

Internet Architecture Board (IAB)
Request for Comments: 6574
Category: Informational
ISSN: 2070-1721

H. Tschofenig
J. Arkko
April 2012

Report from the Smart Object Workshop

Abstract

This document provides an overview of a workshop held by the Internet Architecture Board (IAB) on 'Interconnecting Smart Objects with the Internet'. The workshop took place in Prague on 25 March 2011. The main goal of the workshop was to solicit feedback from the wider community on their experience with deploying IETF protocols in constrained environments. This report summarizes the discussions and lists the conclusions and recommendations to the Internet Engineering Task Force (IETF) community.

Note that this document is a report on the proceedings of the workshop. The views and positions documented in this report are those of the workshop participants and do not necessarily reflect IAB views and positions.

Status of This Memo

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

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. Documents approved for publication by the IAB are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6574>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Constrained Nodes and Networks	5
3. Workshop Structure	6
3.1. Architecture	6
3.1.1. One Internet vs. Islands	6
3.1.2. Domain-Specific Stacks and Profiles	8
3.1.3. Which Layer?	9
3.2. Sleeping Nodes	10
3.3. Security	13
3.4. Routing	14
4. Conclusions and Next Steps	16
5. Security Considerations	19
6. Acknowledgements	20
7. Informative References	20
Appendix A. Program Committee	25
Appendix B. Workshop Materials	25
Appendix C. Accepted Position Papers	25
Appendix D. Workshop Participants	30
Appendix E. IAB Members at the Time of Approval	32

1. Introduction

The Internet Architecture Board (IAB) holds occasional workshops designed to consider long-term issues and strategies for the Internet and to suggest future directions for the Internet architecture. This long-term planning function of the IAB is complementary to the ongoing engineering efforts performed by working groups of the Internet Engineering Task Force (IETF), under the leadership of the Internet Engineering Steering Group (IESG) and area directorates.

Today's Internet is experienced by users as a set of applications, such as email, instant messaging, and services on the Web. While these applications do not require users to be present at the time of service execution, in many cases they are. There are also substantial differences in performance among the various end devices, but in general end devices participating in the Internet are considered to have high performance.

There are, however, a large number of deployed embedded devices, and there is substantial value in interconnecting them with the Internet. The term "Internet of Things" denotes a trend where a large number of devices employ communication services offered by the Internet protocols. Many of these devices are not directly operated by humans, but exist as components in buildings or vehicles, or are spread out in the environment. There is a large variation in the computing power, available memory, (electrical) power, and communications bandwidth between different types of devices.

Many of these devices offer a range of new possibilities or provide additional value for previously unconnected devices. Some devices have been connected using proprietary communication networks in the past but are now migrating to the use of the Internet Protocol suite in order to share the same communication network between all applications and to enable rich communications services.

Much of this development can simply run on existing Internet protocols. For instance, home entertainment and monitoring systems often offer a Web interface to the end user. In many cases the new, constrained environments can benefit from additional protocols and protocol extensions that help optimize the communications and lower the computational requirements. Examples of currently ongoing standardization efforts targeted for these environments include the Constrained RESTful Environments (CoRE), IPv6 over Low power WPAN (6LoWPAN), Routing Over Low power and Lossy networks (ROLL), and the Light-Weight Implementation Guidance (LWIG) working groups of the IETF.

This workshop explored the experiences of researchers and developers when considering the characteristics of constrained devices. Engineers know that many design considerations need to be taken into account when developing protocols and architecture. Balancing between the conflicting goals of code size, economic incentives, power consumption, usability, and security is often difficult, as illustrated by Clark et al. in "Tussle in Cyberspace: Defining Tomorrow's Internet" [Tussle].

Participants at the workshop discussed the experience and approaches taken when designing protocols and architectures for interconnecting smart objects to the Internet. The scope of the investigations included constrained nodes as well as constrained networks.

The call for position papers suggested investigating the area of integration with the Internet in the following categories:

- o Scalability
- o Power efficiency
- o Interworking between different technologies and network domains
- o Usability and manageability
- o Security and privacy

The goals of the workshop can be summarized as follows:

As many deployed smart objects demonstrate, running protocols like the Internet Protocol Version 4 [RFC0791] and Version 6 [RFC2460], the User Datagram Protocol (UDP) [RFC0768], the Transmission Control Protocol (TCP) [RFC0793], the Hypertext Transfer Protocol (HTTP) [RFC2616], etc., on constrained devices is clearly possible. Still, protocol designers, system architects, and developers have to keep various limitations in mind. The organizers were interested to discuss the experience with deploying IETF protocols in different constrained environments.

Furthermore, the organizers were seeking to identify issues either where current implementers do not yet have solutions or where researchers predict potential issues.

The workshop served as a venue to identify problems and to discover common interests that may turn into new work or into changes in direction of already ongoing work at the IETF and or the Internet Research Task Force (IRTF).

2. Constrained Nodes and Networks

The workshop was spurred by the increasing presence of constrained devices on the network. It is quite natural to ask how these limitations impact the design of the affected nodes. Note that not all nodes suffer from the same set of limitations.

Energy constraints:

Since wireless communication can be a large portion of the power budget for wireless devices, reducing unnecessary communication can significantly increase the battery life of a low-end device. The choice of low-power radio can also significantly impact the overall energy consumption, as can sleeping periodically, when the device is not in use. In some cases, these nodes will only wake periodically to handle needed communications. This constraint is quite in contrast to the "always on" paradigm found in regular Internet hosts. Even in the case of non-battery operated devices, power is a constraint with respect to energy efficiency goals.

Bandwidth constraints:

Various low-power radio networks offer only limited bandwidth, and show high packet loss as well as high link quality variability. The data transmission rates vary from 20 to 900 kilobits per second (e.g., in the case of IEEE 802.15.4). Nodes may be used in usually highly unstable radio environments. The physical-layer packet size may be limited (~100 bytes).

Memory constraints:

The amount of volatile and persistent storage impacts the program execution and has important implications for the functionality of the protocol stack. The Arduino UNO board, for example, provides a developer with 2 KB RAM and 32 KB flash memory (without any extensions, such as microSD cards).

