

Independent Submission  
Internet-Draft  
Intended status: Informational  
Expires: August 15, 2021

H. Ayers  
P. Levis  
Stanford University  
February 11, 2021

**Design Considerations for Low Power Internet Protocols**  
**draft-ayers-low-power-interop-03**

Abstract

Low-power wireless networks provide IPv6 connectivity through 6LoWPAN, a set of standards to aggressively compress IPv6 packets over small maximum transfer unit (MTU) links such as 802.15.4.

The entire purpose of IP was to interconnect different networks, but we find that different 6LoWPAN implementations fail to reliably communicate with one another. These failures are due to stacks implementing different subsets of the standard out of concern for code size. We argue that this failure stems from 6LoWPAN's design, not implementation, and is due to applying traditional Internet protocol design principles to low-power networks.

We propose three design principles for Internet protocols on low-power networks, designed to prevent similar failures in the future. These principles are based around the importance of providing flexible tradeoffs between code size and energy efficiency. We apply these principles to 6LoWPAN and show that the modified protocol provides a wide range of implementation strategies while allowing implementations with different strategies to reliably communicate.

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 <https://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 15, 2021.

## Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Background . . . . .	<a href="#">4</a>
<a href="#">3.</a>	6LoWPAN Today . . . . .	<a href="#">6</a>
<a href="#">4.</a>	Traditional Principles: Not Low-Power . . . . .	<a href="#">10</a>
<a href="#">5.</a>	Three Principles . . . . .	<a href="#">10</a>
<a href="#">5.1.</a>	Principle 1: Capability Spectrum . . . . .	<a href="#">10</a>
<a href="#">5.2.</a>	Principle 2: Capability Discovery . . . . .	<a href="#">11</a>
<a href="#">5.3.</a>	Principle 3: Explicit and Finite Bounds . . . . .	<a href="#">12</a>
<a href="#">6.</a>	A Principled 6LoWPAN . . . . .	<a href="#">12</a>
<a href="#">6.1.</a>	Principle 1: Capability Spectrum . . . . .	<a href="#">13</a>
<a href="#">6.2.</a>	Principle 2: Capability Discovery . . . . .	<a href="#">16</a>
<a href="#">6.3.</a>	Principle 3: Provide Reasonable Bounds . . . . .	<a href="#">16</a>
<a href="#">7.</a>	Evaluation . . . . .	<a href="#">17</a>
<a href="#">7.1.</a>	Implementations . . . . .	<a href="#">17</a>
<a href="#">7.2.</a>	Compile-Time Costs . . . . .	<a href="#">18</a>
<a href="#">7.3.</a>	Run-time Performance . . . . .	<a href="#">20</a>
<a href="#">8.</a>	Discussion and Conclusions . . . . .	<a href="#">21</a>
<a href="#">9.</a>	Definitions . . . . .	<a href="#">22</a>
<a href="#">9.1.</a>	Terminology . . . . .	<a href="#">22</a>
<a href="#">10.</a>	Security Considerations . . . . .	<a href="#">22</a>
<a href="#">11.</a>	IANA Considerations . . . . .	<a href="#">22</a>
<a href="#">12.</a>	References . . . . .	<a href="#">23</a>
<a href="#">12.1.</a>	Normative References . . . . .	<a href="#">23</a>
<a href="#">12.2.</a>	Informative References . . . . .	<a href="#">23</a>
	Authors' Addresses . . . . .	<a href="#">25</a>

## [1.](#) Introduction

Interoperability has been fundamental to the Internet's success. The Internet Protocol (IP) allows devices with different software and link layers to communicate. IP provides a basic communication



substrate for many higher layer protocols and applications. In the decades of the Internet's evolution, we have accumulated and benefited from a great deal of wisdom and guidance in how to design, specify, and implement robust, interoperable protocols.

Over the past decade, the Internet has extended to tens of billions of low-power, embedded systems such as sensor networks and the Internet of Things. Hundreds of proprietary protocols have been replaced by 6LoWPAN, a standardized format for IP networking on low-power wireless link layers such as 802.15.4 [[RFC6282](#)] and Bluetooth Low Energy [[RFC7668](#)]. 6LoWPAN was created with the express purpose of bringing interoperable IP networking to low power devices, as stated in the 6LoWPAN WG charter. Many embedded operating systems have adopted 6LoWPAN, [[tinyos](#)] [[riot](#)] [[mbedos](#)] [[contiki-ng](#)] [[contiki](#)] [[lite-os](#)] and [[zephyr](#)], and every major protocol suite uses it (zigbee, openthread). In fact, devices today can communicate with the broader Internet.

However, in many cases 6LoWPAN implementations cannot communicate with each other. We find that no pairing of the major implementations fully interoperates. Despite 6LoWPAN's focus on interoperability, two key features of the protocol -- range extension via mesh networking of devices, and the convenience of different vendors being able to share a gateway -- are largely impossible 10 years later.

Each of the openly available 6LoWPAN stacks, most of which are used in production, implements a subset of the protocol and includes compile-time flags to cut out additional compression/decompression options. As a result, two devices might both use 6LoWPAN, yet be unable to exchange certain IP packets because they use required features the other does not implement. This is especially problematic for gateways, which need to be able to talk to multiple implementations to enable significant scaling in real world applications.

