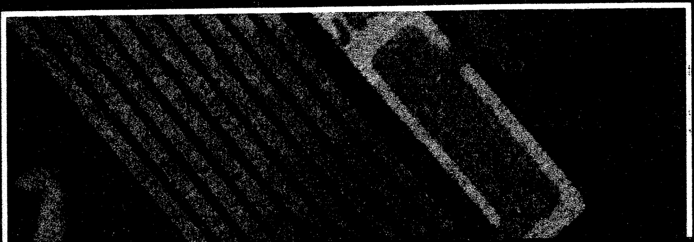

RFIDs, Near-Field Communications, and Mobile Payments

A GUIDE FOR LAWYERS

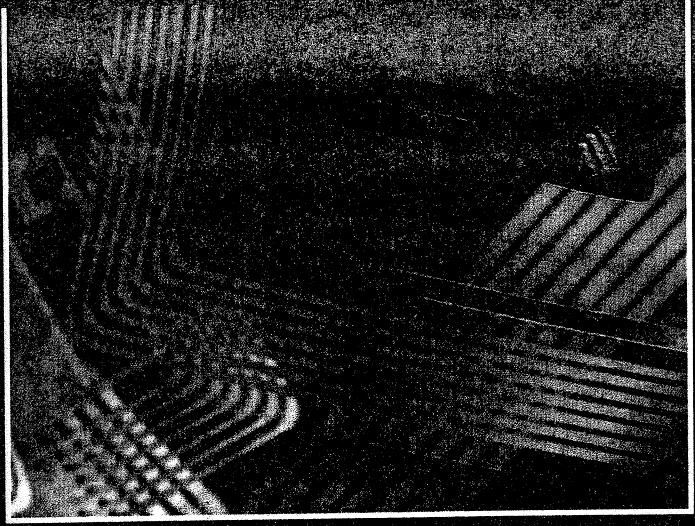

Sarah Jane Hughes, Editor
with Stephen T. Middlebrook
and Candace M. Jones

CYBERSPACE LAW COMMITTEE

AMERICAN BAR ASSOCIATION
Business Law Section

SECTION 6

RFID AND ADDRESS ISSUES

CHAPTER 13

The Handle System and its Application to RFID and the Internet of Things

Patrice A. Lyons and Robert E. Kahn*

Introduction

This essay provides information about the evolution of the Internet, and, in particular, the development of "addressing." As described in Part II of this essay, computer networking began by designating the computers that were communicating with each other and that shared certain protocols that allowed them to communicate. These networks predated the Internet, which is a global information system (and not in itself a network) whose protocols and procedures enable the appropriate information flows among the various computer networks and other computational facilities. Addressing evolved from those early pre-Internet address forms to the open Internet architecture and its Internet Protocol (IP) addresses with which we have become familiar.

*Patrice A. Lyons is General Counsel to the Corporation for National Research Initiatives (CNRI), Reston, VA, President of Law Offices of Patrice Lyons, Chartered in Washington, D.C., and a member of the Cyberspace Law Committee. She can be reached at palyons@bellatlantic.net.

Robert E. Kahn is President and CEO of the Corporation for National Research Initiatives (CNRI), which he founded in 1986 to provide leadership and funding for research and development of the National Information Infrastructure. In addition to many other awards, in 1997, President Clinton awarded Dr. Kahn and Vinton G. Cerf the National Medal of Technology for their joint development of the TCP/IP protocol. Dr. Kahn can be reached at CNRI, www.cnri.reston.va.us.

More recently, physical objects other than computers are being identified in the Internet using RFID technology. The information emanating from an RFID tag can be used to identify and authenticate information about a physical object, but such information also can be useful input to related information resources. The term "Internet of Things" (IoT) has emerged to describe the interoperability between information associated with physical objects and other information systems.

Part III discusses the evolution of Internet addressing and the Digital Object Architecture. Part IV discusses the identification of information related to physical objects using RFID technology. Part V briefly describes the IoT; and Part VI presents several conclusions.

