The Architectural Evolution of the Internet

By: Dr. Robert E. Kahn,
President & CEO, Corporation for National Research Initiatives

1. Early architectural ideas

The idea of computer networking goes back to times before the first such networks were created. Only after several such networks had come into existence did the need for the Internet architecture become a central concern. For several years, individual computer networks were operated independently of one another, but it was quickly recognized within the research community that such packet network interconnections would be worthwhile. In the United States, the Defense Advanced Research Projects Agency (DARPA) had funded the first such computer network, known as ARPANET; and after the development of two additional packet networks (ground radio and satellite) were underway, DARPA undertook to work on the task of internetworking that led to the global information system known as the Internet.

The Internet architecture was designed originally to enable different packet switched networks to intercommunicate and to enable different computers on those networks to work together. Each participating network and their computers were uniquely identified within the architecture. The basic idea was to specify a simple process by which bits sent from one such computer in the form of packets could be reliably and efficiently delivered to any other computer with a valid identifier.

Individual networks of various types were assumed to have different packet sizes, interfaces, error controls, routing disciplines, and administrations. This approach enabled the architecture to accommodate new networks that would be conceived and developed in the future. While the components of the Internet were to be networks, made available by one or more communication service providers, and computing services (initially made available by research participants and their organizations), ultimately they included devices and applications service providers of many different types as well.

The essence of the Internet architecture was the set of protocols and procedures that enabled these components to interoperate. The linchpin of this architecture was the identifier system, as represented by IP addresses and the associated Internet Protocol (IP). Each transmitted internet packet would contain, at a minimum, the IP address of the destination computer. The basic architecture did not specify how any given network would route such packets, but it was clear that a means was needed to route across the overall set of networks. To facilitate that, and to provide a coherent means of interfacing between networks, the notion of a "gateway" was introduced. The gateway, now known as a router, would accept packets from one network, interpret the IP address contained within the packet and determine to which router on that network the packet should be forwarded. This process would be repeated until the packet reached the final destination.

Many other issues were addressed, which are not recounted here, such as how to handle fragmentation in case the network packet limitations were about to be exceeded, or how to put together such fragments at the destination, how to do error control, or deal with duplicate packets. Some of these issues were handled, in part, by the routers but primarily by the participating computers using a protocol that was originally called the Transmission Control Protocol (TCP). Indeed the original notion of TCP included, among other things, the generation of IP packets, as well as their forwarding and/or reception. In due course, the functions of IP processing were separated out and the TCP protocol became known as TCP/IP.

2. Open Architecture

In essence, the Internet was specified by the protocols and procedures that enabled different networks and computers to intercommunicate. Individual organizations could develop and operate their own networks, computer services or use devices of different kinds to participate in the Internet as long as they complied with the appropriate protocols. The growth of the Internet could happen organically and its overall technical management was independent of the management of any constituent network or service. This was a critical aspect of the eventual world-wide propagation of the Internet, since every participant could take responsibility for management of its own resources.

There were, however, certain things for which the management of this global information system had to be coordinated, and, since DARPA's Information Processing Techniques Office (IPTO) was funding the initial work, IPTO undertook to perform this task. One was the assignment of IP and network addresses so that conflicts would not occur. This function fell within my purview at IPTO until the mid 1970s, when Jon Postel (a researcher at USC/ISI) was asked to assume the responsibility with DARPA funding. Shortly thereafter, Vint Cerf, who had been an assistant professor at Stanford University, joined me at DARPA and assumed program responsibility at that time for managing DARPA's internet activities. During this period, we created what became the Internet standards process, and the social structures that could assist DARPA in managing it. In late 1983, responsibility for managing the standards process was transitioned out of government to the then recently formed Internet Activities Board (IAB), and eventually to the current set of structures that now involves bodies such as the Internet Engineering Task Force (IETF). In 1998, oversight of some of the processes (e.g., IP management and DNS, described below) was transitioned into a newly formed organization called the Internet Corporation for Assigned Names and Numbers (ICANN).

The transition from the old ARPANET specific protocol, called NCP, to TCP/IP took place over a (roughly) six month period between January and June of 1983. From that time forward, concerted effort was made to transition as much of the Internet management from the U.S. Government to the private sector as possible. At present, the only remnant has been the U.S. oversight of ICANN. Most countries now make extensive use of the Internet, and many have done so for almost two decades. Yet certain issues remain unresolved and are discussed below.