This draft argues that the failure of 6LoWPAN interoperability stems from applying traditional protocol design principles to low-power networks. Low-power protocols minimize energy consumption via compression. Squeezing every bit out of packet headers, however, requires many different options/operating modes. Principles such as Postel's Law - "Be liberal in what you accept, and conservative in what you send" [[RFC1122](#)] - state that an implementation must be able to receive every feature, even though it only sends some of them. However, code space is tight on many systems. As a result, when an application does not fit, developers cut out portions of the networking stack and stop working with other devices. Put another way, 3kB of unused compression code seems tiny, but when removing it



allows an additional 3kB of useful features, developers cut out parts of 6LoWPAN and devices become part of a custom networked system rather than the Internet.

This draft presents three design principles which resolve this tension between interoperability and efficiency. Protocols following these principles get the best of both worlds: resource-limited devices can implement subsets of a protocol to save code space while remaining able to communicate with every other implementation.

*\*Capability spectrum:* a low-power protocol specifies a linear spectrum of capabilities. Simpler implementations have fewer capabilities and save less energy, while fuller implementations have strictly more capabilities and are able to save more energy. When two devices differ in capability levels, communication can always fall back to the lower one.

*\*Capability discovery:* a low-power protocol provides mechanisms to discover the capability of communicating devices. This discovery can be proactive (advertisements) or reactive (in response to errors).

*\*Explicit, finite bounds:* a low-power protocol specifies explicit, finite bounds on growth during decompression. Without explicit bounds, buffers must be sized too conservatively. In practice, implementations allocate smaller buffers and silently drop packets they cannot decompress.

This draft is not the first to observe poor 6LoWPAN interoperability [[probe-it-plugtest](#)], but it is the first to identify this root cause. It is the first to define design principles for protocols that allow implementations to reduce code size and still interoperate. This draft examines how these principles could be applied to a low power protocol and evaluates the overhead of doing so. It finds that applying these principles to 6LoWPAN promises interoperability across a wide range of device capabilities, while imposing a code size cost of less than 10%. In particular, capability discovery requires an order of magnitude less code than the code size difference at the extremes of our capability spectrum, with minimal runtime overhead.

## **2. Background**

6LoWPAN is a set of standards on communicating IPv6 over low-power wireless link layers such as IEEE 802.15.4 and Bluetooth[RFC4944] [[RFC6282](#)] [[RFC7668](#)]. 6LoWPAN primarily specifies two things: aggressive compression of IPv6 headers and how to fragment/re-assemble packets on link layers whose MTU is smaller than the minimum 1280 bytes required by IPv6. 6LoWPAN also specifies optimized IPv6



Neighbor Discovery. 6LoWPAN is critical to ensuring that IPv6 communication does not exceed the energy budget of low power systems.

The interoperability IP provides is a goal in and of itself, and so many different network stacks, including ZigBee and Thread, have transitioned to supporting IP connectivity with 6LoWPAN. IP connectivity allows systems to easily incorporate new services and applications, and allows applications to build on all of the existing Internet systems, such as network management and monitoring.

#### \*Low Power Hardware and Operating Systems\*

**IoT Platform**	**Code (kB)**	**Year**
-----	-----	-----
EMB-WMB	64	2012
Zolertia Z1	92	2013
TI CC2650	128	2015
NXP MKW40Z	160	2015
SAMR21 XPro	256	2014
Nordic NRF52840 DK	512	2018

Figure 1: Flash Size across IoT Platforms

Figure 1 shows a variety of older and more recent low-power platforms. Modern microcontrollers typically have 128-512 kB of code flash. Applications often struggle with these limits on code flash, and rarely leave code space unused. Embedded systems are application specific, and use their limited available resources toward different ends. Despite this, they still rely on a small number of reusable OSes for basic abstractions.

To support the highly constrained applications and devices for which embedded OSes are used, embedded OSes must be minimal. However, the manner in which they must be minimal varies - some applications require minimal use of radio energy, which can require code size consuming techniques like compression, while others require tiny/low cost MCUs without space for those mechanisms. To support this variety, OSes have compile-time flags to include or exclude parts of the system or networking stack (contiki, mbedos, contiki-ng). Some systems take a more extreme approach, dynamically generating the minimum code to compile from the application itself (tinyc). Part of this minimalist focus is that until application developers demand certain features, OSes are likely to leave them out entirely (a requirement often specified in contribution guides - contiki-ng). These techniques are critical to ensuring a given OS can support a wide range of embedded platforms, and influence the implementation of network protocols.



### 3. 6LoWPAN Today

The original 6LoWPAN working group charter stated "The Working Group will generate the necessary documents to ensure interoperable implementations of 6LoWPAN networks". However looking at implementations today, we find that each includes a different subset of the protocol. Source code shows this is due to concerns with code size. Experiences with a new implementation verify these concerns.

*\*Feature Fail\** A 6LoWPAN receiver has much less flexibility than a sender: it must be able to process any valid compression it receives. Figure 2 shows the receiver features supported by 6 major open-source 6LoWPAN stacks. Some, such as TinyOS, are mostly developed and used in academia. Others, such as ARM Mbed and Nest's OpenThread, are developed and supported commercially. Contiki and Contiki-NG sit somewhere in the middle, having both significant academic and commercial use. Riot is an open-source OS with hundreds of contributors for industrial, academic, and hobbyist use. Two widely used open source stacks excluded from this analysis are LiteOS [[lite-os](#)] and Zephyr [[zephyr](#)] -- they are excluded because LiteOS uses the same 6LoWPAN library (LWIP) used in MBED-OS, and Zephyr simply imports OpenThread.