A system designer also needs to consider CPU constraints, which often relate to energy constraints: a processor with lower performance consumes less energy. As described later in this document, the design of the mainboard may allow certain components to be put to sleep to further lower energy consumption. In general, embedded systems are often purpose built with only the hardware components needed for the given task, while general-purpose personal computers are less constrained with regard to their mainboard layout and typically offer a huge number of optional plug-in peripherals to be connected. A factor that also has to be taken into consideration is the intended usage environment. For example, a humidity sensor

deployed outside a building may need to deal with temperatures from -50 degrees C to +85 degrees C. There are often physical size limitations for smart objects. While traditional mainboards are rather large, such as the Advanced Technology eXtended (ATX) design with a board size of 305 x 244 mm available in many PCs or the mini-ITX design typically found in home theater PCs with 170 x 170 mm, mainboard layouts for embedded systems are typically much smaller, such as the CoreExpress layout with 58 x 65 mm, or even smaller. In addition to the plain mainboard, additional sensors, peripherals, a power adapter/battery, and a case have to be taken into consideration. Finally, there are cost restrictions as well.

The situation becomes more challenging when not only the hosts are constrained but also the network nodes themselves.

While there are constantly improvements being made, Moore's law tends to be less effective in the embedded system space than in personal computing devices: gains made available by increases in transistor count and density are more likely to be invested in reductions of cost and power requirements than into continual increases in computing power.

3. Workshop Structure

With the ongoing work on connecting smart objects to the Internet, there are many challenges the workshop participants raised in more than 70 accepted position papers. With a single workshop day, discussions had to be focused, and priority was given to those topics that had been raised by many authors. A summary of the identified issues are captured in the subsections below.

3.1. Architecture

A number of architectural questions were brought up in the workshop. This is natural, as the architectural choices affect the required technical solutions and the need for standards. At this workshop, questions regarding the separation of traffic, the need for profiling for application-specific domains, the demand for data-model-specific standardization, as well as the design choices regarding the layer at which functionality should be put were discussed and are briefly summarized below.

3.1.1. One Internet vs. Islands

Devices that used to be in proprietary or application-specific networks are today migrating to IP networks. There is, however, the question of whether these smart objects are now on the same IP network as any other application. Controlled applications, like the

fountains in front of the Bellagio hotel in Las Vegas that are operated as a distributed control system [Dolin], probably are not exchanging their control messages over the same network that is also used by hotel guests for their Internet traffic. The same had been argued for smart grids, which are described as follows in [SmartGrid]:

A smart grid is a digitally enabled electrical grid that gathers, distributes, and acts on information about the behavior of all participants (suppliers and consumers) in order to improve the efficiency, reliability, economics, and sustainability of electricity services.

The question that was raised during the workshop is, therefore, in what sense are we talking about one Internet or about using IP technology for a separate, "walled garden" network that is independent of the Internet?

Cullen Jennings compared the current state of smart object deployment with the evolution of Voice over IP (VoIP): "Initially, many vendors recommended to run VoIP over a separate VLAN or a separate infrastructure. Nobody could imagine how to make the type of real-time guarantees, how to debug it, and how to get it to work because the Internet is not ideally suited for making any types of guarantees for real-time systems. As time went on, people got better at making voice work across that type of IP network. They suddenly noticed that having voice running on a separate virtual network than their other applications was a disaster. They couldn't decide if a PC was running a softphone and whether it went on a voice or a data network. At that point, people realized that they needed a converged network and all moved to one. I wouldn't be surprised to see the same happen here. Initially, we will see very separated networks. Then, those will be running over the same hardware to take advantage of the cost benefits of not having to deploy multiple sets of wires around buildings. Over time, there will be strong needs to directly communicate with each other. We need to be designing the system for the long run. Assume everything will end up on the same network even if you initially plan to run it in separate networks."

It is clearly possible to let sensors in a building communicate through the wireless access points and through the same infrastructure used for Internet access, if you want to. Those who want separation at the physical layer can do so as well. What is important is to make sure that these different deployment philosophies do not force loss of interoperability.

The level of interoperability that IP accomplished fostered innovation at the application layer. Ralph Droms reinforced this message by saying: "Bright people will take a phone, build an application and connect it, with the appropriate security controls in place, to the things in my house in ways we have never thought about before. Otherwise, we are just building another telephone network."

3.1.2. Domain-Specific Stacks and Profiles

Imagine a building network scenario where a new light bulb is installed that should, out of the box without further configuration, interoperate with the already present light switch from a different vendor in the room. For many, this is the desired level of interoperability in the area of smart object design. To accomplish this level of interoperability, it is not sufficient to provide interoperability only at the network layer. Even running the same transport protocol and application-layer protocol (e.g., HTTP) is insufficient since both devices need to understand the semantics of the payloads for "Turn the light on" as well.

Standardizing the entire protocol stack for this specific "light switch / light bulb" scenario is possible. A possible stack would, for example, use IPv6 with a specific address configuration mechanism (such as stateless address autoconfiguration), a network access authentication security mechanism such as Protocol for carrying Authentication for Network Access (PANA) [RFC5191], a service discovery mechanism (e.g., multicast DNS with DNS-Based Service Discovery [DNS-SD]), an application-layer protocol, for example, Constrained Application Protocol (CoAP) [CoAP] (which uses UDP), and the syntax and semantic for the light on/off functionality.

As this list shows, there is already some amount of protocol functionality that has to be agreed on by various stakeholders to make this scenario work seamlessly. As we approach more complex protocol interactions, the functionality quickly becomes more complex: IPv4 and IPv6 on the network layer, various options at the transport layer (such as UDP, TCP, the Stream Control Transmission Protocol (SCTP) [RFC4960], and the Datagram Congestion Control Protocol (DCCP) [RFC4340]), and there are plenty of choices at the application layer with respect to communication protocols, data formats and data models. Different requirements have led to the development of a variety of communication protocols: client-server protocols in the style of the original HTTP, publish-subscribe protocols (like the Session Initiation Protocol (SIP) [RFC3261] or Extensible Messaging and Presence Protocol (XMPP) [RFC6121]), and store-and-forward messaging (borrowed from the delay tolerant

networking community). Along with the different application-layer communication protocols come various identity and security mechanisms.

