

Document Title A Most Appropriate Place

Repository ID TI.155

Persistent URL <http://doi.org/10.26869/ti.155>

Authors *Ken Klingenstein*

Publication Date: April 16, 2021

Sponsor: Internet2

A Most Appropriate Place

In Boulder Colorado, where a thousand miles of plains roll in from the east to meet a thousand miles of mountains stretching west, there is a federal science lab built high on a mesa. NCAR (the National Center for Atmospheric Research) is a unique building in a remarkable location, designed by I.M. Pei to resonate with Anasazi inspiration and blend into the Flatiron Mountains rising sharply behind it. It has traditionally been a home for atmospheric science in the US, housing scientists and research instruments, computers and data. It was also the place where, in the summer of 1985, the plan for a scalable and viable Internet was crafted. A vision of a network of networks took shape, connecting supercomputers and scientists. Core protocols and architectures were reassembled from small and specific niches into a larger ensemble that could scale broadly, perhaps globally. Those summer conversations and meetings at NCAR, among a small set of key technical and science leaders, marked a subtle but critical transition from thinking about networks to thinking about internets that connected networks together. It was one of those rare moments where purpose, people and place align and something transformational happens.

Interestingly, that vision of NSFnet (which soon evolved into the Internet) was not even the central part of the conversations that summer. The primary focus of discussions was around alignment of the several NSF Supercomputing Centers that were then emerging: what hardware, operating systems, storage, etc. would each of them use. While those issues were thoroughly chewed on, the topic of how users would connect to these centers was secondary. It had been assumed that dedicated networks would exist for every center, so that each could guarantee the quality of the computing product they were delivering. Moving the supercomputer center leadership from that narrow position toward a vision of a shared network infrastructure was just another issue that summer, even if that decision in turn had the greatest consequence by creating the Internet.

In the arc of technology, some form of global networking was inevitable as the once distinct worlds of telecommunications and computing began to entwine. Indeed the telecommunications companies at the time talked extensively of an “information superhighway” that they might build one day, but the form that the Internet actually took – open, scalable, modular, flexible, resilient, extensible, ubiquitous, free and potentially raucous – was a result of the decisions made in a NCAR conference room poised under the Flatirons, above the plains. Those decisions, in turn, became, over the next few years, blueprints, calls for proposals, grants and contracts, deployments, and the wild invention that ensued. That these

decisions happened in a unique location, where geology and geography combined to lead the eye to a distant horizon, seems most appropriate.

Blame It on the Blue Line

In 1959, Boulder passed a law called the Blue Line, which placed a restriction on building along the escarpment of the Front Range. Designed to protect the remarkable views of the Flatirons, the imaginary Blue Line followed a 5,750-foot altitude contour from Eldorado Springs in the south up to the northern city boundary a few miles along the Front Range. It marked the border west of which the city would no longer provide water. After 1959, nothing has been built above the Blue Line. Except for NCAR.

Indeed, if NCAR was the appropriate place to sketch out the Internet, Boulder thought that high on the mesa, below the Flatirons and above the Blue Line, was the appropriate place for NCAR. Barely several months after the original city charter vote to create the Blue Line, the residents of Boulder overwhelmingly passed the only amendment ever done of the Blue Line, allowing NCAR to be built on the mesa, with the land deeded to NSF, and thereby creating the location that would help converge the forces to assemble a plan for an internet.

The Blue Line, beyond limiting water, had several other impacts. Most notably it meant that the resulting lack of commercial or residential development caused AT&T, the dominant communications operator of the time (dominant as in monopoly) to not run advanced telecommunications infrastructure above the Blue Line to NCAR. Yet NCAR had a large amount of data and computing power to be shared with the atmospheric research community around the country and so, with no land-based alternatives, satellite became the way for specific science projects at NCAR to provide that access. As a shared science site, it was oriented towards open-source software and had experience in satellite-based open network protocols. The unique expertise on how to do that, developed to manage NCAR's location, helped draw the meetings to Boulder.