In almost all cases, each stack's support for features is symmetric for sending and receiving. There are significant mismatches in feature support between stacks. These mismatches lead to deterministic cases when IP communication fails. We verified these failures by modifying existing network applications and testing them on hardware, using Wireshark to verify packets were compressed and formatted as we expected when receivers failed to decode packets. Every implementation pair fails for some type of packet which can be organically generated by one of the stacks. This result may be surprising when compared to prior work which demonstrated successful interoperability, such as [[RPL-interop](#)]. However, this early success preceded the release of [RFC 6282](#), which increased the complexity (and overhead) of 6LoWPAN.



Feature	Stack					
	Contiki	NG	OT	Riot	Arm	TinyOS
Uncompressed IPv6	0	0		0	0	0
6LoWPAN Fragmentation	0	0	0	0	0	0
1280 byte packets	0	0	0	0	0	0
Dispatch_IPHC header prefix	0	0	0	0	0	0
IPv6 Stateless Address Compression	0	0	0	0	0	0
Stateless multicast addr compression	0	0	0	0	0	0
802.15.4 16 bit short address support		0	0	0	0	0
IPv6 Address Autoconfiguration	0	0	0	0	0	0
Stateful Address Compression	0	0	0	0	0	0
Stateful multicast addr compression			0	0	0	
TC and Flow label compression	0	0	0	0	0	0
NH Compression: Ipv6 (tunneled)			0		0	0
IPv6 NH Compression: UDP	0	0	0	0	0	0
UDP port compression	0	0	0	0	0	0
UDP checksum elision						0
Compression + headers past first frag				0	0	
Compression of IPv6 Extension Headers		~	~		0	0
Mesh Header			0		0	~
Broadcast Header						0
Regular IPv6 ND	0	0		0	0	~
<a href="#">RFC 6775</a> 6LoWPAN ND				0	0	
<a href="#">RFC 7400</a> Generic Header Compression						

~ = Partial Support

Figure 2: 6LoWPAN Interoperability Matrix



+-----+-----+-----+-----+-----+					
Stack		Code Size Measurements (kB)			
+-----+-----+-----+-----+-----+					
	6Lo-All	Compression	Frag	Mesh/Bcast Hdr	
+-----+-----+-----+-----+-----+					
Contiki-NG	6.2	3.4	1.9		N/A
Contiki	11.3	6.0	3.3		N/A
OT	26.6	20.0	1.3		4.5
Riot	7.5	>4.7	1.5		N/A
Arm Mbed	22.1	18.0	3.1		1331
TinyOS	16.2	----	----		0.6
+-----+-----+-----+-----+-----+					

Figure 3: 6LoWPAN stack code size for 6 open source stacks. The code size varies by over a factor of 4, in part due to different feature sets. Compression dominated the code requirements of each stack. Individual components do not add up to total as some size cannot be clearly attributed to any subcomponent. TinyOS's whole-program optimization model precluded separating out subcomponents, and only includes incomplete mesh header support.

\*Why?\*

IP communication can consistently fail in low-power networks, despite the presence of succinct standards ([RFC 6282](#) + [RFC 4944](#) is 52 pages) designed for low-power devices. Worse, this failure is silent: the receiver will simply drop the packet. Examining the documentation and implementation of each stack, code size concerns motivated feature elision. Mbed, Riot, and Contiki even provide compile-time flags to remove additional features for particularly constrained use cases.

Figure 3 shows a break down of the code size of each stack. Compression dominates the code size of 6LoWPAN implementations, and in several cases 6LoWPAN's size is comparable to the whole rest of the IPv6 stack. The Contiki and Contiki-NG implementations are significantly smaller than the others in part because they elide significant and complex features. The ARM Mbed IPv6 stack uses 45kB of flash. This is nearly 1/3 of the available space on a CC2650, just for IPv6: it does not include storage, sensors, the OS kernel, cryptography, higher layer protocols, signal processing, or applications.

Can a careful developer implement a leaner, fully-featured stack? To answer this question, we implemented our own 6LoWPAN stack. Our open-source implementation is written in Rust, for Tock, a secure embedded OS [[tock](#)].



Our experiences support the comments and documentation of the other stacks. We surpassed the size of the Contiki-NG and Riot 6LoWPAN code before adding support for recursive compression of IPv6 or the mesh and broadcast headers. We noted several aspects of the protocol required surprisingly complex code to properly handle. For example, 6LoWPAN requires IPv6 tunneled inside a compressed packet to compress interior headers as well. This requires the decompression library to support recursive invocation, which increases minimum execution stack sizes and makes tracking buffer offsets during decompression more difficult. Refusing to support tunneled IPv6 packets (e.g., Contiki) greatly simplified the code. Another example: headers in the first 6LoWPAN sub-IP fragment must be compressed, while headers in subsequent fragments must not be compressed. Given that low-power link layers have variable length headers, correctly determining exactly where to fragment and what should be compressed requires complex interactions between layers of the stack. Finally, 6LoWPAN requires support for out-of-order reception of fragments, potentially from different packets. This forced our receiver to store and track state for a collection of received packets, preventing reliance on a single global receive buffer. The exercise of implementing a 6LoWPAN stack from the ground up affirmed that code size concerns encourage feature elision.

