Smart Object Security Workshop
Internet-Draft
Intended status: Informational
Expires: August 10, 2012

M. Richardson
SSW

February 7, 2012

Challenges in Smart Object Security: too many layers, not enough ram
draft-richardson-smartobject-security-00

Abstract

This is a position paper for the pre-IETF83 Workshop on Smart Object Security. The author contends that layer-2 security solutions are not only in-adequate, but may in fact be harmful when deployed into smart object systems. While layer-2 security services may be valuable, they must be channel bound up to the layer-7 application layer.

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 August 10, 2012.

Table of Contents

1. Introduction

The ROLL RPL specification provides an optional layer-3 security mechanism. The WG did not focus very much on making this security system useable, as most WG participants assumed that layer-2 would provide all the security that most deployments would want.

While the ZigBee 2007 is provides a stack up to the application, and clearly articulates the role of the application in the security system, if ROLL RPL applicability statements specify Zigbee at all (XXX), from Zigbee's point of view the "application" is IPv6. The security provided by Zigbee 2007 does not get translated up to the IPv6 application, and certainly is not leveraged for end-to-end security.

Other specifications, including 6LowPAN (mesh-over) and ISA100 use a network key essentially identical to the 802.11 WEP. While many of these specifications propose to upgrade their mechanisms to include WPA-like usage of EAP, this does not solve the fundamental security problem of *authorization*. Except for auditing purposes, the network does not care who the nodes are, but rather, are they authorized to perform a particular function. In the context of RPL, one of these key functions is routing of packets.

A good example of the lack of security feature is that it is impossible in RPL to create a network where some nodes are authorized to route packets, while other nodes are not. While the specification supports this when doing layer-3 security, it only supports it for asymmetric security methods, widely regarded as too expensive for small devices. If the security is provided by a layer-2, then even if asymmetric methods are used in that layer, they are not available to the RPL (or higher) layers.

2. What we need

What we need is a security service implemented in layer-2 or layer-3, which not only provides for the privacy and integrity that is typically sought, but also can be leveraged by upper layers (including the layer-3 routing layer), to make authorization decisions.

Layer 2 alliances have created detailed and complex security specifications for wireless connectivity of smart objects. The requirements seems to have driven by existing early adopters of building and industrial automation. For many of the participants, security has become magic pixie dust provided by the vendor of the layer 2 MAC/radio.

I had believed that layer 3 security was more appropriate and easier to deploy/update. While requiring possibly more software code space, it might have a lower transitor count as flash is sometimes cheaper than complex logic in a MAC. But, I wondered who would need such flexibility among current industrial and smartgrid users of ROLL? Maybe it just my desire to do ubiquitous l3 networking with strangers on the bus, and I should shut up and believe that l2 security is enough.

Then the ROLL WG came to applicability statements, and it has obvious to me that people installing industrial equipement have much more complex requirements than I could even imagine on the bus. On the bus, I most trust everyone exactly the same: if they get my cryptographically signed packets to my intended destination, then I'm happy. I have really only one level of authorization: I either let you route my packets, or I do not. If I do not trust you to route, I might still trust you to have a cached copy of some data I want, and I have a way to authenticate the data itself.

The home automation users of ROLL was where I figured the most complexity would occur, and this relates mostly to how guests and children in the home will interact with the home system(s). While the lighting and appliance control network in the home looks very much like an commercial building system, how the occupants of the home interact with this system is not well defined as yet. It is likely that initially all interaction will be via hard controls ("light switches"), or via a gateway system that not only connected the 802.11 and 802.15.4 networks, but also provided an authentication and authorization system between the two networks. The ROLL provided security need concern itself only with whether or not a device was part of the home network or not, something that layer-2 security can do.

However, smart phones and personal area networks will begin to get 802.15.4 interfaces, and in some cases home automation is escrewing 802.15.4, claiming that 802.11 is now so cheap (power and bill-of-material) that it makes no sense to assume/require a gateway device. This is where, I thought, the multi-level authorization security would be required, and this would be subject to much innovation, with a number of home automation systems proving inadequate and being upgraded or replaced over time.

I thought that only when we have house guests (consider a teenage child to be an extended stay house guest) that we would run into troubles: we definitely want to authorize our guests to do many things in our homes, but there are many things we do not want them to do.

What I did I not expect was that industrial users would in fact require a multi-level authorization system, and have rekeying and trust issues more complex than that of the home. It was once explained that the militaries aren't into authorization: they care about authentication and auditing. A soldier must always have the ability to exceed their authority, because sometimes it's the right thing to do, and if they did the wrong thing, they have the court martial to solve this problem.

The court martial is not a solution for the industrial ROLL user. The various modes of operation described in [I-D.phinney-roll-rpl-industrial-applicability] require several levels of trust. While layer-2 security has a place, it seems that the installers of the devices (who would have to configure the layer-2 security) are not to be trusted in the long term, and therefore some way to change layer-2 keying material needs to be standardized.

If during any unplanned (i.e. emergency) situation new equipment will be brought into the plant to aid in recovery, that this equipment either will need to be configured (by regular plant personal) with the right security, or it will be necessary to either turn off or revert security to some other more minimal configuration, such that the equipment can be used.

It appears that not only is LAYER 2 security is not only inadequate, but may be actually too difficult to configure, simply because devices can not configure once installed. I am concerned that without a better answer, every building and plant will -- like most BlueTooth heads -- have PIN 0000!

We need something more sophisticated: sophisticated enough to be simple.

3. Security Considerations

This document does not propose any changes to existing or new systems, but rather details limitations of a current security model

4. References

4.1. Normative References

[I-D.phinney-roll-rpl-industrial-applicability]
Phinney, T., Thubert, P., and R. Assimiti, "RPL applicability in industrial networks",

draft-phinney-roll-rpl-industrial-applicability-00 (work in progress), October 2011.

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

4.2. Informative References

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

Appendix A. Additional Stuff

This becomes an Appendix.

Author's Address

Michael C. Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
CA

Email: mcr@sandelman.ca
URI: http://www.sandelman.ca/mcr/