

Networking Working Group
Internet-Draft
Intended status: Informational
Expires: May 12, 2008

K. Pister
R. Enns
Dust Networks
JP. Vasseur
P. Thubert
Cisco Systems, Inc
November 9, 2007

Industrial Routing Requirements in Low Power and Lossy Networks
draft-pister-rl2n-indus-routing-reqs-00

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 12, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Wireless, low power field devices enable industrial users to significantly increase the amount of information collected and the number of control points that can be remotely managed. The deployment of these wireless devices will significantly improve the productivity and safety of the plants while increasing the efficiency

Internet-Draft [draft-pister-rl2n-indus-routing-reqs-00](#) November 2007

of the plant workers. For wireless devices to have a significant advantage over wired devices in an industrial environment the wireless network needs to have three qualities: low power, high reliability, and easy installation and maintenance. The aim of this document is to analyze the requirements for the routing protocol used for low power and lossy networks (L2N) in industrial environments.

Requirements Language

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

Table of Contents

1.	Terminology	3
2.	Introduction	4
2.1.	Applications and Traffic Patterns	5
3.	Quality of Service (QoS) Routing requirements	7
3.1.	Configurable Application Requirement	9
4.	Network Topology	9
5.	Device-Aware Routing Requirements	10
6.	Broadcast/Multicast	11
7.	Route Establishment Time	11
8.	Mobility	11
9.	Manageability and Ease Of Use	12
10.	Security	12
11.	Informative Reference	13
12.	IANA Considerations	13
13.	Acknowledgements	13
14.	References	13
14.1.	Normative References	13
14.2.	Informative References	13
	Authors' Addresses	14
	Intellectual Property and Copyright Statements	15

Internet-Draft [draft-pister-rl2n-indus-routing-reqs-00](#) November 2007

1. Terminology

Access Point: The access point is an infrastructure device that connects the low power and lossy network system to a plant's backbone network.

Actuator: a field device that moves or controls plant equipment.

Channel Hopping: An algorithm by which field devices synchronously change channels during operation.

Channel: RF frequency band used to transmit a modulated signal carrying packets.

Closed Loop Control: A process whereby a host application controls an actuator based on information sensed by field devices.

Downstream: Data direction traveling from the host application to the field device.

Field Device: physical devices placed in the plant's operating environment (both RF and environmental). Field devices include sensors and actuators as well as network routing devices and access points in the plant. **Superframe:** A collection of timeslots repeating at a constant rate.

HART: "Highway Addressable Remote Transducer", a group of specifications for industrial process and control devices administered by the HART Foundation (see [[HART](#)]). The latest version for the specifications is HART7 which includes the additions for WirelessHART.

Host Application: The host application is a process running in the plant that communicates with field devices to perform tasks on that may include control, monitoring and data gathering.

ISA: "International Society of Automation". ISA is an ANSI accredited standards making society. ISA100 is an ISA working group whose charter includes defining a family of standards for industrial automation. ISA100.11a is a work group within ISA100 that is working on a standard for non-critical process and control applications.

L2N: Low Power and Lossy Network Slotted-Link: a data structure that is associated with a superframe that contains a connection between field devices that comprises a timeslot assignment, and channel and usage information.

Open Loop Control: A process whereby a plant technician controls an

actuator over the network where the decision is influenced by information sensed by field devices.

Timeslot: A fixed time interval that may be used for the transmission or reception of a packet between two field devices. A timeslot used for communications is associated with a slotted-link. Upstream: Data direction traveling from the field device to the host application.

RF: Radio Frequency Sensor: A field device that measures data and/or detects an event.

RL2N: Routing in Low power and Lossy Networks

[2.](#) Introduction

Wireless, low power field devices enable industrial users to significantly increase the amount of information collected and the number of control points that can be remotely managed. The deployment of these wireless devices will significantly improve the productivity and safety of the plants while increasing the efficiency of the plant workers.

Wireless field devices enable expansion of networked points by appreciably reducing cost of installing a device. The cost reductions come from eliminating cabling costs and simplified planning. Cabling for a field device can run from \$100s/ft to \$1,000s/ft. depending on the safety regulations of the plant. Cabling also carries an overhead cost associated with planning the

installation, where the cable has to run, and the various organizations that have to coordinate its deployment. Doing away with the network and power cables reduces the planning and administrative overhead of installing a device.