\*Why does it matter?\*

For any 6LoWPAN implementation, there exists a border router implementation that can connect it to the broader Internet. But this status-quo model of connectivity forces vertical integration and fails to meet the original design goals of 6LoWPAN, for two reasons.

First, a 6LoWPAN gateway can not know how to safely compress packets for different nodes, unless it communicates only with devices produced by the same vendor. As a result, for a coverage area containing devices produced by 5 different vendors, 5 gateways are required. If not for feature mismatches, a single gateway would suffice. Second, the current situation significantly limits the potential for range extension via mesh topologies. Most existing 6LoWPAN meshes rely on "route-over" mesh routing at the network layer, which requires that each node can at least partially decompress and recompress IP headers when forwarding. Mesh-under routing is no better, as implementation of the mesh header is not universal (see Figure 2). Poor interoperability worsens the usability, range, cost, and efficiency of these networks.



#### **4. Traditional Principles: Not Low-Power**

Over the past 45 years, the Internet community has coalesced on a small number of design principles. Connectivity through interoperability is a key premise of the Internet. Principles such as layering and encapsulation support composing protocols in new ways (e.g., tunneling), while the end-to-end principle [[end-to-end](#)] allows building a robust network out of an enormous and complex collection of unreliable parts. The robustness principle asserts that implementations should make no assumptions on packets they receive: bugs, transmission errors, and even memory corruption can cause a device to receive arbitrarily formatted packets. It also asserts that an implementation must be ready to receive any and all properly formatted packets. This aspect of the principle is often attributed to John Postel as Postel's Law, first written down in the initial specification of IPv4: "In general, an implementation should be conservative in its sending behavior, and liberal in its receiving behavior."

Protocols often have optional features ("MAY, SHOULD, and OPTIONAL" in RFC language). Implicitly, due to Postel's Law, a receiver needs to handle either case. This scenario creates an asymmetry, where sender code can have reduced complexity but receiver code must be large and complex.

We need to think about low-power protocols differently. They need new principles to help guide their design. These principles need to embrace that there is no "one size fits all" design, while defining how devices choosing different design points interoperate. Flexibility needs to exist not only for senders, but also receivers, without harming interoperability.

#### **5. Three Principles**

This section describes three protocol design principles which prevent failures in low-power protocols. These principles are absolutely necessary to ensure interoperable implementations in this space, and should be closely observed.

##### **5.1. Principle 1: Capability Spectrum**

A low power protocol should be deployable/installable on devices which are at the low end of code and RAM resources. Rather than require every device pay the potential energy costs of fewer optimizations, a protocol should support a linear spectrum of device capabilities.



This may seem familiar -- the IP [[RFC0791](#)] and TCP [[RFC0793](#)] specifications provide optional fields which can be used by endpoints at their leisure; many non-standard HTTP headers will be ignored unless both client and server support them; TLS ciphersuite support is often asymmetrical. But this principle is different. For those examples, no linear spectrum exists -- support for any particular capability is generally unrelated to support for any other. Checking for support of any feature requires explicit enumeration of each, making it impossible to effectively compress such options. A non-linear spectrum requires storing feature support for every neighbor in RAM, or re-discovering capabilities on every exchange.

Low power protocols require simpler capability management. A low power protocol should define a capability spectrum with a clear ordering in which especially resource constrained devices can reduce code size or RAM use by eliding features. Such a spectrum makes a protocol usable by extremely low resource devices without forcing more resourceful devices to communicate inefficiently.

This capability spectrum should be a linear scale. For a device to support capability level  $N$ , it must also support all lower capability levels. More complex configuration approaches (e.g., a set of independent options) might allow for a particular implementation to be more efficient, picking the features that give the most benefit for the least added complexity. However, this sort of optimization makes interoperability more difficult, as two devices must negotiate each specific feature to use.

## **5.2. Principle 2: Capability Discovery**

The second principle follows from the first: if different capability levels exist, there should be a mechanism for two devices to determine what level to communicate with.

The capability negotiation we propose here differs from capability discovery mechanisms built for traditional systems, such as IP Path MTU discovery or the Link Layer Discovery Protocol (LLDP). IP Path MTU discovery relies on continual probing until an acceptable value is discovered. LLDP requires regular, detailed capability advertisements at a fixed interval. The energy overhead of network probing or advertising is unacceptable in most low power environments. Capability discovery in low power networks should require no more than one failure between any two neighbors, even if this slightly increases the overhead per error. Proactive capability discovery should be built into baseline communication required for tasks like neighbor discovery or route maintenance. Further, assumptions for traditional systems that prohibit storing per-endpoint state do not apply, as nodes store information about link-



layer neighbors, `_not IP endpoints_`. This is needed because low-power networks with route over topologies frequently involve decompression and re-compression at each hop to enable forwarding. Low power nodes have few neighbors, so storing a few bits of state for each is feasible and can significantly reduce the amount of radio energy needed for communication. The code size cost of storing state is small compared to the cost of complex compression mechanisms.

In a low power network with capability discovery, if two devices wish to communicate, they default to the lower of their supported capability levels. E.g. a level 2 and a level 4 device should communicate at level 2. One offshoot of this principle is that it requires implementations have symmetric capabilities for send and receive - no benefits can be realized from an asymmetric implementation.

### **5.3. Principle 3: Explicit and Finite Bounds**

Protocols must specify explicit and reasonable bounds on recursive or variable features so implementations can bound RAM use. This allow implementations to safely limit their RAM use without silent interoperability failures. This also ensures that capability discovery is sufficient for interoperability.

