



IP for Smart Objects

Internet Protocol for Smart Objects (IPSO) Alliance

White paper #1

Version 1.1

Adam Dunkels, PhD, Senior scientist, Swedish Institute of Computer Science
JP Vasseur, Distinguished Engineer, Cisco Systems

July 2010 (Original version September 2008)

Executive Summary

The emerging application space for smart objects requires scalable and interoperable communication mechanisms that support future innovation as the application space grows. IP has proven itself a long-lived, stable, and highly scalable communication technology that supports both a wide range of applications, devices, and underlying communication technologies. The IP stack is lightweight (see [IPSO-WP2]) and runs on tiny, battery operated embedded devices. IP therefore has all the qualities to make “The Internet of Things” a reality, connecting billions of communicating devices (see [IP-Smart-Objects]).

Introduction

Smart objects are small computers with a sensor or actuator and a communication device, embedded in objects such as thermometers, car engines, light switches, and industry machinery. Smart objects enable a wide range of applications in areas such as home automation, building automation, factory monitoring, smart cities, structural health management systems, smart grid and energy management, and transportation. Until recently, smart objects were realized with limited communication capabilities, such as RFID tags, but the new generation of devices has bidirectional wireless communication and sensors that provide real-time data such as temperature, pressure, vibrations, and energy measurement. Smart objects can be battery-operated, but not always, and typically have three components: a CPU (8-, 16- or 32-bit micro-controller), memory (a few tens of kilobytes) and a low-power wireless communication device (from a few kilobits/s to a few hundreds of kilobits/s). The size is small and the price is low: a few square mm and few dollars.

The technical development in low-cost sensors and actuators combined with low-power communication technologies such as IEEE 802.15.4, low-power WiFi, and power line communication has been rapid (IEEE 1901.2, ITU G.hnem). Nevertheless, the emergence of smart object applications has not been as fast because the large number of proprietary or semi-closed systems has led to partial and non-interoperable solutions.

The current situation for smart objects is similar to what computer networks looked like about two decades ago: islands of computers communicating with their own protocol, for example SNA, IPX, and Vines, interconnected by complex multi-protocol gateways. Subsequently, these architectures evolved to IP-based tunneling mechanisms such as DLSw or XOT. Today, these networks operate on fully end-to-end IP-based architectures.

Many of today's non-IP-based sensor architectures are evolving toward a protocol-translation gateway model, similar to the path computer networks went through before quickly moving to fully IP-based architectures. *Have we not learnt from the past?* Protocol gateways are inherently complex to design, manage, and deploy. The network fragmentation leads to non-efficient networks because of inconsistent routing, QoS, transport and network recovery techniques.

By contrast, end-to-end IP architectures are widely accepted as the only alternative to design scalable and efficient networks of large numbers of communicating devices.

The Internet of Things: IP for Smart Objects

To support the large number of emerging applications for smart objects, the underlying networking technology must be inherently scalable, interoperable, and have a solid standardization base to support future innovation as the application space grows.

IP has proven itself a long-lived, stable, and highly scalable communication technology that supports both a wide range of application, a wide range of devices, and a wide range of underlying communication technologies. The layered architecture of IP

provides a high level of flexibility and innovation. IP already supports a plethora of applications, such as email, the World Wide Web, Internet telephony, video streaming, and collaborative tools. Over the past 20 years, IP has evolved to support new mechanisms for high availability, enhanced security, support of Quality of Service (QoS), real-time transport, and Virtual Private Networks (VPNs).

IP has a long history as a communication mechanism for general-purpose PC computers and network servers. It was therefore long believed that IP was too heavy weight to run on highly constrained devices. Several recent lightweight IP stacks have shown, however, that lightweight IP stacks can be designed to meet the requirements of light footprint devices with a few kilobytes of RAM and ROM, limited processing power and severe energy constraints.

IP provides standardized, lightweight, and platform-independent network access to smart objects and other embedded networked devices. The use of IP makes devices accessible from anywhere and from anything; general-purpose PC computers, cell phones, PDAs as well as database servers and other automated equipment such as a temperature sensor or a light bulb.

