Published June 1, 2002. Distribution restricted to Sponsors until September 1, 2002.

AUTO-ID CENTER

TECHNICAL MEMO

A Comparison of the Electronic Product Code Identification Scheme & the Internet Protocol Address Identification Scheme

Daniel W. Engels

AUTO-ID CENTER MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 77 MASSACHUSETTS AVENUE, BLDG 3-449, CAMBRIDGE, MA 02139-4307, USA

ABSTRACT

This memo presents a summary comparison of the Electronic Product Code (EPC) identification scheme and the Internet Protocol version 6 (IPv6) Address identification scheme. The EPC identification scheme is designed to uniquely identify all physical objects, and the EPC identifiers are designed to be used as pointers to information about their associated object. The IPv6 addressing scheme is designed to identify network interfaces with aggregatable unicode addresses providing globally unique identifiers. IPv6 addresses are designed to be used by the Internet Protocol (IP) packet communication/routing system that provides efficient communication between networked computing systems. The primary conclusion from this comparison is that neither scheme may be used to replace the other without compromising its ability to be used within its originally intended application.

TECHNICAL MEMO

A Comparison of the Electronic Product Code Identification Scheme & the Internet Protocol Address Identification Scheme

Biography

by Daniel W. Engels
Associate Director

Daniel W. Engels received his B.S. from the University at Buffalo, his M.S. from the University of California, Berkeley, and his Ph.D. from the Massachusetts Institute of Technology all in Electrical Engineering and Computer Science. His master's thesis is in the area of computer-aided design for electronic systems, and his doctoral thesis is in the field of theoretical computer science. Dr. Engels joined the Auto-ID Center after obtaining his doctoral degree where he leads the day-to-day research activities of the Center. Dr. Engels' research interests include scheduling theory and applications, real-time system design, distributed and mobile computing, and computer-aided design for embedded systems.

TECHNICAL MEMO

A Comparison of the Electronic Product Code Identification Scheme & the Internet Protocol Address Identification Scheme

Contents

1. INTRODUCTION

This memo presents a summary comparison of the Electronic Product Code (EPC) identification scheme and the Internet Protocol version 6 (IPv6) address identification scheme. The EPC identification scheme is designed to uniquely identify all physical objects [1] [2]. Furthermore, EPC identifiers are designed to be used as pointers to information about their associated object. In contrast, the IPv6 addressing scheme is designed to identify network interfaces [3] [4] [5]. The aggregatable global unicode addresses (IPv6 addresses containing the binary Format Prefix of 001) provide the set of globally unique IPv6 identifiers. IPv6 addresses are designed to be used by the Internet Protocol (IP) packet communication/routing system that provides efficient communication between networked computing systems.

The EPC and the IPv6 address identification schemes are designed for different sets of objects and distinctly different applications. However, the EPC identifiers and the IPv6 aggregatable global unicode (AGU) addresses are very similar in structure and allocation management. The primary questions that this memo seeks to answer are:

“Can the EPC scheme be used to replace the AGUs (or more generically the IPv6 address scheme) while maintaining its original use?” and, symmetrically, “Can the IPv6 addressing scheme (or more specifically, the AGUs) be used to replace the EPC scheme while maintaining its original use?”

We summarize the EPC identification scheme in Section 2 and the IPv6 addressing scheme in Section 3. A comparison of their key characteristics is presented in Section 4. We conclude with answers to our primary questions in Section 5.

2. EPC SUMMARY

An Electronic Product Code is either a 64-bit or 96-bit identifier (additional versions of greater length are anticipated in the future). The Electronic Product Code namespace is segmented into four hierarchically encapsulated partitions, version, domain manager, object class, and serial number, regardless of the total number of bits in the identifier. This segmentation gives the EPC namespace a singly rooted out-tree topology with a height of exactly four: depth 0 – root node, depth 1 – version number, depth 2 – domain manager, depth 3 – object class, and depth 4 – serial number. Each node in the namespace topology is labeled with a zero padded binary number (the total number of bits in this number is determined by the version number). Sibling nodes may not have identical labels, but labels may be identical between non-sibling nodes. The null label is reserved for the root node of the tree.

A path from the root node to a leaf node corresponds to a valid EPC identifier. The value of the EPC identifier is determined by concatenating the labels of the nodes on the path from the root node to the leaf node. EPC identifiers that differ only in their version number, that is their domain manager, object class, and serial numbers are identical (ignoring leading zeroes), are defined to be identical; thus, they correspond to the same physical object. Therefore, the EPC effectively has a three level hierarchy in its namespace. An object may be assigned multiple EPC identifiers, and, once assigned, the EPC will permanently be associated with the object.

The out-tree topology of the EPC namespace also indicates the identifier allocation and management authority relationships within the namespace. The Auto-ID Center defines the valid version numbers which are the parents to the domain manager numbers. Thus, the Auto-ID Center allocates the domain manager numbers. In turn, the owner of a domain manager number allocates the object class numbers that are its children in the namespace topology, and similarly for the allocation of serial numbers.