Background of "Addressing in the Internet"

The first packet network, the ARPANET, addressed computers by addressing (in essence) the wires on that network to which they were connected. Within that one network, routing to specific wires was easily managed. A global scheme of assigning numbers to wires around the globe could conceivably have been devised; however, a means of routing from network to network would be required as well as a means of incorporating other new communication networks in the future. The subsequent development of packet radio and packet satellite networks (sponsored by the Defense Advanced Research Projects Agency (DARPA)), and the growing need for an infrastructure to enable information to be accessed and used across the various networks and other associated computational facilities, led to the design and implementation of the Internet.

The original Internet architecture focused on enabling intercommunication between individual packet networks (each of which was assumed to have its own unique characteristics such as packet size, data rates, error conditions, and interfaces) and the computational resources connected to them. The approach taken was to develop a protocol for the reliable transport of packets from source to destination, transiting one or more networks along the way.[1] Part of the protocol that is now called the IP dealt with

[1] Robert E. Kahn and Vinton G. Cerf, *A Protocol for Packet Network Intercommunication*,

best-efforts communication over the various networks; and a host protocol, known as Transmission Control Protocol (TCP) (which originally included the IP functions) ensured reliable communication on an end-to-end basis. Although both protocols were originally integrated, subsequently they were separated to facilitate certain real-time traffic such as speech that did not require perfectly reliable end-to-end communication. Other protocols such as UDP—a datagram protocol—were developed to be used in place of TCP for applications that did not require the reliability of virtual circuit sequencing provided by most implementations of TCP.

A global addressing mechanism based on IP addresses was developed. Each machine participating in the Internet was assigned its own unique 32 bit IP address consisting of a network number and a host number on that network. All communication in the Internet was to use IP addresses for delivery purposes. Gateways (now called routers) were to be placed between the networks to facilitate global routing; the gateways used the network portion of the IP address for this purpose. All activities within and between computational resources, other than the chosen method of delivery of packets (i.e., TCP or UDP) were assumed to be under the direction of the users.

With the introduction of workstations and personal computers, it eventually became necessary to have a simpler means for users to remember IP addresses. The Domain Name System (DNS) was selected for this purpose during the 1980s. Its architecture attributed a name, known as a "domain name," to each computer. The DNS provided a means of mapping domain names into IP addresses. The names were all contained in a host file, publicly accessible from a known machine that was typically downloaded once a day to get the updated list until the file size started to increase substantially with the advent of local area networks (LANs) and personal computers (PCs). Initially, all domain names in the Internet were of the form host. arpa, reflecting the role that the funding agency ARPA (by then it was

IEEE Transactions on Communications, vol. Com-22, No. 5, at 627 (May 5, 1974), *available at*http://ece.ut.ac.ir/Classpages/F86/ECE571/Papers/CK74.pdf; *see also* Robert E. Kahn, *The Architectural Evolution of the Internet,* CNRI (Nov. 17, 2010), at http://www.cnri.reston.va.us/publications.html.

known as DARPA) had played in creating the Internet. "Dot ARPA" was the first "Top Level Domain" (TLD), but today plays a more limited role.

As the number of machines increased beyond a few hundred to many thousands, the host file was partitioned and each partition was associated with its own TLD. Seven more TLDs ("dot" net, org, com, edu, gov, mil and int) were created for this purpose. These became known as generic TLDs or gTLDs. Country code domains (i.e., ccTLDs) were later added to allow individual countries to manage the assignment and mapping of domain names.[2]

Typical early applications of the Internet were to send and receive e-mail and to transfer files. These uses have continued to the present. Over the years, e-mail has been improved and upgraded to accommodate such advances as graphics, attachments, stylized fonts, and expanded character sets. File transfer was simplified by use of the World Wide Web (Web), which turned a procedural method of transfer driven by user commands into an automated procedure activated by clicking on a structured citation known as a uniform resource locator (URL).

