ranges from neighbor states to derived routes will allow unauthorized sources to influence the operation and convergence of the routing protocols and ultimately impact the forwarding decisions made in the network.

Manipulation of topology and reachability information will allow unauthorized sources to influence the nodes with which routing information is exchanged and updated. The consequence of manipulating routing exchanges can thus lead to suboptimality and fragmentation or partitioning of the network by restricting the universe of routers with which associations can be established and maintained.

A suboptimal network may use too much power and/or may congest some routes leading to premature failure of a node and a denial of service (DoS) on the entire network.

In addition, being able to attract network traffic can make a black-hole attack more damaging.

The forms of attack that allow manipulation to compromise the content and validity of routing information include:

- o falsification, including overclaiming and misclaiming (claiming routes to devices that the device cannot in fact reach);
- o routing information replay;
- o Byzantine (internal) attacks that permit corruption of routing information in the node even when the node continues to be a validated entity within the network (see, for example, [RFC4593] for further discussions on Byzantine attacks); and
- o physical device compromise or remote device access attacks.

6.3.2. Node Identity Misappropriation

Falsification or misappropriation of node identity between routing participants opens the door for other attacks; it can also cause incorrect routing relationships to form and/or topologies to emerge. Routing attacks may also be mounted through less-sophisticated node identity misappropriation in which the valid information broadcasted or exchanged by a node is replayed without modification. The receipt of seemingly valid information that is, however, no longer current can result in routing disruption and instability (including failure to converge). Without measures to authenticate the routing participants and to ensure the freshness and validity of the received information, the protocol operation can be compromised. The forms of attack that misuse node identity include:

- o Identity attacks, including Sybil attacks (see [Sybil2002]) in which a malicious node illegitimately assumes multiple identities.
- o Routing information replay.

6.4. Threats and Attacks on Availability

The assessment in Section 4.2 indicates that the process and resource assets are exposed to threats against availability; attacks in this category may exploit directly or indirectly information exchange or forwarding (see [RFC4732] for a general discussion).

6.4.1. Routing Exchange Interference or Disruption

Interference is the threat action and disruption is the threat consequence that allows attackers to influence the operation and convergence of the routing protocols by impeding the routing information exchange.

The forms of attack that allow interference or disruption of routing exchange include:

- o routing information replay;
- o ACK spoofing; and
- o overload attacks (Section 7.3.2).

In addition, attacks may also be directly conducted at the physical layer in the form of jamming or interfering.

6.4.2. Network Traffic Forwarding Disruption

The disruption of the network traffic forwarding capability will undermine the central function of network routers and the ability to handle user traffic. This affects the availability of the network because of the potential to impair the primary capability of the network.

6.4.3. Communications Resource Disruption

Attacks mounted against the communication channel resource assets needed by the routing protocol can be used as a means of disrupting its operation. However, while various forms of DoS attacks on the underlying transport subsystem will affect routing protocol exchanges and operation (for example, physical-layer Radio Frequency (RF) jamming in a wireless network or link-layer attacks), these attacks cannot be countered by the routing protocol. As such, the threats to the underlying transport network that supports routing is considered beyond the scope of the current document. Nonetheless, attacks on the subsystem will affect routing operation and must be directly addressed within the underlying subsystem and its implemented protocol layers.

6.4.4. Node Resource Exhaustion

A potential threat consequence can arise from attempts to overload the node resource asset by initiating exchanges that can lead to the exhaustion of processing, memory, or energy resources. The establishment and maintenance of routing neighbors opens the routing process to engagement and potential acceptance of multiple neighboring peers. Association information must be stored for each peer entity and for the wireless network operation provisions made to periodically update and reassess the associations. An introduced proliferation of apparent routing peers can, therefore, have a negative impact on node resources.

Node resources may also be unduly consumed by attackers attempting uncontrolled topology peering or routing exchanges, routing replays, or the generating of other data-traffic floods. Beyond the disruption of communications channel resources, these consequences may be able to exhaust node resources only where the engagements are able to proceed with the peer routing entities. Routing operation and network forwarding functions can thus be adversely impacted by node resources exhaustion that stems from attacks that include:

- o identity (including Sybil) attacks (see [Sybil2002]);
- o routing information replay attacks;
- o HELLO-type flood attacks; and
- o overload attacks (Section 7.3.2).