Satellite based networking had significant technical challenges but offered compelling economic and deployment options. The technical challenges stemmed largely from the time delays in sending signals 22000 miles into space and back. Compared to the conventional landlines at the time, the signal delay to satellites was problematic to everything from echoing a character being typed on a keyboard to keeping the pipes full doing large data transfers. (Research at that time showed that the half second delay between typing on a keyboard and seeing the character appear on the screen was very difficult for scientists to handle. Ultimately, workarounds began to emerge to help users manage the time delays.) Despite those challenges, satellite-based networking had several benefits around the "last-mile" problem. The last mile refers to providing connectivity from the larger network to the actual location where the computer was housed. Last miles were often very expensive if in fact a solution could be found; often the local telco had no expertise and couldn't provide any service offering. The last mile for satellites was often just a few hundred feet of institutional wiring from the dish to the campus network. And so, if the approach would be to use satellites as part of the NSFnet

telecommunications infrastructure, and run open network protocols on those satellite links, then NCAR, above the Blue Line, was the place to meet and test the ideas.

From networks to internetworks

The initial research activities that created the concept of networks and TCP/IP protocols are well documented. Beginning in the late 1960's, a small core group of researchers, funded by DARPA, built a 4 node testbed network that ran NCP (a precursor to the TCP and IP protocols). Not only were the technologies well-crafted, but in their principled approach set an architectural paradigm for much of the development that followed. A few design principles were paramount: Keep things simple. Use a layered approach. Allow complexity at the edges. Design for robustness. Plan for change. From those principles flowed a modular design with replaceable elements and distinct interfaces with other modules. At the same time, a process was established to iteratively build using those designs. The famous Internet phrase "we believe in rough consensus and running code" became a shibboleth about the development process.

These prototypes grew in an hoc but constrained fashion over a decade. From those first packet-switched experiments in 1969 through early 1985, a set of mission-driven networks moved the art of TCP/IP packet-switched networking along. Central to that was the ARPAnet, a significantly expanded version of the original testbeds funded by ARPA. Other networks such as Milnet (the military R&D network), DoEnet (connecting energy research labs) and CSnet (computer science departments at universities), as well as dedicated networks connected to some federal supercomputer sites added to the landscape of mission-driven TCP/IP networks which increasingly had a few common touchpoints. The Internet as a concept grew from this connection of dedicated networks. Before the summer of 1985, the emphasis had been on the network; after the summer the story had moved to the internet word. It was a critical revisioning of the world, with broader implications. Through these summer meetings, a rough beast was being born, with the catalyst of NSF.

NSF was an appropriate organization to be interested in an internetwork infrastructure. NSF had traditionally supported a few strategic supercomputer centers in key computational sciences to provide computer power and data storage for those research domains. Each of these supercomputer centers in turn were looking to connect their resources to key researchers across a diversity of US universities. NSF began to realize the scaling issue and wondered if a more generalized infrastructure was possible. NSF also had a flexible leadership approach that allowed campus "rotators" to come into NSF positions for 1-2 years, do mischief and ruffle feathers, and then safely rotate back to their regular positions. A reasonable way to encourage vision and risk-taking. And into a rotator role in scientific computing stepped a most appropriate visionary risk-taker.

Dennis Jennings is a tall and charismatic Irishman, whose height and rich brogue voice allows him to command a room. That combination was helpful to sell the vision that Dennis and a few others had of an Internet not as a purpose-limited and dedicated set of connections but as a

more ubiquitous utility. To get there would need an alignment of technology, deployment, and business model. It would take both vision and sales.

Validating the vision

In that summer of 1985, Jennings convened a first meeting in Boulder to explore the technical and business dimensions of what an Internet could look like. Boulder was chosen for several reasons, including its relatively central US location, NCAR's unique expertise in satellite-based networking, and a place where the views might encourage sweeping ideas.

Three ideas needed to be verified.

The first was that a general purpose shared network infrastructure could also sustain the specific science performance requirements that dedicated networks could provide. At the time, network traffic tended to be generated by two scientific use cases with very different requirements and it was unclear that a single architecture could serve both needs, especially at scale. The first use case was sending single characters, as in typing a command to a remote computer across the network, with a reasonable response time. (Early experiments showed that if the response time was too long to echo a character back on screen after pressing a key, it was a major disruption to cognitive processes.) The second use case was sending large chunks of scientific data across the network in a bulk file transfer, where the critical network characteristic was good throughput. This classic networking tradeoff – good response time vs high throughput – was now going to be accommodated across a much larger shared infrastructure. One that might blossom over time with rich new uses. Was that possible?