For wireless devices to have a significant advantage over wired devices in an industrial environment the wireless network needs to have three qualities: low power, high reliability, and easy installation and maintenance. The routing protocol used for low power and lossy networks (L2N) is important to fulfilling these goals.

Industrial automation is segmented into two distinct application spaces, known as "process" or "process control" and "discrete manufacturing" or "factory automation". In industrial process control, the product is typically a fluid (oil, gas, chemicals ...). In factory automation or discrete manufacturing, the products are individual elements (screws, cars, dolls). While there is some overlap between products and systems between these two segments, they

are surprisingly separate communities. The specifications targeting industrial process control tend to have more tolerance for network latency than what is needed for factory automation.

Both application spaces desire battery operated networks of hundreds of sensors and actuators communicating with wired access points. In an oil refinery, the total number of devices is likely to exceed one million, but the devices will be clustered into smaller networks reporting to existing wired infrastructure.

Existing wired sensor networks in this space typically use communication protocols with low data rates - from 1,200 baud (wired HART) into the one to two hundred kbps range for most of the others. The existing protocols are often master/slave with command/response.

Note that the total low power and lossy network system capacity for devices using the IEEE802.15.4-2006 2.4 GHz radio is at most 1.6 Mbps when spatial reuse of channels is not utilized.

[2.1.](#) Applications and Traffic Patterns

The industrial market classifies process applications into three

broad categories and six classes.

- o Safety

- * Class 0: Emergency action - Always a critical function Control
- * Class 1: Closed loop regulatory control - Often a critical function
- * Class 2: Closed loop supervisory control - Usually non-critical function
- * Class 3: Open loop control - Operator takes action and controls the actuator (human in the loop)

- o Monitoring

- * Class 4: Alerting - Short-term operational effect (for example event-based maintenance)
- * Class 5: Logging and downloading / uploading - No immediate operational consequence (e.g., history collection, sequence-of-events, preventive maintenance)

Critical functions are effect the basic safety or integrity of the plant. Timely deliveries of messages are more important as class

numbers decrease.

Industrial customers are interested in deploying wireless networks for the monitoring classes 4 and 5 and in the non-critical portions of classes 3 through 1.

Classes 4 and 5 also include equipment monitoring which is strictly speaking separate from process monitoring. An example of equipment monitoring is the recording of motor vibrations to detect bearing wear.

Most low power and lossy network systems in the near future will be for low frequency data collection. Packets containing samples will be generated continuously, and 90% of the market is covered by packet rates of between 1/s and 1/hour, with the average under 1/min. In

industrial process these sensors include temperature, pressure, fluid flow, tank level, and corrosion. There are some sensors which are bursty, such as vibration monitors which may generate and transmit tens of kilo-bytes (hundreds to thousands of packets) of time-series data at reporting rates of minutes to days.

Almost all of these sensors will have built-in microprocessors which may detect alarm conditions. Time crucial alarm packets are expected to have lower latency than sensor data, often requiring substantially more bandwidth.

Some devices will transmit a log file every day, again with typically tens of Kbytes of data. For these applications there is very little "downstream" traffic coming from the access point and traveling to particular sensors. During diagnostics, however, a technician may be investigating a fault from a control room and expect to have "low" latency (human tolerable) in a command/response mode.

Low-rate control, often with a "human in the loop" or "open loop" is implemented today via communication to a centralized controller, i.e. sensor data makes its way through the access point to the centralized controller where it is processed, the operator sees the information and takes action, and control information is sent out to the actuator node in the network.

In the future, it is envisioned that some open loop processes will be automated (closed loop) and packets will flow over local loops and not involve the access point. These closed loop controls for non-critical applications will be implemented on L2Ns. Non-critical closed loop applications have a latency requirement that can be as low as 100 ms but many control loops are tolerant of latencies above 1 s.

