# 17 AN ARCHITECTURAL ALTERNATIVE FOR THE INTERNET

A common thread and an inherent weakness among the existing alternatives for bringing OSI and other protocols into the Internet is that they approach the problem in a patchwork fashion. A potential alternative to the four methods described in Chapter 16 for accommodating multiple protocols in the Internet has emerged from a series of workshops sponsored by the Internet Architecture Board. Consideration of the way in which the concept of "the Internet" has evolved since its inception has led to the formulation of a proposal that might be called "a new multiprotocol architecture for the Internet." In this chaper, the authors present a brief and highly speculative look into the possible future of the Internet.[1]

## What Is "the Internet"?

It is very difficult to deal constructively with the issue of "the multiprotocol Internet" without first determining what "the Internet" is (or should be). In this context, it is important to distinguish the Internet (a set of communicating systems) from the Internet community (a group of people and organizations). Most people would accept a loose definition of the latter as "the group of people who regard themselves as being part of the Internet community"; however, no such "sociological" definition of the Internet itself is likely to be useful.

1. These ideas emerged at the Internet architecture workshop held at the San Diego Supercomputer Center in June 1991. The participants in that workshop deserve much of the credit—such as it may be—for the contents of this chapter; see RFC 1287.

The scope of the Internet has traditionally been defined by the existence of IP connectivity. IP and ICMP were and remain the only "required" Internet protocols. It has been relatively easy to identify a host or router as a participant in the Internet; namely, if my host (router) could ping your host (router) and yours could ping mine, then we were both on the Internet. From this relatively simple test, a very satisfying, and eminently workable, definition of the Internet—transitive closure of IP-speaking systems (see Figure 17.1)—is thus constructed.

Until recently, the IP-connectivity model clearly distinguished systems that were "on the Internet" from those that were not. As the Internet has grown, and the technology on which it is based has gained widespread commercial acceptance, the notion of what it means for a system to be on the Internet has expanded to include the following:

- Any system that runs the TCP/IP protocol suite, whether or not it is actually accessible from other parts of the Internet
- Any system that can exchange RFC 822 mail (without the intervention of mail gateways or mail object transformations)
- Any system with E-mail connectivity to the Internet, whether or not a mail gateway or mail object transformation is required

These definitions of "the Internet," however, remain limited, for they are still based on the original concept of connectivity and merely percolate the addressing upon which connectivity is based up the stack to a "host" level rather than an IP level (e.g., if I can somehow convey enough information in an E-mail address to identify my host unambiguously to the recipient of my mail, there's a good chance the recipient will find a reciprocal path, be it through BITNET, UUCP, RFC 822 mail, or whatever).

A new definition of "the Internet" has been proposed, one that is

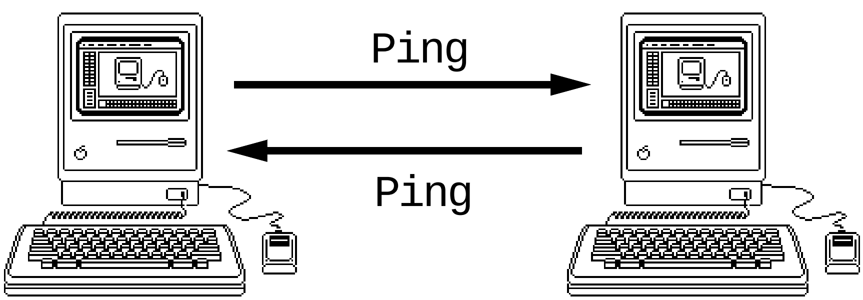


FIGURE 17.1 The "ping" Model of the Internet

based on a more powerful and *unifying* concept. Whereas the "old Internet" concept relied on IP connectivity—in which the IP address served as the fundamental, organizing principle and provided a common network address space—the "new Internet" is based on the notion of application name space. Specifically, the organizing principle is the existence of the domain name system and directories; the ping test for connectivity becomes a name game.

## A Naming-based Concept of Internet Connectivity

In the "new Internet," the idea of "connected status"—traditionally bound to the IP address—is coupled to the names (and related identifying information) contained in the distributed Internet directory (see Figure 17.2).