The idea of imposing bounds is, on its own, not unique to this space. TCP enforces a finite limit of 40 bytes for TCP options which may be appended to the TCP header, as does IPv4. DHCP allows for the communication of maximum DHCP message sizes. In the space of low power Internet protocols, however, this idea `_must be pervasive_`. Notably, the original designers of a specification may not know exactly what these values should be. This is not a new problem: TCP congestion control, for example, had to specify initial congestion window values. In this space, bounds should initially be very conservative. Over time, if increasing resources or knowledge suggests they should grow, then future devices will have the onus of using fewer resources to interoperate with earlier ones. The capability spectrum defined in the previous two principles can be helpful in this regard.

## **6. A Principled 6LoWPAN**

This section proposes how to apply the three principles in the previous section to 6LoWPAN through specific modifications to the protocol. These modifications ensure that two 6LoWPAN devices can communicate even if they choose different code size/energy efficiency tradeoffs. We refer to this modified protocol as Principled 6LoWPAN (P6LoWPAN).



This application of our principles is not intended as a suggestion that these changes be made immediately to 6LoWPAN, as modifying an established protocol is a complex task very different from constructing new protocols. Instead, this application is a tool for evaluating these principles, and an example for how they should be applied.

### 6.1. Principle 1: Capability Spectrum

We propose replacing the large collection of "MUST" requirements -- the features in Figure 2 -- into 6 levels of functionality. These "Capability Levels" are depicted in Figure 4.

**Capability**	**Basic Description / Added Features**
Level 0	Uncompressed IPv6 + ability to send ICMP errors
	<ul style="list-style-type: none"> <li>- Uncompressed IPv6</li> <li>- 6LoWPAN Fragmentation (Fragment Header)</li> <li>- 1280 Byte Packets</li> <li>- Stateless decompression of source addresses</li> </ul>
Level 1	IPv6 Compression Basics + Stateless Addr Compression
	<ul style="list-style-type: none"> <li>- Support for the Dispatch\_IPHC Header Prefix</li> <li>- Correctly handle elision of IPv6 length and version</li> <li>- Stateless compression of all unicast addresses</li> <li>- Stateless compression of multicast addresses</li> <li>- Compression + 16 bit link-layer addresses</li> </ul>



	- IPv6 address	
	autoconfiguration	
+-----+	+-----+	+-----+
Level 2	Stateful IPv6 Address Compression	
+-----+	+-----+	+-----+
	- Stateful compression of	
	unicast addresses	
	- Stateful compression of	
	multicast addresses	
+-----+	+-----+	+-----+
Level 3	IPv6 Traffic Class and Flow Label	
	Compression	
+-----+	+-----+	+-----+
	- Traffic Class compression	
	- Flow Label Compression	
	- Hop Limit Compression	
+-----+	+-----+	+-----+
Level 4	Next Header Compression + UDP	
	Port Compression	
+-----+	+-----+	+-----+
	- Handle Tunneled IPv6	
	correctly	
	- Handle the compression of the	
	UDP Next Header	
	- Correctly handle elision of	
	the UDP length field	
	- Correctly handle the	
	compression of UDP ports	
	- Handle headers past the first	
	fragment, when first fragment	
	compressed.	
+-----+	+-----+	+-----+
Level 5	Entire Specification	
(all routers)		
+-----+	+-----+	+-----+
	- Support the broadcast header	
	and the mesh header	
	- Support compression of all	
	IPv6 Extension headers	
+-----+	+-----+	+-----+



Figure 4: Capability Spectrum

These levels prioritize features which provide the greatest energy savings per byte of added code size, based off our code size measurements and the number of bits saved by each additional compression mechanism. They allow for a wide range of code size/efficiency tradeoffs.

For example, addresses dominate an uncompressed IPv6 header. Level 0 devices only support compressed source addresses, while level 1 devices support all stateless address compression. In one early design of this spectrum, Level 0 supported only uncompressed packets. However, this raises a problem with ICMP error generation. If a node cannot decompress the source address of a received packet, it cannot send ICMP errors. ICMP errors are required for capability discovery. Stateful compression depends on an out-of-band signal to set up state, such that nodes only send statefully compressed packets to nodes who also support it. Therefore decompressing stateless source addresses is a minimum requirement.

The classes in this scale do not precisely reflect the current feature support of the implementations described in [Section 3](#). For example, Contiki supports UDP port compression (level 4) but does not support 802.15.4 short addresses (level 2) or stateful multicast compression (level 3): following this formulation, Contiki only provides level 1 support. If Contiki supported 16-bit addresses, it would provide level 2 support. A concrete spectrum such as the one above gives stack designers a structure and set of guidelines on the order in which to implement features. Based on our experiences developing a 6LoWPAN stack, we believe that if this scale existed as part of the initial specification, implementations would have made an effort to adhere to it.

One additional advantage of this spectrum is that it allows for some future additions to the P6LoWPAN specification without breaking interoperability between new and old implementations. For example, our scale does not include support for Generic Header Compression [[RFC7400](#)] because none of the open-source stacks we analyzed implement it. Despite this, support for this RFC could easily be added as a new class on this linear scale (as Class 6), and devices supporting it would know to not use it when communicating with lower class implementations.