Over time, the nature of the Internet has become blurred in a number of important ways. Many people do not distinguish the Internet from the Web, especially if all they do is use the Web to access information. To others, it appears as if e-mail is provided by the Web, without regard to the underlying technology. Some think the Internet is based on the DNS, since it is embedded into many applications including e-mail and the Web. The DNS has been a helpful application in the Internet, but the Internet actually uses IP addresses for routing over the Internet. Even the Internet itself, which is a global information system defined by its protocols and procedures for intercommunication, is often confused with the various components that have been interconnected.[3]

[2] For general overview of the origins and early development of the Internet, including DNS, see Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff, *A Brief History of the Internet, available at* http://www.isoc.org/internet/history/brief.shtml (last visited Oct. 5, 2012).

[3] *See Patrice A. Lyons, The End-End Principle and the Definition of Internet, Working Grp. on Internet Governance, available at http://www.wgig.org/docs/CNRInovember.pdf (last visited*

Evolution of Internet Addressing: Digital Object Architecture

In the 1980s, Corporation for National Research Initiatives (CNRI) developed an innovative approach for managing information in the Internet. It assumed there were repositories of information and computational resources accessible in the Internet, made use of existing Internet capabilities, and extended them in several important ways. This work started with the notion of mobile programs in the Internet that would carry out tasks on behalf of a user.[4] The user, perhaps through an application program, would create programs that could be dispatched in the Internet. Such programs could visit multiple sites to access or collect information, or carry out computational tasks making use of such information.

In the early 1990s, this approach was split into two parts. One part concerned the management of mobile programs. The other part was concerned with the management of information in digital form: this latter part resulted in the development by CNRI of the Digital Object Architecture (DO Architecture).[5]

The *lingua franca* of the DO Architecture is the "digital object" (DO), and the protocol for accessing a DO is known as the digital object interface protocol (DOIP). A DO is a data structure that is machine-and-platform independent and that has a unique persistent identifier known as a "handle" (or, generically, a "digital object identifier"). A resolution system maps these identifiers into "state information" about the DO being identified. The state information could be the various locations where the DO may be accessed in the Internet (e.g., a list of several IP addresses), terms and conditions for

Oct. 5, 2012); *see also* Patrice A. Lyons, *Some Myths about the Internet*, *Soc'y for Scholarly Publ'g*, at http://www.sspnet.org/community/news/some-myths-about-the-internet.

[4] Robert E. Kahn and Vinton G. Cerf, *The Digital Library Project, Volume 1: The World of Knowbots* (March 1988), Corp. for Nat'l Research Initiatives, http://www.cnri.reston.va.us/kahn-cerf-88.pdf (last visited Oct. 5, 2012).

[5] Robert Kahn and Robert Wilensky, *A Framework for Digital Object Services*, 6 Int'l J. on Digital Libraries 115, *available at* http://doi.info/topics/2006_05_02_Kahn_Framework.pdf ; *see also* Corp. for Nat'l Research Initiatives, *Overview of the Digital Object Architecture* (July 28, 2012), http://www.cnri.reston.va.us/papers/OverviewDigitalObjectArchitecture.pdf.

access to and use of the DO, or authentication information and the like. The creator of the DO, or an authorized administrator, typically supplies this state information under a public key regime that is integrated within the DO Architecture.

Public key (PK) technology emerged in the late 1970s and has provided one way to deal with signatures and other forms of authentication. Unlike traditional encryption schemes, where a single key is used to encrypt and decrypt, a PK system uses two keys, one of which is kept private and the other made public. A PK Infrastructure (PKI) is needed to enable the widespread use of this technology. The DO Architecture enables a PKI by providing a straightforward means of making public keys available. Existing repository technology implements PKI based on the DO Architecture; other applications would have to do so, as well, in order to take advantage of it.