In critical control, 10's of milliseconds of latencies are typical. In many of these systems, if a packet does not arrive within the specified interval, the system will enter an emergency shutdown state, often with substantial financial repercussions. For a 1 second control loop in a system with a mean-time between shutdowns target of 30 years, the latency requirement implies nine 9s of reliability.

Thus, the routing protocol for L2Ns MUST support multi-topology routing (e.g especially critical for critical control applications). The routing protocol MUST provide the ability to color slotted-links (where the color corresponds to a user defined slotted-link attribute) and can be used to include/exclude slotted-links from a logical topology.

For all but the most latency-tolerant applications, route discovery is likely to be too slow a process to initiate when a route failure is detected.

The routing protocol MUST support multiple paths (a tree-based solution is not sufficient).

3. Quality of Service (QoS) Routing requirements

The industrial applications fall into four large service categories:

1. Published data. Data that is generated per periodically and has a well understood data bandwidth requirement. The end-to-end latency of this data is not as important as regularity with which it is presented to the host application.
2. Event data. This category includes alarms and aperiodic data reports with bursty data bandwidth requirements
3. Client/Server. Many industrial applications are based on a client/server model and implement a command response protocol. The data bandwidth required is often bursty. The round trip latency for some operations can be 200 ms.
4. Bulk transfer. Bulk transfers involve the transmission of blocks of data in multiple packets where temporary resources are assigned to meet a transaction time constraint. Bulk transfers assign resources for a limited period of time to meet the QoS requirements.

For industrial applications QoS parameters include:

- o Data bandwidth - periodic, burst statistics

- o Latency - the time taken for the data to transit the network from the source to the destination. This may be expressed in terms of a deadline for delivery
- o Transmission phase - process applications can be synchronized to wall clock time and require coordinated transmissions. A common coordination frequency is 4 Hz (250 ms).
- o Reliability - the end-to-end data delivery statistic. All applications have latency and reliability requirements. In industrial applications, these vary over many orders of magnitude. Some non-critical monitoring applications may tolerate latencies of days and reliability of less than 90%. Most monitoring latencies will be in seconds to minutes, and industrial standard such as HART7 has set user reliability expectations at 99.9%. Regulatory requirements are a driver for some industrial applications. Regulatory monitoring requires high data integrity because lost data is assumed to be out of compliance and subject to fines. This can drive reliability requirements to higher than 99.9%.
- o QoS contract type - revocation priority. L2Ns have limited network resources that can vary with time. This means the system can become fully subscribed or even over subscribed. System policies determine how resources are allocated when resources are over subscribed. The choices are blocking and graceful degradation.
- o Transmission Priority - within field devices there are limited resources need to be allocated across multiple services. For transmissions, a device has to select which packet in its queue will be sent at the next transmission opportunity. Packet priority is used as one criterion for selecting the next packet. For reception a device has to decide how to store a received packet. The field devices are memory constrained and receive buffers may become full. Packet priority is used to select which packets are stored or discarded.

In industrial wireless L2Ns a time slotting technology is used. A time slotted media access protocol synchronizes channel hopping which is one of the means that is used to make the wireless network reliable. Timeslots also are employed to reduce the power by minimizing the active duty cycle for field devices. Communications between devices are assigned a combination of timeslot and channel assignment called a slotted-link.

The routing protocol MUST also support different metric types for each slotted-link used to compute the path according to some objective function (e.g. minimize latency, maximize reliability, ...).

Industrial application data flows between field devices are not necessarily symmetric. The routing protocol MUST be able to set up routes that are directional.

[3.1.](#) Configurable Application Requirement

Time-varying user requirements for latency and bandwidth will require changes in the provisioning of the underlying L2 protocols. The wireless worker may initiate a bulk transfer to configure or diagnose a field device. A level sensor device may need to perform a calibration and send a bulk file to a host. The routing protocol MUST route on paths which are changed to appropriately provision the application requirements. The routing protocol MUST support the ability to recompute paths based on slotted-link characteristics that may change dynamically.

[4.](#) Network Topology

Network topology is very tough to generalize, but networks of 10 to 200 field devices and maximum number of hops from two to twenty covers the majority of existing applications. It is assumed that the field devices themselves will provide routing capability for the network, and in most cases additional repeaters/routers will not be required.