7. Countermeasures

By recognizing the characteristics of LLNs that may impact routing, this analysis provides the basis for understanding the capabilities within RPL used to deter the identified attacks and mitigate the threats. The following subsections consider such countermeasures by grouping the attacks according to the classification of the ISO 7498-2 model so that associations with the necessary security services are more readily visible.

7.1. Confidentiality Attack Countermeasures

Attacks to disclosure routing information may be mounted at the level of the routing information assets, at the points of access associated with routing exchanges between nodes, or through device interface access. To gain access to routing/topology information, the attacker may rely on a compromised node that deliberately exposes the information during the routing exchange process, on passive wiretapping or traffic analysis, or on attempting access through a component or device interface of a tampered routing node.

7.1.1. Countering Deliberate Exposure Attacks

A deliberate exposure attack is one in which an entity that is party to the routing process or topology exchange allows the routing/topology information or generated route information to be exposed to an unauthorized entity.

For instance, due to misconfiguration or inappropriate enabling of a diagnostic interface, an entity might be copying ("bridging") traffic from a secured ESSID/PAN to an unsecured interface.

A prerequisite to countering this attack is to ensure that the communicating nodes are authenticated prior to data encryption applied in the routing exchange. The authentication ensures that the LLN starts with trusted nodes, but it does not provide an indication of whether the node has been compromised.

Reputation systems could be used to help when some nodes may sleep for extended periods of time. It is also unclear if resulting datasets would even fit into constrained devices.

To mitigate the risk of deliberate exposure, the process that communicating nodes use to establish session keys must be peer-to-peer (i.e., between the routing initiating and responding nodes). As is pointed out in [RFC4107], automatic key management is critical for good security. This helps ensure that neither node is exchanging routing information with another peer without the

knowledge of both communicating peers. For a deliberate exposure attack to succeed, the comprised node will need to be more overt and take independent actions in order to disclose the routing information to a third party.

Note that the same measures that apply to securing routing/topology exchanges between operational nodes must also extend to field tools and other devices used in a deployed network where such devices can be configured to participate in routing exchanges.

7.1.2. Countering Passive Wiretapping Attacks

A passive wiretap attack seeks to breach routing confidentiality through passive, direct analysis and processing of the information exchanges between nodes.

Passive wiretap attacks can be directly countered through the use of data encryption for all routing exchanges. Only when a validated and authenticated node association is completed will routing exchange be allowed to proceed using established session keys and an agreed encryption algorithm. The mandatory-to-implement CCM mode AES-128 method, described in [RFC3610], is believed to be secure against a brute-force attack by even the most well-equipped adversary.

The significant challenge for RPL is in the provisioning of the key, which in some modes of RFC 6550 is used network wide. This problem is not solved in RFC 6550, and it is the subject of significant future work: see, for instance, [AceCharterProposal], [SolaceProposal], and [SmartObjectSecurityWorkshop].

A number of deployments, such as [ZigBeeIP] specify no Layer 3 (L3) / RPL encryption or authentication and rely upon similar security at Layer 2 (L2). These networks are immune to outside wiretapping attacks but are vulnerable to passive (and active) routing attacks through compromises of nodes (see Section 8.2).

Section 10.9 of [RFC6550] specifies AES-128 in CCM mode with a 32-bit Message Authentication Code (MAC).

Section 5.6 of ZigBee IP [ZigBeeIP] specifies use of CCM, with PANA and EAP-TLS for key management.

7.1.3. Countering Traffic Analysis

Traffic analysis provides an indirect means of subverting confidentiality and gaining access to routing information by allowing an attacker to indirectly map the connectivity or flow patterns (including link load) of the network from which other attacks can be

mounted. The traffic-analysis attack on an LLN, especially one founded on a shared medium, is passive and relies on the ability to read the immutable source/destination L2 and/or L3 routing information that must remain unencrypted to permit network routing.

One way in which passive traffic-analysis attacks can be muted is through the support of load balancing that allows traffic to a given destination to be sent along diverse routing paths. RPL does not generally support multipath routing within a single DODAG. Multiple DODAGs are supported in the protocol, and an implementation could make use of that. RPL does not have any inherent or standard way to guarantee that the different DODAGs would have significantly diverse paths. Having the diverse DODAGs routed at different border routers might work in some instances, and this could be combined with a multipath technology like Multipath TCP (MPTCP) [RFC6824]. It is unlikely that it will be affordable in many LLNs, as few deployments will have memory space for more than a few sets of DODAG tables.