Individual DOs are stored in repositories in the Internet and accessed by their unique identifiers. Thus, changes in technology for storage, operating systems, languages, and platforms can be easily handled in this way. The DOIP allows commands to be represented as DOs specified by their unique identifiers, permits optional verification of the repository by the user, and enables verification of the user by the repository for purposes of access control. Resolving a DO's identifier allows the DO to be authenticated, if the state information contains authentication information such as a fingerprint or hash of the DO. Network communications can also be signed or encrypted for security.

In addition, registries that contain metadata about DOs can be used to build collections across multiple repositories and to search for identifiers of DOs based on the use of metadata schemas tailored for the DOs. Thus, searching for a book represented in digital form would be based on one type of metadata schema that would be quite different from that used for a chip design, or for a song, or movie, or biological structure.

Because a DO may contain executable programs, passive data, or both, access involves the possibility of deriving more dynamic actions than just the transfer of prestored and prestructured information from one place to another. Indeed, distributed applications comprised of multiple interacting DOs may be structurally interconnected using identifiers that allow dynamic reconfiguration of both the underlying application and the presentations generated for the participants.

Identification of RFID Information

Identifiers can be obtained in many ways. Typical means might involve conveying them in an e-mail, or posting them via one or more applications obtained from an "apps" outlet or from a Web page. They might be obtained as citations from a journal in an Internet accessible archive, or from information passed between components of a distributed simulation. More recently, near-field radio frequency techniques have attracted attention, due to the use of cell phones or other devices to convey information very locally, or physical entities, such as a passport or access card from which certain information can be accessed wirelessly.

Radio frequency identification (RFID) devices are small passive or active transceivers that can emit internal information, such as an identifier over an "air interface." A typical active transceiver would require a power source such as a battery, solar cell, or other source of power. A typical passive transceiver embodies a very small antenna with fixed internal information such as an identifier that is radiated by reflecting an input burst of energy back to the user suitably modulated with the fixed internal information. Because the input energy to a passive transceiver is usually very small, the reflected energy is even smaller and, thus, only works over very short distances, such as in an aisle of a grocery store or at the checkout counter, in warehouse inventory management, or in cargo management while in transportation.

The amount of information in an RFID device (RFID Information) is usually quite small, but larger devices could contain somewhat more information. All that is really required, however, is a unique identifier for the "state information" associated with a physical object or entity being identified where the state information represented in digital form can be accessed by a resolution mechanism. If the RFID device were used subcutaneously, such as in a pet, its internal information (i.e., the identifier) would be resolved into the relevant state information about the entity (i.e., information in digital form pertaining to the pet). If the RFID device were attached to a supermarket product, it would be resolved to relevant information about the particular product, or a means to access such information.

RFID Information is often confused with identifying the physical resource to which the RFID tag is physically attached. This issue also shows up with identity management (IdM). If a person is allotted an identifier, that identifier can be associated with the person or with managed information about the person. Most often, the identifier for IdM is used to obtain a public key, or other credentials, to be used for authentication and/or access control, although the identifier can be used for other purposes such as addressing e-mail, determining coordinates, or other contact information for the individual.

Security plays an important role here. If an RFID tag can easily be removed from one physical object and attached to another, then the relevance of the RFID Information and/or the additional information would be questionable. So, the permanence of the method of attachment is one consideration. It is conceivable for an RFID tag, assuming it is small enough, to be implanted subcutaneously in an individual. A human could presumably find a way to extract such a tag, but an animal would have more difficulty doing so.

Implementation of RFID technology has led in recent years to adoption of DNS addressing technology, in particular the Object Naming Service (ONS) based on the DNS, for the management of RFID Information.[6] The ONS was derived from work originally done by the Auto-ID group at MIT with leadership from EPCglobal,[7] along with support from a number of companies such as Procter & Gamble and Walmart.