IP runs over virtually any underlying communication technology, ranging from high-speed wired Ethernet links to low-power 802.15.4 radios, 802.11 (WiFi) and low power PLC technologies. For long-haul communication, IP data is readily transported through encrypted channels over the global Internet.

IP: and Open and Flexible Standard

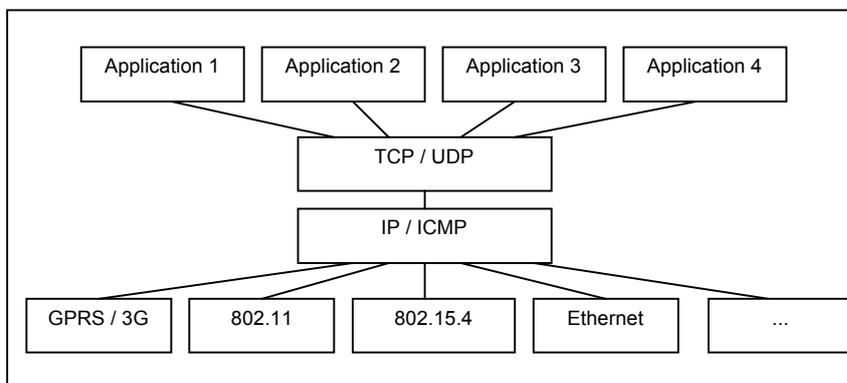


Figure 1. The layered IP architecture,

IP is based on a layered architecture as shown in Figure 1. IP runs over a wide variety of Physical and MAC layers and is therefore media-independent, ranging from low-power radios such as 802.15.4, 802.11 (WiFi), low power PLC technologies through long-range radio technology such as GPRS and 3G, to high-speed wired Ethernet, and multi-gigabit links.

The IP layer is responsible for packet delivery across the network: addressing of hosts, packet forwarding, and routing. Every IP packet carries the address of the destination, and

the IP network forwards the packet through the network until it reaches its destination according to the paths computed by the routing engine (dynamically recomputed upon network failures). IP also supports differentiated services where IP packets can support a range of Class of Service. IP packets are marked (colored) typically at the edge of the network and if congestion occurs along the forwarding path differentiated forwarding treatments such as queuing and congestion avoidance mechanisms are triggered according to the required QoS

ICMP does error reporting on behalf of IP. In IPv6, ICMP is also used for address auto configuration and neighbor discovery.

UDP provides a best-effort datagram service, where applications send chunks of data to other applications, with no delivery guaranty and flow control. Applications running on top of UDP therefore must detect and recover from lost datagrams, if 100% reliability is necessary. UDP is typically used for data where immediate delivery is more important than 100% reliability, such as IP telephony and real-time data streaming.

TCP provides a reliable packet transport layer. The data is reliably delivered, as long as there is a network path between the endpoints. TCP support sophisticated flow control mechanisms thanks to adaptive windowing techniques.

The IP packets consist of application data and a set of headers. The packet headers contain protocol data such as addresses, sequence numbers, and flags. Each protocol in the stack adds its own header to the packets.

For low-power and low-speed links, the headers can be compressed in a way that avoids the transmission of redundant bits of header data. A recent example of how header compression enables low-power operation is the 6lowpan effort which provides IPv6 transport over a low-power 802.15.4 radio. 6lowpan compresses the standard 48 bytes of the IPv6 and UDP headers down to 6 bytes for the common case (see [IP-HC]).

IP is Open

The Internet Engineering Task Force (IETF) is the open recognized International Standards Organization (ISO) in charge of standardizing the IP protocol. The IETF was formed in 1986 and is organized into working groups, each with a specific charter and set of milestones. The working groups are organized into several areas, such as routing, transport, and security.

The IP protocol suite is standardized in the form of documents, called RFCs. Each RFC document specifies a part of the protocol suite. For example, RFC791 specifies IPv4, RFC792 specifies UDP, RFC793 specifies TCP, and RFC2460 specifies IPv6.