This spectrum requires that a node store 3 bits of state for each neighbor. Given that low-power nodes often store ten or more bytes for each entry in their link table (link quality estimates, addresses, etc.), this cost is small. 6LoWPAN already assumes that



routers are more resourceful devices, P6LoWPAN routers are required to be level 5.

## **6.2. Principle 2: Capability Discovery**

We propose two mechanisms by which P6LoWPAN performs capability discovery: neighbor discovery (ND) and ICMP\@. Neighbor discovery [[RFC4861](#)] is analogous to ARP in IPv4: it allows IPv6 devices to discover the link layer addresses of neighboring addresses as well as local gateways. Devices use neighbor discovery to proactively discover capability levels and ICMP to detect when incompatible features are used. Of the two, only ICMP is required. Neighbor discovery simply allows a pair of differing nodes to avoid an initial ICMP error, and allows for optimization of host-router communication during neighbor discovery.

*\*ICMP:* We propose adding a new ICMPv6 message type--P6LoWPAN Class Unsupported--which a device sends in response to receiving 6LoWPAN features it does not understand. This error encodes the device's capability level. A node receiving such an error updates its link table entry with the capability level. In the future, any packets sent to that address use at most the supported level.

*\*Neighbor discovery:* We propose adding an IPv6 ND option that allows a device to communicate its capability class during network association. This option would be included in Router Solicitations and Neighbor Advertisements, and would allow all devices that obtain link-layer addresses via ND to also know how to send packets which that neighbor can receive. When a node uses ND to resolve an IP address to a link layer address, it learns the supported capability level as well as the link layer address. This option minimizes the energy cost of communicating capabilities. It is worth noting that [RFC 7400](#) already employs a similar method for communicating whether devices implement General Header Compression: adding such an option is clearly viable [[RFC7400](#)].

## **6.3. Principle 3: Provide Reasonable Bounds**

[Section 3](#) discussed two missing bounds which affect 6LoWPAN interoperability: limits on header decompression and bounds on recursion when decompressing tunneled IPv6.

For P6LoWPAN, we propose that header decompression be bounded to 51 bytes. This bound allows for significant RAM savings in implementations that decompress first fragments into the same buffer in which the fragment was originally held. 51 bytes is a good tradeoff between RAM savings and how frequently we expect such a bound would force packets to be sent uncompressed. A 51 byte limit



allows for transmission of a packet containing a maximally compressed IP header (+38 bytes), a maximally compressed UDP header (+6 bytes), and one maximally compressed IPv6 extension header (+7 bytes). This allows saving hundreds of bytes of RAM, without jeopardizing interoperability. Packets requiring more decompression than this are extremely rare, and could be sent uncompressed. How rare? It is only possible to surpass this limit if tunneled IPv6 is used or multiple IPv6 extension headers are present. As of 2014, a real-world study of IPv6 extension header use found that 99% of packets with `_multiple_` extension headers were dropped in the real Internet, as published at IETF 90.

Second, we propose that headers for tunneled IPv6 should not be compressed. The primary motivation for this feature was from the RPL protocol [[RFC6550](#)], as discussed in [Section 3](#). However, the fact that RPL must tunnel IPv6 in this way is generally agreed to be a problem and a wart in its design that should be avoided when possible. This change allows implementations to avoid recursive functions to decompress these headers, and instead use simple if/else statements.

## [7. Evaluation](#)

This section evaluates the costs of applying our principles to 6LoWPAN. The principles are written such that interoperability comes by construction, and thus interoperability of the modified protocol cannot be directly evaluated without observing implementations written by different stakeholders. But indirect evaluation is possible. Can a reasonable set of capability levels provide a good range of implementation complexity from which a developer can choose? Is the overhead of the proposed mechanisms low enough to make them viable? Are the savings afforded by a linear capability spectrum worth the associated limitations? We find the incremental costs of capability discovery mechanisms is small, adding 172-388 bytes of code in the worst case. We find that the capability spectrum allows meaningful savings in code size and memory usage. Finally, we find capability discovery has a low run-time performance cost when a linear spectrum is used.

### [7.1. Implementations](#)

First, we implemented the proposed P6LoWPAN on the Contiki-NG 6LoWPAN stack, modifying it such that a compile-time option determines which features of 6LoWPAN are compiled. We selected Contiki-NG because it has the smallest 6LoWPAN stack of those tested, so any overheads the mechanisms introduce would be most pronounced. Our changes required modifying 500 lines of code relative to the head of the 4.2 release of Contiki-NG. We did not add additional 6LoWPAN features which were



absent from the original Contiki-NG 6LoWPAN stack. Our code size numbers therefore represent a conservative lower bound of the total possible savings. All code sizes provided in this section are compiled with the Texas Instruments CC2650 as the target.

We also added ICMP and ND support for capability discovery. The updated stack responds to incompatible 6LoWPAN messages with an ICMP error, and communicates its capability level in Router Solicitation messages using the 6CIO prefix originally defined in [\[RFC7400\]](#). It stores the capability class of each node in its link table, and compresses IPv6 packets by the maximum amount supported by the destination.

