On Routing and Forwarding

Vinton G. Cerf
Google

Everything on the Internet need not be connected.

In the earliest days of Internet design, I was largely fixated on the idea that everything should be able to talk to everything else on the network. A recipient of an Internet packet could reject or ignore it but senders were free to send. As the Internet penetrated into the commercial private sector, enterprises looked for ways to isolate their computing equipment from the global network through the use of firewalls that were not part of the original design.

In subsequent years, routers have become extremely elaborate systems and routing methods have become more complex. The Border Gateway Protocols grew more complex as routing choices collided with economic considerations (e.g. near-end vs. far-end hop off). Multiprotocol Label Switching (MPLS) was used to groom traffic onto alternative optical streams to manage channel occupancy. Virtual LANs were developed to group devices together in a common communication channel even when they were on distinct physical LANs. OpenFlow expanded the basis for routing and forwarding decisions from simple destination address lookup to use of any bits in a packet on which to base forwarding choices. It also demonstrated the feasibility of centralizing forwarding table production for systems of suitable scale. The core, inter-data-center network of Google adopted this practice to very good effect.

It has finally dawned on me that not everything has to be connected to everything and that routing and forwarding can be deliberately constructed to confine connectivity to a desired cohort of devices. I have been so consumed with "everything has to be connected" that this recognition has been slow in coming. Duh. The forwarding tables in routers can be constructed in many ways and more than one forwarding table can readily be imagined. The implication is that one can use conventional routers or more recent Software Defined Networks (SDNs) to isolate groups of devices from the rest of the Internet. Such an implementation strikes me as a different way to realize the concept of Virtual Private Network without the encapsulation and potential hazards of conventional VPN implementation. The forwarding table defines the constituents of the virtual network. A flat table of specific 32 bit or 128 bit IP addresses could create an isolated group of devices able to communicate only with each other, for example.

Interestingly, the redefinition of forwarding tables also leads to concepts such as Information Centric Networks or Content Centric Networks that route on content indicators rather than addresses. Moreover, one can extend this line of thinking in other ways. The existing Domain Name System (DNS), for example, maps domain names into IP addresses which are then routed by the common routing system(s) of the Internet. It is quite possible to imagine a different set of identifiers, other than domain names, that could be mapped into IP addresses. One example is the Digital Object Architecture developed by Robert Kahn at the Corporation for National Research

Published by the IEEE Computer Society
1089-7801/18/$33.00 ©2018 IEEE

Initiatives (CNRI). Every digital object gets a unique digital identifier (handle) that can be looked up to find the location or locations at which the object may be found. David D. Clark's forthcoming book, Designing an Internet (Information Policy), explores a variety of potential new designs for Internet-like functionality.

As memory becomes less a barrier and backbone link capacities increase, more elaborate routing mechanisms may prove feasible. New SDN designs, with switches that are programmable, may provide a basis for more refined and sophisticated routing and forwarding mechanisms. In the coming era of the Internet of Things, the ability to isolate groups of devices for protective reasons may prove to be an essential step towards improving the security of the Internet. Such ideas will place significant demand on configuration tools to cope with scale and the use of cryptography to protect against the risks of misrouting.

I am sure that there are among the readers of this column, many who are much more cognizant of advanced thinking about routing and forwarding, so I hope they will take time to draw attention to their ideas for future evolution of this all-important function of the Internet.

BIO

Vinton G. Cerf is vice president and chief Internet evangelist at Google, and past president of ACM. He's widely known as one of the "fathers of the Internet." He's a Fellow of IEEE and ACM. Contact him at vgcerf@gmail.com.