Autumn 2016 Volume 13, Number 3

iHOMES and BUILDINGS

THE MAGAZINE OF THE CONTINENTAL AUTOMATED BUILDINGS ASSOCIATION

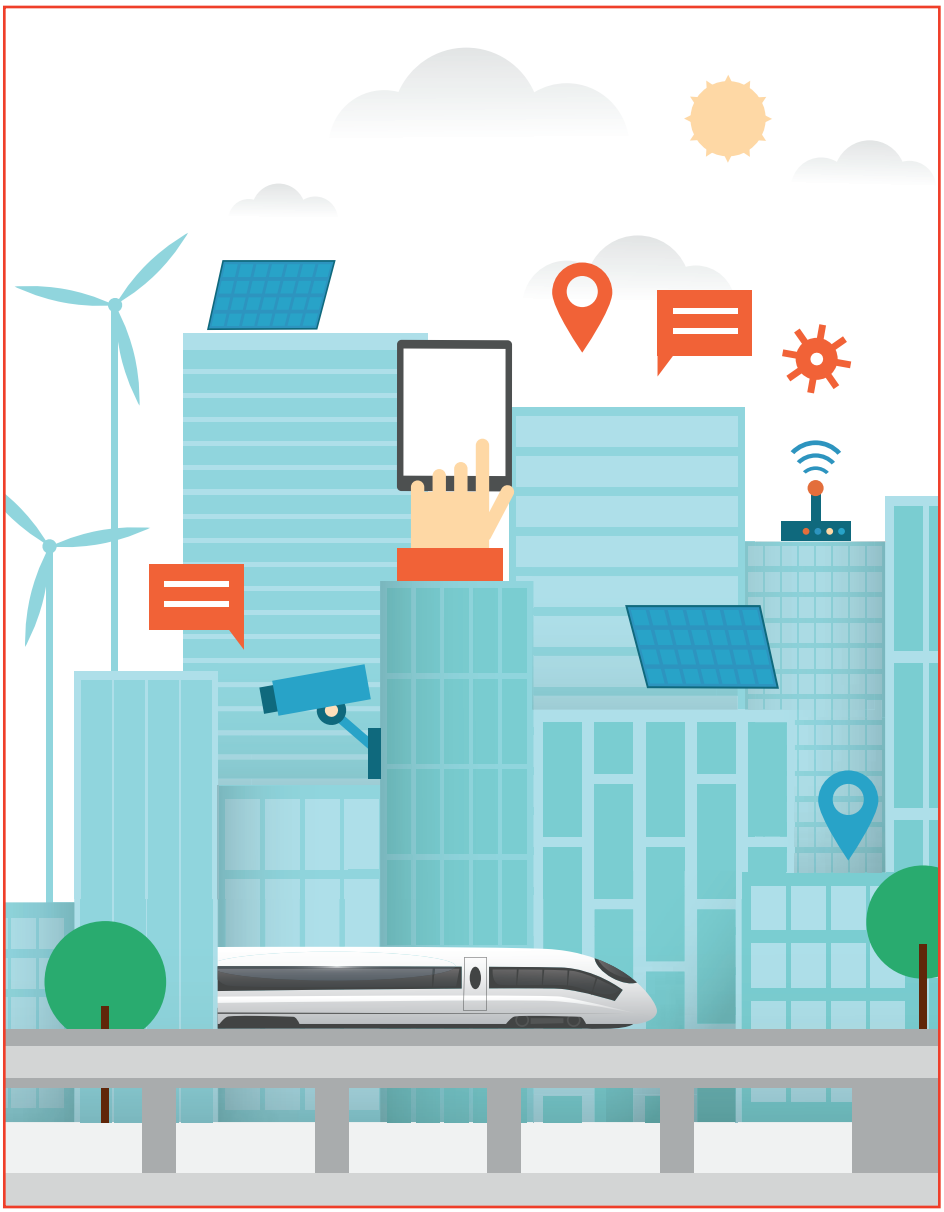

CABA™ Connect to what's next™

www.caba.org/ihomesandbuildings

iHOMES and BUILDINGS

THE MAGAZINE OF THE CONTINENTAL AUTOMATED BUILDINGS ASSOCIATION

Autumn 2016, Volume 13, Number 3

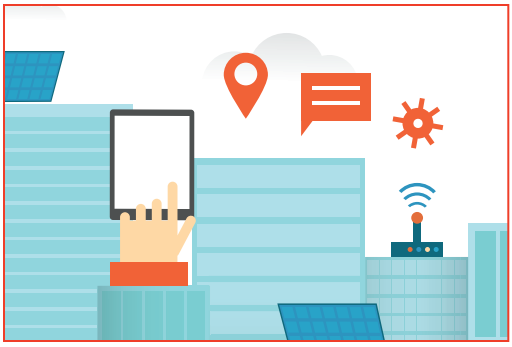

Contents

Features

Columns

Departments

Ken Wacks' Perspectives

Avoiding Chaos with the Internet of Things

By Ken Wacks

In the spring 2016 issue of *iHomes and Buildings* I explained that the "Internet of Things" (IoT) describes machine-to-machine communications. In home and building systems these devices are typically sensors, actuators, controllers, and user interfaces.

Some IoT proponents propose that all network devices communicate via the Internet Protocol (IP). However, in practice home and building devices communicate via a local network using a variety of non-IP communication protocols specialized for the operating environment, speed, and product cost-targets.

In this paper, I explain how a world of IP in all devices might operate and consider the benefits and risks. Key risks impact data security and customer privacy resulting in chaos if boundaries and priorities are not established.

Internet addresses

To appreciate a world of interconnected devices all using IP where each device could communicate with every other IoT device, we need to start with some fundamentals about Internet addressing. Every node on the public Internet is assigned a universally unique address.

Since the creation of the Internet, this address (called the IP address) has been encoded as 32 bits (1s and 0s). Rather than writing a string of 32 ones and zeros, four groups of eight bits each are expressed as four decimal numbers ranging from 0 to 255, such as 96.230.106.20. A 32-bit address accommodates a theoretical maximum of 2^{32} (4,294,967,296) unique addresses. This addressing scheme is called IP version 4 (IPv4).

Each country is assigned a pool of IPv4 addresses for government use and for Internet Service Providers (ISPs, usually cable or telephone companies). ISPs then assign IP addresses to customers from the national pools. Canada has about 80 million addresses; the U.S. has about 1.5 billion. The smallest pool is 256 IP addresses for the island of Saint Lucia. Some of the founding institutions of the Internet and early participants were each assigned 2^{24} (16,777,216) IPv4 addresses when IP addresses seemed limitless. Companies and universities such as Apple Computer, General Electric, Hewlett-Packard, IBM, and MIT still own such large blocks of addresses.