The second idea was that a viable and scalable implementation could be approached by running a single robust high-speed national backbone with regional networks covering several states hanging off that backbone. This deployment model needed to be tested for both technical and financial merits. The technical approach had to provide performance, offer some robustness, and be manageable. And while the model being sketched in the room was going to be funded initially by scientific grants from federal agencies, most notably NSF, it was important to have the possibility of a business model. Could the regional networks being proposed find paths to sustainability? Financially, did the economies of scale in running a regional hub and then sharing longer distance costs work? And was the loss of control that users of a dedicated network might feel be balanced with the improved operations that a regional staff could provide?

The last issue was one of network protocols. The two major computer companies of the time – IBM and DEC – had developed proprietary mechanisms to connect their computers together. These protocols were licensed and expensive and optimized for the vendors computer architectures, which in turn were oriented toward the businesses they marketed to. (IBM to a corporate marketplace; DEC to a science one.) Effective in those niches, such networks had problems serving the diverse set of requirements that NSF had identified. Moreover, there was a higher-order principle in the NSF user community that the network protocols be open and not proprietary. That driver led to consideration of TCP/IP as the network protocol. Moreover, the network should not be optimized to support a particular set of transactions, but instead offer a core infrastructure that all use cases could leverage. There

was yet another major network protocol capability to consider. The network would operate over a variety of communications media, including land lines of various flavors and satellite based communication. Could TCP/IP could be run at adequate performance across a satellite?

At one point in the meeting, a step ladder was used to place the satellite numbers and data high on the ten foot white board that lined the conference room wall, allowing room at a lower level to be used for more dynamic writing and calculations. One by one, the answers to the above questions were confirmed. Moreover, there was an interplay among the issues that reinforced the way they could leverage each other. For example, one of the ways to improve the performance of a satellite based telecommunications path was to modify software to allow an exceptionally large number of packets in transit. Because of the open source code being used, those modifications could be made and shipped as a high performance version of the software. Campuses confirmed their interest to build a campus-wide network to enable all type of NSF researchers. On occasion, some of those involved in the conversations would leave to the room to walk among the rocks on the mesa and gaze out across the high plains at the distant horizon. Something consequential was taking shape.

Sales

That the vision was validated at the design meeting in Boulder then led to a second meeting that summer at NCAR with the NSF supercomputing center leadership teams. They had to be convinced on several points: that there should be a single network connecting them instead of dedicated networks per site; that this single network would be general purpose for all the NSF research interests; and then critically that the supercomputing centers themselves should build and operate this inter-network of networks.

It was not an easy sell. The core business of the centers was supercomputers, not networks. Their primary interest was in serving their select sets of users, not the broader research community. They were not funded to do the work. The resistance to the idea was a shared opinion among normally competitive facilities. Over the course of the meeting, however, the compelling potential of the network changed their attitude. Finally, sitting around a picnic table at lunch the last day, one of the centers took the lead and asked Dennis if a backbone network proposal was submitted by them to NSF, would he fund it? Dennis said he would do everything in his power. Standing up, they said if that was a guarantee, they would call and place the circuit orders right then. Dennis reached in his pocket and took out a dime for them to use the pay phone and make the call. With that guarantee of a dime, it started.



The room where it happened

A Place in Time

And so it happened. Under the primary focus of serving the science community, NSF would catalyze a powerful general purpose Internet capability, one whose potential stretched off like the vastness of the Plains below. In those meetings at NCAR in Boulder, it wasn't that those

present knew what specifically was coming. There was no notion of the web, rich media, streaming, etc. The applications then were almost all email and telnet – a command line interface for logging onto a remote computer. But even with that primitive set of network applications, the power was obvious. The most powerful technology since the printing press took shape that summer.

One can speculate what might have taken shape if the meetings in Boulder in 1985 hadn't happened, if Dennis Jennings hadn't catalyzed the Internet we know today. Almost certainly, some integration of computing and telecommunications would have happened, as the leverage the two technologies offer each other is clear. But likely the character of the integration would have been quite different. If the telecommunications companies had actually implemented their vision, the result might be a "managed" global network, with rigid points of entry and limited capabilities. If some other national government initiative would have happened, it's possible that the policy aspects of the Internet would not be the free agora it is today. If proprietary technologies had been chosen, the wild innovation that ensued would have been constrained.

There are several places to sit among the rocks and under the ponderosa pines at NCAR and get perspective. Satellite dishes still sit discreetly next to the building. The picnic table still looks out over a vast landscape, a horizon as seemingly unbounded as the invention that began there. For contemplation on what was started here, it remains a most appropriate place.