

Networking Working Group	K. Pister	
Internet-Draft	Dust Networks	
Intended status: Informational	P. Thubert	
Expires: October 25, 2008	Cisco Systems, Inc	
	April 23, 2008	

[TOC](#)

## **Industrial Routing Requirements in Low Power and Lossy Networks draft-pister-roll-indus-routing-reqs-01**

### **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 October 25, 2008.

### **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 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 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

---

## Table of Contents

- [1.](#) Terminology
- [2.](#) Introduction
  - [2.1.](#) Applications and Traffic Patterns
  - [2.2.](#) Network Topology of Industrial Applications
    - [2.2.1.](#) The Physical Topology
- [3.](#) Service Requirements
  - [3.1.](#) Configurable Application Requirement
  - [3.2.](#) Different Routes for Different Flows
- [4.](#) Reliability Requirements
- [5.](#) Device-Aware Routing Requirements
- [6.](#) Broadcast/Multicast
- [7.](#) Route Establishment Time
- [8.](#) Mobility
- [9.](#) Manageability
- [10.](#) Security
- [11.](#) IANA Considerations
- [12.](#) Acknowledgements
- [13.](#) References
  - [13.1.](#) Normative References
  - [13.2.](#) Informative References
  - [13.3.](#) External Informative References
- [§](#) Authors' Addresses
- [§](#) Intellectual Property and Copyright Statements

---

## 1. Terminology

[TOC](#)

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

Closed Loop Control: A process whereby a device controller controls an actuator based on information sensed by one or more field devices.

Downstream: Data direction traveling from the plant 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 L2N access points in the plant.

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. 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 Access Point: The L2N access point is an infrastructure device that connects the low power and lossy network system to a plant's backbone network.

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.

Plant Application: The plant 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.

Upstream: Data direction traveling from the field device to the plant application.

RL2N: Routing in Low power and Lossy Networks.

---

## 2. Introduction

[TOC](#)

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, determining where the cable has to run, and interfacing with the various organizations required 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 of 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 L2N 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 that report to an existing plant network infrastructure.

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

---

## 2.1. Applications and Traffic Patterns

[TOC](#)

The industrial market classifies process applications into three broad categories and six classes.

### \*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)

### \*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 effect the basic safety or integrity of the plant. Timely deliveries of messages becomes more important as the class number decreases.

Industrial users 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 asset monitoring and tracking which include equipment monitoring and are essentially separate from process monitoring. An example of equipment monitoring is the recording of motor vibrations to detect bearing wear.

In the near future, most low power and lossy network systems 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. Some sensors are bursty, such as vibration monitors that 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 that may detect alarm conditions. Time-critical alarm packets are expected to have lower latency than sensor data.

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 L2N 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" (also referred to as "open loop"), is implemented today via communication to a centralized controller. The sensor data makes its way through the L2N access point to the centralized controller where it is processed, the operator sees the information and takes action, and the control information is then 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 L2N 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, tens of milliseconds of latency is typical. In many of these systems, if a packet does not arrive within the specified interval, the system enters an emergency shutdown state, often with substantial financial repercussions. For a one-second control loop in a

system with a mean-time between shutdowns target of 30 years, the latency requirement implies nine 9s of reliability.

---

## **2.2. Network Topology of Industrial Applications**

[TOC](#)

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

For most industrial applications, a manager, gateway or backbone router acts as a sink for the wireless sensor network. The vast majority of the traffic is real time publish/subscribe sensor data from the field devices over a L2N towards one or more sinks. Increasingly over time, these sinks will be a part of a backbone but today they are often fragmented and isolated.

The wireless sensor network is a Low Power and Lossy Network of field devices for which two logical roles are defined, the field routers and the non routing devices. It is acceptable and even probable that the repartition of the roles across the field devices change over time to balance the cost of the forwarding operation amongst the nodes.

The backbone is a high speed network that interconnects multiple WSNs through backbone routers. Infrastructure devices can be connected to the backbone. A gateway / manager that interconnects the backbone to the plant network of the corporate network can be viewed as collapsing the backbone and the infrastructure devices into a single device that operates all the required logical roles. The backbone is likely to always become an important function of the industrial network. The Internet at large is not considered as a viable option to perform the backbone function.