Another approach to countering passive traffic analysis could be for nodes to maintain a constant amount of traffic to different destinations through the generation of arbitrary traffic flows; the drawback of course would be the consequent overhead and energy expenditure.

The only means of fully countering a traffic-analysis attack is through the use of tunneling (encapsulation) where encryption is applied across the entirety of the original packet source/destination addresses. Deployments that use L2 security that includes encryption already do this for all traffic.

7.1.4. Countering Remote Device Access Attacks

Where LLN nodes are deployed in the field, measures are introduced to allow for remote retrieval of routing data and for software or field upgrades. These paths create the potential for a device to be remotely accessed across the network or through a provided field tool. In the case of network management, a node can be directly requested to provide routing tables and neighbor information.

To ensure confidentiality of the node routing information against attacks through remote access, any local or remote device requesting routing information must be authenticated and must be authorized for that access. Since remote access is not invoked as part of a routing protocol, security of routing information stored on the node against remote access will not be addressable as part of the routing protocol.

7.2. Integrity Attack Countermeasures

Integrity attack countermeasures address routing information manipulation, as well as node identity and routing information misuse. Manipulation can occur in the form of a falsification attack and physical compromise. To be effective, the following development considers the two aspects of falsification, namely, the unauthorized modifications and the overclaiming and misclaiming content. The countering of physical compromise was considered in the previous section and is not repeated here. With regard to misuse, there are two types of attacks to be deterred: identity attacks and replay attacks.

7.2.1. Countering Unauthorized Modification Attacks

Unauthorized modifications may occur in the form of altering the message being transferred or the data stored. Therefore, it is necessary to ensure that only authorized nodes can change the portion of the information that is allowed to be mutable, while the integrity of the rest of the information is protected, e.g., through well-studied cryptographic mechanisms.

Unauthorized modifications may also occur in the form of insertion or deletion of messages during protocol changes. Therefore, the protocol needs to ensure the integrity of the sequence of the exchange sequence.

The countermeasure to unauthorized modifications needs to:

- o implement access control on storage;
- o provide data integrity service to transferred messages and stored data; and
- o include a sequence number under integrity protection.

7.2.2. Countering Overclaiming and Misclaiming Attacks

Both overclaiming and misclaiming aim to introduce false routes or a false topology that would not occur otherwise, while there are not necessarily unauthorized modifications to the routing messages or information. In order to counter overclaiming, the capability to determine unreasonable routes or topology is required.

The counter to overclaiming and misclaiming may employ:

- o Comparison with historical routing/topology data.

- o Designs that restrict realizable network topologies.

RPL includes no specific mechanisms in the protocol to counter overclaims or misclaims. An implementation could have specific heuristics implemented locally.

7.2.3. Countering Identity (including Sybil) Attacks

Identity attacks, sometimes simply called spoofing, seek to gain or damage assets whose access is controlled through identity. In routing, an identity attacker can illegitimately participate in routing exchanges, distribute false routing information, or cause an invalid outcome of a routing process.

A perpetrator of Sybil attacks assumes multiple identities. The result is not only an amplification of the damage to routing but extension to new areas, e.g., where geographic distribution is explicitly or implicitly an asset to an application running on the LLN, for example, the LBR in a P2MP or MP2P LLN.

RPL includes specific public key-based authentication at L3 that provides for authorization. Many deployments use L2 security that includes admission controls at L2 using mechanisms such as PANA.

7.2.4. Countering Routing Information Replay Attacks

In many routing protocols, message replay can result in false topology and/or routes. This is often countered with some kind of counter to ensure the freshness of the message. Replay of a current, literal RPL message is, in general, idempotent to the topology. If replayed, an older (lower DODAGVersionNumber) message would be rejected as being stale. If the trickle algorithm further dampens the effect of any such replay, as if the message was current, then it would contain the same information as before, and it would cause no network changes.

Replays may well occur in some radio technologies (though not very likely; see [IEEE.802.15.4]) as a result of echos or reflections, so some replays must be assumed to occur naturally.

Note that for there to be no effect at all, the replay must be done with the same apparent power for all nodes receiving the replay. A change in apparent power might change the metrics through changes to the Expected Transmission Count (ETX); therefore, it might affect the routing even though the contents of the packet were never changed. Any replay that appears to be different should be analyzed as a selective forwarding attack, sinkhole attack, or wormhole attack.