A naming-based definition of "the Internet" implies a much larger Internet community and a much more dynamic (and unpredictable) operational Internet. This argues for an Internet architecture based on adaptability to a broad spectrum of possible future developments. Rather than specifying a particular "multiprotocol Internet," comprised of a predetermined number of specific protocol architectures, the proposal instead presents a process-oriented model of the Internet, which accommodates different protocol architectures—new members—according to the traditional "protocols and combinations that do useful work" principle that permeates the Internet community and keeps the Internet itself from falling to pieces.

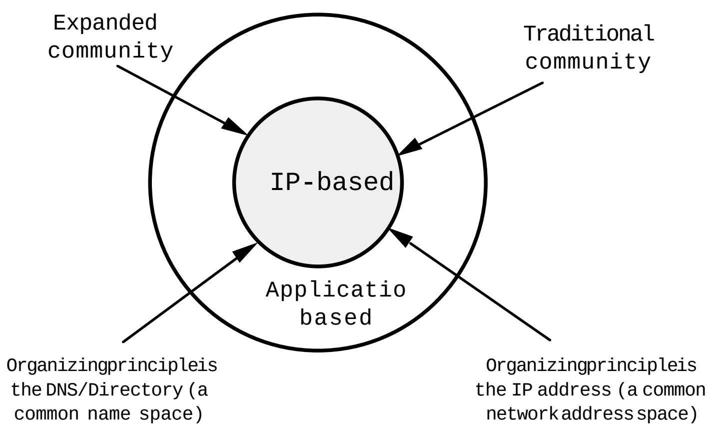


FIGURE 17.2 Expanded Definition of the Internet

The process-oriented model of the Internet asserts that there is no *steady-state* "multiprotocol Internet." Protocols and combinations may be added to the Internet as their usefulness, openness, and widespread use are demonstrated (in this model, it is also conceivable that a protocol or combination that ceases to be useful may be weeded from the Internet). A major assumption here is that everyone would agree that it is better to adapt and be connected than to be isolated or partitioned from the global community.

This seems to suggest that the forces driving the evolution of the Internet are pushing it toward multiprotocol diversity. Not so. Although we may never adopt a single "pure" stack, we may achieve a catholic notion of protocol-stack uniformity by adaptation and hybridization of stacks. Current wisdom suggests that the tendency of the Internet is to evolve toward what may be described as a "thermodynamically stable" state, defined through four components of the process-based Internet architecture (see Figure 17.3).

**Component 1: The "Core" Internet Architecture**

The traditional TCP/ IP-based architecture will remain the "magnetic center" of Internet evolution. This is not a statement of superiority but merely a recognition that (1) having something that people can point to as the common thread that holds the Internet together is good, and (2) IP connectivity remains that common thread. The argument here is as follows. The success of the Internet has stretched IP to its limits, and IP by itself can no longer be the lone measure of connectivity. For the Internet to extend further, other protocols must coexist; however, regardless of whether or not the actual state of IP ubiquity can be achieved in practice

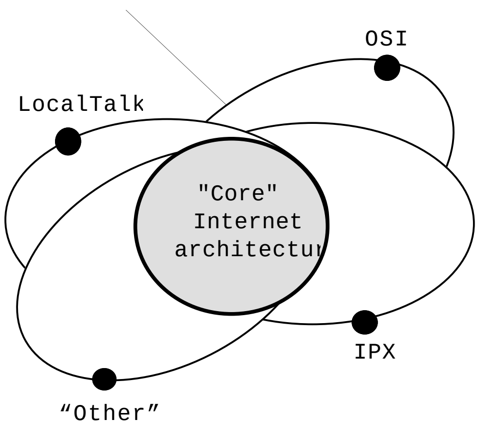


FIGURE 17.3 The Process Model of the Internet

in a global operational Internet, it is generally recognized that "getting there using IP" is a good thing.

Two additional components of the new Internet architecture express the ways in which the scope and extent of the Internet have been expanded.

### Component 2: "Link Sharing"

Physical resources (transmission media, network interfaces) will be shared by multiple, noninteracting protocol suites. The necessity and convenience of coexistence at this level is not at issue; issues of how common resources are to be managed, how to minimize the degree to which protocol suites interfere with each other, and how to create "equal access" in real Internet systems are critical to the continued success of the Internet. The Internet cannot grow without new controls or the means to extend existing controls across a multiprotocol platform; without these, prospects for a multiprotocol Internet would be grim.

