

The Route Less Taken: The Homegrown Los Alamos Integrated Computer Network

Nicholas Lewis

► **To cite this version:**

Nicholas Lewis. The Route Less Taken: The Homegrown Los Alamos Integrated Computer Network. IFIP International Conference on the History of Computing (HC), May 2016, Brooklyn, NY, United States. pp.1-22, 10.1007/978-3-319-49463-0_1 . hal-01620138

HAL Id: hal-01620138

<https://hal.inria.fr/hal-01620138>

Submitted on 20 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



The Route Less Taken: The Homegrown Los Alamos Integrated Computer Network*

Nicholas Lewis

History of Science, Technology, and Medicine Program, University of Minnesota,
Minneapolis, Minnesota USA, Charles Babbage Institute
lewi0740@umn.edu

Abstract. Between the 1970s and 1990s, Los Alamos National Laboratory built and utilized a largely custom computer network for the Lab's supercomputers. Designed to support the unusual performance, storage, and security requirements of an American weapons lab, the Los Alamos Integrated Computer Network, as the focus of historical study, complicates and enriches the history of computer networking development, exploring the approaches and contributions to computer networking of an institution outside the better-known worlds of industry, academia, and the military. For example, the Lab's reticence to adopt TCP/IP due to performance and security concerns further complicates the narrative of the ARPANET/Internet protocol suite's adoption among advanced networking sites in the 1980s and 90s.

Keywords: Los Alamos National Laboratory, supercomputing, computer networking, network standardization, ANSI, TCP/IP, history of computing.

1. Introduction

Since the Second World War, the Laboratory at Los Alamos has engaged in advanced scientific computing as part of its primary mission as a nuclear weapons lab, and increasingly for other fields of computing-intensive scientific research, such as climate modelling. Because of its unusual computational demands, Los Alamos, during and after the Cold War, often sought levels of computing performance, and hardware and software capabilities, before high-tech vendors were capable or interested in meeting those requirements. As a consequence, Los Alamos had a history of developing custom solutions to computing-related problems, particularly where new technologies were not yet available on the market, or existing systems lacked key features or capabilities Lab supercomputer users desired. Homegrown solutions were costlier in terms of development and maintenance, but were typically well-tailored to the Lab's technical and security requirements. The custom approach also provided in-house knowledge and expertise for technologies important to the

* LA-UR-16-20103

Lab. The Los Alamos Integrated Computer Network of the 1970s and 80s was a prime example of the Lab's homegrown computing solutions.¹

The Lab's custom approach toward computer networking, followed by its adoption of outside networking technologies, such as TCP/IP, from the late 1980s through the mid-1990s, offers a window onto several important facets of networking history. First, it traces the maturation of computer networking over the course of two decades from the perspective of an American weapons laboratory, with its unusual networking requirements interacting with the capabilities and business interests of an emerging industry. Second, it highlights the alternative approaches and contributions toward computer networking of an institution outside academia, industry, or the Department of Defense. Finally, it builds upon historiographical works, such as Janet Abbate's *Inventing the Internet* and Andrew Russell's *Open Standards and the Digital Age*, which argue against previous common conceptions of how and why TCP/IP spread as a networking standard, demonstrating that the DoD's protocol suite was not always adopted readily at computer networking sites. Indeed, Los Alamos considered TCP/IP to be a technical step backward, in terms of security and performance, over the custom Lab protocols, and initially rejected using the DoD suite in the early 1980s. Networking development at Los Alamos followed a largely unique path during its first two decades, attempting to integrate slowly with outside standards and technologies as the high-performance computer networking market and industry evolved appreciably in the 1980s and 90s. However, various internal, but primarily external pressures, brought a far more rapid end to the homegrown integrated computer network than Los Alamos had anticipated.²

2. Early Lab Computing

Founded in 1943 as the primary R&D site of the Manhattan Project, the Laboratory at Los Alamos emerged from World War II as a major center of scientific computing in the United States. Using, building, and buying some of the earliest digital computers,

¹ For an overview of Los Alamos and its computing history, see Donald MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," *Annals of the History of Computing* 13, No. 2 (1991).

² For detailed discussions of how and why the TCP/IP suite developed and became the standard set of networking protocols in use today, see Janet Abbate, *Inventing the Internet* (Cambridge, MA: MIT Press, 1999), Andrew L. Russell, *Open Standards and the Digital Age: History, Ideology, and Networks* (New York: Cambridge University Press, 2014), and Laura DeNardis, *The Global War for Internet Governance* (New Haven: Yale University Press, 2014). These works challenge previous depictions of the history of ARPANET and the Internet that depicted the development and spread of TCP/IP as open and uncontested processes, when, in fact, the opposite was true. The protocol suite resulted from closed-door planning and implementation decisions, and entered widespread usage largely as a result of the US Department of Defense exerting financial and other forms of coercion on networked institutions and vendors to adopt the suite.

Los Alamos played a role in the maturation of computer hardware and software in the first decades of the Cold War, going back to its use of ENIAC on that computer's very first calculation. As part of its collaboration with IBM on the design of the IBM 7030 "Stretch," Los Alamos helped to create and define the concept of supercomputing. Research and development on operating systems, programming languages, and compilers also accompanied the Lab's work on hardware, creating the software tools which vendors like IBM, Control Data Corporation, and Cray Research either could not, or would not, provide for the small (but growing) niche of high-performance scientific computing.³

Beginning in the early 1950s, when digital computing was still in its infancy, Los Alamos actively engaged with the broader, though still small, scientific computing community, exchanging correspondence and software with sites like the RAND Corporation, and with other researchers, taking part in conferences and user groups, like IBM SHARE, of which Los Alamos was a founding member. The purpose was to share in the community's pool of knowledge and experience concerning a new technological field. These exchanges continued and became more elaborate and important over time, with Los Alamos directly funding R&D at universities and among vendors, including IBM, RCA, and Burroughs, on computing technologies, such as high-speed storage and memory, which weren't feasible or cost-effective to research at the Lab. Participating in the scientific computing community and encouraging R&D activities outside the Lab were considered to be investments that would reduce the need for costlier in-house development. However, in-house development was a common practice in Los Alamos computing, particularly along new technological branches, like high-performance and high-security networking in the 1970s.⁴

³ MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories," 186; Memorandum from Walter H. Brummet, Jr., Chief Contracts and Procurement Branch, to George Udell, Director, Supply Division, Nov. 17, 1955, "Report of Conferences in Washington, D. C. Regarding Proposed Purchase of Computing Machine for LASL," Edward A. Voorhees Papers, Box 16, Folder 6 Stretch Solicitation and Selection, LANL Archives; Charles J. Bashe, et al., *IBM's Early Computers* (Cambridge, MA: MIT Press, 1986), 430, 432; For a detailed assessment of the intersection between computing and the development of nuclear weapons at Los Alamos, see Anne Fitzpatrick's Ph.D. dissertation, "Igniting the Light Elements: The Los Alamos Thermonuclear Weapon Project, 1942-1952" (The George Washington University, 1999); Thomas Haigh, Mark Priestley, and Crispin Rope, "Los Alamos Bets on ENIAC: Nuclear Monte Carlo Simulations, 1947-1948," *IEEE Annals of the History of Computing* 36, No. 3 (July-Sept. 2014) examines specifically Los Alamos' use of ENIAC for the first simulations of nuclear weapons on a digital computer.

⁴ Correspondence from Bengt Carlson to Carson Mark, May 23, 1960, "UHS Computer Study Contracts," Edward A. Voorhees Papers, Box 32, Folder 1, Series 25, SPARC Correspondence, LANL Archives; Correspondence from Edward A. Voorhees to Paul Armer, Head, Computer Sciences Department, The RAND Corporation, March 17, 1960, Edward A. Voorhees Papers, Box 4, Folder 10, Series 9, File 1, Voorhees Correspondence 1960, LANL Archives; Correspondence between T. L. Jordan and Donald C. Cashman,

3. C Division and Custom Lab Networking

In 1968, Los Alamos formed its first dedicated Computing or “C” Division, its personnel drawn from the three computing-related groups within the Theoretical Physics or “T” Division. T Division had led the original Los Alamos hydrogen-bomb project, and was still the most computing-intensive division in the 1960s, occupying 40% of the Lab’s available computing time. Transferring oversight of the Lab’s supercomputers to a dedicated computing division was intended to unify the Lab’s computing efforts, as multiple science divisions were purchasing their own machines, and were pursuing differing technical paths from one another. A dedicated computing division was also intended to mitigate the potential conflict of interest found with the most computer-dependent division allocating access to all other divisions.⁵

Between the mid-1960s and the formation of the new C Division, T Division faced a rapidly growing demand from users on Lab computing resources. The actual computing power at the Lab’s Central Computing Facility (CCF) quadrupled every two years between 1966 and 1974, but this increase still did not meet the user demand. With the mission-critical nature of Lab computing resources, particularly amid the escalating demands from users, T Division was reluctant to adopt new technologies that had the potential to disrupt vital services or security, at least until those technologies had undergone extensive testing before entering widespread use at the Lab. Computer users from other divisions who were dissatisfied with T Division’s conservative approach did not find significant change with the formation of C Division under its first leader, Roger Lazarus. A physicist and Los Alamos veteran since 1951, Lazarus was an expert on the application of digital computing for nuclear weapons development. Lazarus likewise resisted the adoption of technologies that, among other factors, might impinge on the efficiency or availability of the Lab’s supercomputers for weapons work.⁶

SHARE Distribution Agent, January 13, 1961 through June 27, 1961, Edward A. Voorhees Papers, Box 4, Folder 18, Series 9, File 2, Jordan Correspondence 1960-61, LANL Archives; MacKenzie, “The Influence of the Los Alamos and Livermore National Laboratories,” 189-193.

⁵ “Use of Computing Time by Division,” Edward A. Voorhees Papers, Box 67, Folder 3, Series 48, Monthly Summary Sheets 1961-1964, LANL Archives; “New Division Formed at LASL,” *The Atom* 5, no. 5 (May 1968), 17-18

⁶ The increasing user demand upon computer resources at Los Alamos between the mid-1960s and mid-70s was due to a variety of factors: First, the Partial Test-Ban Treaty of 1963 moved all nuclear testing underground, which made weapons tests more expensive, and increased the reliance upon computer simulation. Second, after having diversified in their research, the US weapons labs experienced a reducing budget year-over-year beginning in 1960, spurring Los Alamos to transfer large numbers of personnel back toward its core mission, which increased the number of heavy computer users, despite the Lab population remaining stable. Third, the weapons themselves became more sophisticated and design tolerances became tighter starting in the mid-1960s, as US nuclear strategic policy shifted toward a counter-force (targeting military assets) strategy. A mix of more complex weapons

For example, Lazarus resisted the development or adoption of time-sharing operating systems at Los Alamos, believing that time sharing, in its early stages in the late 1960s and early 1970s, posed security risks, and wasted clock cycles at a time when the Lab still could not keep up with the growing demand for computing power. Computer users from heavily computer-dependent divisions, such as P (Physics), MP (Material Physics), and W (Weapons), found this approach to time sharing to be too cautious, and successfully petitioned for Lazarus to be replaced in 1973. His successor, Frank McGirt, quickly entered into collaboration with the Lawrence Livermore Lab to utilize its LTSS (Livermore Time-Sharing System) operating system on the CDC (Control Data Corporation) 7600 supercomputers at Los Alamos. Livermore's Computation Division, under Sidney Fernbach from 1955 to 1982, typically approached new technical solutions more readily than did Los Alamos' C Division under Lazarus. Curiously, Fernbach was less open than Los Alamos to the potential of smaller computer systems in the 1970s, as he remained focused on large machines. He usually delegated responsibility for smaller systems to others in the Computation Division.⁷