7.2.5. Countering Byzantine Routing Information Attacks

Where a node is captured or compromised but continues to operate for a period with valid network security credentials, the potential exists for routing information to be manipulated. This compromise of the routing information could thus exist in spite of security countermeasures that operate between the peer routing devices.

Consistent with the end-to-end principle of communications, such an attack can only be fully addressed through measures operating directly between the routing entities themselves or by means of external entities accessing and independently analyzing the routing information. Verification of the authenticity and liveness of the routing entities can, therefore, only provide a limited counter against internal (Byzantine) node attacks.

For link-state routing protocols where information is flooded with, for example, areas (OSPF [RFC2328]) or levels (IS-IS [RFC7142]), countermeasures can be directly applied by the routing entities through the processing and comparison of link-state information received from different peers. By comparing the link information from multiple sources, decisions can be made by a routing node or external entity with regard to routing information validity; see Chapter 2 of [Perlman1988] for a discussion on flooding attacks.

For distance vector protocols, such as RPL, where information is aggregated at each routing node, it is not possible for nodes to directly detect Byzantine information manipulation attacks from the routing information exchange. In such cases, the routing protocol must include and support indirect communications exchanges between non-adjacent routing peers to provide a secondary channel for performing routing information validation. S-RIP [Wan2004] is an example of the implementation of this type of dedicated routing protocol security where the correctness of aggregate distance vector information can only be validated by initiating confirmation exchanges directly between nodes that are not routing neighbors.

RPL does not provide any direct mechanisms like S-RIP. It does listen to multiple parents and may switch parents if it begins to suspect that it is being lied to.

7.3. Availability Attack Countermeasures

As alluded to before, availability requires that routing information exchanges and forwarding mechanisms be available when needed so as to guarantee proper functioning of the network. This may, e.g., include the correct operation of routing information and neighbor state information exchanges, among others. We will highlight the key

features of the security threats along with typical countermeasures to prevent or at least mitigate them. We will also note that an availability attack may be facilitated by an identity attack as well as a replay attack, as was addressed in Sections 7.2.3 and 7.2.4, respectively.

7.3.1. Countering HELLO Flood Attacks and ACK Spoofing Attacks

HELLO Flood [Karlof2003], [HELLO], and ACK spoofing attacks are different but highly related forms of attacking an LLN. They essentially lead nodes to believe that suitable routes are available even though they are not and hence constitute a serious availability attack.

A HELLO attack mounted against RPL would involve sending out (or replaying) DODAG Information Object (DIO) messages by the attacker. Lower-power LLN nodes might then attempt to join the DODAG at a lower rank than they would otherwise.

The most effective method from [HELLO] is bidirectional verification. A number of L2 links are arranged in controller/spoke arrangements and are continuously validating connectivity at layer 2.

In addition, in order to calculate metrics, the ETX must be computed, and this involves, in general, sending a number of messages between nodes that are believed to be adjacent. One such protocol is [MESH-LINK].

In order to join the DODAG, a Destination Advertisement Object (DAO) message is sent upwards. In RPL, the DAO is acknowledged by the DAO-ACK message. This clearly checks bidirectionality at the control plane.

As discussed in Section 5.1 of [HELLO], a receiver with a sensitive receiver could well hear the DAOs and even send DAO-ACKs as well. Such a node is a form of wormhole attack.

These attacks are also all easily defended against using either L2 or L3 authentication. Such an attack could only be made against a completely open network (such as might be used for provisioning new nodes) or by a compromised node.

7.3.2. Countering Overload Attacks

Overload attacks are a form of DoS attack in that a malicious node overloads the network with irrelevant traffic, thereby draining the nodes' energy store more quickly when the nodes rely on batteries or energy scavenging. Thus, it significantly shortens the lifetime of

networks of energy-constrained nodes and constitutes another serious availability attack.

With energy being one of the most precious assets of LLNs, targeting its availability is a fairly obvious attack. Another way of depleting the energy of an LLN node is to have the malicious node overload the network with irrelevant traffic. This impacts availability since certain routes get congested, which:

- o renders them useless for affected nodes; hence, data cannot be delivered;
- o makes routes longer as the shortest path algorithms work with the congested network; and
- o depletes battery and energy scavenging nodes more quickly and thus shortens the network's availability at large.