What controls are needed and for what resources? Currently, there are communities interested in providing type of service routing, effective implementation of the notions of IP precedence, CLNP priority, and OSI's embarrassingly vague QOS within the framework of the existing "best-effort delivery," datagram environment of the Internet. There is a desire to introduce alongside these improvements an effective means of allocating or reserving resources so that applications such as packet voice, video, and teleconferencing might be supported across the Internet. These imply capabilities that are not currently accommodated in the Internet (there is a tendency to conclude that applications such as these require "connections," but it isn't necessarily the case that these connections must adhere to the traditional model of the "telephone call").

Beyond these controls, there are additional levels of control that are equally important: policy, access control, and accounting must all be improved.

### Component 3: Application Interoperability

In those circumstances in which it is not possible to achieve interoperability among (all) the underlying protocol stacks, it is still possible to achieve application ubiquity, through the deployment of gateways. The ultimate role of the Internet is to serve as a basis for communication among applications. At a minimum, the Internet should enable communities of end users sharing a common application "objective"—among these, the exchange of electronic mail, files, documents, images, and access to remote resources—but having diverse applications as "tools" to operate and migrate from one underlying protocol suite to another without unacceptable loss of functionality.

The scope of the Internet and the size of the Internet community are expanded considerably by these changes, but at a cost. The Internet may become harder to administer from backbone to individual hosts. The cost of multiprotocol hardware, software, and support and administrative staff increases as well. Accounting, accountability, and security are no longer merely important but essential. The changes are unavoidable, but they are not painless, and each time functionality is lost, or additional system complexity and costs are endured in order to expand the scope of the Internet, the decision to grow out from the "warm and comfy" world of IP connectivity will be subjected to renewed harsh scrutiny. In a perfect world, however, the Internet would evolve and expand without these penalties; there is a tendency, therefore, for the Internet to evolve in favor of the homogeneous architecture represented by component 1, and away from the compromised architectures of components 2 and 3. Component 4 expresses this tendency.

## Component 4: Hybridization/ Integration

This component expresses the tendency of the Internet, as a system, to attempt to return to the original "state of grace" represented by the uniform architecture of component 1. It is a force acting on the evolution of the Internet rather than a process whereby the Internet actually returns to a uniform state at some point in the future. Component 4 recognizes the desirability of integrating similar elements from different Internet protocol architectures to form hybrids that reduce the variability and complexity of the (central) Internet system. It also recognizes the desirability of leveraging the existing Internet infrastructure to facilitate the absorption of "new stuff" into the Internet, applying to "new stuff" the established Internet practice of test, evaluate, adopt.

The result is something akin to bombarding the nucleus of an atom with neutrons (the neutrons here being complementary or replacement protocols for the existing "core" stack). According to this dynamic model, running X.400 mail over RFC 1006 on a TCP/IP stack, integrated IS-IS routing, transport bridges, and the development of a single common internetwork protocol successor to IP and CLNP are all examples of "good things"[2]—they represent movement away from the nonuniformity of components 2 and 3 in the direction of greater homogeneity, under the influence of the "magnetic field" asserted by component 1, following the hybridization dynamic of component 4. Thus, homogeneity is accommodated not through adoption of a single, "pure" stack but through widespread implementation

2. The authors wish to emphasize that these are merely examples and are not intended to suggest that the very fine work done to produce MIME and OSPF should be set aside!

of the protocols and combinations that compose the "core."

What protocols and combinations will (initially) comprise the "core" remains highly speculative. Certainly, the existing TCP/IP "stack" will remain in the core (including ARP, OSPF, and BGP-something); these are likely to be complemented by CLNP and integrated IS-IS routing, the OSI directory and message-handling applications, and the RFC 1006 transport service bridges. Because of its virtual ubiquity among IP-based systems, the simple network management protocol (SNMP) is likely to remain the fundamental management component; it now runs over an OSI stack, as will its heir apparent, SNMP version 2.

This is, of course, highly speculative and considers only the complementing items from OSI. Depending upon stakeholders' interest in making other networking technologies "open," and their willingness to do so, the "core" may be complemented by applications and services of current, proprietary, or *de jure* networking architectures.