A plant or corporate network is also present on the factory site. This is the typical IT network for the factory operations beyond process control. That network is out of scope for this document.

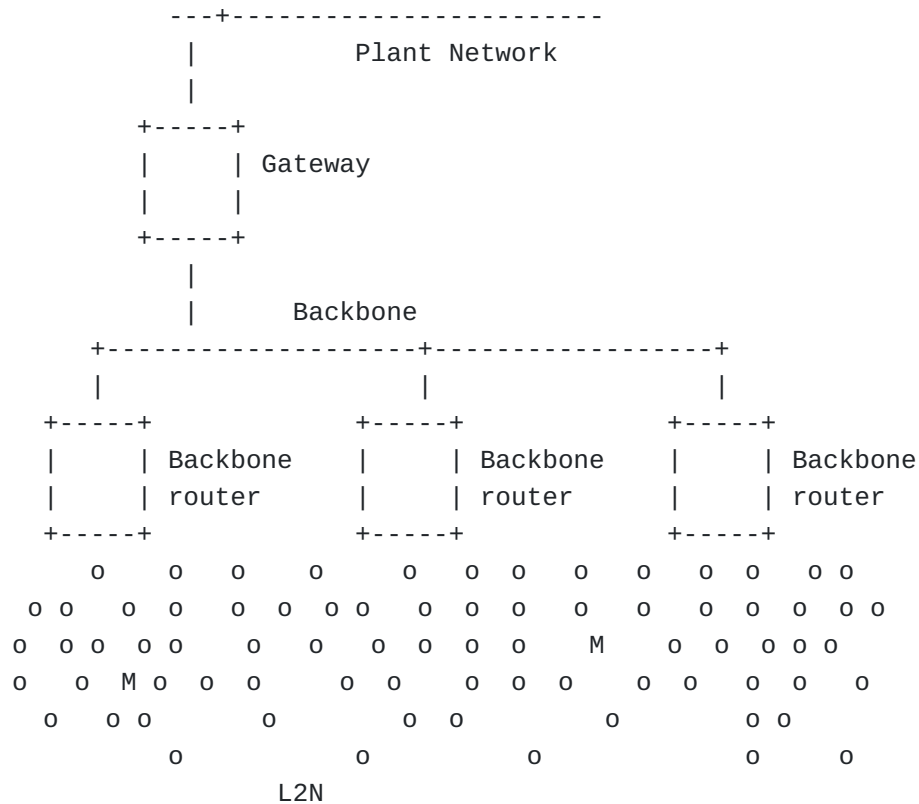
---

### **2.2.1. The Physical Topology**

[TOC](#)

There is no specific physical topology for an industrial process control network. One extreme example is a multi-square-kilometer refinery where isolated tanks, some of them with power but most with no backbone connectivity, compose a farm that spans over of the surface of the plant. A few hundred field devices are deployed to ensure the global coverage using a wireless self-forming self-healing mesh network that might be 5 to 10 hops across. Local feedback loops and mobile

workers tend to be only one or two hops. The backbone is in the refinery proper, many hops away. Even there, powered infrastructure is also typically several hops away. So hopping to/from the powered infrastructure will in general be more costly than the direct route. In the opposite extreme case, the backbone network spans all the nodes and most nodes are in direct sight of one or more backbone router. Most communication between field devices and infrastructure devices as well as field device to field device occurs across the backbone. From afar, this model resembles the WIFI ESS (Extended Service Set). But from a layer 3 perspective, the issues are the default (backbone) router selection and the routing inside the backbone whereas the radio hop towards the field device is in fact a simple local delivery.



Considering that though each field device to field device route computation has specific constraints in terms of latency and availability it can be expected that the shortest path possible will often be selected and that this path will be routed inside the LLN as opposed to via the backbone. It can also be noted that the lifetimes of the routes might range from minutes for a mobile workers to tens of years for a command and control closed loop. Finally, time-varying user requirements for latency and bandwidth will change the constraints on the routes, which might either trigger a constrained route recomputation, a reprovisioning of the underlying L2 protocols, or both in that order. For instance, a 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 plant.

For these reasons, the ROLL routing infrastructure MUST be able to compute and update constrained routes on demand (that is reactively), and it can be expected that this model will become more prevalent for field device to field device connectivity as well as for some field device to Infrastructure devices over time.