Overload attacks can be countered by deploying a series of mutually non-exclusive security measures that:

- o introduce quotas on the traffic rate each node is allowed to send;
- o isolate nodes that send traffic above a certain threshold based on system operation characteristics; and
- o allow only trusted data to be received and forwarded.

As for the first one, a simple approach to minimize the harmful impact of an overload attack is to introduce traffic quotas. This prevents a malicious node from injecting a large amount of traffic into the network, even though it does not prevent the said node from injecting irrelevant traffic at all. Another method is to isolate nodes from the network at the network layer once it has been detected that more traffic is injected into the network than allowed by a prior set or dynamically adjusted threshold. Finally, if communication is sufficiently secured, only trusted nodes can receive and forward traffic, which also lowers the risk of an overload attack.

Receiving nodes that validate signatures and sending nodes that encrypt messages need to be cautious of cryptographic processing usage when validating signatures and encrypting messages. Where feasible, certificates should be validated prior to use of the associated keys to counter potential resource overloading attacks. The associated design decision needs to also consider that the validation process requires resources; thus, it could be exploited for attacks. Alternatively, resource management limits can be placed

on routing security processing events (see the comment in Section 6, paragraph 4, of [RFC5751]).

7.3.3. Countering Selective Forwarding Attacks

Selective forwarding attacks are a form of DoS attack that impacts the availability of the generated routing paths.

A selective forwarding attack may be done by a node involved with the routing process, or it may be done by what otherwise appears to be a passive antenna or other RF feature or device, but is in fact an active (and selective) device. An RF antenna/repeater that is not selective is not a threat.

An insider malicious node basically blends in neatly with the network but then may decide to forward and/or manipulate certain packets. If all packets are dropped, then this attacker is also often referred to as a "black hole". Such a form of attack is particularly dangerous if coupled with sinkhole attacks since inherently a large amount of traffic is attracted to the malicious node, thereby causing significant damage. In a shared medium, an outside malicious node would selectively jam overheard data flows, where the thus caused collisions incur selective forwarding.

Selective forwarding attacks can be countered by deploying a series of mutually non-exclusive security measures:

- o Multipath routing of the same message over disjoint paths.
- o Dynamically selecting the next hop from a set of candidates.

The first measure basically guarantees that if a message gets lost on a particular routing path due to a malicious selective forwarding attack, there will be another route that successfully delivers the data. Such a method is inherently suboptimal from an energy consumption point of view; it is also suboptimal from a network utilization perspective. The second method basically involves a constantly changing routing topology in that next-hop routers are chosen from a dynamic set in the hope that the number of malicious nodes in this set is negligible. A routing protocol that allows for disjoint routing paths may also be useful.

7.3.4. Countering Sinkhole Attacks

In sinkhole attacks, the malicious node manages to attract a lot of traffic mainly by advertising the availability of high-quality links even though there are none [Karlof2003]. Hence, it constitutes a serious attack on availability.

The malicious node creates a sinkhole by attracting a large amount of, if not all, traffic from surrounding neighbors by advertising in and outwards links of superior quality. Hence, affected nodes eagerly route their traffic via the malicious node that, if coupled with other attacks such as selective forwarding, may lead to serious availability and security breaches. Such an attack can only be executed by an inside malicious node and is generally very difficult to detect. An ongoing attack has a profound impact on the network topology and essentially becomes a problem of flow control.

Sinkhole attacks can be countered by deploying a series of mutually non-exclusive security measures to:

- o use geographical insights for flow control;
- o isolate nodes that receive traffic above a certain threshold;
- o dynamically pick up the next hop from a set of candidates; and
- o allow only trusted data to be received and forwarded.

A canary node could periodically call home (using a cryptographic process) with the home system, noting if it fails to call in. This provides detection of a problem, but does not mitigate it, and it may have significant energy consequences for the LLN.

Some LLNs may provide for geolocation services, often derived from solving triangulation equations from radio delay calculation; such calculations could in theory be subverted by a sinkhole that transmitted at precisely the right power in a node-to-node fashion.

While geographic knowledge could help assure that traffic always goes in the physical direction desired, it would not assure that the traffic is taking the most efficient route, as the lowest cost real route might match the physical topology, such as when different parts of an LLN are connected by high-speed wired networks.

7.3.5. Countering Wormhole Attacks

In wormhole attacks, at least two malicious nodes claim to have a short path between themselves [Karlof2003]. This changes the availability of certain routing paths and hence constitutes a serious security breach.