3. Pragmatism

In the early 1970s, when the original work on the Internet protocols began, it was assumed by many that there would be only limited interest in such interconnected networks. Commercial time sharing systems were still very expensive (typically a million dollars or more) and the personal computer had not been introduced into the marketplace. Indeed, the first microprocessors were still in their infancy. Yet, within a very short time, the first workstations became available in the research community along with local area networks such as the Ethernet. The idea that there might be only a small number of computers and networks to interconnect was quickly overtaken by events. Fortunately, the basic Internet architecture was able to scale to handle these new technologies, but not without taking into consideration the realities of the new situation.

The IP Address (now referred to as IPv4) is 32 bits long. Originally, only 8 of those bits were to identify the network, and the remaining 24 bits were to identify the host on that network. This seemed like an enormous amount of address space at the time, but it was soon apparent that we could have thousands (or more) of Ethernets alone, and the possibility of even more personal computers and workstations loomed large.

Within the research community, plans were drawn up to reassign the interpretation of the IPv4 addresses to contemplate large, medium and small size networks, which we called class A, B and C respectively. This entailed a means of determining the class of an IP address and interpreting the bits appropriately. There could be a small number of class A networks, each with a large number of IP addresses, and a large number of class C networks, each with a small number of IP addresses. Class B networks fell in the middle.

Finally, it had become clear that names would be important so that individuals would not have to remember individual IP addresses (there is a corresponding need to map to individual ports or sockets on those machines, but this aspect is not discussed further here in the interest of simplicity). This naming approach had been taken in the early ARPANET days, when individual computers, called hosts, were given 16-bit ARPANET addresses. Names were selected by DARPA in consultation with the participating sites and the mapping between names and the 16-bit addresses was maintained by SRI. With regularity, the network sites would download a file called “host.txt” from SRI to get the most up-to-date mapping. The NCP protocol was used to process requests sent to those addresses. When the transition to the TCP/IP protocol occurred, a “domain name” construct “dot arpa” was used to distinguish the two cases. For example, the name ucla was meant to indicate the 16-bit address in the host.txt file, and the name ucla.arpa was used to mean a 32-bit address (that basically included the 16-bit address to which was added the network identifier).

When the number of computers had grown significantly, it was clear that the host.txt file was getting too large to manage as a singly downloadable file; and a scheme that USC/ISI had developed called the “domain name system” was selected by DARPA for use on the Internet. Instead of a single top level domain (i.e. “dot arpa”), seven more generic domains (i.e. gTLDs), namely “dot com, net, org, edu, mil, gov, and int”, were introduced, and each of these domains would be interactively accessed to learn about individual domain names under a specific gTLD. In due course, country code domains were introduced, and under ICANN administration, additional gTLDs have since been authorized.

The World Wide Web ("web") was later introduced and has become an important application that makes use of the Internet. By combining domain names with file names, to produce what are known as Universal Resource Locators (URLs), the process of accessing files containing structured data over the Internet was simplified. Instead of having to know the procedure by which remote files can be accessed, the procedures were (in effect) automatically invoked behind the scenes when the URL was clicked. Further, the data structuring enabled the information to be displayed effectively and to adapt to the user's presentation window size. Various advances in the web have improved its functionality, and, for the time being, it remains the main vehicle for accessing information on the Internet.

4. Realism

By the late 1990s, the impact of the Internet was being felt across the globe and many countries were making plans to make use of it for many purposes ranging from health, education and information access, to strategic communications and business. As a potential critical infrastructure, a basic question was asked, namely "who is in charge?" This question was raised in the context of the millennium development goals of the United Nations.

A series of consultations ensued leading to two meetings of the World Summit on the Information Society (WSIS; Geneva 2003 and Tunis 2005) to deal with such concerns. The question of "Internet Governance" was high on the list of topics for consideration. While much of the initial focus was on ICANN and the oversight role that the U.S. government played, a result of the WSIS at Tunis was to embark upon a series of meetings, called "the Internet Governance Forum (IGF)," at which topics of interest with regard to the Internet and its perceived governance could be discussed. To date, there have been five such IGF meetings, run under UN auspices, all of which have served to engage a large set of interested stakeholders from governments and the private sector in discussions. Although there have been many views of governance expressed in the context of the IGF, the general concept has served to focus the ensuing discussions around many of the most important internet management considerations still unresolved.