Cautious toward the new arena of computer networking, Roger Lazarus oversaw the construction of an experimental remote-access system initiated under T Division. Called "MUX," the Multiple User eXperiment was meant to determine the utility and viability of remote computer connections at Los Alamos. MUX provided up to 75 users the ability to run batch jobs on one of the Lab's CDC 6600 supercomputers via remote terminals at 300bits/s. Its specifications completed in 1967, and the system entering use in 1968, MUX consisted of about \$50,000 worth of hand-built hardware and custom application software, using 8,000 integrated circuits and two miles of twisted-pair cable. The custom MUX application software resided on the 6600 and

designs and more complex codes used for development and simulation also markedly increased the processing and storage demands of the heaviest users of Lab computing resources between the 1960s and 70s. Part of T and then C Division's conservatism with new technologies was the need to prevent disruptions in the already overbooked CCF's production schedules.

⁷ H. Butler, MP-1, and T. Gardiner, P-1, to Edward Voorhees, CADP, 11-29-66, "Delay in Procurement of Time-Share Service," Edward A. Voorhees Papers, Box 35, Folder 2, Series 27, LASL Computer Needs, LANL Archives; Bill Collins, interview by Nicholas Lewis, Minneapolis, MN, February 19, 2015; Michael, George, and Marilyn Ghausi. "An Interview with Sidney Fernbach." An Interview with Sidney Fernbach. Accessed March 20, 2016. <http://www.computer-history.info/Page1.dir/pages/Fernbach.html>. Information regarding Sidney Fernbach's potential blindspot for smaller computer systems is derived from George Michael's introduction to Marilyn Ghausi's February 1989 interview with Fernbach. Fernbach stepped down as the Livermore Computation Division leader in 1982, with Fernbach arguing that he had been removed. Fernbach at Livermore had heavily backed the development of the Control Data STAR-100 supercomputer in the early 1970s, but the poor performance of that system (of which Livermore purchased two, while Los Alamos backed out of its contract to purchase a single STAR) had consequences for Fernbach's position, and for the relative freedom of the two labs in their selection of new computing systems.

one of its Peripheral Processing Units (PPU), which handled I/O operations for the CPU. The application placed remotely submitted jobs into the 6600's job queue, then returned completed jobs to users via the terminal interface built by the MUX team. MUX was not a time-sharing system, as users could not interact with their jobs as they ran on the 6600, but it offered a platform for introducing and testing time-sharing-like capabilities in the Lab setting. Ron Christman, one of MUX's three original developers, recalled that the system was reliable to the point that, when it did fail on rare occasions, its maintenance staff had to "re-learn" how MUX worked in order to repair it. The experience with MUX, as a tentative first step into networking technology, characterized C Division's cautious approach to networking development between the 1960s and 70s.⁸

MUX was enormously popular with users, and remained in use for over a decade. While quite limited in its capabilities, MUX indicated to Lazarus and C Division that a demand for computer networking existed at Los Alamos, and that the technology could be implemented securely and reliably in the Lab setting. While Lab supercomputer users were, indeed, enthusiastic over the potential of larger-scale networking at Los Alamos, MUX itself did not fully prepare C Division, despite its usual caution, for the technical challenges of building a computer network. Even with inside knowledge of a similar networking effort at another American weapons lab, Los Alamos wrestled with many of the same problems facing other early computer networking sites, but with the added complications introduced with the niche requirements of supercomputing at a secure facility.⁹

4. Hydra

In 1970, with the success of MUX, Lazarus authorized development of the Lab's first computer network. The network came online in 1974, one year later than planned. Ron Christman, also a designer of the new network, indicated that the delay was largely due to an overly ambitious early design, and insufficient coordination within the development team for such a complex system. The relative simplicity and small scale of MUX did not prepare the C Division networking team for the challenges of interconnecting multiple supercomputers in the Central Computing Facility (CCF). The resulting network was, by necessity, greatly scaled-back from the original plan, with proposed features like graphical interface capabilities stripped from the final design. Called "Hydra", the network initially connected each of the Lab's four CDC 7600 supercomputers to a shared pool of short-term, long-term, and archival storage systems through a central front-end machine, itself a Control Data 6600

⁸ Ronald D. Christman, "MUX, a Simple Approach to On-Line Computing," *Communications of the ACM* 15, No. 5 (May 1972), 319-329; Jeffery L. Peterson, "MUX, adieu, faithful friend," *The Atom* (October 1979), 14-16; Ronald D. Christman, interview by Nicholas Lewis, Los Alamos, NM, July 2, 2014.

⁹ Peterson, "MUX, adieu, faithful friend," 14-16; Christman, interview by Nicholas Lewis, 2014.

supercomputer. The front-end machine, doing no production work of its own, also managed remote terminal connections to the four 7600s. Although the original plans called for a new computer to serve as the front-machine, C Division opted to convert one of its production 6600s for use on Hydra due to an initial lack of funds. The one 6600, itself, represented only a small portion of the Lab's overall computing capacity. Christman argued at the time of Hydra's development that the benefits of more flexible, shared storage options would outweigh the small loss in processing power. Roger Lazarus agreed.¹⁰