With the smart object constraints, it feels natural to develop these stacks since each application domain (e.g., healthcare, smart grids, building networking) will have their unique requirements and their own community involved in the design process. How likely are these profiles going to be the right match for the future, specifically for the new innovations that will come? How many of these stacks are we going to have? Will the differences in the profiles purely be the result of different requirements coming from the individual application domains or will these mismatches reflect the spirit, understanding, and preferences of the community designing them? How many stacks will multipurpose devices have to implement?

Standardizing profiles independently for each application is not the only option. Another option is to let many different applications utilize a common foundation, i.e., a protocol stack that is implemented and utilized by every device. This, however, requires various application domains to be analyzed for their common characteristics and to identify requirements that are common across all of them. The level of difficulty for finding an agreement of how such a foundation stack should look depends on how many layers it covers and how lightweight it has to be.

From the discussions at the workshop, it was clear that the available options are not ideal and further discussions are needed.

3.1.3. Which Layer?

The end-to-end principle states that functionality should be put into the end points instead of into the networks. An additional recommendation, which is equally important, is to put functionality higher up in the protocol stack. While it is useful to make common functionality available as building blocks to higher layers, the wide range of requirements by different applications led to a model where lower layers provide only very basic functionality and more sophisticated features were made available by various applications. Still, there has been the desire to put application-layer functionality into the lower layers of the networking stack. A common belief is that performance benefits can be gained if functionality is placed at the lower layers of the protocol stack. This new functionality may be offered in the form of a gateway, which bridges different communication technologies, acts on behalf of other nodes, and offers more generic functionality (such as name-based routing and caching).

Two examples of functionality offered at the network layer and discussed during the workshops were location and name-based routing:

Location:

The notion of location gives each network node the understanding of proximity (e.g., 'I am a light bulb and in the same room as the light switch.'). Today, a router may provide information as to whether other nodes belong to the same subnet or they may provide location information (for example, in the form of GPS-based coordinates). However, providing a sense of the specific environment (e.g., a room, building, campus, etc.) is not possible without manual configuration. While it has been a desirable feature for many ubiquitous computing applications to know this type of information and to use it for richer application-layer interactions, widespread deployment has not happened yet.

Name-Based Routing:

With the work on recent "clean slate" architecture proposals, such as named data networking, flexible naming concepts are being researched to allow application developers to express their application-layer concepts in a way that they can be used natively by the underlying networking stack without translation. For example, Jeff Burke provided the example of his work in a theater with a distributed control system of technical equipment (such as projectors, dimmers, and light systems). Application developers name their equipment with human-readable identifiers, which may change after the end of a rehearsal, and address them using these names. These naming concepts based on variable-length strings raise questions regarding scalability.

The workshop participants were not able to come to an agreement about what functionality should be moved from the application layer to the network layer.

3.2. Sleeping Nodes

One limitation of smart objects is their available energy. To extend battery life, for example, of a watch battery or single AAA battery for months, these low-power devices have to sleep 99% to 99.5% of their time. For example, a light sensor may only wake up to check whether it is nighttime to turn on light bulbs. Most parts of the system, particularly communication components, are put into a sleeping state (e.g., WLAN radio interface) and selected components,

such as sensors, periodically check for relevant events and, if necessary, turn on other hardware modules. Every bit is precious, as is every round trip and every millisecond of radio activity.

Many IETF protocols are implicitly designed to be always on, i.e., the protocol implementation waits for and reacts to incoming messages, and may maintain session state (at various layers of the stack). These protocols work well when energy consumption is not a concern and when receiving and sending messages happen at a low cost. It should be understood that energy is consumed both in transmitting messages (and often more importantly) in keeping the receiver awake. Allowing devices to sleep most of the time preserves energy but creates challenges for protocol designs.

The inherent issue encountered by a stationary node resuming from sleep is that another node may have chosen the same address while the node was asleep. A number of steps have to be taken before sending a packet. A new address may have to be obtained, for example using the Dynamic Host Configuration Protocol (DHCP) or stateless address autoconfiguration. Optionally, Detecting Network Attachment (DNA) procedures (see [RFC4436] and [RFC6059]) can be used to shorten the setup time by noticing that the node is using the same default gateway.

The issue with a mobile node that is sleeping is that the node may have been moved to another network (e.g., a sleeping laptop being transported to a new environment) where on resumption it may discover that its address has become invalid.

The following design considerations should be taken into account when energy efficiency is a concern:

1. Rethink the Always-On Assumption

When designing a protocol that assumes that the nodes are always on, alternatives need to be considered. This could involve eliminating functionality (e.g., not implementing DNA or duplicate address detection) or considering the use of a sleep proxy. While sleep proxies (e.g., proxZzzy(TM) [proxZzzy]) enable periodic messages to be sent on behalf of sleeping nodes, this approach assumes that energy management constraints do not apply to the sleep proxy, which may not always be the case (e.g., if the entire network is deployed in the field without access to power). Yet another option is for devices to explicitly communicate sleep cycles so that they can only check for messages periodically (be it measured in milliseconds, seconds, or hours).

This is the approach taken in IEEE 802.11, which supports multiple energy conservation mechanisms designed to enable a station to spend a large fraction of the time sleeping.

2. Reduce Network Attachment Costs

As noted above, the procedures for obtaining an address and assuring its uniqueness can be costly. In a network where nodes spend a large fraction of the time sleeping, but are not necessarily mobile, are all of these procedures really necessary?

Can we find ways to reduce the number of protocol interactions without sacrificing correctness? The main focus of past investigations has been on IPv6 and ND, but other protocols do also deserve a deeper investigation, such as DNS and DHCP.

3. Avoid Verbose Protocols

Protocols involving multiple roundtrips and/or lengthy messages with verbose encodings (e.g., XML) are not always well-suited for constrained environments. Low-power nodes utilizing verbose protocols have to use more energy in sending messages and have to stay powered on for a longer period of time as they wait for the multi-roundtrip protocol exchanges to complete.

The best way to address these problems is to ensure that the simplest possible protocol exchanges are used when the protocols in question are designed. In some cases, alternative encoding formats and compression may also help.

4. Think about System-Wide Efficiency

While energy efficiency is critical for low-power devices running on batteries, it is also beneficial for other devices as well, including notebook computers, desktop computers, and smartphones. However, if the goal is energy efficiency as opposed to battery life extension for low-power devices, then it is important to consider the total energy consumption of all the elements of the system.