5. Clean-Slate

Several organizations have asserted that existing problems in the Internet require a fresh start and that a "clean slate" approach to the future internet needs to be taken. I have always been an ardent supporter of research support for good ideas, as one never knows what will emerge or where the ideas will take us. Good ideas about the how best to evolve the Internet should be supported. However, there is a fundamental question as to whether a clean slate approach is really required, and, more fundamentally, whether a clean slate approach is really feasible or even desirable.

There is a large installed base of Internet users, as well as providers of communications and other services. Any clean slate approach could serve to disenfranchise these users and begin to fragment the user community. Clean slate users will either have to maintain connectivity to two different systems or split off. Yet, to adopt a clean-slate approach, even if coherent "clean slate" architecture can be developed, may not be the wisest choice. Rather, an approach that bootstraps on the current Internet capabilities, or integrates new functionality in the current Internet, would seem to be a far better approach.

For example, although the NSF-supported GENI project began as a search for new architectures and a means to test and evaluate them, my view is that the most important contributions from that program are about exploring how to integrate multi-stack routers into parallel operation in the current Internet. This would allow different strategies to be explored by the research community in parallel with the provision of existing Internet services.

For many years, concerns have been expressed about enabling communication service providers to integrate new functionality within their networks. Since they play a critical role in interfacing users to the rest of the Internet components, the opportunity for anti-competitive behavior is apparent. If sufficient competition exists in a given country, this problem becomes primarily a business issue. I say primarily, since even with transparency and full disclosure by a service provider, there will still be need for oversight. How best to handle such oversight will be an important issue in each country.

6. Digital Object Architecture

The Digital Object Architecture (DOArch) is a new architectural approach to the Internet that originated in work done on mobile programs by my organization (CNRI) during the 1980s. In the early 1990s, we separated the mobility aspects from the information management aspects and the result was the DOArch (brief overview available at http://hdl.handle.net/4263537/5041). It represents a coherent meta-level architecture that advanced the Internet from the original objective of just moving bits (in the form of packets) from one machine and network to another, to the more important objective of managing digital information.

The DOArch is fully compatible with the existing Internet and can greatly augment its capabilities along several lines. It deals with information in the form of "digital objects" (DOs) which are data structures that have unique persistent identifiers. CNRI calls such identifiers "handles" (or generically "digital object identifiers"); others may brand them differently. By managing information in the form of DOs, one injects architectural infrastructure into the information management process. Further, the information is potentially portable from system to system so that it can be parsed and understood by multiple systems, if necessary.

DOs consist of multiple elements, each of which consists of a type-value pair. Each of the types is represented by identifier and can thereby be interrogated individually. Identifying the data structure itself, instead of a specific file or folder that may contain it, or perhaps the machine on which it was first made available, enables persistent information access that is decoupled from most aspects of the underlying technology.

The DOArch has three basic components and operates in an Internet environment. One component is the DORepository (http://www.dorepository.org), which relies on commercial storage systems and provides a means of storing and accessing DOs. The DORepository relies on a standard securable meta-level interface based on the use of identifiers. Thus, identifiers can be used to define operations to be performed, targets of those operations, individuals using the system, as well as the systems themselves. Every instance of a DORepository is potentially interoperable with every other DORepository, subject to administrative controls. Indeed, although we describe the architecture in terms of its components, and while software implementations of the technology are available for download on the Internet under a CNRI open source license, as with the Internet, the DOArch consists primarily of the protocols and procedures used to enable these components to work together.

If the user knows the relevant identifiers (or, more generally, if the user's system does), access to the relevant information only requires that the identifiers be resolvable into useful state information about the DOs to which they pertain. Such a resolution system, called the Handle System (http://www.handle.net), was developed by CNRI and is in wide-spread use around the globe. It has been in service on the Internet for almost 17 years and in 7x24 operation for more than a decade. It has been demonstrated to support the DNS system as well as more general types of resolution. Each handle resolves into a "handle record, that contains, for example, multiple IP addresses where a given DO resides, authentication information for DOs, as well as public keys, terms and conditions for use and other relevant information all entered, as appropriate, by the administrator for the given DO under PKI security. Many local handle services around the globe are run by individual organizations that manage their own handle records. Resolution times are usually measured in fractions of a second.