Lawrence Livermore, beginning development in the mid-1960s, constructed a similar network before Hydra, called "Octopus," which used two DEC PDP-6 computers as its front-end machines. Despite the historical rivalry between Los Alamos and Lawrence Livermore, the two nuclear weapons design labs of the AEC complex, their computing divisions maintained (and still maintain) collaborative relationships on the development, purchase, and usage of HPC systems, as both labs had similar computing requirements and faced similar technical problems. However, particularly between the 1950s and 1980s, the two often followed somewhat different technological paths, as in the case of networking development, due to differences in funding, lab-specific goals, particular demands from users, and styles of leadership. Livermore was founded in the 1950s largely to stimulate a competitive environment with Los Alamos, with either lab pursuing similar lines of R&D, but from slightly different philosophical angles. These tendencies were reflected in their approaches to computing. Los Alamos C-Division staff remained apprised of the problems with Livermore's Octopus through publications and inter-lab technical meetings, and attempted to avoid similar issues with Hydra. Hydra had fewer initial performance and reliability issues than did Octopus, partly due to the higher performance of the CDC 6600 as the Hydra front-end machine, and the lessons learned from Livermore's earlier foray into centralized-storage networking. However, both custom networks had similar drawbacks for maintenance, operability, and reliability. The total dependence on a central front-end machine was the single largest drawback for network reliability. C Division prepared one of the production 6600s in the Central Computing Facility (CCF) as a backup for the Hydra front-end machine, allowing the other CDC computer to take over if the primary 6600 failed, but that solution could not completely address the inherent deficiencies in Hydra's configuration, only minimize their impact.¹¹

The purpose of Hydra, a "star" configuration network, was to reduce the redundancy of data storage purchased for each of the Lab's supercomputers, and to allow users to access their files on any of the machines in the CCF, regardless of

¹⁰ Ron Christman, "A Review of the Hydra Project," (Los Alamos, NM: Los Alamos Scientific Laboratory, 1974); "The Computer Science and Services Division Annual Report, January—December 1975" (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-6228-PR, 1976), 5.

¹¹ "The Computer Science and Services Division Annual Report, January—December 1975," 5; William J. Worlton, Lab Notebook (Los Alamos, NM: Los Alamos Scientific Laboratory, 4/6/71), 27.

whichever machine they had used previously. Aside from the Control Data channel couplers, the point-to-point interconnects which C Division had repurposed to link each of the supercomputers to the Hydra front-end machine at 12Mbits/s, the network was entirely custom to Los Alamos, including the Hydra OS, which replaced the CDC Scope operating system on the 6600 front-end machine. After the MUX development team found the Scope OS to be unreliable and difficult to adapt to the needs of the project, C Division opted to pursue its custom OS for Hydra, despite the greater development costs. Data storage on Hydra consisted of CDC disk units for short- and long-term storage, and the IBM 1360 “Photostore”, a massive archival system that stored data on photographic film chips. The Photostore, one of only six IBM constructed, provided about one-terabit of storage in total.¹²

4.1 The Homegrown Tradition

Retired Lab networking experts Don Tolmie and John Morrison recalled that C Division opted to develop its custom networking and storage system because the commercial market for supercomputer networking did not yet exist in the early 1970s. Los Alamos and Livermore were among the few civilian facilities at that time pursuing the kinds of data speeds, volumes, and security needed for interconnecting supercomputers for classified work. This was not an unusual situation at the Laboratory. In the early 1950s, the digital computing industry was still in its infancy, with very few computers in existence worldwide. In addition to its modified IBM accounting machinery and a Hand Calculation Group of human computers, Los Alamos had spent the immediate post-war years utilizing whatever available computer resources existed outside the Lab, such as the ENIAC in both Pennsylvania and Maryland, and IBM’s electromechanical SSEC machine in New York. The most demanding of these calculations were usually in support of the hydrogen bomb project. However, traveling great distances for computer access, particularly to machines that were not typically well suited to the types of calculations Los Alamos performed in the early Cold War, was not a practical solution.¹³

In the 1950s, Los Alamos cultivated a native capacity for developing and maintaining its own computing solutions. For example, the MANIACs I and II, from 1952 and 1957, respectively, were first-generation digital computers constructed at Los Alamos. These machines provided computing capacity for the Lab when vendor

¹² “The Computer Science and Services Division Annual Report, January—December 1975,” 5-6, 9; Christman, “A Review of the Hydra Project”; Christman, interview by Nicholas Lewis, 2014; John Morrison, interview by Nicholas Lewis, Los Alamos, NM, 7/1/2014.

¹³ Morrison, interview by Nicholas Lewis, 2014; Don Tolmie, interview by Nicholas Lewis, Los Alamos, NM, 7/14/2014; The first ever calculation on the ENIAC was a Los Alamos feasibility study of the hydrogen bomb. The calculation was less than ideal, because ENIAC’s limited memory forced Los Alamos to simplify the variables of the problem to the point that the results were mostly inconclusive. However, Los Alamos would use ENIAC repeatedly in the 1940s, due to the scarcity of digital computers in the early Cold War years.

solutions were still relatively scarce. Just as important, the intimate knowledge of digital computing gleaned from developing the MANIACs also provided the Lab community with in-depth knowledge of a mission-critical technology, which continued forward as Los Alamos transitioned to the maturing array of vendor solutions that appeared from the mid-1950s through the early 60s. The Lab continued to customize vendor hardware as needed, and to create bespoke devices when no others were available. T and then C Divisions of the 1950s and 60s developed much of their own software in-house, due to the relative dearth of software for the niche scientific-computing market, and to maintain a local capacity for producing software tailored to the Lab's requirements. Much was the case in the 1970s with computer networking. C Division opted to pursue its own networking solutions to meet the Lab's unusual requirements while the supercomputer-networking industry was young and provided few ready-made options. Pursuing technically challenging custom solutions, according to the C Division leadership of the 1970s, also fostered local networking expertise, and aided the division in attracting new talent.¹⁴

5. The Integrated Computer Network

C Division began to modify the centralized Hydra network soon after it first entered service, replacing the 6600 with a CDC Cyber 73 as the front-end machine in 1975, when the funding became available. C Division also took advantage of the growing range of minicomputers available on the market to augment network services. Minicomputers became the Swiss Army knives of Lab networking, with C Division customizing UNIX-powered DEC PDP-11s, and later DEC VAX, SEL, and Gould minicomputers, to serve in a variety of roles. For example, in order to connect remote users to Hydra, Los Alamos developed the KCC (Keyboard Communications Concentrator), a custom device based on a PDP-11 and DEC interface unit to provide a single connection to the Hydra front-end machine for up to 96 keyboard-based terminals. The CCC (Computer Communications Concentrator) provided a similar connection for up to 18 computer-based terminals, which were usually other DEC minicomputers.¹⁵

In 1975, when the number of keyboard-based terminals in use exceeded the capacity of a single KCC, C Division developed the PDP-11-based Synchronous Concentrator (SYNC). The SYNC served as a concentrator and message switcher to interconnect the KCCs with the Hydra front-end machine. More importantly, the

¹⁴ Christman, interview by Nicholas Lewis, 2014; MacKenzie, "The Influence of the Los Alamos and Livermore National Laboratories," 186.