For example, consider energy consumption in a home environment. In these scenarios it is important to evaluate the energy usage of the entire system. A light bulb utilizing Internet technology described in this document may use less power but there is also the device that controls the bulb that may consume a lot of energy. If the goal is to reduce overall energy usage, then it is important to consider these two devices (and potentially many others) together.

3.3. Security

In the development of smart object applications, as with any other protocol application solution, security has to be considered early in the design process. As such, the recommendations currently provided to IETF protocol architects, such as RFC 3552 [RFC3552] and RFC 4101 [RFC4101], apply also to the smart object space.

While there are additional constraints, as described in Section 2, security has to be a mandatory part of the solution. The hope is that this will lead to implementations that provide security features, deployments that utilize them, and finally use of better security mechanisms. It is important to point out that the lack of direct user interaction will place hard requirements on deployment models, configuration mechanisms, and software upgrade / crypto-agility mechanisms.

Since many of the security mechanisms allow for customization, particularly with regard to the cryptographic primitives utilized, many believe that IETF security solutions are usable without modifications in a large part of the smart object domain. Others call for new work on cryptographic primitives that make use of a single primitive (such as the Advanced Encryption Standard (AES) [AES]) as a building block for all cryptographic functions. The benefit would be a smaller footprint of the overall solution, specifically with respect to hardware limitations (e.g., the hardware architecture of certain embedded devices prevents pipelining from being used). In the excitement for new work on optimizations of cryptographic primitives, other factors have to be taken into consideration that influence successful deployment, such as widespread support in libraries, as well as intellectual property rights (IPR). As an example of the latter aspect, the struggle of Elliptic Curve Cryptography (ECC)-based cryptographic algorithms [ECC] to find deployment can partially be attributed to its IPR situation. The reuse of libraries providing cryptographic functions is clearly an important way to use available memory resources in a more efficient way. To deal with the performance and footprint concerns, investigations into offloading certain resource-hungry functions to parties that possess more cryptographic power have been considered. For example, the ability to delegate certificate validation to servers has been standardized in the IETF before, for the support of the Online Certificate Status Protocol (OCSP) in the Internet Key Exchange protocol version 2 (IKEv2) and in Transport Layer Security (TLS); see [RFC4806] and [RFC5246], respectively.

Focusing only on the cryptographic primitives would be shortsighted; many would argue that this is the easy part of a smart object security solution. Key management and credential enrollment,

however, are considered a big challenge by many, particularly when usability requirements have to be taken into account. Another group of challenges concern privacy; with smart grids, for example, some have voiced concerns regarding the ability of third parties to keep track of an individual's energy consumption (and draw associated conclusions). As another example, it is easy to see how a scale that is connected to the Internet for uploading weight information to a social network could lead to privacy concerns. While security mechanisms that are used to offer protection of the communication between different parties also provide a certain degree of privacy protection, they are clearly not enough to address all concerns. Even with the best communication security and access control mechanisms in place, one still needs additional safeguards against the concerns mentioned in the examples.

While better deployment of security protocols on the entire Internet would be very desirable, practical considerations regarding usability and the incentives of the stakeholders involved have often lead to slower adoption.

3.4. Routing

A smart object network environment may also employ routers under similar constraints as the end devices. Currently two approaches to routing in these low-power and lossy networks are under consideration -- namely, mesh-under and route-over. The so-called "mesh-under" approach places routing functions at the link layer, and consequently all devices appear as immediate neighbors at the network layer. With the "route-over" approach, routing is done in the IP layer and not at all in the link layer. Each physical hop appears as a single IP hop (ignoring devices that just extend the physical range of signaling, such as repeaters). Routing in this context means running a routing protocol. The IPv6 Routing Protocol for Low-power and Lossy Networks (RPL) [RPL], for example, belongs to the route-over category.

From an architectural point of view there are several questions that arise from where routing is provided, for example:

- o Protocols often assume that link characteristics are predictable when communicating with any device attached to the same link. Latency, throughput, and reliability may vary considerably between different devices that are multiple link-layer hops away. What timeout should be used? What happens if a device is unreachable? In case of default router selection, two advertised routers may be a different number of hops away. Should a device have visibility into the path to make a decision, and what path characteristics would be useful to have?

- o Scoped message delivery to a neighboring IP hop (via link-local addressing) allows certain types of IP protocols to communicate with their immediate neighbors and to therefore scope their recipients. A link-local multicast message will be received by all nodes in the same IP link-local realm unless some special optimizations are provided by the link layer.
- o When path computations are done at the link layer as well as on the network layer, then what coordination needs to take place?

When multiple different link-layer technologies are involved in a network design, routing at layer 3 has to be provided in any case. [IoT-ARCH] talks about these tradeoffs between route-over and mesh-under in detail. Furthermore, those who decide about the deployment have to determine how to connect smart objects to the Internet infrastructure, and a number of wired and wireless technologies may be suitable for a specific deployment. Depending on the chosen technologies the above-mentioned mesh-under vs. route-over approach will have to be decided, and further decisions will have to be made about the choice of a specific routing protocol.

In 2008, the IETF formed the Routing Over Low power and Lossy networks (ROLL) working group to specify a routing solution for smart object environments. During its first year of existence, the working group studied routing requirements in detail (see [RFC5867], [RFC5826], [RFC5673], and [RFC5548]), and it worked on a protocol survey comparing a number of existing routing protocols, including Ad hoc On-Demand Distance Vector (AODV)-style protocols [RFC3561], against the identified requirements. The protocol survey [PROT-SURVEY] was inconclusive and abandoned without giving rise to publication of an RFC.

The ROLL WG concluded that a new routing protocol satisfying the documented requirements has to be developed and the work on RPL was started as the IETF routing protocol for smart object networks. Nevertheless, controversial discussions at the workshop about which routing protocols is best in a given environment are still ongoing. Thomas Clausen, for example, argued for using an AODV-like routing protocol in [Clausen].

4. Conclusions and Next Steps

The workshop allowed the participants to be exposed to interesting applications and their requirements (buildings, fountains, theater, etc.), to have discussions about radically different architectures and their issues (e.g., information centric networking), to look at existing technology from a new angle (sleeping nodes, energy consumption), to focus on some details of the protocol stack (neighbor discovery, security, routing) and to learn about implementation experience.

