CLOUDS VERSUS STRINGS: WHY IP WILL CONTINUE TO PROVIDE THE FOUNDATION OF THE INTERNET

By Judy Estrin, President and CEO
Packet Design, Inc.
Menlo Park, California

As millions of new users flock to the Internet, the ability to scale reliably -- to grow big enough and fast enough to meet demand -- will be the single most important factor for building an effective infrastructure. The Internet Protocol (IP), as the foundation of that infrastructure, is the key to providing it with continued scalability. Other solutions have garnered interest because they have held out promises of telephony-like manageability -- promises that, without scalability, are meaningless. These solutions gained a foothold because they coincided with a shift in allocation of development resources to switching in the mid-90s, which in turn led to a sharp drop in industry investment in the evolution of the routing technologies that had originally established IP as the basis of the Internet.

The reason the so-called Internet Economy came about -- and is thriving today -- is that IP was designed in such a way that it could scale indefinitely in an environment of unpredictable growth, allowing "networks of networks" to be constructed. IP's distributed design has allowed first universities and government labs, and subsequently many others, to be able to keep adding to the network "cloud" without negatively impacting what was already there.

IP has come to dominate the networking market for several reasons:

It is open and available to everyone, encouraging rapid, competitive innovation.

It is application-independent, requiring no proprietary application-layer gateways.

And, probably most critical, its designers made three key decisions about its distributed architecture. First, because IP provides for separation of the control plane and the forwarding path, major advances in these two areas could be made independently: forwarding speeds could be pushed to the limit in silicon, while routing control functionality could be improved concurrently in software. Second, services are placed at the edges of the network rather than integrated into the network itself; this allows services to evolve without impacting the network and keeps complexity out of the network core. Third, and perhaps most important, the ability of packets to carry globally meaningful addresses enables network nodes to make autonomous decisions in processing each packet. This allows for the distribution of work throughout the nodes, both providing redundancy and improving scalability. (This is analogous to accepted theories on how to best scale organizations: by distributing information throughout those organizations, so that individuals are empowered to be better decision-makers.)

The distributed nature of IP's architecture makes it very different from the public telephone network, or PSTN, whose design is centralized and more complex. Telephony networks, which are based on connections, or circuits, work more like organizations with rigid hierarchies, where individuals (in this case network switches) are told what to do by a central controller. Circuit-based network architectures are inherently limited in scalability because new connections ("strings") must be set up, and consequently managed and accounted for, every time a new element is added.

Over the years the PSTN world has enthusiastically backed such technologies as Asynchronous Transfer Mode (ATM) and, more recently, Multi-Protocol Label Switching (MPLS). These technologies have major gut appeal to the telephony industry because they give the network a familiar circuit-oriented look, which network operators understand well and believe they can manage effectively. Being deterministic in nature (i.e., permitting the establishment of logical circuits), these approaches have claimed unique abilities in several areas: first, quality of service (QoS); later, traffic engineering; today, integration into emerging optical networks.

However, the convergence of data, voice and video we are seeing today is driven by the dramatic increase in data traffic now being pushed across the public infrastructure. The traffic patterns associated with new data applications are very different from those of phone conversations. Yes, the Internet needs to maintain the manageability of the telephony world -- but not at the expense of scalability. MPLS, a string-oriented technology, was developed to solve a point problem: integrating local IP and ATM environments. It works well for that use, but its proponents have positioned it as a panacea for all sorts of other problems. At first glance, MPLS seems like the perfect answer to a converged Internet. But it's really just a quick fix. Because its architecture is based on strings rather than clouds, it has all the disadvantages of strings and, in the long run, it creates more problems than it solves.

If IP was so great in the first place, why did technologies like ATM and MPLS even come about? The answer goes back to the early nineties, when the use of routing as the primary basis for data networks was giving way to Ethernet switching -- a technology that was faster, cheaper and easier to understand. Most companies and universities turned their research and development efforts away from routing during this period. MPLS arose out of a research vacuum, a void of routing expertise. But while circuit-switching may work fine in a relatively static local network, a connection-oriented approach that can't scale to hundreds of nodes, much less millions, cannot be applied to create an effective long-term end-to-end solution. And the Internet is the least static and least localized network of all time.

To create the Internet that we will want and need in the future, the industry must return to its routing roots, and work on evolving IP, particularly the IP control plane and its routing algorithms (e.g., BGP, IS-IS).

IP has already proven it can grow in functionality and speed, incorporating capabilities once thought to be the exclusive domain of circuit networks. For example, the DiffServ (differentiated services) protocol (IETF RFCs 2474 and 2475) has brought QoS functionality to IP networks, and IP Multicast has enabled streaming audio and video to scale to large numbers. IP traffic engineering will certainly be the next step.

Similarly, IP can and will evolve to support optical network integration. Think about it: what are optical networks all about if not scaling to support massive numbers of users, extreme data-transfer speeds, and ubiquitous services? There has been a recent effort to try to adapt the connection-oriented MPLS protocol for this purpose -- so-called Multi-Protocol Lambda Switching. Once again, take a peek behind the curtain and the strings are all too visible. Evolving IP -- not replacing it -- is the answer.

Article draft
Special to E.E. TIMES

CLOUDS VERSUS STRINGS: WHY IP WILL CONTINUE TO
PROVIDE THE FOUNDATION OF THE INTERNET

Revised
August 26, 2000

APPROVAL:

Judy Estrin ________________________

PD02