As will be discussed further below, addressing and naming in the RFID context are now viewed as basic elements in what is sometimes called the

[6] Rolf. H. Weber and Romana Weber, INTERNET OF THINGS: LEGAL PERSPECTIVE 6 (2010); *see also* S. S. Chawathe, et.al, *Managing RFID Data, 2004 PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON VERY LARGE DATA BASES 1189* , *available at* http://www.vldb.org/conf/2004/IND6P2.PDF.

[7] *See generally* EPCGLOBAL, EPCglobal Object Name Service (ONS) 1.0.1, GS1, (MAY 29, 2008), *available at* http://www.gs1.org/gsmp/kc/epcglobal/ons/ons_1_0_1-standard-20080529.pdf (last visited Oct. 5, 2012).

Internet of Things (IoT). As noted in a report on a 2006 meeting in Brussels on the IoT:

> [i]f one wants to connect with something, one must know where it is. In the Internet, a hierarchy of **domain name servers** (DNS) allows one to do this. The root server is queried, which redirects the query to another server, and so on, until the physical address is found. An extension of this system, **object name servers** (ONS), is expected to serve RFIDs and the Internet of Things.[8]

The restrictions of the DNS in this context are not mentioned in the above report.

Internet of Things

Efforts are under way to create what appears on the surface to be a completely new capability, namely an "Internet of Things."[9] In reality, it may be viewed as just another Internet application. Historically, the Internet has been about communication of information, typically in the form of packets, from one computational facility to another over an interconnected set of networks. More recently, advances in technology have introduced new architectures, such as the DO Architecture, that assume the historical Internet capabilities and integrate information management along with them.

[8] John Buckley, *From RFID to the Internet of Things*, *Networks and Commc'n Techs, Directorate*, at 26 (March 2006), *available at* ftp://ftp.cordis.europa.eu/pub/ist/docs/ka4/au_conf670306_buckley_en.pdf (last visited Oct. 5, 2012).

[9] *See, e.g.*, Joint Coordination Activity on Network Aspects of Identification Systems (including RFID), JCA-NID, Doc. O-044 (Jul. 22, 2010) (activity renamed to focus on IoT). Information on this JCA effort and other IoT-related activities at the International Telecommunications Union (ITU) is available at http://www.itu.int/en/ITU-T/techwatch/Pages/internetofthings.aspx). *See also ITU Internet Reports 2005: The Internet of Things*, Executive Summary, Int'l Telecomm. Union, http://www.itu.int/dms_pub/itu-s/opb/pol/S-POL-IR.IT-2005-SUM-PDF-E.pdf (last visited Oct. 5, 2012). *Internet of Things* was the subject of a panel discussion at ITU Telecom World 2011 (http://forum.world2011.itu.int/sessions/f22-internet-of-things).

Although the IoT may include sensors and actuators as "things," it is really about managing information about things and, thus, there is no compelling reason to treat this area differently than any other area that involves information management in the Internet. This view finds some support in work being carried out in Europe. The final report on a 2009 meeting held at Brussels on the IoT observed that

> [t]he growth of the Internet is an ongoing process: only twenty-five years ago it was connecting about a thousand hosts and has grown ever since to link billions of people through computers and mobile devices. One major next step in this development is to progressively evolve from a network of interconnected computers to a network of interconnected objects, from books to cars, from electrical appliances to food, and thus create an 'Internet of things' (IoT).[10]

Although still in its early stages of formation, stated plans for the IoT appear tied very closely to certain aspects of the current Internet infrastructure, specifically the DNS, yet they also seem to assume everything else is new and different. This is, perhaps, a Hobson's choice that severely limits the potential for interoperability with other types of information systems deployed in the Internet now or in the future.

This is also particularly true where such systems move beyond current end-to-end limitations to focus on the management of discrete units of information at finer levels of granularity than currently may be contemplated. An identifier associated with a digital object can be resolved in two steps to get to the actual digital object. Step one is to resolve the identifier in the Handle System[11] to obtain the relevant state information for the digital object, such as where it is accessible in the Internet. Step two is to access the digital object. Unlike the DO Architecture, use of the DNS and ONS requires three successive steps. Step one is to resolve the ONS domain name

[10] *Internet of Things – An Action Plan for Europe,* COM (2009) 278 final, at 2 (June 18, 2009), *available at* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2009:0278:FIN:EN:PDF.