One goal of the workshop was to identify areas that require further investigation. The list below reflects the thoughts of the workshop participants as expressed on the day of the workshop. Note that the suggested items concern potential work by the IETF and the IRTF, and the order does not imply a particular preference.

Security:

A discussion of security is provided in Section 3.3. In general, security-related protocol exchanges and the required amount of computational resource requirements can contribute significantly to the overall processing. Therefore, it remains a challenge to accomplish performance improvements without sacrificing the overall security level, taking the usability of the entire system into consideration.

Another challenge is how to balance the security and performance needs of smart objects with the interoperability requirements of hosts on the Internet since different suites of algorithms tend to be chosen for these different environments. This involves trade-offs between performance on the smart objects versus end-to-end security. Solutions that mandate a "trusted" middlebox whose only role is to terminate security associations tend to be frowned upon from the security perspective, especially since end users of challenged devices (where those exist) are unlikely to understand the security consequences of such middleboxes.

The discussion concluded with the following recommendations:

- * Investigate usability in cryptographic protocol design with regard to credential management. As an example, the Bluetooth pairing mechanism was mentioned as a simple and yet reasonably secure method of introducing devices into a new environment. In fact, the IETF working group Credential and Provisioning (ENROLL) was established years ago to deal with residential

networks but was terminated prematurely due to lack of progress. The email archive still exists and is available [ENROLL] and may provide useful historical information.

- * Standardized authentication and key exchange mechanisms should be surveyed for suitability in smart object environments with respect to message size, computational performance, number of messages, code size, and main memory requirements. A starting point of such an investigation (in the case of IKEv2) was provided by Tero Kivinen with [MINIMAL-IKEv2], and a suitable venue for discussion could be the recently established Light-Weight Implementation Guidance (LWIG) working group [LWIG].
- * Research has to be done in the area of lightweight cryptographic primitives -- namely, block ciphers, stream ciphers, and cryptographic hash functions. It's worthwhile to mention the early work with the National Institute of Standards and Technology (NIST) new cryptographic hash algorithm 'SHA-3' competition [SHA3]. A suitable forum for discussion is the IRTF Crypto Forum Research Group (CFRG) [CFRG].

The difficulty and impact of choosing specialized algorithms for smart objects should not be underestimated. Issues that arise include the additional specification complexity (e.g., TLS already has hundreds of ciphersuites defined, most of which are unused in practice), the long latency in terms of roll out (many hosts are still using deprecated algorithms 5-10 years after those algorithms were deprecated), and the barriers that IPR-encumbered schemes present to widespread deployment. While research on this topic within CFRG and the cryptographic research community is a very worthwhile goal, any such algorithms will likely have to offer very significant benefits before they will be broadly adopted. 20% less CPU usage is unlikely to be a winning argument no matter what an algorithm inventor believes.

Energy Design Considerations:

One part of the workshop was focused on the discussion of energy implications for IETF protocol design with proposals being made about how to extend protocols to better support nodes that are mostly sleeping. Discussions are encouraged to take place on the RECIPE mailing list [RECIPE]. The workshop position paper [Wasserman] by Margaret Wasserman provides a good starting point for further investigations.

Information-/Content-Centric Networking:

Information/Content Centric Networking is about accessing named content, and a number of research projects have emerged around this theme. At this point in time, the work is not yet ready for standardization in the IETF. Instead, the formation of an IRTF research group has been proposed, and more details are available on the IRTF DISCUSS mailing list [irtf-discuss].

Architectural Guidelines:

Participants suggested developing an architectural write-up about what can be done with the IETF protocols we have today and how these different elements may be combined to offer an explanation for the broader community. This would be a task for the IAB. An example of prior work that serves as input is [RFC6272].

Network Management:

While this topic did not make it onto the workshop agenda, it was considered useful to start a discussion about how to implement network management protocols, such as Network Configuration Protocol (NETCONF) [RFC6241], on smart objects. The following position papers may be useful as a starting point for further discussions: [Ersue] and [Schoenwaelder]. An IETF draft document is also available: [SNMP-OPT].

Congestion Control:

When smart objects transmit sensor readings to some server on the Internet, these protocol interactions often carry a small amount of data and happen infrequently, although regularly. With the work on new application protocols, like CoAP [CoAP], the question of whether a congestion control mechanism should be used at the underlying transport protocol or by the application protocol itself arises. [CoAP-CC] is a starting point for CoAP, but further work is needed.

Data Models:

Standardization of application-layer protocols is important but does not ensure that, for example, a light switch and a light bulb are able to interact with each other. One area where participants saw the need for further work was in the area of data models. While prior IETF standardization work on, for example, location [GEOPRIV] can be reused, the question was raised where the IETF

should focus its standardization efforts since domain expertise in each area will be needed. As a potential example, energy pricing was discussed, with an example provided by [ENERGY-PRICING].

Building Networking:

Network architectures in residential as well as commercial buildings should take the presence of smart objects and dedicated subnetworks focusing on smart objects into account. A new working group, Home Networking (HOMENET) [HOMENET], was created after the workshop to look at this topic.

Discovery:

Additional extensions to developed discovery protocols, such as multicast DNS (mDNS), may be needed for the smart object environment. For instance, the HOMENET working group wants to extend current discovery protocols to work across multiple subnets. Smart object networks are often organized in separate subnets, so these extensions may be useful in that environment as well.

Routing:

Changing radio conditions and link fluctuation may lead to the need for renumbering. Workshop participants argued that work should be started on the multi-link subnetworks to mitigate this problem, i.e., to extend the notion of a subnet to be able to span multiple links. With the publication of RFC 4903 [RFC4903], the Internet Architecture Board had looked into this topic already and provided pros and cons.

The applicability of specific routing protocols designed for smart object networks needs further investigation. The ROLL working group is chartered with the task of constructing an applicability document for RPL, for instance.

5. Security Considerations

The workshop discussions covered a range of potential engineering activities, each with its own security considerations. As the IETF community begins to pursue specific avenues arising out of this workshop, addressing relevant security requirements will be crucial.