Finally, we implemented a second modified 6LoWPAN stack in Contiki-NG, which does not follow the recommendation of using a linear capability spectrum. In this modified implementation, each node can select any of the 6LoWPAN features it chooses. We refer to this implementation as FLEX-6LoWPAN. For this alternative policy, we isolated 26 features of 6LoWPAN as single bit flags in a 32 bit bitfield. Thus, FLEX-6LoWPAN stores and communicates capabilities using 4 byte objects. FLEX-6LoWPAN also supports the added granularity required to maximally compress outgoing messages intended for a device supporting any specific combination of features. We did not add back in any 6LoWPAN features which the Contiki-NG stack did not originally support. This second implementation required modifying about 300 additional lines of code from the P6LoWPAN implementation.

## [7.2.](#) Compile-Time Costs

Figure 5 shows the size of the original Contiki-NG 6LoWPAN stack compiled at each possible capability level. Each capability level adds between 0.25 and 1.05 kB of code, and the spectrum enables implementations to cut the size of the 6LoWPAN stack by up to 45%. The code size cost of adding capability discovery, using the P6LoWPAN implementation with the linear capability spectrum, is shown in Figure 6. Capability discovery adds 178-388 bytes, a fraction of the size which implementations can save by supporting lower capability levels. The code added for communication varies across capability levels because the number of code paths for ICMP error generation and compression changes.



Capability	Code Size (kB)	Increase (kB)
Level 0	3.2	\-
Level 1	4.2	1.0
Level 2	4.8	0.6
Level 3	5.1	0.3
Level 4	5.6	0.5
Level 5	6.2	0.6

Figure 5: 6LoWPAN code size of different capabilities levels in Contiki-NG. The spectrum spans a nearly 100% increase in code size.

Capability	Base	w/Discovery	Increase
Level 0	3.2	3.4	188 bytes
Level 1	4.2	4.4	260 bytes
Level 2	4.8	5.2	388 bytes
Level 3	5.1	5.4	340 bytes
Level 4	5.6	5.9	296 bytes
Level 5	6.2	6.3	172 bytes

Figure 6: The cost of implementing capability discovery in Contiki-NG is on average less than 5% of the total 6LoWPAN size; the maximum size reduction from choosing a lower capability level is 10x the discovery cost.

Figure 7 presents the compile-time costs of using an arbitrary bitfield instead of a linear capability spectrum by comparing our P6LoWPAN implementation with our FLEX-6LoWPAN implementation. The bitfield approach requires 32 bits per neighbor to store capabilities, instead of 3 bits. More importantly, it complicates determining the allowable compression between two nodes, as demonstrated by the code size increase. The important takeaway here is that opting for a less restrictive set of feature combinations mitigates much of the savings provided by implementing capabilities. For example, a FLEX-6LoWPAN device with the equivalent of level 4 capabilities requires more code space than a level 5 P6LoWPAN device - the linear capability spectrum makes a difference. The code size addition for FLEX-6LoWPAN is a conservative lower bound, as we did not need to add checks for handling 6LoWPAN compression features that Contiki-NG does not support.



### 7.3. Run-time Performance

**\*ND Cost\*** 6LoWPAN ND communication [[RFC6775](#)] is host-initiated and flows through routers (which must be level 5), and nodes store neighbor capability levels alongside link layer addresses: thus there is no possibility of communication failures due to capability mismatches. Therefore the cost of capability discovery in networks that use IPv6 ND is exclusively that certain ND messages become longer (router solicitations and neighbor advertisements are sent with an added capability option). To put this added cost in perspective, The equation below shows the total link-layer payload bytes sent/received for ND by a node in its initial wake-up period. This equation assumes the configuration described in [RFC 6775](#) as the "Basic Router Solicitation Exchange" - route over topology, 1 6LoWPAN context, 1 on-link prefix, and the host requires address registration. All variables not affected by the use of capability discovery are assigned the minimum possible value for the scenario discussed, so that the overhead of capability discovery represents a worst case.

If  $C$  = total link layer payload sent/recieved for ND, and  $N$  = endpoints requiring address resolution:

$$C = \text{Router Solicitation \{RS\}} + \text{Min. IP Hdr \{2\}} + \text{Router Advertisement \{104\}} + \text{Min. IP Hdr \{2\}} + (\text{Neighbour Solicitation \{24\}} + \text{Min. IP Hdr \{2\}}) * N + (\text{Neighbour Advertisement \{24\}} + \text{Min. IP Hdr \{2\}}) * N + \text{Address Registration Options in first NS \{24\}} + \text{Address Registration Options in first NA \{16\}}.$$

Figure 8 shows the values of RS and NA for each 6LoWPAN implementation, and the resulting total ND cost. Notably, use of an arbitrary bitfield increases the size of the capability option by 4 bytes, making use of existing ND options like the 6CIO option impossible. In both cases the additional bytes added for capability discovery are small compared to the total cost of ND (  $\leq 8\%$  linear spectrum /  $\leq 16\%$  arbitrary bitfield).

**\*ICMP Cost\*** In networks that do not use IPv6 ND the cost of capability discovery is the energy/latency required for one ICMP packet per failure between any two nodes. For P6LoWPAN capability based failures can only happen in one direction, so the size of this link-layer payload is:

$$C_{\text{icmp}} = \text{Compressed IP Header Size} + 4$$

For FLEX-6LoWPAN  $C_{\text{icmp}} = 48$ , because the recipient does not know the capabilities of the sender, and thus must send an uncompressed packet to ensure successful reception of its own capabilities. This example



reveals why use of an arbitrary bitfield is so undesirable - the ability to compress headers in ICMP errors can reduce overhead by a factor of 4 or more (in the common case of 8 byte compressed headers).

