# Looking back at the ARPANET effort, 34 years later



By David C. Walden (http://www.walden-family.com)

In late 1968, the Advanced Research Projects Agency (ARPA) of the US Department of Defense announced the award to Bolt Beranek and Newman Inc. (BBN) of a competitively procured contract to develop the backbone of packet switches for the ARPANET.

The ARPANET initially consisted of packet switches called Interface Message Processors (IMPs), long-distance leased telephone circuits between pairs of IMP, and host computers connected to directly to an IMP at the host computer site. Network users were connected to their local host computers and could specify for their local host computer to communicate with a host computer connected to an IMP at another site. A set of communication protocols specified the interface between the hosts and the IMPs and conventions for communicating between a host connected to one IMP and a host connect to another IMP.

Eventually, I think it is fair to say, the ARPANET evolved into what has become the Internet. Thus, the ARPANET activities in 1969 were an important step in the development of what is arguably the most significant societal and technology development since the development of the computer itself. The massive distribution of computation to the desktop that we know today was directly related (a) to the decreasing price and size of computation, and (b) the computer networking advances that followed and to a significant extent grew out of the ARPANET, e.g., Local Area Networking and the connection of these networks together into the internetwork of LANs (https://broadbandnow.com/guides/optimize-wifi-network-faster-speeds) and WANs that covers the world today.

Many people and institutions were involved in the ARPANET activities in 1969 including ARPA itself, the ARPANET host sites (the first four being at University of California at Los Angeles, Stanford Research Institute, University of California at Santa Barbara and University of Utah), Network

Analysis Corporation, and others. By virtue of its contract to develop the packet switches, BBN was a central player during the first ARPANET year, 1969.

BBN's contract from ARPA was a one-year contract to develop and deliver a backbone network of four IMPs, and on August 30, 1969, the first IMP was delivered to UCLA. By the end of calendar 1969 the first four IMPs had been delivered. Naturally, the year of effort at BBN was intense, and I had the good fortune to be part of this invigorating activity.

While I don't remember the exact dates, I think by 1968 *Robert Kahn* (https://www.livinginternet.com/i/ii_kahn.htm), who was aware that there would be a competitive procurement for the ARPANET backbone of IMPs, had convinced BBN management that BBN should prepare itself to bid on this procurement when it came out. BBN pulled a team together under the supervision of Frank Heart that included Bob, Severo Ornstein, me and perhaps others to think about what we would bid once the Request for Proposal came out from the government. In 1968 the RFP did come out and a number of people from throughout BBN helped draft and review the proposal. Before the proposal was submitted, Will Crowther was recruited from MIT Lincoln Laboratory (where Frank, Severo and I had previously worked) to join our team. By roughly the time the contract was awarded Ben Barker join us from Harvard where Severo had met him in a class Severo taught. Also, soon after contract award, Bernie Cosell, who had been working on another project at BBN was added to the team. Thus, the primary team that worked on developing the IMP for the first year consisted of

- Frank Heart (team leader)
- Robert Kahn (communications theorist and the person who did much of the communication with the external world)
- Severo Ornstein and Ben Barker (hardware development)
- Will Crowther, Bernie Cosell and I (software development).

Others from BBN contributed or got added full-time to the project later in the year.

In retrospect, our approach to bidding on the RFP was pretty smart. We had decided to submit a fairly detailed design, including initial hardware designs, a software architecture and fairly detailed initial timing analysis, principles of system operation, and so forth as part of our bid. This level of detail helped us (I suspect) win the procurement. It also left us in a fairly advanced starting position at the beginning of the actual contract. This helped us finish in the specified one year, removed a lot of uncertainty from the beginning of the implementation period (we were sure we could do the design and development on time), and enabled us to begin the actual development period with what was in effect a second design cycle. Often, people talk about developing good designs before starting implementation on complex or difficult implementation projects. It was the inclination of our team to actually do so.

Regardless of how well prepared we were to begin the design and development effort, the year was an intense one, and one that was under intense scrutiny with lots of collaboration from ARPA and the members of the host community. In particular, BBN had contractual commitments to deliver a series of design and specification documents reports during that first year.

Upon rereading some of these reports from the first year (BBN Reports 1763, 1783, 1837, 1890, 1822, and 1928) a few years ago, I was struck by a number of design characteristics.

> 1. Many features to make the IMPs run reliably and with minimal on-site assistance and with cross-network diagnosis, debugging, and new releases
>
> 2. Considerable facilities for network monitoring and measurement
>
> 3. No constraints put on the data hosts could exchange over the network
>
> 4. Highly successful initial algorithms for IMP-to-IMP communications and network routing (both were changed over time, especially the latter, but they did an excellent job in their time, and provided in the initial implementations in the Internet of a system of positive acknowledgments and time-outs and a distributed algorithm for routing)
>
> 5. Much less successful initial algorithms for Host-to-IMP and source-IMP-to-destination-IMP communications-the former was too limited because of the assumption of a direct electrical connection rather than a remote communications interface, and the latter was simply inadequate to the congestion control and multiplexing task it was designed for
>
> 6. A design and implementation that was very high performance in terms of use of memory and machine cycles and very reliable in terms of the IMPs not crashing because of coding bugs.

Although there were some missteps, the initial IMP design and implementation was quite robust and provided good support for the host experiments and a powerful mechanism for releasing incremental improvements as they were needed. In fact, aspects of our original design can still be seen in how the Internet works today.

As I reflect back to that first year, now so many years ago, I am also struck by the general competence of the effort and team certainty of successful completion. Today, nearly thirty-five years later, people often ask me whether I was worried to be a member of a team that had so much to accomplish in only one year. Of course developing that first IMP system was a relatively small project compared to the massive extent of what people think of today when they think of the Internet. We also knew we had a tight schedule, and we worked very hard. However, I didn't see any real worry from any member of the team at anytime. We were a small team of highly motivated and, on average, highly experienced people that worked well together during that first year. We were one of

those “hot teams” that sometimes get written up in management books. We were very focused — the team was enormously pragmatic and concentrated on getting a system delivered on time that worked “well enough.”

Of course, much of the ARPANET development effort happened at other places than BBN — at ARPA, UCLA, SRI, UCSB, the University of Utah, and Network Analysis Corporation. While we at BBN (and I personally) participated intensely in the broader efforts of the ARPANET community, others are better able than I to reflect on the non-BBN parts of the story. I can only say what a joy it was for me to be part of the ARPANET effort.

Dave Walden
East Sandwich, Massachusetts
February 2003