As described in this report, part of the agenda was focused on the discussion of security; see Section 3.3.

6. Acknowledgements

We would like to thank all the participants for their position papers. The authors of the accepted position papers were invited to the workshop.

Big thanks to Elwyn Davies for helping us to fix language bugs. We would also like to thank Andrei Robachevsky, Ulrich Herberg, Thomas Clausen, Bruce Nordman, Alissa Cooper, Dave Thaler, Bernard Aboba, and Henning Schulzrinne for their review comments.

Additionally, we would like to thank Ericsson and Nokia Siemens Networks for their financial support.

7. Informative References

- [AES] Wikipedia, "Advanced Encryption Standard", December 2011, <http://en.wikipedia.org/w/index.php?title=Advanced_Encryption_Standard&oldid=481153988>.
- [CFRG] McGrew (Chair), D., "IRTF Crypto Forum Research Group (CFRG)", June 2011, <<http://irtf.org/cfrg>>.
- [Clausen] Clausen, T. and U. Herberg, "Some Considerations on Routing in Particular and Lossy Environments", IAB Interconnecting Smart Objects with the Internet Workshop, Prague, Czech Republic, March 2011, <<http://www.iab.org/wp-content/IAB-uploads/2011/03/Herberg.pdf>>.
- [CoAP] Shelby, Z., Hartke, K., Bormann, C., and B. Frank, "Constrained Application Protocol (CoAP)", Work in Progress, October 2011.
- [CoAP-CC] Eggert, L., "Congestion Control for the Constrained Application Protocol (CoAP)", Work in Progress, January 2011.
- [DNS-SD] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", Work in Progress, December 2011.
- [Dolin] Dolin, B., "Application Communications Requirements for 'The Internet of Things'", IAB Interconnecting Smart Objects with the Internet Workshop, Prague, Czech Republic, March 2011, <<http://www.iab.org/wp-content/IAB-uploads/2011/03/Dolin.pdf>>.

- [ECC] Wikipedia, "Elliptic Curve Cryptography", December 2011, <http://en.wikipedia.org/w/index.php?title=Elliptic_curve_cryptography&oldid=480053167>.
- [ENERGY-PRICING] Jennings, C. and B. Nordman, "Communication of Energy Price Information", Work in Progress, July 2011.
- [ENROLL] "The ietf-enroll Archives", <<http://mailman.mit.edu/pipermail/ietf-enroll/>>.
- [Ersue] Ersue, M. and J. Korhonen, "Position Paper on 'Interconnecting Smart Objects with the Internet'", IAB Interconnecting Smart Objects with the Internet Workshop, Prague, Czech Republic, February 2011, <<http://www.iab.org/wp-content/IAB-uploads/2011/03/Ersue.pdf>>.
- [GEOPRIV] IETF, "Geographic Location/Privacy (geopriv) Working Group", <<http://datatracker.ietf.org/wg/geopriv/>>.
- [HOMENET] "Home Networking (homenet) Working Group", <<http://datatracker.ietf.org/wg/homenet>>.
- [IoT-ARCH] Hui, J. and J. Vasseur, "Routing Architecture in Low-Power and Lossy Networks (LLNs)", Work in Progress, March 2011.
- [LWIG] IETF, "Light-Weight Implementation Guidance (lwig) Working Group", June 2011, <<http://datatracker.ietf.org/wg/lwig/charter/>>.
- [MINIMAL-IKEv2] Kivinen, T., "Minimal IKEv2", Work in Progress, February 2011.
- [PROT-SURVEY] Tavakoli, A., Dawson-Haggerty, S., and P. Levis, "Overview of Existing Routing Protocols for Low Power and Lossy Networks", Work in Progress, April 2009.
- [RECIPE] "Reducing Energy Consumption with Internet Protocols Exploration (RECIPE) Mailing List", <<https://www.ietf.org/mailman/listinfo/recipe>>.
- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [RFC3561] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC 3561, July 2003.
- [RFC4101] Rescorla, E. and IAB, "Writing Protocol Models", RFC 4101, June 2005.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.
- [RFC4436] Aboba, B., Carlson, J., and S. Cheshire, "Detecting Network Attachment in IPv4 (DNav4)", RFC 4436, March 2006.
- [RFC4806] Myers, M. and H. Tschofenig, "Online Certificate Status Protocol (OCSP) Extensions to IKEv2", RFC 4806, February 2007.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, June 2007.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, April 2010.
- [RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.
- [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, November 2010.
- [RFC6121] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence", RFC 6121, March 2011.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC6272] Baker, F. and D. Meyer, "Internet Protocols for the Smart Grid", RFC 6272, June 2011.
- [RPL] Brandt, A., Vasseur, J., Hui, J., Pister, K., Thubert, P., Levis, P., Struik, R., Kelsey, R., Clausen, T., and T. Winter, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", Work in Progress, March 2011.

- [SHA3] NIST, "Cryptographic Hash Algorithm Competition", December 2010, <<http://csrc.nist.gov/groups/ST/hash/sha-3/index.html>>.
- [SNMP-OPT] Schoenwaelder, J., Mukhtar, H., Joo, S., and K. Kim, "SNMP Optimizations for Constrained Devices", Work in Progress, October 2010.
- [Schoenwaelder] Schoenwaelder, J., Tsou, T., and B. Sarikaya, "Protocol Profiles for Constrained Devices", IAB Interconnecting Smart Objects with the Internet Workshop, Prague, Czech Republic, February 2011, <<http://www.iab.org/wp-content/IAB-uploads/2011/03/Schoenwaelder.pdf>>.
- [SmartGrid] Wikipedia, "Smart Grid", December 2011, <http://en.wikipedia.org/w/index.php?title=Smart_grid&oldid=479750548>.
- [Tussle] Clark, D., Wroclawski, J., Sollins, K., and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", In Proc. ACM SIGCOMM, 2002, <<http://conferences.sigcomm.org/sigcomm/2002/papers/tussle.html>>.
- [Wasserman] Wasserman, M., "It's Not Easy Being 'Green'", IAB Interconnecting Smart Objects with the Internet Workshop, Prague, Czech Republic, March 2011, <<http://www.iab.org/wp-content/IAB-uploads/2011/03/Wasserman.pdf>>.
- [irtf-discuss] Ohlman, B., "Draft ICN RG Charter", message to IRTF DISCUSS Mailing List, May 2011, <<http://www.ietf.org/mail-archive/web/irtf-discuss/current/msg00041.html>>.
- [proxZzzy] ECMA, "proxZzzy(TM) for sleeping hosts", Standard ECMA-393, February 2010, <<http://www.ecma-international.org/publications/standards/Ecma-393.htm>>.