As an example of this allocation and management, consider a manufacturer of goods, such as Gillette, that obtains a domain manager number. Each product produced by that manufacturer, such as Mach3 4-pack blades, will be assigned an object class, or product, identifier, and each instance of that product will be assigned a unique serial number. By choosing an appropriate version number, each instance of the product will be assigned a unique EPC identifier.

The public Object Name Service (ONS), used to map EPCs to IP addresses, is logically deployed in a topology that mimics the EPC namespace topology. This topology enables ONS to provide efficient services.

3. IPv6 SUMMARY

An Internet Protocol version 6 (IPv6) address is a 128-bit identifier. There are three types of IPv6 addresses: unicast, anycast, and multicast. Unicast addresses are typically assigned to one and only one network interface; however, when multiple network interfaces are logically treated as a single interface, a single unicast address may be assigned to all of these network interfaces. Anycast and multicast addresses may be assigned to multiple network interfaces. A network interface may be assigned multiple, unicast, anycast, and multicast IPv6 addresses. IPv6 uses a variable length Format Prefix to distinguish between the multiple types of IPv6 addresses. Unlike its MAC address[6], the IPv6 address(es) assigned to a network interface may be changed over its lifetime. IPv6 unicast addresses are designed assuming that the Internet routing system makes forwarding decisions based on a “longest prefix match” algorithm on arbitrary bit boundaries and does not have any knowledge of the internal structure of IPv6 addresses. The structure in IPv6 addresses is for assignment and allocation.

The IPv6 aggregatable global unicode (AGU) addresses provide globally unique identifiers for network interfaces. The AGU address namespace is segmented into three hierarchically encapsulated partitions, public topology, site topology, and interface identifier. This segmentation gives the AGU namespace a singly rooted out-tree topology with a height of at least three: depth 0 – root node, depth 1 – public topology, depth 2 – site topology, and depth 3 – interface identifier. Additional partitions may exist within the public topology identifier, the site topology identifier, and the interface identifier. The format prefix (equal to 001 in binary) corresponds to the three most significant bits of the public topology identifier. Each node in the namespace topology is labeled with a binary number. The public topology identifier is 48 bits in length. The site topology identifier is 16 bits in length, and the interface identifier is 64 bits in length. Sibling nodes may not have identical labels, but labels may be identical between non-sibling nodes. The null label is reserved for the root node of the tree.

A path from the root node to a leaf node corresponds to a valid IPv6 AGU address. The IPv6 AGU address is determined by concatenating the labels of the nodes on the path from the root node to the leaf node.

The out-tree topology of the IPv6 AGU address namespace also indicates the identifier allocation and management authority relationships within the space. The Internet Assigned Numbers Authority (IANA) is the registration authority for the assignment of unique public topology identifiers (also referred to as network identifiers) for IP addresses. The IANA is chartered by the Internet Society (ISOC) and the Federal Network Council (FNC) to act as the clearinghouse to assign and coordinate the use of numerous Internet protocol parameters including IP public topology identifiers. The owner of a unique public topology identifier is responsible for the allocation and assignment of site topology identifiers. Similarly, the owner of the site topology identifier is responsible for the allocation and assignment of interface identifiers under the namespace defined by the given unique public topology identifier and site topology identifier.

The effect of the out-tree topology is that an IPv6 AGU address details the hierarchical set of networks and subnetworks that contain the network interface. That is, the AGU address provides directions on how to reach a network interface by detailing the set of networks and subnetworks to which the network interface is attached. If a network interface is connected to a different network or subnetwork without changing its AGU address, all communication to it would be sent to an incorrect network or subnetwork. Thus, the interface would not receive any communications sent to its AGU address.

The network routing system relies upon the logical tree structure topology created by the network hierarchy described in assigned IPv6 AGU addresses for efficient operation.

4. COMPARISON AND COMPATIBILITY EVALUATION

Both the EPC identification scheme and the IPv6 addressing scheme are designed to identify objects. The EPC identification scheme is designed to uniquely identify all physical objects. The IPv6 addressing scheme is designed to identify network interfaces, a specific type of object.

In addition to identifying a network interface, the IPv6 addressing scheme is designed such that a particular IPv6 address identifies the type of communication that it will be used for: either unicast, anycast, or multicast. The communication type is determined from the format prefix of an IPv6 address. An EPC identifier contains no additional information about its associated object or pertaining to its use. There is no provision for embedding additional information within an EPC identifier. Therefore, the EPC identification scheme cannot be used to replace the general IPv6 addressing scheme.

The IPv6 AGU addressing scheme is designed to uniquely identify network interfaces on a global scope. An AGU address is used primarily for unicast communication, but may also be used for anycast communication. The type of communication is determined by the communication/routing system and not specifically denoted by an AGU address. The IPv6 AGU addressing scheme is very similar to the EPC identification scheme. Table 1 summarizes the key properties of an IPv6 AGU address and an EPC identifier. The namespace structure of these two schemes are very similar.