---

### 3. Service Requirements

[TOC](#)

The industrial applications fall into four large service categories [ISA100.11a]:

1. Periodic data (aka buffered). Data that is generated periodically and has a well understood data bandwidth requirement, both deterministic and predictable. Timely delivery of such data is often the core function of a wireless sensor network and permanent resources are assigned to ensure that the required bandwidth stays available. Buffered data usually exhibits a short time to live, and the newer reading obsoletes the previous. In some cases, alarms are low priority information that gets repeated over and over. The end-to-end latency of this data is not as important as the regularity with which the data is presented to the plant application.
2. Event data. This category includes alarms and aperiodic data reports with bursty data bandwidth requirements. In certain cases, alarms are critical and require a priority service from the network.
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 acceptable round-trip latency for some legacy systems was based on the time to send tens of bytes over a 1200 baud link. Hundreds of milliseconds is typical. This type of request is statistically multiplexed over the L2N and cost-based fair-share best-effort service is usually expected.
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. Transient resources are assigned for a limited period of time (related to file size and data rate) to meet the bulk transfers service requirements.

For industrial applications Service parameters include but might not be limited to:

\*Data bandwidth – the bandwidth might be allocated permanently or for a period of time to a specific flow that usually exhibits well defined properties of burstiness and throughput. Some bandwidth will also be statistically shared between flows in a best effort fashion.

\*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. Most monitoring latencies will be in seconds to minutes.

\*Transmission phase - process applications can be synchronized to wall clock time and require coordinated transmissions. A common coordination frequency is 4 Hz (250 ms).

\*Service 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.

\*Transmission priority - the means by which limited resources within field devices are 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.

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

Industrial application data flows between field devices are not necessarily symmetric. In particular, asymmetrical cost and unidirectional routes are common for published data and alerts, which represent the most part of the sensor traffic. The routing protocol MUST be able to set up unidirectional or asymmetrical cost routes that are composed of one or more non congruent paths.

---

### 3.1. Configurable Application Requirement

[TOC](#)

Time-varying user requirements for latency and bandwidth will require changes in the provisioning of the underlying L2 protocols. A

technician may initiate a query/response session or bulk transfer to diagnose or configure a field device. A level sensor device may need to perform a calibration and send a bulk file to a plant. The routing protocol MUST route on paths that are changed to appropriately provision the application requirements. The routing protocol MUST support the ability to recompute paths based on underlying link characteristics that may change dynamically.

---

### 3.2. Different Routes for Different Flows

[TOC](#)

Because different services categories have different service requirements, it is often desirable to have different routes for different data flows between the same two endpoints. For example, alarm or periodic data from A to Z may require path diversity with specific latency and reliability. A file transfer between A and Z may not need path diversity. The routing algorithm MUST be able to generate different routes for different flows.

---

## 4. Reliability Requirements

[TOC](#)

Another critical aspect for the routing is the capability to ensure maximum disruption time and route maintainance. The maximum disruption time is the time it takes at most for a specific path to be restored when broken. Route maintainance ensures that a path is monitored to be restored when broken within the maximum disruption time. Maintenance should also ensure that a path continues to provide the service for which it was established for instance in terms of bandwidth, jitter and latency.

In industrial applications, reliability is usually defined with respect to end-to-end delivery of packets within a bounded latency. Reliability requirements vary over many orders of magnitude. Some non-critical monitoring applications may tolerate a reliability of less than 90% with hours of latency. Most industrial standards, such as HART7, have 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%.

Hop-by-hop path diversity is used to improve latency-bounded reliability. Additionally, bicasting or pluricasting may be used over multiple non congruent / non overlapping paths to ensure that at least one instance of a critical packet is actually delivered.

Because data from field devices are aggregated and funneled at the L2N access point before they are routed to plant applications, L2N access

point redundancy is an important factor in overall reliability. A route that connects a field device to a plant application may have multiple paths that go through more than one L2N access point. The routing protocol MUST support multiple L2N access points and load distribution among L2N access points. The routing protocol MUST support multiple L2N access points when L2N access point redundancy is required. Because L2Ns are lossy in nature, multiple paths in a L2N route MUST be supported. The reliability of each path in a route can change over time. Hence, it is important to measure the reliability on a per-path basis and select a path (or paths) according to the reliability requirements.