Appendix A. Program Committee

The following persons are responsible for the organization of the associated workshop and are responsible also for this event: Jari Arkko, Hannes Tschofenig, Bernard Aboba, Carsten Bormann, David Culler, Lars Eggert, JP. Vasseur, Stewart Bryant, Adrian Farrel, Ralph Droms, Geoffrey Mulligan, Alexey Melnikov, Peter Saint-Andre, Marcelo Bagnulo, Zach Shelby, Isidro Ballesteros Laso, Fred Baker, Cullen Jennings, Manfred Hauswirth, and Lukas Kencl.

Appendix B. Workshop Materials

Main page:

<http://www.iab.org/about/workshops/smartobjects/>

Position papers:

<http://www.iab.org/about/workshops/smartobjects/papers/>

Slides:

<http://www.iab.org/about/workshops/smartobjects/agenda/>

Minutes:

<http://www.iab.org/activities/workshops/smartobjects/smartobjectworkshopmeetingminutes/>

Appendix C. Accepted Position Papers

1. "Deployment Experience with Low Power Lossy Wireless Sensor Networks" by C. Adjih, E. Baccelli, P. Jacquet, P. Minet, M. Philipp, and G. Wittenburg
2. "CoAP improvements to meet embedded device hardware constraints" by Davide Barbieri
3. "Unified Device Networking Protocols for Smart Objects" by Daniel Barisic and Anton Pfefferseder
4. "Simplified neighbour cache implementation in RPL/6LoWPAN" by Dominique Barthel
5. "Home Control in a consumer's perspective" by Anders Brandt
6. "Authoring Place-based Experiences with an Internet of Things: Tussles of Expressive, Operational, and Participatory Goals" by Jeff Burke

7. "A Dynamic Distributed Federated Approach for the Internet of Things" by Diego Casado Mansilla, Juan Ramon Velasco Perez, and Mario Lopez-Ramos
8. "Quickly interoperable Internet of Things using simple transparent gateways" by Angelo P. Castellani, Salvatore Loreto, Nicola Bui, and Michele Zorzi
9. "Position Paper of the Brno University of Technology Department of Telecommunications" by Vladimir Cervenka, Lubomir Mraz, Milan Simek, and Karel Pavlata
10. "Secure Access to IOT Network: An Application-based Group Key Approach" by Samita Chakrabarti and Wassim Haddad
11. "Domain-Insulated Autonomous Network Architecture (DIANA)" by W. Chun
12. "Yet Another Definition on Name, Address, ID, and Locator (YANAIL)" by W. Chun
13. "The Challenge of Mobility in Low Power Networks" by E. Davies
14. "If the routing protocol is so smart, why should neighbour discovery be so dumb?" by Nicolas Dejean
15. "Making Smart Objects IPv6 Ready: Where are we?" by M. Durvy and M. Valente
16. "Position Paper on 'Interconnecting Smart Objects with the Internet'" by Mehmet Ersue and Jouni Korhonen
17. "The Real-time Enterprise: IoT-enabled Business Processes" by Stephan Haller and Carsten Magerkurth
18. "Making Internet-Connected Objects readily useful" by Manfred Hauswirth, Dennis Pfisterer, and Stefan Decker
19. "Some Considerations on Routing in Particular and Lossy Environments" by Thomas Clausen and Ulrich Herberg
20. "A Security Protocol Adaptation Layer for the IP-based Internet of Things" by Rene Hummen, Tobias Heer, and Klaus Wehrle
21. "Simplified SIP Approach for the Smart Object" by Ryota Ishibashi, Takumi Ohba, Arata Koike, and Akira Kurokawa

22. "Mobility support for the small and smart Future Internet devices" by Antonio J. Jara and Antonio F. G. Skarmeta
23. "The Need for Efficient Reliable Multicast in Smart Grid Networks" by J. Jetcheva, D. Popa, M.G. Stuber, and H. Van Wyk
24. "Lightweight Cryptography for the Internet of Things" by Masanobu Katagi and Shiho Moriai
25. "Thoughts on Reliability in the Internet of Things" by James Kempf, Jari Arkko, Neda Beheshti, and Kiran Yedavalli
26. "IKEv2 and Smart Objects" by Tero Kivinen
27. "Position Paper on Thing Name Service (TNS) for the Internet of Things (IoT)" by Ning Kong and Shuo Shen
28. "Connecting Smart Objects to Wireless WANS" by Suresh Krishnan
29. "Towards an Information-Centric Internet with more Things" by D. Kutscher and S. Farrell
30. "Application of 6LoWPAN for the Real-Time Positioning of Manufacturing Assets" by Jaacan Martinez and Jose L. M. Lastra
31. "Lighting interface to wireless network" by Jaroslav Meduna
32. "Social-driven Internet of Connected Objects" by Paulo Mendes
33. "Protocols should go forward that are required by non IP based protocols" by Tsuyoshi Momose
34. "Web services for Wireless Sensor and Actuator Networks" by Guido Moritz
35. "Connecting BT-LE sensors to the Internet using Ipv6" by Markus Isomaki, Kanji Kerai, Jari Mutikainen, Johanna Nieminen, Basavaraj Patil, Teemu Savolainen, and Zach Shelby
36. "Building Networks" by Bruce Nordman
37. "Issues and Challenges in Provisioning Keys to Smart Objects" by Yoshihiro Ohba and Subir Das
38. "Challenges and Solutions of Secure Smart Environments" by Eila Ovaska and Antti Evesti