Three IETF working group are currently standardizing IP protocols for smart objects: 6lowpan (IPv6 over IEEE 802.15.4 – see [6LoWPAN]), ROLL (Routing Over Low power and Lossy networks – see [ROLL]) that specifies a new routing protocol designed for IP smart objects networks ([RPL]) and CoRE ([CoRE]) devoted to defining a framework for resource-oriented applications intended to run on constrained IP networks.

Note that these three IETF working groups are exclusively relying on IP version 6, the successor of IP version 4, which is fully compliant with the IP architecture.

IP is Lightweight

IP used to be believed to be heavyweight, but the small footprint of many recent lightweight IP stacks has successfully refuted this belief. A small memory footprint is essential for operation on the low-power, low-cost microcontrollers of smart objects. The available memory of such microcontrollers typically is on the order of a few kilobytes. Power budgets are constrained, requiring low-power communication solutions.

Memory-efficient implementations of the IP stack show that IP can successfully work in as little as a few kilobytes of RAM, and require less than 10 kilobytes of ROM. Figure 2 shows the memory footprint of five embedded TCP/IP stacks: the open source uIP stack from the Contiki operating system, one commercially available TinyOS-based IPv6 stack, the commercially available NanoStack, and the open source lwIP stack. Their footprint is around 10 kilobytes, except for lwIP that is around 20 kilobytes. Recent implementations of RPL (the routing protocol) show that the requirements in terms of RAM and ROM have been increased by a few Kbytes, which is far below the limit imposed by current hardware.

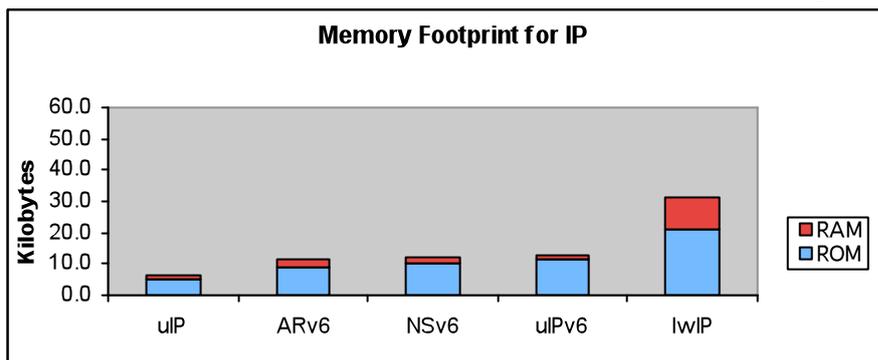


Figure 2. Memory footprint for five embedded TCP/IP stacks..

For power-constrained devices, recent standardization work has made IP power-efficient enough to run over sub-milliwatt radio links such as 802.15.4. Such low-power operation enables years of lifetime on typical AA batteries, even for multi-hop routing nodes.

IP is Versatile

IP is used for a wide range of applications, on a wide range of devices, and over a wide range of underlying network technologies. IP provides an open platform that allows for future innovation as the application space for smart objects evolves beyond today's visions.

IP supports virtually any type of application, including low-data rate applications such as remote device control delay-sensitive applications such as Internet telephony, large-size bulk data transfer such as file downloads, and high-data rate with stringent QoS requirements applications such as high definition TV. This diversity is enabled by the well-tested and well-engineered design of IP and its flexible, layered architecture.

Virtually any device supports IP networking, from high-end server clusters and high-speed disk racks, to cell phones and low-cost embedded devices. Although IP was originally designed for communication between computers, the flexibility inherent to IP has made it possible to interconnect very different types of devices and support a wide variety of applications.

IP is Ubiquitous

IP is available in most, if not all, operating systems for general-purpose computers and servers, and there is an ever-growing body of software available for IP networking for embedded systems. Both commercially licensed and open source implementations are generally available: general-purpose operating systems such as Microsoft Windows and Linux or micro-controller operating systems such as Contiki, TinyOS, and FreeRTOS. Most software packages also provide the necessary device drivers for the underlying communication hardware.

Most, if not all, networks provide IP access. It is possible today to get low-cost IP access both through wired connections such as DSL and through wireless offerings over 3G wireless networks. Last-hop IP access through technologies such as 802.11 (WiFi) is also available at hotspots.