Essentially, two malicious insider nodes use another, more powerful, transmitter to communicate with each other and thereby distort the would-be-agreed routing path. This distortion could involve shortcutting and hence paralyzing a large part of the network; it

could also involve tunneling the information to another region of the network where there are, e.g., more malicious nodes available to aid the intrusion or where messages are replayed, etc.

In conjunction with selective forwarding, wormhole attacks can create race conditions that impact topology maintenance and routing protocols as well as any security suits built on "time of check" and "time of use".

A pure wormhole attack is nearly impossible to detect. A wormhole that is used in order to subsequently mount another kind of attack would be defeated by defeating the other attack. A perfect wormhole, in which there is nothing adverse that occurs to the traffic, would be difficult to call an attack. The worst thing that a benign wormhole can do in such a situation is to cease to operate (become unstable), causing the network to have to recalculate routes.

A highly unstable wormhole is no different than a radio opaque (i.e., metal) door that opens and closes a lot. RPL includes hysteresis in its objective functions [RFC6719] in an attempt to deal with frequent changes to the ETX between nodes.

8. RPL Security Features

The assessments and analysis in Section 6 examined all areas of threats and attacks that could impact routing, and the countermeasures presented in Section 7 were reached without confining the consideration to means only available to routing. This section puts the results into perspective, dealing with those threats that are endemic to this field, that have been mitigated through RPL protocol design, and that require specific decisions to be made as part of provisioning a network.

The first part of this section, Sections 8.1 to 8.3, presents a description of RPL security features that address specific threats. The second part of this section, Section 8.4, discusses issues of the provisioning of security aspects that may impact routing but that also require considerations beyond the routing protocol, as well as potential approaches.

RPL employs multicast, so these alternative communications modes MUST be secured with the same routing security services specified in this section. Furthermore, irrespective of the modes of communication, nodes MUST provide adequate physical tamper resistance commensurate with the particular application-domain environment to ensure the confidentiality, integrity, and availability of stored routing information.

8.1. Confidentiality Features

With regard to confidentiality, protecting the routing/topology information from unauthorized disclosure is not directly essential to maintaining the routing function. Breaches of confidentiality may lead to other attacks or the focusing of an attacker's resources (see Section 6.2) but does not of itself directly undermine the operation of the routing function. However, to protect against and reduce consequences from other more direct attacks, routing information should be protected. Thus, to secure RPL:

- o Implement payload encryption using L3 mechanisms described in [RFC6550] or
- o Implement L2 confidentiality

Where confidentiality is incorporated into the routing exchanges, encryption algorithms and key lengths need to be specified in accordance with the level of protection dictated by the routing protocol and the associated application-domain transport network. For most networks, this means use of AES-128 in CCM mode, but this needs to be specified clearly in the applicability statement.

In terms of the lifetime of the keys, the opportunity to periodically change the encryption key increases the offered level of security for any given implementation. However, where strong cryptography is employed, physical, procedural, and logical data access protection considerations may have a more significant impact on cryptoperiod selection than algorithm and key size factors. Nevertheless, in general, shorter cryptoperiods, during which a single key is applied, will enhance security.

Given the mandatory protocol requirement to implement routing node authentication as part of routing integrity (see Section 8.2), key exchanges may be coordinated as part of the integrity verification process. This provides an opportunity to increase the frequency of key exchange and shorten the cryptoperiod as a complement to the key length and encryption algorithm required for a given application domain.

8.2. Integrity Features

The integrity of routing information provides the basis for ensuring that the function of the routing protocol is achieved and maintained. To protect integrity, RPL must run either using only the secure versions of the messages or over a L2 that uses channel binding between node identity and transmissions.

Some L2 security mechanisms use a single key for the entire network, and these networks cannot provide a significant amount of integrity protection, as any node that has that key may impersonate any other node. This mode of operation is likely acceptable when an entire deployment is under the control of a single administrative entity.

Other L2 security mechanisms form a unique session key for every pair of nodes that needs to communicate; this is often called a per-link key. Such networks can provide a strong degree of origin authentication and integrity on unicast messages.

However, some RPL messages are broadcast, and even when per-node L2 security mechanisms are used, the integrity and origin authentication of broadcast messages cannot be as trusted due to the proliferation of the key used to secure them.