¹⁵ "The Computer Science and Services Division Annual Report, January—December 1975," 5; Fred W. Dorr, "Computer Science and Services Division of the Los Alamos Scientific Laboratory" (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-UR-74-1566, 1974), 2; Leo Romero and Bill Buzbee, "Overview of the Los Alamos Integrated Computing Network" (Los Alamos, NM: Los Alamos Scientific Laboratory, LALP-85-34, 1985), 23-24.

SYNC provided direct terminal connections with each of the networked supercomputers in the CCF. Providing direct terminal access to the “worker” supercomputers reduced the centralization of the Hydra system, which C Division increasingly recognized as an unreliable network model. The CDC channel couplers proved to be the single most unreliable component of Hydra, as they were never intended to connect more than two CDC computers together. Troubleshooting a connection error required maintenance staff to take both the originating machine and the Hydra front-end computer offline.¹⁶

As frustrating as Hydra’s shortcomings were to users and to C Division, including its problematic interconnect technologies, a front-end machine that CDC had not intended to be a storage or terminal-access controller, and an interface that was unfriendly to users, they offered a hands-on learning experience with the complexities of networking in the Lab environment, just as the Lab’s home-built computers had with digital computing in the 1950s. Having a clearer understanding of what did and did not work well at Los Alamos, and in supercomputer networking in general, C Division’s networking team developed a network plan for the remainder of the 1970s. The plan emphasized a reduced dependency on Hydra and the eventual phasing-out of the Hydra storage system completely. The plan also called for an increase in network speeds to support supercomputers with faster I/O capabilities, such as the new Cray-1, and growing the network to accommodate a variety of new services and users. The plan meant adding new transport, switching, and security systems, all of which would be custom to Los Alamos, primarily due to the relative infancy of the high-performance computer networking industry as the 1970s came to a close.¹⁷

One of the first products of C Division’s new networking plan was the Los Alamos File Transport machine, introduced in 1977. While the SYNCs provided complete terminal access to the supercomputers of the CCF independent of Hydra, Los Alamos lacked similar direct connectivity for file access. The File Transport (FT) machine used the same architecture and low-level protocols as the SYNC to provide high-speed data links of up to 40Mbit/s between the CCF supercomputers. The first FT linked together the Cray-1 and two CDC 7600s. Additional FT machines formed a File-Transport Network (FTN), creating a high-speed backbone for direct file transfers, which reduced the need for routing data through the Hydra front-end machine. The FTN was also important for serving the rapidly growing number of high-speed graphics terminals at the Lab. The new terminals, once connected to the FTN, allowed users to stream color graphical data at 150kbits/s from the supercomputers of the Central Computing Facility directly to their offices.

¹⁶ Romero, “Overview of the Los Alamos Integrated Computing Network,” 23; “The Computer Science and Services Division Annual Report, January—December 1975,” 5-6; Christman, interview by Nicholas Lewis, 2014.

¹⁷ Christman, interview by Nicholas Lewis, 2014; Fred W. Dorr, “Computer Science and Services Division CCF Five-Year Plan” (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-6921-MS, 1977), 3-4, 13, 43-45, 79; Fred W. Dorr, “Computer Science and Services Division Activities and Plans” (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-7093, 1978), 4, 12, 19.

Conventional terminals operated at only about 9.6kbits/s. Improving support for a growing range of small computing systems would become an important theme in Los Alamos network development over the following decade. While the Lab network would grow far beyond its original size and complexity as the 1980s began, reorienting to accommodate small-scale computing, in addition to adapting to more powerful and data-intensive supercomputers in the CCF, would pose a significant challenge to C Division, as it often meant rethinking how and why computer networks were used at Los Alamos. Those challenges were only emerging during the ICN's first great expansion, taking place in the late 1970s.¹⁸

5.1 ICN Security in the Late 1970s

Increasing the range of services and users on the new Integrated Computer Network in the late 1970s required a significant reconfiguration of network security. The original Hydra network had a simple approach to the complex issue of security at a weapons laboratory; only classified users were allowed on the network. In order to accommodate unclassified supercomputer users, and to facilitate administrative Lab functions across the network, C Division developed a custom, minicomputer-based Network Security Controller (NSC) to monitor and restrict file and resource access based upon user passwords. The NSCs separated the ICN into three partitions: Secure, Administrative, and Open. Several of the supercomputers in the CCF were moved to the Open partition for unclassified use, while the Administrative partition users had dedicated CDC Cyber 73 mainframes. Users of the Administrative or Open partitions could share files with users on the Secure partition, but not the other way around, in order to prevent a "write down" of classified data to a lower classification level. Each data packet ran through two NSC checks, providing redundancy to guard against error. The NSCs did not run any user code, and prevented the need to rely upon worker machines to enforce security rules. Using the NSC to monitor and control network security on the single physical network was far less costly than developing separate networks for each security level, allowing classified and unclassified users to share the same resources.¹⁹

5.2 The Los Alamos High-Speed Parallel Interface

The ICN jumped from about 2,200 mostly local users in 1977 to over 6,000 by the mid-1980s, with an additional 2,000 users remote to Los Alamos. As networked resources became more widely used and more mission-critical to the Laboratory, the demands of connecting larger numbers of machines and services to the ICN grew more difficult and time consuming. The various supercomputers, minicomputers,

¹⁸ Romero, "Overview of the Los Alamos Integrated Computing Network," 24; Morrison, interview by Nicholas Lewis, 2014.

