Requirements for the Internet

Vinton G. Cerf

Senior Vice President, Architecture and Technology
MCI Communications Corporation
Ashburn, VA 20147
U.S.A.
vinton.g.cerf@mci.com

The requirements for the Internet emerged in part from the practical experiences drawn from the development, deployment and use of the ARPANET. Bob Kahn was one of the principal design architects of the ARPANET packet switch (Interface Message Processor) while he worked at Bolt Beranek and Newman. Bob joined the Defense Advanced Research Projects Agency in late 1972 and outlined his requirements for network evolution that year [3] (quoted with minor punctuation correction):

> Four ground rules were critical to Kahn's early thinking:
>
> - Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.
> - Communications would be on a best effort basis. If a packet didn't make it to the final destination, it would shortly be retransmitted from the source.
> - Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be no information retained by the gateways about the individual flows of packets passing through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
> - There would be no global control at the operations level.
>
> Other key issues that needed to be addressed were:
>
> - Algorithms to prevent lost packets from permanently disabling communications and enabling them to be successfully retransmitted from the source.
> - Providing for host-to-host "pipelining" so that multiple packets could be enroute from source to destination at the discretion of the participating hosts, if the intermediate networks allowed it.
> - Gateway functions to allow it to forward packets appropriately. This included interpreting IP headers for routing, handling interfaces, breaking packets into smaller pieces if necessary, etc.
> - The need for end-end checksums, reassembly of packets from fragments and detection of duplicates, if any.
> - The need for global addressing
> - Techniques for host-to-host flow control.
> - Interfacing with the various operating systems
> - There were also other concerns, such as implementation efficiency, internetwork performance, but these were secondary considerations at first.

In our discussions that started in early 1973 on the question of networking, Bob and I considered specifically the problem of interconnecting a mobile packet radio network (PRNET), a multi-access packet satellite network (SATNET) and the wireline ARPANET. Bob's Open Architecture ideas clearly took root in the design of the TCP (later TCP/IP) protocol [1, 2].

The requirement that the networks comprising the Internet NOT be modified led directly to the need for gateways (now called routers) and to the need for a global address space orthogonal to any intra-network addressing structure. The need for end-to-end reliability led to the retransmission mechanisms of TCP. Flow control was a known requirement because not all hosts would be equal in capacity to send or receive data, and the network could

Proceedings of the 11th IEEE International Requirements Engineering Conference
1090-705X/03 $17.00 © 2003 IEEE

potentially be congested. A great deal of attention was paid to the problem of long-delayed packets that might arrive at a most inconvenient time for the TCP protocol. Implementation experience with the December 1974 TCP specification [1] led directly to the addition of a "three-way handshake" to establish the initial sequence numbers for each direction of the bisymmetric TCP flow.

It was understood, if dimly, that classes of service would need support, hence the TCP/IP packet header contained flags for "reliability," "speed," or both. With regard to speed (timeliness of delivery), the motivation for breaking out the IP protocol from TCP was precisely the carriage of real-time speech over the Internet. This application favored rapid transport over absolute reliability. Eventually UDP and higher level real-time protocols emerged that avoided the use of TCP for streaming audio, video and real-time gaming.

As the number of networks comprising the Internet increased. it became clear that the original 256 network design was inadequate so the 32-bit IP address structure was altered to allow for a much larger number of networks (about 2 million) through the introduction of Class A, B, C and D networks. Eventually even this method proved inadequate and Classless Inter-Domain Routing (CIDR) was introduced to create additional flexibility in subnetwork sizing.

Not long after the roll out of TCP/IP, it became clear that the routing table (the so-called Host.Txt file) would not scale to the tens to hundreds of thousands of hosts on the Internet so the table was replaced by a distributed design for host naming called Domain Name System (DNS) and was developed beginning about 1984 by Paul Mockapetris and Jon Postel. This hierarchical system has scaled to tens of milliions of hosts in the network. The commercialization of the Internet brought new requirements to the DNS, particularly with regard to competition, intellectual property protection and dispute resolution.

The growing Internet created its own requirements for scalable routing and that led to the development and evolution of the so-called Border Gateway Protocol. This system was designed to be used in conjunction with some care in IP address assignment so as to allow consolidation of the global routing tables to keep their absolute size and update requirements under control. Intranet routing protocols such as IS-IS and OSPF reflected similar needs for scalability and hierarchical structure.

The evolution of email brought a new set of requirements, some of which were reflected in extensions to the DNS—in particular, the creation of the MX record to allow a single MTA to serve multiple apparent email domains.

The advent of voice over IP has created new requirements for interlinking of the public switched telephone network and/or private branch networks with the public Internet and with private IP networks. The mechanisms proposed to support this requirement are found in the ENUM extensions to the DNS. In effect, ENUM maps international E.164 telephone numbers into Universal Resource Identifiers through a potentially iterative system of regular expression evaluations embedded in the DNS. A new DNS record type, the Naming Authority Pointer (NAPTR) is the active component of this design.

The creation of the World Wide Web has underscored the value of streaming audio and video and real-time protocols for interactive services including multi-party role playing games.

The notion of virtualizing computing and storage resources on the Internet has led to the concept of the global GRID and the so-called GLOBUS protocols. Among these are SOAP, XML, UDDI and others.

A major missing requirement from the initial Internet design was security and authenticity. Although much effort went into making the system robust against random failures, it was not proof against deliberate attacks. There was a version designed for use by the military that made heavy use of end-to-end packet encryption and that version was intended to provide considerable security. However the details were and are classified and the public system does not have these features, yet.

References

[1] Cerf, V.G. and Kahn, R.E., "A Protocol for Packet Network Interconnection", *IEEE Transactions on Communication Technology* **COM-22**(5), pp. 627–641 (May 1974).

[2] Cerf, V.G., Dalal, Y., and Sunshine, C., "Specification of Internet Transmission Control Program", Request for Comments, RFC675, Network Working Group, (December 1974), http://www.faqs.org/rfcs/rfc675.html.

[3] Leiner, B.M., Cerf, V.G., Clark, D.D., Kahn, R.E., Kleinrock, L., Lynch, D.C., Postel, J., Roberts, L.G., and , S. Wolff, "A Brief History of the Internet", Internet Society (ISOC), Reston, VA (2000), http://www.isoc.org/internet/history/brief.shtml#Introduction.

Proceedings of the 11th IEEE International Requirements Engineering Conference
1090-705X/03 $17.00 © 2003 IEEE