[11] Information on the Handle System may be found at *http://www.handle.net*.

in the DNS. Step two is to obtain the necessary state information from ONS to access the product information, wherever it may be stored. Step three is to obtain that information.

The IoT is not yet a defined conceptual entity[12] although efforts such as those involving RFID tags, which were undertaken by EPCglobal, are assumed to be part of it. It would be a mistake simply to adopt the existing RFID system approach for the IoT, rather than to find a way to accommodate it within the evolving Internet. Indeed, the IoT should not differ substantially in its infrastructure from the system used for information management in the Internet more generally.[13]

The beauty of the Internet, as it currently exists, is its open architecture and ability to accommodate future innovations in communication networks and application services, without making fundamental changes to the basic protocols and procedures that comprise the Internet as a global information system. The same approach should be adopted for information management. Fortunately, there is one reasonably well-known approach, the DO Architecture, that is widely used in several contexts and can be used in this context as well. Indeed, some components of that architecture have been available in the Internet for more than a decade.

The notion of a uniform resource name (URN) and a URI has been explored in related efforts by the Internet Engineering Task Force (IETF) and other organizations. Yet, the Handle System,[14] which is a key part of the DO Architecture, predated these efforts. At a minimum, the Handle System provides the same key functionality and offers the ability to evolve the

[12] *See* Stephan Haller, *The Things in the Internet of Things, SAP Research Center Zurich, http://www.iot-a.eu/public/news/resources/TheThingsintheInternetofThings_SH.pdf; but see* Elgar Fleisch, *What is the Internet of Things,* Auto-ID Labs White Paper WP-BIZAPP-053 (Jan. 2010), *http://www.autoidlabs.org/uploads/media/AUTOIDLABS-WP-BIZAPP-53.pdf (last visited Oct. 5, 2012).*

[13] For additional analysis of the issues in this section of this essay, see Dugie Standeford, *IoT Naming System Said "Critical" for Network of Connected Devices But Which One Unclear,* 13 Wash. Internet Daily 1 (Aug. 27, 2012) (reporting on prospects for a European Commission recommendation to European Union governments on the IoT and on a new competitor "Handle" to the current object-naming service (ONS) that is similar to the domain-name system (DNS) for information addressing), *available at* http://www.cnri.reston.va.us/papers/wwid082712.pdf [hereinafter Standeford].

[14] *See supra* note 11; *See also* Standeford, *supra* note 13 at 3–4.

capability by customizing the results of the resolution mechanism. Information about various physical objects is identified using this technology, and the information about them is accessible from their identifiers by means of metadata registries.[15]

If a physical object were to have an RFID tag attached to it, the ideal situation would be for that tag to contain its digital object identifier and make that available to a "reader" to resolve via the Handle System, which can refer to IP addresses, the DNS, and other state information, as appropriate. The information provided from the tag is already in the form to be resolved, only in this case it would be supplied to the Handle System instead of to the ONS. The Handle System would directly resolve the identifier to the resolution information requested without the need to access a third system. The administrator of the digital object identifier would simply make that information available in its local handle service.

Resolution in the Handle System is carried out in a distributed fashion by a collection of local handle services. There is no need for a single service provider as the function can be distributed among many relevant providers. A common thread to this approach is the part of the Handle System, called the Global Handle Registry (GHR),[16] which maps an identifier resolution request to the local service that can resolve it. This GHR capability is also distributed and can be managed by multiple organizations in the future. Indeed, multiple nations and private sector organizations can participate in providing such GHR functions. In this approach, there is no need to determine a single provider to operate the overall service. However, effective administration and coordination globally is essential to maintaining the integrity and long-term reliability of GHR services and the evolution of the DO Architecture more generally. Interoperability of heterogeneous information systems, including identity management systems such as that enabled by RFID technology, is an important area for technical research

[15] CrossRef, operated by Publishers International Linking Association, Inc., is a leading provider of metadata services using DO Architecture technology, CrossRef, *http://www.crossref.org* (last visited Oct. 5, 2012).