¹⁹ Romero, "Overview of the Los Alamos Integrated Computing Network," 24; Robert H. Ewald, et al., "Computing Division Two-Year Operational Plan, FY 1981-1982" (Los Alamos, NM: Los Alamos National Laboratory, 1981), 41-42.

storage systems, printing services, and other devices employed at the Lab were never intended to function together on a single network. A variety of incompatible interconnect technologies and proprietary communication protocols meant that C Division had to develop expensive hardware and software conversion systems that could translate the native I/O of one machine into that of every other device with which it needed to connect on the network. This engineering task grew increasingly impractical as the network became larger and more heterogeneous as the 1970s drew to a close. C Division's solution was one of the first efforts toward network interface standardization in the supercomputing arena.²⁰

Beginning development in 1978, the Los Alamos High-Speed Parallel Interface (HSPI) converted the native I/O of a computer or another networked device into the standard ICN higher-level protocol suite, which was also developed at Los Alamos. The suite primarily consisted of SIMP (Simple Inter-Machine Protocol), PTP (process-to-process), and AFT (File Transport), which would remain in use for well over a decade, much like the HSPI. The custom Lab protocols were designed to accommodate large data blocks with low latency, high burst rates, extensive error detection and correction, and minimal message passing. C Division had to design and construct a HSPI, including the hardware and drivers, for each new device placed on the network, but each HSPI only had to convert the native I/O interface and protocols of the new device into the standard ICN protocols. The HSPI of each receiving device then converted the ICN protocols back into that device's native I/O. Once connected to its HSPI, a device had full access to the entire Los Alamos network, which significantly reduced the time required to integrate new equipment onto the ICN.²¹

HSPIs provided full-duplex point-to-point transfers to the various machines of the CCF, with error correction integrated into the HSPI hardware. The HSPI supported data transfers of up to 60Mbits/s, which was intended to exceed the I/O capabilities of the fastest devices on the network, so that the network interfaces would not act as performance bottlenecks. In order to accommodate different data rates (for example, 3Mbits/s for the PDP-11, and 50Mbits/s for the Cray-1), a handshake between HSPIs determined the maximum data rates of the exchanging machines. The HSPI employed hardware error detection and correction while data was on the fly, so as not to slow the data transfer. When double-bit errors were detected, the HSPI hardware aborted

²⁰ Dorr, "Computer Science and Services Division Activities and Plans," 5; Ewald "Computing Division Two-Year Operational Plan, FY 1981-1982," 36; Romero, "Overview of the Los Alamos Integrated Computing Network," 43; Tolmie, interview by Nicholas Lewis, 2014; Andrew and David Dubois, interview by Nicholas Lewis, 7/17/2014.

²¹ Donald E. Tolmie, et al., "Interconnecting Computers with the High-Speed Parallel Interface" (Los Alamos, NM: Los Alamos National Laboratory, LA-9503-MS, 1982), 2-12; Ewald "Computing Division Two-Year Operational Plan, FY 1981-1982," 43; Granville Chorn, et al., "The Standard File Transport Protocol" (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-7388-MS, 1978).

the transfer. The first completed HSPI was for the ubiquitous PDP-11, with others following in short succession in the 1970s and 80s.²²

5.3 CFS and the ICN of the 1980s

The late 1970s' introduction of higher networking speeds, new security partitions, and standardized interfaces and protocols on the Integrated Computer Network coincided with the end of Hydra in 1979, completing the Lab's transition to a modular network configuration. Replacing Hydra as the networked storage system on the ICN was the Common File System (CFS), which represented a middle ground between vendor-sourced networked storage systems and the entirely custom Hydra system. The networked-storage market had improved significantly in the years after Hydra's deployment, and C Division was able to use off-the-shelf IBM storage hardware and operating systems, rather than developing its own. IBM disc units replaced older Control Data storage devices, and the tape-cartridge-based IBM 3850 replaced the Photostore. While CFS still required a central storage controller, unlike the repurposed CDC front-end computers of Hydra, the CFS controller was designed and used specifically for storage. The CFS hardware offered a maximum data transfer rate of about 5Mbits/s, and stored about 384 gigabytes of combined online and offline data on the network as of 1980, with a growth rate of about 175 gigabytes per year. While the hardware and operating system were vendor-sourced, saving C Division considerable development time and cost, the division had to develop in-house the CFS application and interface software, as vendors in the HPC space still did not offer a complete package needed for tying together their storage devices with the complex, high-performance network environment of a national laboratory.²³

After the introduction of CFS, the C Division leader, Robert "Bo" Ewald, noted a drop in the number of tapes at the Central Computing Facility, as the storage system reduced the number of redundant, partially filled tapes in circulation. Much of the CFS development team originally worked on the Hydra project, and had feared a repeat of the design difficulties that had delayed the rollout of Hydra. As a consequence, the team leaders invested in a "structured" design process of allocating and reviewing programming work. Individual programmers on the project were assigned specific parts of the CFS coding effort, with their code being checked at least three times by other programmers for errors and efficiency before the code was added to the CFS application suite. The process itself was time-consuming, but team

²² Norman R. Morse and Joseph L. Thompson, "Bringing A Large Computer Network Into Focus" (Los Alamos, NM: Los Alamos Scientific Laboratory, LA-UR-82-83, 1982), 8-11; Tolmie, "Interconnecting Computers with the High-Speed Parallel Interface," 1-5.