--	Linear Spectrum	Arbitrary Bitfield
-----	-----	-----
6LoWPAN Code Size	5.9 kB	6.5 kB
RAM per neighbor	19 Bytes	22 Bytes
-----	-----	-----

Figure 7: Resource requirements for a 6LoWPAN stack in Contiki-NG using a linear capability spectrum vs. using an arbitrary capability bitfield.

--	RS	NA	C (Total ND Cost)
-----	-----	-----	-----
6LoWPAN	20	24	$168 + 52 \cdot N$
P6LoWPAN	24	28	$172 + 56 \cdot N$
FLEX-6LoWPAN	28	32	$176 + 60 \cdot N$
-----	-----	-----	-----

Figure 8: Total ND cost for each implementation

## 8. Discussion and Conclusions

A new generation of low-power devices face a connectivity dilemma: Internet protocols are not designed for energy efficiency, but compression and other energy saving adaptations takes up precious code space. Device deployments specialized for single-vendor local networks make trade-offs specific to their application requirements. As a result, IP communication between IP enabled devices fails. This problem is not specific to 6LoWPAN -- Iova et. al. recently noted similar issues in the RPL protocol: "RPL has too large of a footprint for resource-constrained devices, and requires all devices in a network to run the same mode of operation, limiting heterogeneity" [[iova](#)].

Part of the challenge is that some traditional protocol design principles do not apply well to the low-power setting. We present three design principles for low-power protocols that attempt to remedy this. These principles explicitly acknowledge the unique code space/energy tradeoffs of low-power devices.

Looking forward, considering this tension is critical for protocol designers in this ecosystem of diverse hardware capabilities and



application tradeoffs. 6LoWPAN is not the only low power Internet protocol -- the low power space uses its own routing protocols, address discovery protocols, and application layer protocols [[RFC6550](#)] [[RFC7252](#)]. Additional protocols will follow as the space matures. Many of these protocols will be initially developed outside the IETF -- Jonathan Hui was a graduate student when he presented the first complete IPv6-based network architecture for sensor nets [[hui](#)], as was Adam Dunkels when he created Contiki. While we have also presented these principles outside of the context of the IETF [[dcoss-design-considerations](#)], there is no guarantee they will be followed. When protocols in this class are adopted by the IETF, these principles represent important considerations. We present a roadmap for how these principles can reframe the discussion of how to connect the next hundred billion devices to the Internet.

## **[9.](#) Definitions**

### **[9.1.](#) Terminology**

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[RFC2119](#)].

## **[10.](#) Security Considerations**

This informational document does have some implications for security if followed.

First, capability advertisements of the type recommended in this document are liable to leak some information regarding the type of device sending those advertisements. In any situation for which this information is privileged, such advertisements must be suppressed.

Second, implementations should be careful not to take for granted that the suggestions in this document will be implemented by all other transmitting devices. Accordingly, though this document recommends reasonable bounds, receivers still must be careful to prevent buffer overflows in the event these bounds are not followed.

## **[11.](#) IANA Considerations**

This document has no actions for IANA.



## **12. References**

### **12.1. Normative References**

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", [RFC 6775](#), DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.

### **12.2. Informative References**

- [contiki] Dunkels, A., "Contiki OS", n.d., <<http://www.contiki-os.org/>>.
- [contiki-ng] Duquennoy, S., "Contiki-NG", n.d., <<https://github.com/contiki-ng/contiki-ng>>.
- [dcoss-design-considerations] Levis, P., "Design Considerations for Low Power Internet Protocols", n.d., <<https://ieeexplore.ieee.org/abstract/document/9183440>>.
- [end-to-end] Clark, D., "End-to-end Arguments in System Design", n.d..
- [hui] Culler, D., "IP is Dead, Long Live IP for Wireless Sensor Networks", n.d..
- [iova] Kiraly, C., "RPL The Routing Standard for the Internet of Things...Or Is It?", n.d..



- [lite-os] Huawei, ., "LiteOS", n.d., <<https://liteos.github.io/>>.
- [mbedos] ARM, ., "ARM MbedOS", n.d., <<https://os.mbed.com/>>.
- [probe-it-plugtest]  
Huang, X., "Tehcnical Interoperability 6LoWPAN-CoAP Report from Interop Event", n.d..
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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](#), DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", [RFC 7400](#), DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", [RFC 7668](#), DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.



- [riot] Berlin, F., "riot OS", n.d., <<https://www.riot-os.org/>>.
- [RPL-interop] Ko, J., "Contikierpl and tinyrpl - Happy Together", n.d..
- [tinyos] Levis, P., "TinyOS An Operating System for Sensor Networks", n.d., <[https://link.springer.com/chapter/10.1007/3-540-27139-2\\_7](https://link.springer.com/chapter/10.1007/3-540-27139-2_7)>.
- [tock] Levy, A., "Multiprogramming a 64kB Computer Safely and Efficiently", n.d..
- [zephyr] Foundation, T., "zephyrOS", n.d., <<https://www.zephyrproject.org/>>.

#### Authors' Addresses

Hudson Ayers  
Stanford University  
350 Serra Mall  
Stanford, CA  
United States

Email: [hayers@stanford.edu](mailto:hayers@stanford.edu)

Philip Levis  
Stanford University  
350 Serra Mall  
Stanford, CA  
United States

Email: [pal@cs.stanford.edu](mailto:pal@cs.stanford.edu)