The Handle System has a Global Handle Registry, which is currently administered by CNRI, and three independently run mirrors. At present, in addition to the primary registry service provided by CNRI, there is one mirror in China, one mirror in Europe, and one mirror in the US.

Users/programs that need to find the relevant identifiers make use of more general DORegistries that are designed to support searching across collections of DOs using metadata supplied by the administrator for each DO (or perhaps extracted automatically from the DO in some cases). It is assumed that some (or perhaps all) of this information is private and security is built in.

The DORegistry technology (http://www.doregistry.org) assumes that the authorized administrator of a given DORegistry will supply a "metadata schema" for each class of DOs and convey, or arrange to have conveyed, the relevant metadata in each case. The metadata schema for each class of DOs may be learned from the DORegistry. Users/programs will search over specified collections based on allowable terms in the metadata schema. The DORegistry uses the DORepository to store its metadata records, and thus inherits the security provisions of the DORepository.

8. Opportunities

The integration of the DOArch into the Internet offers several important benefits. First, any two instances of the DORepository technology are automatically interoperable, if this is desirable in a specific application. Second, if the DORepository implementation is fast enough (i.e., real-time), and allows a clean separation between the processing of DO requests and the actual storage of DOs, it can provide an effective means of accessing legacy systems. The challenge here is to map the inputs and outputs of the legacy system into a digital object framework. This is primarily a task in understanding the semantics of each legacy system and the internal implementation choices, which may be a relatively easy or difficult task, depending on the legacy system. However, once done, it is then accessible from within the DOArch by other participants, if they have the appropriate access rights.

If this latter approach is adopted, then a major advance is possible in the way interactions take place across the Internet. From its earliest days, the Internet assumed, and, indeed, required users to direct what happens over the Internet. In certain pre-planned cases, a portion of the task could be automated and relegated to the appropriate programs. For example, email was preplanned in format and delivery mode so that on arrival it was understood to be email and placed in the appropriate user inbox. From there on, the user had to take charge. But for many new applications, it would be desirable for the destination system to carry out new functions based only on the contents of the arriving DO and without intervention by a user. Integrating medical inputs automatically into a patient's medical record with security is one such example.

Another opportunity is the possibility of gaining international acceptance of an approach whereby administration of the Global Handle Registry (GHR) is shared among a set of trusted participants. At present, CNRI provides the GHR administration and provides regular updates to the mirrors. A very small charge is collected from each participating organization to maintain their entry in the GHR and to defray the costs of running the GHR. In the shared administration, called a "Multi-Primary GHR" some number of independent organizations would be designated as primaries; and each could take responsibility for entering organizational information into the GHR. There would be need to coordinate the actions of the primaries, but only at the level of insuring they do not give out the same identifier to different organizations. For this purpose, and similar to what the International Telecommunication Union (ITU) does in maintaining country codes for the telephone system, one possibility would be for the ITU to hand out blocks of identifiers on an infrequent basis (e.g. once a year or once a decade) to each primary for purposes of unique assignments by the various primaries. This coordination could be achieved without the need for the primaries to coordinate directly on this block allocation.

Many of the details of this Multi-Primary GHR remain to be worked out, but the initial implementation of the system has been completed and testing is scheduled to begin next year. In this context, administrative issues and other global coordination matters will continue to be important to consider going forward.

Finally, for secure operation of the system, some form of identity management (IdM) will be required. If we assume multiple IdM systems and trust frameworks will be developed and deployed, the Handle System as well as other parts of the DOArch may be useful in achieving interoperability between them.

9. Conclusions & Observations

The evolution of the Internet is a work in progress and, hopefully, will continue for the indefinite future. In my view, ICANN has been working reasonably well and should continue to play the role originally given to it, i.e., primarily the coordination of the IP Address and DNS systems. At the same time, emphasis needs to be placed on allowing the Internet to grow and evolve in ways that allow innovation to flourish, but that are not disruptive to existing users.

We need to understand how best to incorporate fundamentally new capabilities into the Internet. To date most of these efforts have been incremental, and even some that erroneously appear to be merely parameter changes (i.e., going from IPv4 to IPv6) have substantive technological implications and have proven to be much more difficult to accomplish than originally envisioned. In many ways, a transition to a Digital Object Architecture seems a relatively easy and useful next step to adding important universal functionality. It bootstraps on the current Internet services and integrates well with the existing capabilities.

November 17, 2010