²³ Bill Collins, interview by Nicholas Lewis, Minneapolis, MN, 2/19/2015; Christman, interview by Nicholas Lewis, 2014; Bill Collins, Marjorie Devaney, and David Kitts, "Profiles in Mass Storage: A Tale of Two Systems," *Ninth IEEE Symposium on Mass Storage Systems, 1988. Storage Systems: Perspectives* (Oct. 31, 1988 – Nov. 3, 1988), 61-67.

members recalling the effort cite it as a key reason for the on-time, unusually smooth roll-out of CFS, and its long-term reliability.²⁴

To the astonishment and delight of the CFS development team, Los Alamos network users, who were never shy about voicing their dissatisfaction over the problems and perceived shortcomings of the Lab's network services, wrote letters of praise to CFS for its on-time delivery, reliability, and intuitive, platform-agnostic interface. The transition to CFS concluded with the Photostore, which IBM no longer supported, being placed into a read-only mode until 1980, so that users could move their remaining files before Hydra was finally shutdown. With its success at Los Alamos, General Atomics Corporation sold CFS commercially as DataTree. The CFS development team aided in creating or reconfiguring components of the CFS application suite for use in a variety of supercomputing environments. Over the next decade CFS, as DataTree, entered use at several dozen data-intensive computing sites worldwide. Los Alamos networking personnel often visited sites, including those overseas, that had purchased the DataTree system, and aid in its installation, optimization, and personnel training for that site's unique requirements. CFS, and its success at and beyond Los Alamos, during the 1980s represented a transitional period in networked data storage, when the number of data-intensive computing sites had grown markedly since the development of Hydra in the early-to-mid 1970s. CFS also marked the important role that Los Alamos would play in networked storage development, with the Lab collaborating with partners in industry and academia on many other high-speed and high-capacity storage systems, such as HPSS and Panasas, from the 1990s to the present.²⁵

By 1981, the Los Alamos ICN consisted of several modules. The modules included the "worker" machines (the nine supercomputers in the CCF, which performed the bulk of the scientific computing work at the Lab), the Common File System (which allowed the supercomputers to store and share files via the network), the Terminal Network (connecting terminal users from the Secure, Administrative,

²⁴ Christman, interview by Nicholas Lewis, 2014; Ewald "Computing Division Two-Year Operational Plan, FY 1981-1982", 37.

²⁵ M. Blood, R. Christman, and B. Collins, "Experience with the LASL Common File System," *Fourth IEEE Symposium on Mass Storage Systems* (New York: The Institute of Electrical and Electronics Engineers, Inc., 1980); Bob Ewald, "Overview," *User Interface* 67 (May 1981), 3-4; Collins, interview by Nicholas Lewis, 2015; Gary Grider, interview by Nicholas Lewis, Los Alamos, NM, 8/21/2015; Fred McClain, "DataTree and UniTree: Software for File and Storage Management," *Digest of Papers. Tenth IEEE Symposium on Mass Storage Systems: Crisis in Mass Storage* (1990): 126-28. Lawrence Livermore National Laboratory also developed its own networked-storage system, which sold commercially as UniTree. As the name indicates, UniTree ran atop the Unix operating system, while Los Alamos' DataTree/CFS relied on IBM's MVS OS. While MVS was a more mature platform in the late 1970s, when both networked-storage solutions were under development, it left DataTree/CFS tied to IBM operating system development and support at a time when the national labs and supercomputing in general were moving toward Unix in the late 1980s and 1990s.

and Open security partitions to the ICN), and PAGES (a service center that allowed users of the ICN to print to shared laser printers, plotters, or other graphical media). The various modules were interconnected through the File Transport Network, which provided the ICN's high-speed data transfer capability. Two additional modules catered to differing levels of production control on the network. FOCUS, a VAX-11-based control interface, allowed operators to monitor and manage the Lab's Cray and CDC 7600 supercomputers via the ICN. The final module, XNET (eXtended Network Access System), marked the beginning of the Lab's incorporation of outside protocols onto portions of the ICN. Beginning in 1979, the XNET module brought "distributed processing" to Los Alamos, that is, support for user-operated minicomputers on the network, which a growing number of ICN users were requesting as the 1970s drew to a close. Text-based and graphical terminals performed no production work of their own, and simply connected remote users to the supercomputers of the CCF. VAX-11 minicomputer "distributed processors" allowed users to perform production work locally, or on a remote VAX machine, while retaining full access to the resources of the ICN. Most XNET users were local to Los Alamos, but some were located in other states. Smaller-scale computing complicated the structure and philosophy of computer networking at Los Alamos, as it did elsewhere, with small, but powerful systems performing an increasing amount of production work that normally would have been done on the Lab's larger computers. C Division invested significant time and research energies into supporting smaller computing systems in a variety of capacities on the ICN over the 1980s, while attempting to maintain the overall performance and security of the larger network.²⁶

The VAX distributed processors initially used the proprietary DECnet protocol to communicate with the XNET concentrators, forming a smaller network of minicomputers on a module of the larger ICN. The use of DECnet on the XNET module offered the Lab several advantages. For one, it allowed Los Alamos ICN users access to the latest tools available from vendors, such as smaller-scale computing and networking, without requiring the Lab to develop its own, costlier solutions. Los Alamos already followed a similar, but custom path with its terminal access, as it continually upgraded the terminal network module to coincide with improvements in vendor-supplied terminals over the 1970s. However, upgrading the custom system took longer, and offered no vendor support. The XNET module also allowed Los Alamos to experiment and gain experience with outside networking advancements, such as the proprietary DECnet system, while retaining its custom, mission-critical portions of the ICN.²⁷

²⁶ Robert H. Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984" (Los Alamos, NM: Los Alamos National Laboratory, LA-9712-MS, 1983), 43-52; Ewald, "Computing Division Two-Year Operational Plan, FY 1981 - 1982," 27, 31-38; Romero, "Overview of the Los Alamos Integrated Computer Network," 24.

²⁷ Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984," 67; Ewald, "Computing Division Two-Year Operational Plan, FY 1981 - 1982," 39-40; Romero, "Overview of the Los Alamos Integrated Computer Network," 24.