IP addresses are traditionally assigned to computers that request data (called clients) and computers (called servers) that retrieve and deliver the requested data from a repository. Typically these data constitute the text and graphics for a Web page.

For convenience, a scheme for assigning text-based names (called URLs – Uniform Resource Locators) was developed so users could find Web information with descriptive names such as www.caba.org. A worldwide network of servers forms the Domain Name System (DNS) that translates URLs into IP addresses.

Port addresses

To facilitate and expedite Internet services such as browsing and email, messages are associated with specific services via port numbers included in the message header data. A port number facilitates the dispatch of a message at a client and server to the appropriate process software for the intended service. The Internet messaging protocol (TCP[1] or UDP[2]) assigns one or more 16-bit port numbers to each service, which is written following the IP address as a number preceded by a colon: 96.230.106.20:80 (Port 80). Of the 2^{16} (65,536) port addresses available, a few thousand are pre-assigned for specific services such as browsing, mail, file transfers, and manufacturer-specific applications.

All traffic on the Internet consists of digital data messages (a string of 0s and 1s called bits, meaning "binary digits"). An Internet message is divided into packets (like

1 TCP = Transmission Control Protocol (used for web browsing and email)

2 UDP = User Datagram Protocol (used for time-sensitive applications such as real-time audio (VoIP) and video (IPTV))

telegrams) labeled with the sender IP address (called the Source Address) and the recipient IP address (called the Destination Address). Each packet may take a different path from the sender through various routers to the recipient where the packets are reassembled into the message, which is then passed to the assigned port for processing by the associated service. The combined TCP and IP headers containing the port and IP addresses are illustrated in Figure 1 for IPv4.

Figure 1 – Format of addresses and ports in a TCP/IP data packet (IPv4) (One Byte = 8 bits)

Byte #	Function	
1-12	[Initial header data]	
13-16	Source Address	
17-20	Destination Address	
21-24	[other header data]	
25-28	Source port	Destination port
29-44	[other header data]	
45+	Data (email, web graphics, etc.)	

Addresses for the Internet of Things

The Internet of Things concept supplements the Web and e-mail servers with devices that offer smart amounts of data, such as the local temperature or whether a light is on or off. Each device on the Internet needs an address. The more than four billion possible IPv4 addresses could soon be exhausted if these IoT devices were each as-signed unique IP addresses.