IP is Scalable

With the global Internet, IP has proven itself to be inherently scalable. No other networking technology has ever been deployed and tested at such an immense scale and with such a large number of devices. As smart objects will connect an even larger number of devices than that of the existing Internet, scalability is a primary concern.

The next generation Internet protocol, IPv6, expands the address space of IP to 2^{128} . Such a large address space has been said to be enough to provide every grain of sand on the planet with an IP address. Specific protocols such as RPL in charge of routing in IP smart object networks have been designed to handle an extremely large number of devices.

IP is Manageable

Management in IP networks is done with a suite of well-understood network management protocols and mechanisms, for which a wide range of commercial and open source toolsets are available. Within the IP architecture, naming is done with the DNS protocol. Dynamic address assignment and network setup is done with the DHCP protocol. Network management is done with the SNMP protocol. And DNS, DHCP, SNMP are few a few examples of protocols that smart objects can reuse at no cost. Leveraging existing mechanisms, IP for smart objects do not need to reinvent the wheel in each of these instances.

IP is Stable

The IP standard has existed for nearly 30 years. Although the standard has been updated several times over the years, its foundation as a packet-based communication technology

has remained firm. Because IP forms the basis of the Internet, IP will continue to exist well into the future.

Because of its prevalence, IP has a large, well-established knowledge base. Network administrators, network architects, and developers have learned IP over decades of training and education.

IP is End-to-End

IP provides end-to-end communication between devices, without intermediate protocol translation gateways. Protocol gateways are inherently complex to design, manage, and deploy. The objective of a gateway is to translate or map between two or more protocols. Such translation, however, typically requires significant semantic and functional translation for the protocols to work together. Mechanisms on both sides usually differ significantly, thus requiring the adoption of a “least common denominator” approach that leads to non-efficient networks because of inconsistent routing, QoS, transport and network recovery techniques. With the end-to-end architecture of IP, there are no protocol translation gateways involved.

With the IP end-to-end architecture, there is no single point of failure. Intermediate routers may fail, but the end-to-end communication will choose alternate paths through the network. In contrast, if a protocol translation gateway fails, the entire network fails.

In the IP architecture, protocols can change without affecting the underlying network. Routers operate independently of the protocols running over them. In contrast, a protocol translation gateway needs to be updated every time a protocol changes, no matter how small the change.

With the success of today’s global Internet, the end-to-end architecture of IP has proven itself scalable, stable, and efficient. For the future Internet of things, scalability, stability, and efficiency is even more important than ever. IP therefore is the future-proof choice for the Internet of things.

Conclusions

Smart objects enable a wide range of applications that will improve our lives in many areas such as energy management, healthcare, and safety. The recent progress in low-cost embedded devices is about to make the Internet of Things a reality. For this to come true, we must learn from the lessons of the past and adopt a flexible, scalable, efficient and open based networking technology. IP has proven itself to fulfill these requirements and it is now a fact that IP can meet the strict requirements of highly constrained smart object networks.

References

[6LoWPAN] IETF 6LoWPAN Working Group:

<http://datatracker.ietf.org/wg/6lowpan/charter/>

[CoRE] IETF CoRE Working Group: <http://datatracker.ietf.org/wg/core/charter/>

[IPSO-WP2] IPSO White Paper #2: “Lightweight IPv6 Stacks for Smart Objects: the Experience of Three Independent and Interoperable Implementations”

[IP-HC] J. Hui, P. Thubert, “Compression Format for IPv6 Datagrams in 6LoWPAN Networks”, draft-ietf-6lowpan-hc-07, work in progress.

[IP-Smart-Objects] “Interconnecting Smart Objects with IP: The Next Internet”, JP Vasseur and Adam Dunkels, Morgan Kaufmann, June 2010.

[ROLL] IETF ROLL Working Group: <http://datatracker.ietf.org/wg/roll/charter/>

[RPL] T Winter (Ed), P. Thuber (Ed), JP. Vasseur, T. Clausen, P. Levis, R. Kelsey, J. Hui, K. Pister, A. Brandt, R. Struik, “RPL: IPv6 Routing Protocol for Low power and Lossy Networks”, draft-ietf-roll-rpl-10, work in progress.