[16] *See supra* note 11. The DOI System developed by the International DOI Foundation is based on the DO Architecture, including the Handle System. It is a branded version of this technology. *Id.*

and development. This technology holds great promise in many sectors. Physical items such as paper still play a key role in legal matters and business operations. Where information having legal significance is represented in digital form and structured as a digital object, it is possible to introduce new capabilities that were not possible to achieve in paper-based data structures. For example, the concept of a bill of lading may be reconceptualized to incorporate RFIDs or other identifiers affixed to physical objects, such as packing crates or containers, and combined with identifiers for related insurance and banking information to form a more flexible and secure legal instrument. Authentication of such an instrument and the establishment of change of "control" when it moves in the conduct of commerce in the Internet are important aspects that can be enabled by the DO Architecture technology.[17]

To take another example: in the case of a will represented in digital form, it may be useful to uniquely identify, as part of the will, information about a specific bequest, for example, a painting held in a bank vault, particularly where the specific location of the painting may change over time. The same is true with real property references in, for example, deeds of trust or other mortgage-related documentation, or notarial certificates or apostilles[18] associated with such documents.

[17] *See, e.g.*, Robert E. Kahn and Patrice A. Lyons, *Representing Value as Digital Objects: A Discussion of Transferability and Anonymity,* 5 J. TELECOMM. & HIGH TECH. L. 189, *available at* http://www.jthtl.org/content/articles/V5I1/JTHTLv5i1_KahnLyons.PDF; *see also* U.N. Comm. on Int'l Trade Law (UNCITRAL), *Comm on Int'l Trade Law Possible future work on electronic commerce – Recommendations for future work on electronic commerce, Proposal of the United States of America on electronic transferable records, 42nd Sess.*, *June 29-July 17, 2009, ¶¶ 17-18, U.N. Doc. A/CN.9/681/Add. 1 (June 18, 2009), (June 18, 2009), available at* http://daccess-dds-ny.un.org/doc/UNDOC/GEN/V09/845/21/PDF/V0984521.pdf?OpenElement.

[18] *See generally* Hague Conference on Private International Law, *6th International Forum on the electronic Apostille Pilot Program (e-APP)* (June 29-30, 2010) at http://www.hcch.net/index_en.php?act=events.details&year=2010&varevent=197; for information on a specific plug-in to associate an apostille with a given document that was prepared by CNRI for this project, see http://www.handle.net/hs-tools/adobe/index.html.

Conclusion

The term "Internet of Things" has been introduced to deal with arbitrary "things" of indefinite scope and range, but, in reality, the IoT is really dealing with information about such things. It has been suggested that this formulation represents something quite different from the Internet as we know it today. This view is quite misplaced. Although one can perform different functions in the Internet when one is addressing arbitrary things, the architecture of the Internet as we know it need not change substantially. The main difference is in the nature of the specific application, including, in particular, the approach used for embedding and resolving identifiers.

In light of the limitations inherent in the DNS, it appears ill-advised to restrict future Internet-based systems only to this approach. This is evident when the attributes of more advanced information management systems such as that provided by the DO Architecture components are considered. Where the DO Architecture is implemented, information about each "thing" would have its own unique, persistent identifier, which could be obtained by resolving the identifier. That resolution information, in turn, could contain what may be needed to perform a given operation, or it could refer to one or more locations in the Internet where a DO containing a larger set of information about the thing could be accessed. This is similar to what a user would do for other applications in the Internet. The DO Architecture represents one such innovative system, but other such systems will surely be designed and developed in the future. The Internet needs to accommodate such innovations as they occur over time.