Two approaches to deal with the size of the IP-address space are available: sub-addresses and longer IP addresses. Most home and building networks are actually operating as local sub-networks where one IP address is assigned to the entire home or building (or to a company in the building) and a locally generated address is assigned to each computer and networked device.

In 1998 the Internet Engineering Task Force (IETF) proposed increasing the IP address from 32 to 128 bits, thereby expanding the number of nodes from 4.3 billion to 3.4 x 10^{38}. This scheme is called IP version 6 or IPv6. IPv4 is still the universal addressing scheme for the Internet.

Local address and public addresses

IPv4 routers within homes and buildings assign local IP addresses to attached computers and devices. Local addresses are not directly accessible by outside devices or computers and are not unique beyond the home, building, or company. The following blocks of IP addresses have been reserved for local assignment:

10.0.0.0 – 10.255.255.255
172.16.0.0 – 172.31.255.255
192.168.0.0 – 192.168.255.255

The ISP assigns one IP address to the home, building, or company router. The router uses one of the local IP addresses to identify each connected computer, smart phone, smart thermostat, smart TV, etc., whether connected via an Ethernet cable or Wi-Fi. This local address is mapped to the public IP address with a unique port number assigned by the router. Thus, a home computer with a local address of 192.168.0.3 may be assigned a public *source:port* address of 96.230.106.20:5000. The router is responsible for this address substitution in all outbound packets from this computer. Packets received with this address in the destination address and 5000 in the port field are sent to the local address of this computer.

Expanding the address space to IPv6

With 3.4 x 10^{38} IPv6 addresses available, it is possible to assign a unique address to practically every physical object needing to communicate. Even with IPv6 addressing we could retain the mapping function of a router, which would be assigned an IPv6 address by the ISP and still use port numbers to differentiate among devices in the home, building, company.

The network interface in Windows includes code for both IPv4 and IPv6 addresses. These addresses may be assigned once and retained indefinitely by the router as a *static IP address*. Alternatively addresses can be assigned (or reassigned) upon request of the router, such as when the router is turned on. This is called a dynamic IP address. Dynamic address assignment is preferred by ISPs since they have a limited pool of addresses to share among their customers, who may occasionally disconnect from the network, thereby releasing an IP address for reassignment.

Static address assignment is usually a premium service. However, static addresses facilitate machine-

Figure 2 – Features of the international gateway standard

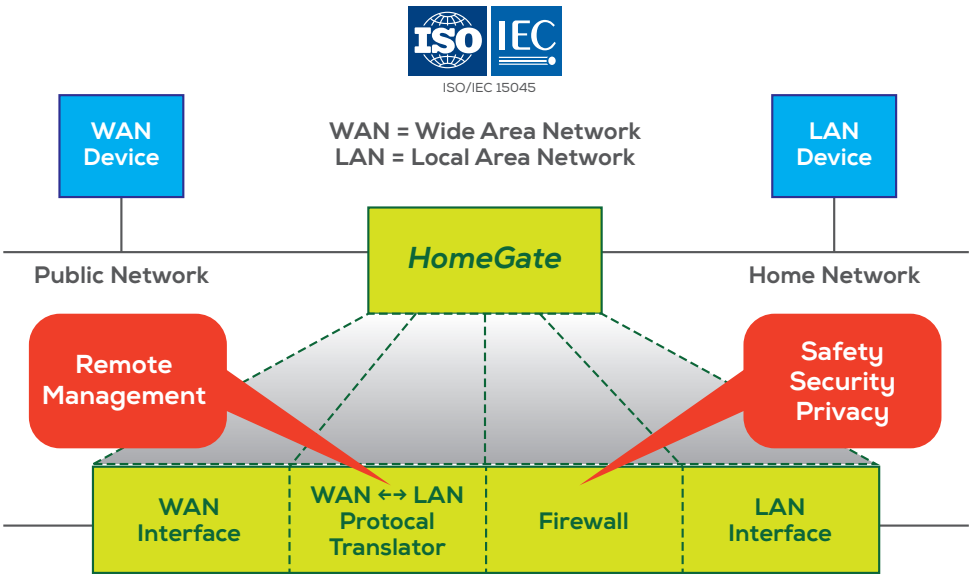

to-machine (M2M) communications since the sender can be programmed with the address of the recipient. M2M communications is possible with dynamic address. A server might be responsible for tracking the dynamic addresses of M2M devices, identified by a descriptive term. Each device would contact this server periodically to report its current dynamic address. Senders would ask this server for the IP address of the recipient.