Table 1: Summary Comparison of EPC and IPv6 AGU properties.

EPC AND IPv6 PROPERTIES

	EPC	IPV6 AGU
OBJECTS TO IDENTIFY	all physical	network interfaces
PRIMARY APPLICATION	information pointer	routing address
IDENTIFIER SCOPE	global	global
IDENTIFIER UNIQUENESS	unique	unique
SAME IDENTIFIER ASSIGNED TO MULTIPLE OBJECTS	no	yes
IDENTIFIER PERMANENTLY ASSIGNED	yes	no
IDENTIFIER CODING	binary	binary
IDENTIFIER LENGTH (BITS)	64, 96, other	128
NAMESPACE PARTITIONS ENABLE	allocation	allocation
NAMESPACE MAXIMUM EFFECTIVE DEPTH	3	3

The IPv6 AGU address namespace contains a three level hierarchy structure for allocation purposes. The EPC namespace effectively contains a three level hierarchy structure for allocation purposes. Both of these schemes have a structure designed to enable efficient allocation and management of their respective namespaces.

However, they are designed for use by very different applications. Thus, the partition identifiers in the two identification schemes have different interpretations. The EPC partitions effectively identify the manufacturer of the object, its product class (as determined by its manufacturer), and the unique instance of the product. The IPv6 AGU partitions effectively identify the hierarchy of networks and subnetworks to which a network interface belongs. This hierarchy is interpreted by the communication/ routing system as an address (or path, or directions) to the network interface. Due to their interpretations, the EPC namespace is effectively managed by the manufacturers of products and the IPv6 namespace is effectively managed by network administrators.

Because of its interpretation by the communication/routing system, IPv6 addresses are designed for fixed location (logically) network interfaces. A mobile network interface requires either its primary IP address to change as it changes location or a secondary IP address to be assigned. The secondary IP address must be changed as the interface moves, and the primary IP address must be used as a pointer to the secondary IP address. An EPC identifier is designed for mobile objects and for pointing to information, such as the currently associated IP address, about that object. Furthermore, once assigned to an object, an EPC identifier is forever more associated with that object.

5. CONCLUSIONS

The Electronic Product Code identification scheme and the Internet Protocol version 6 identification scheme are very similar in structure and appearance. However, they are not interchangeable in the sense that neither the EPC identification scheme nor the IPv6 addressing scheme can be used for both permanently assigned unique identifiers and addresses for IP communication.

The inability to encode information beyond its naming authorities prevents the EPC from being used as an address in the IP communication/routing system. An EPC identifier does not provide information on its associated object beyond the encoding of its naming authorities. An IPv6 address does provide information on the type of communication to be performed with that address in addition to encoding its naming authorities.

The requirement to interpret an IPv6 identifier as an address for IP communication prevents its use as a permanently assigned identifier on mobile objects. An IPv6 identifier must have its value assigned such that it may be correctly interpreted as an address to its associated object by the IP communication/routing system. Thus, as soon as an object moves from one subnet to another, its assigned identifier must be modified to reflect its new address. An identifier that acts as a pointer to information about an object must be assigned at the time the object is manufactured to prevent the loss of information about that object. If information about an object is gathered and stored in multiple locations, then its originally assigned identifier must be permanently associated with it. Physical objects are mobile and will have data collected and stored on them in a multitude of locations. Thus, the initially assigned identifier for an object must be an identifier such as an EPC.

6. REFERENCES

1. **David Brock. The compact Electronic Product Code — a 64-bit representation of the Electronic Product Code.**
   Technical Report MIT-AUTOID-WH-008, MIT Auto-ID Center,
   Massachusetts Institute of Technology, November 2001.

2. **David Brock. The Electronic Product Code — a naming scheme for physical objects.**
   Technical Report MIT-AUTOID-WH-002,
   MIT Auto-ID Center, Massachusetts Institute of Technology, January 2001.

3. **R. Hinden and S. Deering. IP Version 6 Addressing Architecture.**
   IETF, RFC 2373, July 1998.
   http://www.ietf.org/rfc/rfc2373.txt.

4. **R. Hinden, M. O'Dell, and S. Deering. An Aggregatable Global Unicast Address Format.**
   IETF, RFC 2374, July 1998.
   http://www.ietf.org/rfc/rfc2374.txt.

5. **Internet Protocol: DARPA Internet Program Protocol Specification.**
   University of Southern California Information Sciences Institute.
   IETF, RFC 791, September 1981. http://www.ietf.org/rfc/rfc0791.txt.

3. **David C. Plummer. An Ethernet Address Resolution Protocol.**
   IETF, RFC 826, November 1982. http://www.ietf.org/rfc/rfc0826.txt.

Designed and produced by Foxner. www.foxner.com