RPL has two specific options that are broadcast in RPL Control Messages: the DIO and the DODAG Information Solicitation (DIS). The purpose of the DIS is to cause potential parents to reply with a DIO, so the integrity of the DIS is not of great concern. The DIS may also be unicast.

The DIO is a critical piece of routing and carries many critical parameters. RPL provides for asymmetric authentication at L3 of the RPL Control Message carrying the DIO, and this may be warranted in some deployments. A node could, if it felt that the DIO that it had received was suspicious, send a unicast DIS message to the node in question, and that node would reply with a unicast DIS. Those messages could be protected with the per-link key.

8.3. Availability Features

Availability of routing information is linked to system and network availability, which in the case of LLNs require a broader security view beyond the requirements of the routing entities. Where availability of the network is compromised, routing information availability will be accordingly affected. However, to specifically assist in protecting routing availability, nodes MAY:

- o restrict neighborhood cardinality;
- o use multiple paths;
- o use multiple destinations;
- o choose randomly if multiple paths are available;
- o set quotas to limit transmit or receive volume; and

- o use geographic information for flow control.

8.4. Key Management

The functioning of the routing security services requires keys and credentials. Therefore, even though it's not directly an RPL security requirement, an LLN MUST have a process for initial key and credential configuration, as well as secure storage within the associated devices. Anti-tampering SHOULD be a consideration in physical design. Beyond initial credential configuration, an LLN is also encouraged to have automatic procedures for the revocation and replacement of the maintained security credentials.

While RPL has secure modes, some modes are impractical without the use of public key cryptography, which is believed to be too expensive by many. RPL L3 security will often depend upon existing LLN L2 security mechanisms, which provide for node authentication but little in the way of node authorization.

9. Security Considerations

The analysis presented in this document provides security analysis and design guidelines with a scope limited to RPL. Security services are identified as requirements for securing RPL. The specific mechanisms to be used to deal with each threat is specified in link-Land deployment-specific applicability statements.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, June 2005, <<http://www.rfc-editor.org/info/rfc4107>>.
- [RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.
- [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, September 2012, <<http://www.rfc-editor.org/info/rfc6719>>.

- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [ZigBeeIP] ZigBee Alliance, "ZigBee IP Specification", Public Document 15-002r00, March 2013.

10.2. Informative References

- [AceCharterProposal] Li, Kepeng., Ed., "Draft Charter V0.9c - Authentication and Authorization for Constrained Environment Charter", Work in Progress, December 2013, <http://trac.tools.ietf.org/wg/core/trac/wiki/ACE_charter>.
- [HELLO] Park, S., "Routing Security in Sensor Network: HELLO Flood Attack and Defense", Work in Progress, draft-suhopark-hello-wsn-00, December 2005.
- [IEEE.802.11] IEEE, "IEEE Standard for Information Technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Std 802.11-2012, March 2012, <<http://standards.ieee.org/about/get/802/802.11.html>>.
- [IEEE.802.15.4] IEEE, "IEEE Standard for Local and metropolitan area networks - Specific requirements - Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANS)", IEEE Std 802.15.4-2011, September 2011, <<http://standards.ieee.org/getieee802/802.15.html>>.
- [ISO.7498-2.1989] International Organization for Standardization, "Information processing systems - Open Systems Interconnection -- Basic Reference Model - Part 2: Security Architecture", ISO Standard 7498-2, 1989.

- [Karlof2003] Karlof, C. and D. Wagner, "Secure Routing in Wireless Sensor Networks: Attacks and Countermeasures", Elsevier Ad Hoc Networks Journal, Special Issue on Sensor Network Applications and Protocols, 1(2):293-315, September 2003, <<http://nest.cs.berkeley.edu/papers/sensor-route-security.pdf>>.
- [MESH-LINK] Kelsey, R., "Mesh Link Establishment", Work in Progress, draft-kelsey-intarea-mesh-link-establishment-06, May 2014.
- [Myagmar2005] Myagmar, S., Lee, A.J., and W. Yurcik, "Threat Modeling as a Basis for Security Requirements", in Proceedings of the Symposium on Requirements Engineering for Information Security (SREIS'05), Paris, France pp. 94-102, August 2005.
- [Perlman1988] Perlman, R., "Network Layer Protocols with Byzantine Robustness", MIT LCS Tech Report, 429, August 1988.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [RFC3067] Arvidsson, J., Cormack, A., Demchenko, Y., and J. Meijer, "TERENA'S Incident Object Description and Exchange Format Requirements", RFC 3067, February 2001, <<http://www.rfc-editor.org/info/rfc3067>>.
- [RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, September 2003, <<http://www.rfc-editor.org/info/rfc3610>>.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, October 2006, <<http://www.rfc-editor.org/info/rfc4593>>.
- [RFC4732] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006, <<http://www.rfc-editor.org/info/rfc4732>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.

- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008, <<http://www.rfc-editor.org/info/rfc5191>>.
- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, March 2008, <<http://www.rfc-editor.org/info/rfc5216>>.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009, <<http://www.rfc-editor.org/info/rfc5548>>.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009, <<http://www.rfc-editor.org/info/rfc5673>>.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010, <<http://www.rfc-editor.org/info/rfc5751>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, April 2010, <<http://www.rfc-editor.org/info/rfc5826>>.
- [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, <<http://www.rfc-editor.org/info/rfc5867>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, March 2011, <<http://www.rfc-editor.org/info/rfc6192>>.
- [RFC6574] Tschofenig, H. and J. Arkko, "Report from the Smart Object Workshop", RFC 6574, April 2012, <<http://www.rfc-editor.org/info/rfc6574>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.

- [RFC7142] Shand, M. and L. Ginsberg, "Reclassification of RFC 1142 to Historic", RFC 7142, February 2014, <<http://www.rfc-editor.org/info/rfc7142>>.
- [RFC7397] Gilger, J. and H. Tschofenig, "Report from the Smart Object Security Workshop", RFC 7397, November 2014, <<http://www.rfc-editor.org/info/rfc7397>>.
- [SmartObjectSecurityWorkshop] Klausen, T., Ed., "Workshop on Smart Object Security", March 2012, <<http://www.lix.polytechnique.fr/hipercom/SmartObjectSecurity>>.
- [SolaceProposal] Bormann, C., Ed., "Notes from the SOLACE ad hoc at IETF 85", November 2012, <<http://www.ietf.org/mail-archive/web/solace/current/msg00015.html>>.
- [Sybil2002] Douceur, J., "The Sybil Attack", First International Workshop on Peer-to-Peer Systems, March 2002.
- [Wan2004] Wan, T., Kranakis, E., and PC. van Oorschot, "S-RIP: A Secure Distance Vector Routing Protocol", in Proceedings of the 2nd International Conference on Applied Cryptography and Network Security, pp. 103-119, June 2004.
- [Yourdon1979] Yourdon, E. and L. Constantine, "Structured Design: Fundamentals of a Discipline of Computer Program and Systems Design", Yourdon Press, New York, Chapter 10, pp. 187-222, 1979.

Acknowledgments

The authors would like to acknowledge the review and comments from Rene Struik and JP Vasseur. The authors would also like to acknowledge the guidance and input provided by the ROLL Chairs, David Culler and JP Vasseur, and Area Director Adrian Farrel.

This document started out as a combined threat and solutions document. As a result of a series of security reviews performed by Steve Kent, the document was split up by ROLL Co-Chair Michael Richardson and Security Area Director Sean Turner as it went through the IETF publication process. The solutions to the threats are application and L2 specific and have, therefore, been moved to the relevant applicability statements.

Ines Robles and Robert Cragie kept track of the many issues that were raised during the development of this document.

Authors' Addresses

Tzeta Tsao
Eaton's Cooper Power Systems Business
910 Clopper Rd., Suite 201S
Gaithersburg, Maryland 20878
United States
EMail: tzetatsao@eaton.com

Roger K. Alexander
Eaton's Cooper Power Systems Business
910 Clopper Rd., Suite 201S
Gaithersburg, Maryland 20878
United States
EMail: rogeralexander@eaton.com

Mischa Dohler
CTTC
Parc Mediterrani de la Tecnologia, Av. Canal Olímpic S/N
Castelldefels, Barcelona 08860
Spain
EMail: mischa.dohler@kcl.ac.uk

Vanesa Daza
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 308
Barcelona 08003
Spain
EMail: vanesa.daza@upf.edu

Angel Lozano
Universitat Pompeu Fabra
P/ Circumval.lacio 8, Oficina 309
Barcelona 08003
Spain
EMail: angel.lozano@upf.edu

Michael Richardson (editor)
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z5V7
Canada
EMail: mcr+ietf@sandelman.ca