Timeslot size is about 10 ms and timeslot synchronization requirements are on the order of +/-1 ms for non-critical process and control data (some L2 protocols provide/require tighter synchronization). Wall clock time *accuracy* requirements vary substantially, but are generally about 100ms. Some applications that time stamp data require 1 ms accuracy to determine the sequence of events reported. (Note that data time stamping does not translate to a latency requirement.)

In low power and lossy network systems using the IEEE802.15.4-2006 2.4 GHz radio the total raw throughput per radio is 250 kbps. 10 ms timeslots reduces this to 101.6 Kbps for maximum sized packets. This constrains the typical throughput of a single radio access point to less 100 kbps. Therefore an access point with one IEEE 802.15.4 radio has a maximum aggregate throughput 100 packets per second and

no more than about 100 Kbps.

A graph that connects a field device to a host application may have more than one access point. The routing protocol MUST support multiple access points and load distribution when aggregate network throughputs need to exceed 100 kbps. The routing protocol MUST support multiple access points when access point redundancy is required.

5. Device-Aware Routing Requirements

Wireless L2N nodes in industrial environments are powered by a variety of sources. Battery operated devices with lifetime requirements of at least 5 years are the most common. Battery operated devices have a cap on their total energy, and typically can report some estimate of remaining energy, and typically do not have constraints on the short term average power consumption. Energy scavenging devices are more complex. These systems contain both a power scavenging device (such as solar, vibration, or temperature difference) and an energy storage device, such as a rechargeable battery or a capacitor. Therefore these systems have limits on both the long term average power consumption (which cannot exceed the average scavenged power over the same interval) as well as the short-term limits imposed by the energy storage requirements. For solar-powered systems, the energy storage system is generally designed to provide days of power in the absence of sunlight. Many industrial sensors run off of a 4-20mA current loop, and can scavenge on the order of mW from that source. Vibration monitoring systems are a natural choice for vibration scavenging, which typically only provides tens or hundreds of microwatts. Due to industrial temperature ranges and desired lifetimes, the choices of energy storage devices can be limited, and the resulting stored energy is often comparable to the energy cost of sending or receiving a packet rather than the energy of operating the node for several days. And of course some nodes will be line-powered.

Example 1: solar panel, lead-acid battery sized for two weeks of rain.

Example 2: vibration scavenger, 1mF tantalum capacitor

Field devices have limited resources. Low power, low cost devices have limited memory for storing route information. Typical field devices will have a finite number of routes they can support for their embedded sensor/actuator application and for forwarding other devices packets in a mesh network slotted-link.

Users may have strong preferences on lifetime that is different for the same device in different locations. A sensor monitoring a non-

critical parameter in an easily-accessed location may have a lifetime requirement that is shorter and tolerate more statistical variation than a mission-critical sensor in a hard to reach place that requires shutdown of the plant to replace.

The routing algorithm MUST support node constrained routing (e.g. taking into account the existing energy state as a node constraint). Node constraints include power and memory as well as constraints placed on the device by the user such as battery life.

[6.](#) Broadcast/Multicast

Existing industrial host applications do not use broadcast or multicast addressing to communicate to field devices. Unicast address support is sufficient. However wireless field devices with communication controllers and protocol stacks will require control and configuration such as firmware downloading that may benefit from broadcast and multicast addressing.

The routing protocol SHOULD support broadcast and multicast addressing.

[7.](#) Route Establishment Time

Network connectivity in real deployments is always time-varying, with time constants from seconds to months. Optimization is perhaps not the right word to use, in that network optimization will need to run continuously, and single-slotted-link failures that cause loss of connectivity are not likely to be tolerated. Once the network is formed, it should never need to "optimize" to a new configuration in

response to a lost slotted-link. The routing algorithm SHOULD not have to re-optimize in response to the loss of a slotted-link. The routing algorithms SHOULD always be in the process of plesio-optimizing the system for the changing RF environment. The routing algorithm MUST re-optimize the path when field devices change due to insertion, removal or failure.