39. "Research Experiences about Internetworking Mechanisms to Integrate Embedded Wireless Networks into Traditional Networks" by Jose F. Martinez, Ivan Corredor, and Miguel S. Familiar
40. "Scalable Video Coding in Networked Environment" by Naeem Ramzan, Tomas Piatrik, and Ebroul Izquierdo
41. "Challenges in Opportunistic Networking" by Mikko Pitkaenen and Teemu Kaerkkaeinen
42. "Position Statement" by Neeli R. Prasad and Sateesh Addepalli
43. "A Gateway Architecture for Interconnecting Smart Objects to the Internet" by Akbar Rahman, Dorothy Gellert, Dale Seed
44. "Routing Challenges and Directions for Smart Objects in Future Internet of Things" by T. A. Ramrekha, E. Panaousis, and C. Politis
45. "6LoWPAN Extension for IPsec" by Shahid Raza, Thiemo Voigt, and Utz Roedig
46. "Connected Vehicle as Smart Object(s)" by Ryuji Wakikawa
47. "Problem and possible approach of development and operation of Smart Objects" by Shoichi Sakane
48. "Connecting Passive RFID Tags to the Internet of Things" by Sandra Dominikus and Joern-Marc Schmidt
49. "Protocol Profiles for Constrained Devices" by Juergen Schoenwaelde, Tina Tsou, and Behcet Sarikaya
50. "Multihoming for Sensor Networks" by J. Soininen
51. "Internet Objects for Building Control" by Peter van der Stok and Nicolas Riou
52. "Optimal information placement for Smart Objects" by Shigeya Suzuki
53. "Multi-National Initiative for Cloud Computing in Health Care (MUNICH)" by Christoph Thuemmler
54. "The time of the hourglass has elapsed" by Laurent Toutain, Nicolas Montavont, and Dominique Barthel

55. "Sensor and Actuator Resource Architecture" by Vlasios Tsiatsis, Jan Hoeller, and Richard Gold
56. "IT'S NOT EASY BEING 'GREEN'" by Margaret Wasserman
57. "Trustworthy Wireless Industrial Sensor Networks" by Markus Wehner, Thomas Bartzsch, Dirk Burggraf, Sven Zeisberg, Alexis Olivereau, and Oualha Nouha
58. "Interconnecting smart objects through an overlay networking architecture" by Anastasios Zafeiropoulos, Athanassios Liakopoulos and Panagiotis Gouvas
59. "Building trust among Virtual Interconnecting Smart Objects in the Future Internet" by Theodore Zahariadic, Helen C. Leligou, Panagiotis Trakadas, and Mischa Dohler
60. "Experience and Challenges of Integrating Smart Devices with the Mobile Internet" by Zhen Cao and Hui Deng
61. "The 'mesh-under' versus 'route over' debate in IP Smart Objects Networks" by JP. Vasseur and Jonathan Hui
62. "Identification and Authentication of IoT Devices" by Alper Yegin
63. "Security Challenges For the Internet of Things" by Tim Polk and Sean Turner
64. "Application Communications Requirements for 'The Internet of Things'" by Bob Dolin
65. "Interoperability Concerns in the Internet of Things" by Jari Arkko
66. "Privacy in Ubiquitous Computing" by Ivan Gudymenko and Katrin Borcea-Pfitzmann
67. "The 10 Laws of Smart Object Security Design" by Hannes Tschofenig and Bernard Aboba
68. "Position Paper on 'In-Network Object Cloud' Architecture and Design Goals" by Alex Galis and Stuart Clayman
69. "Temptations and Difficulties of Protocols for Smart Objects and the Internet" by Alexandru Petrescu

70. "The Internet of Things - Challenge for a New Architecture from Problems" by Gyu Myoung Lee and Noel Crespi

71. "Garrulity and Fluff" by Carsten Bormann and Klaus Hartke

Appendix D. Workshop Participants

We would like to thank the following workshop participants for attending the workshop:

Adrian Farrel
Akbar Rahman
Alissa Cooper
Alper Yegin
Anastasios Zafeiropoulos
Anders Brandt
Angelo P. Castellani
Antonio F. G. Skarmeta
Antonio Jara
Arvind Ramrekha
Behcet Sarikaya
Bernard Aboba
Bruce Nordman
Carsten Bormann
Cullen Jennings
Daniel Barisic
Dave Thaler
Davide Barbieri
Diego Casado Mansilla
Dirk Kutscher
Dominique Barthel
Dorothy Gellert
Elwyn Davis
Emmanuel Baccelli
Fred Baker
Guido Moritz
Gyu Myoung Lee
Hannes Tschofenig
Hui Deng
Ivan Gudymenko
Jaacan Martinez
Jari Arkko
Jaroslav Meduna
Jeff Burke
Johanna Nieminen
Jonathan Hui
Jonne Soininen
Jouni Korhonen

JP. Vasseur
Karel Pavlata
Klaus Hartke
Lars Eggert
Laura Gheorghe
Laurent Toutain
Lukas Kencl
Marcelo Bagnulo
Marco Valente
Margaret Wasserman
Markus Isomaki
Markus Wehner
Masanobu Katagi
Mathilde Durvy
Mehmet Ersue
Mikko Pitkaenen
Milan Simek
Neeli R. Prasad
Nicolas Dejean
Ning Kong
Pascal Thubert
Paulo Mendes
Pete Resnick
Peter van der Stok
Ralph Droms
Rene Hummen
Ross Callon
Ruediger Martin
Russ Housley
Ryota Ishibashi
Ryuji Wakikawa
Samita Chakrabarti
Sandra Dominikus
Sean Shen
Sean Turner
Shahid Raza
Shoichi Sakane
Spencer Dawkins
Stephan Haller
Stephen Farrell
Stewart Bryant
Subir Das
Suresh Krishnan
Tea Sang Choi
Tero Kivinen
Theodore Zahariadis
Thomas Clausen
Tim Polk

Tina Tsou
Tsuyoshi Momose
Vladimir Cervenka
Wassim Haddad
Woojik Chun
Zach Shelby

Appendix E. IAB Members at the Time of Approval

Bernard Aboba
Ross Callon
Alissa Cooper
Spencer Dawkins
Joel Halpern
Russ Housley
David Kessens
Olaf Kolkman
Danny McPherson
Jon Peterson
Andrei Robachevsky
Dave Thaler
Hannes Tschofenig

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
EMail: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Jari Arkko
Ericsson
Jorvas 02420
Finland

EMail: jari.arkko@piuha.net

Internet Architecture Board

EMail: iab@iab.org