---

## 5. Device-Aware Routing Requirements

[TOC](#)

Wireless L2N nodes in industrial environments are powered by a variety of sources. Battery operated devices with lifetime requirements of at least five years are the most common. Battery operated devices have a cap on their total energy, and typically can report an 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. These systems, therefore, have limits on both 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 4-20 mA current loop, and can scavenge on the order of milliwatts 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 strongly prefer that the same device have different lifetime requirements 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 a plant shutdown in order 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**

[TOC](#)

Existing industrial plant 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 or multicast addressing.

The routing protocol **SHOULD** support broadcast or multicast addressing.

---

## **7. Route Establishment Time**

[TOC](#)

During network formation, installers with no networking skill must be able to determine if their devices are “in the network” with sufficient connectivity to perform their function. Installers will have sufficient skill to provision the devices with a sample rate or activity profile. The routing algorithm **MUST** find the appropriate route(s) and report success or failure within several minutes, and **SHOULD** report success or failure within tens of seconds.

Network connectivity in real deployments is always time varying, with time constants from seconds to months. So long as the underlying connectivity has not been compromised, this link churn should not substantially affect network operation. The routing algorithm **MUST** respond to normal link failure rates with routes that meet the Service requirements (especially latency) throughout the routing response. The routing algorithm **SHOULD** always be in the process of optimizing the system in response to changing link statistics. The routing algorithm **MUST** re-optimize the paths when field devices change due to insertion, removal or failure, and this re-optimization **MUST** not cause latencies greater than the specified constraints (typically seconds to minutes).

---

[TOC](#)

## 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 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, and low command and response latencies to the plant behind the L2N access points, to applications, and to field devices. The routing protocol SHOULD also support the bandwidth allocation for bulk transfers between the field device and the handheld device of the wireless worker. The routing protocol SHOULD 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

[TOC](#)

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 have the installation process not require any new skills for the plant personnel. The person would install the wireless sensor or wireless actuator the same way the wired sensor or wired actuator is installed, except the step to connect wire is eliminated.

The routing protocol for L2Ns is expected to be easy to deploy and manage. Because the number of field devices in a network is large, provisioning the devices manually would not make sense. Therefore, the routing protocol MUST support auto-provisioning of field devices. The protocol also MUST support the distribution of configuration from a centralized management controller if operator-initiated configuration change is allowed.

---

## 10. Security

[TOC](#)

Given that 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 that 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 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. IANA Considerations

[TOC](#)

This document includes no request to IANA.

---

## 12. Acknowledgements

[TOC](#)

Many thanks to Rick Enns and Chol Su Kang for their contributions.

---

## 13. References

[TOC](#)

### 13.1. Normative References

[TOC](#)

[RFC2119]	<a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
-----------	--

---

### 13.2. Informative References

[TOC](#)

[I-D.culler-rl2n-routing-reqs]	Vasseur, J. and D. Cullerot, " <a href="#">Routing Requirements for Low Power And Lossy Networks</a> ," draft-culler-rl2n-routing-reqs-01 (work in progress), July 2007 ( <a href="#">TXT</a> ).
--------------------------------	--

---

### 13.3. External Informative References

[TOC](#)

[HART]	<a href="http://www.hartcomm.org">www.hartcomm.org</a> , "Highway Addressable Remote Transducer", a group of specifications for industrial process and control devices administered by the HART Foundation."
[ISA100.11a]	ISA, "SP100.11 Working Group Draft Standard, Version 0.1," December 2007.

---

### Authors' Addresses

[TOC](#)

	Kris Pister
	Dust Networks
	30695 Huntwood Ave.
	Hayward, 94544
	USA
Email:	<a href="mailto:kpister@dustnetworks.com">kpister@dustnetworks.com</a>
	Pascal Thubert
	Cisco Systems, Inc
	Village d'Entreprises Green Side - 400, Avenue de Roumanille
	Sophia Antipolis, 06410
	FRANCE
Email:	<a href="mailto:pthubert@cisco.com">pthubert@cisco.com</a>

---

## Full Copyright Statement

[TOC](#)

Copyright © The IETF Trust (2008).

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](mailto:ietf-ipr@ietf.org).