8. Mobility

Various economic factors have contributed to a reduction of trained workers in the plant. The industry as a whole appears to be trying to solve this problem with what is called the "wireless worker". Carrying a PDA or something similar, this worker will be able to accomplish more work in less time than the older, better-trained workers that he or she replaces. Whether the premise is valid, the

Pister, et al.

Expires May 12, 2008

[Page 11]

Internet-Draft [draft-pister-rl2n-indus-routing-reqs-00](#) November 2007

use case is commonly presented: the worker will be wirelessly connected to the plant IT system to download documentation, instructions, etc., and will need to be able to connect "directly" to the sensors and control points in or near the equipment on which he or she is working. It is possible that this "direct" connection could come via the normal L2Ns data collection network. This connection is likely to require higher bandwidth and lower latency than the normal data collection operation.

The routing protocol SHOULD support the wireless worker with fast network connection times of a few of seconds, low latency command and response latencies to host behind the access points and to applications and to field devices. The routing protocol SHOULD also support configuring graphs for bulk transfers. The routing protocol MUST support walking speeds for maintaining network connectivity as the handheld device changes position in the wireless network.

Some field devices will be mobile. These devices may be located on moving parts such as rotating components or they may be located on vehicles such as cranes or fork lifts. The routing protocol SHOULD support vehicular speeds of up to 35 kmph.

9. Manageability and Ease Of Use

The process and control industry is manpower constrained. The aging demographics of plant personnel are causing a looming manpower problem for industry across many markets. The goal for the industrial networks is to make the installation process not require any new skills for the plant personnel. The industrial customers do not even want to require the current level of networking knowledge needed for do-it-yourself home network installations.

The routing protocol for L2Ns must be easy to deploy and manage. In a further revision of this document, metrics to measure ease of deployment for the routing protocol will be detailed.

10. Security

Wireless sensor networks in industrial automation operate in systems that have substantial financial and human safety implications, security is of considerable concern. Levels of security violation which are tolerated as a "cost of doing business" in the banking industry are not acceptable when in some cases literally thousands of lives may be at risk.

Industrial wireless device manufactures are specifying security at

the MAC layer and the Transport layer. A shared "Network Key" is used to authenticate messages at the MAC layer. At the transport layer, commands are encrypted with unique randomly-generated end-to-end Session keys. HART7 and ISA100.11a are examples of security systems for industrial wireless networks.

Industrial plants may not maintain the same level of physical security for field devices that is associated with traditional network sites such as locked IT centers. In industrial plants it must be assumed that the field devices have marginal physical security and the security system needs to have limited trust in them. The routing protocol SHOULD place limited trust in the field devices deployed in the plant network.

The routing protocol SHOULD compartmentalize the trust placed in field devices so that a compromised field device does not destroy the security of the whole network. The routing MUST be configured and managed using secure messages and protocols that prevent outsider

attacks and limit insider attacks from field devices installed in insecure locations in the plant.

11. Informative Reference

[HART] "Highway Addressable Remote Transducer", a group of specifications for industrial process and control devices administered by the HART Foundation, www.hartcomm.org.

12. IANA Considerations

This document includes no request to IANA.

13. Acknowledgements

14. References

14.1. Normative References

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

14.2. Informative References

[I-D.culler-rl2n-routing-reqs]
Vasseur, J. and D. Cullerot, "Routing Requirements for Low

Pister, et al.

Expires May 12, 2008

[Page 13]

Internet-Draft [draft-pister-rl2n-indus-routing-reqs-00](#) November 2007

Power And Lossy Networks",
[draft-culler-rl2n-routing-reqs-01](#) (work in progress),
July 2007.

Authors' Addresses

K. Pister
Dust Networks
30695 Huntwood Ave.
Hayward, Denmark 94544

USA

Email: kpister@dustnetworks.com

Rick Enns
Dust Networks
30695 Huntwood Ave.
Hayward, 94544
USA

Email: renns@dustnetworks.com

JP Vasseur
Cisco Systems, Inc
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Email: jpv@cisco.com

Pascal Thubert
Cisco Systems, Inc
Village d'Entreprises Green Side - 400, Avenue de Roumanille
Sophia Antipolis, 06410

Email: pthubert@cisco.com

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).