Consumer electronics in IPv6

The Consumer Technology Association (CTA, formerly called the Consumer Electronics Association) developed an American National Standard for consumer electronics using IPv6. This standard was published in 2015 as ANSI/CEA-2048, "Host and Router Profiles for IPv6." I was a member of the CTA committee that wrote this standard.

The assignment of IPv6 addresses to consumer products such as televisions elicited considerable discussion at CTA about consumer privacy. We finally decided not to allow permanent static IP address assignment.

The CTA decision is explained in the standard [Section 3]:

> This standard includes methods for assigning addresses to consumer devices as required for these devices to communicate via the Internet using the Internet Protocol version 6 (IPv6) for an expanded address space (compared to IPv4). However, these IPv6 addresses are not intended to become permanent identifiers for these devices. In order to protect consumer privacy, various techniques, some of which are based on established IETF RFCs (specifications of the Internet Engineering Task Force), are required in this standard so an IPv6 address cannot be permanently associated with specific devices, and in effect then be associated with the individual(s) using the device. Such privacy protections may need to be strengthened via amendments or future versions of this standard as cyber-security technology evolves. Therefore, deployed devices may need to be upgraded to maintain privacy protections.

NOTE: The United States HIPAA law (Health Insurance Portability and Accountability Act) lists "Internet Protocol (IP) address numbers" as a specific type of Protected Health Information (PHI) identifier if it is used within records that support health care treatments, payments, or operations. The association of a permanent IPv6 address with a device that is subject to HIPAA health regulations could then be expressly prohibited unless the IP address is safeguarded in full accordance with all associated HIPAA requirements. HIPAA requirements for protecting PHI identifiers are detailed and complex.

Establishing a privacy perimeter

If all devices had static addresses, a world map of M2M-enabled devices could be determined. There are Internet Web sites that search for accessible M2M devices and publish maps of M2M networks for curiosity or malice. Many of these devices were never intended to process messages from random sources. For example, messages might include commands for an actuator to open a valve in a chemical or water processing plant.

This raises the fundamental dilemma of IoT. Should every device be accessible via the Internet? The answer is YES for convenience and NO for security. Clearly we don't want the chaos of everything talking to everything with the potential for malicious control or inadvertent remote operation. If we want the convenience of universal access, we must incur the considerable expense and complexity of cybersecurity protection for every IoT device. Furthermore, each IoT device would need a processor capable of handling IPv6 with the requisite speed while achieve a cost target for a marketable product.

Another approach to protecting valuable assets is to create a fortress with sentries that allow only authorized access. In a cyber-physical world this could be accomplished by a gateway that assigns local addresses and monitors message flows between a network outside a home or building (called the wide area network) and the network inside (called the local area network).

The world standard for the gateway was developed by the international committee that I chair. This series of standards, informally called HomeGate, was published in two parts in 2004 and 2012 as ISO/IEC 15045-1 and ISO/IEC 15045-2. A third part is under development to specify features in the gateway for security and privacy. Figure 2 illustrates this gateway with these added features. The gateway would provide cybersecurity features rather than burdening each IoT device in the home or building.

The gateway could be programmed to filter messages arriving from unknown sources or messages that initiate possible dangerous actions. For example, if you have a smart phone app to start your oven, the gateway might check that the command originated from your phone. It could also be linked to a smart appliance that checks for a closed oven door.

Forethought for IoT devices

The need for universal connectivity with remote access and control should be critically examined for each application. For example, should every light switch be an IoT device connected to a server in the cloud, which then communicates with an actuator in the house to turn on a lamp? Imagine if communications were interrupted during a storm so the switch could not speak to the lamp. You would be sitting in the dark even if you had backup power such as a generator or a storage battery fed by renewable sources (wind or solar).

IoT is all-inclusive term for the automation of homes, buildings, and factories. Some market prognosticators paint a rosy future of low-cost sensors and actuators enabling a myriad of clever applications. However, if the cybersecurity and privacy issues discussed in this paper are ignored, there is the potential for unintended consequences where IoT descends into chaos. ●

Dr. Kenneth Wacks has been a pioneer in establishing the home systems industry. He advises manufacturers and utilities worldwide on business opportunities, network alternatives, and product development in home and building systems. In 2008, the United States Department of Energy appointed him to the GridWise Architecture Council. For further information, please contact Dr. Wacks at 781.662.6211; kenn@alum.mit.edu; www.kenwacks.com.