In the early 1980s, Los Alamos became part of the broader networking community, establishing links in 1983 with multiple external networks, including a purchased link to GTE's Telenet, a commercial packet-switching network. Connecting with external networks brought resources to Los Alamos from beyond the Lab, and allowed easier collaboration and information access for remote users. The Telenet connection, for example, allowed Open partition users access to globally distributed computing and information resources. Telenet also offered remote dialup access to the ICN for a substantially lower cost than a dedicated phone line. 1983 also marked the Lab's first link with ARPANET, the Department of Defense's packet-switching network of universities, research institutions, and unclassified DoD traffic. C Division connected an ARPANET IMP (Interface Message Processor) to a custom VAX gateway, which translated ARPANET's TCP/IP traffic into the ICN's local protocols. For security purposes, traffic from external networks connected only to the Open partition of ICN.²⁸

Although TCP was being forcibly rolled-out onto ARPANET in 1983, and was spreading elsewhere due to ARPA and the Department of Defense incentivizing its adoption among outside networks and computing vendors, Los Alamos had little interest in TCP/IP in the early 1980s. The ARPANET suite had little to offer in terms of performance, security, or compatibility with systems in use on the Lab network. Indeed, maintaining a custom-networking environment offered the Lab distinct security advantages, as few outsiders were familiar with the Los Alamos network or its largely bespoke systems and protocols. However, as with changes in supercomputing technologies, the Computing Division actively studied new networking technologies that were emerging from the maturing computer networking community in the 1980s.²⁹

While Los Alamos remained committed to its locally developed protocol suite for mission-critical parts of the ICN, such as the worker machine and File Transport modules, C Division was open to expanding its use of external protocols where they offered greater compatibility with the growing range of networked devices entering the market. Such protocols and devices had the potential to reduce the high development and maintenance costs associated with custom solutions. In 1983, C Division began investigating the potential of the X.25 networking protocol for its distributed processor network, as DEC was slated to incorporate X.25 support into its VAX VMS operating system. As an international standard, X.25 was also promising as the Lab's preferred protocol for interfacing with commercial packet-switching networks, such as Telenet and Tymnet. By 1985, C Division had converted XNET to run the X.25 protocol over DECnet, with over 75 VAX distributed processors on XNET connecting to the ICN via three, and soon four, gateways. After extensive security testing, the Lab allowed distributed XNET processors access to the Secure

²⁸ Morrison, interview by Nicholas Lewis, 2014; Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984," 67.

²⁹ Abbate, *Inventing the Internet*, 140-143; Morrison, interview by Nicholas Lewis, 2014; Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984," 76.

network partition, bringing small-scale computing to all network users at Los Alamos.³⁰

The Lab's connections with the wider networked world changed rapidly in the mid-1980s. In 1984, the DoD separated MILNET from ARPANET. ARPANET retained the university and research traffic, while MILNET linked together government, Defense Department, and Department of Energy facilities, including Los Alamos. Los Alamos established direct links to three other MILNET nodes. Each node connected to MILNET via a gateway host computer. The Lab's host allowed minicomputers on DECnet to exchange e-mail, and enabled TCP/IP access to the open partition of the larger Department of Defense Network. In 1986, Los Alamos added a fourth security partition to the ICN, the National Security partition, so that Department of Defense users who lacked Department of Energy (DOE) clearances could have access to secure computing resources separate from secure DOE computing work. Los Alamos moved one of its Cray-1A supercomputers into the new security partition for the benefit of military users and Defense contractors, like the Defense Nuclear Agency (DNA), SDI, and DARPA.³¹

As a site once selected for its geographical isolation, the Lab's experience with computer networking quickly brought Los Alamos into the center of collaborative information processing, distribution, and management for the Department of Energy in the 1980s. The Department of Energy's WBCN (Wide-Band Communications Network) and NWCNET (Nuclear Weapons Complex Network) projects exemplified the Lab's place in bridging remote government computing sites. Begun in 1985, the DOE initiated the two related networking projects to link together the computing resources of the national Nuclear Weapons Complex, which consisted of Los Alamos, Lawrence Livermore National Laboratory, Sandia National Laboratory, and eleven other research and development sites. WBCN was the hardware component that allowed for secure transmission of data among the sites of the nuclear weapons complex, and consisted of DEC VAX-11/750 minicomputers that served as gateway controllers at each location, along with accompanying encryption and communication hardware.³²

WBCN initially ran over terrestrial circuits, which were replaced with satellite circuits in the late 1980s. NWCNET comprised the basic tools used at each site for the storage and transmission of weapons-related data across WBCN. The first

³⁰ Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984," 47, 67; Robert H. Ewald, "Computing Division Two-Year Operational Plan, FY 1984-1985" (Los Alamos, NM: Los Alamos National Laboratory, LA-9978-MS, 1984), 100-101; Romero, "Overview of the Los Alamos Integrated Computer Network," 25.

³¹ Abbate, *Inventing the Internet*, 142-145; Ewald, "Computing Division Two-Year Operational Plan, FY 1984-1985," 101-102; N. R. Morse, "C-Division Annual Review and Operating Plan, January 1987" (Los Alamos, NM: Los Alamos National Laboratory, LA-10932-MS, 1987), 84.

³² Morse, "C-Division Annual Review and Operating Plan, January 1987," 84; N. R. Morse and B. L. Buzbee, "C-Division Annual Review and Operating Plan" (Los Alamos, NM: Los Alamos National Laboratory, LA-10631-MS, 1986), 110-111.

NWCNET tool was CFS, the Common File System developed at Los Alamos, which stored and managed the secure data at each location in the complex. Another tool was the software that interfaced CFS with the WCBN, while a third tool was the software that allowed the various NWCNET devices to communicate via Hyperchannel. A fourth tool allowed for secure e-mail capability between the different locations. The system first linked Los Alamos and the Rocky Flats weapons production facility in Colorado in 1986, with the other sites following soon thereafter. Many of the components and services driving the WCBN and NWCNET systems originated at Los Alamos. C Division produced a modified version of CFS that could be used at the other weapons complex locations, and created a network interface for the IBM MVS operating system (the OS platform for the CFS application suite), which allowed users and worker machines access to CFS across a Hyperchannel link. Because of its extensive networking experience, Los Alamos functioned as both the development site and the Tech Control for WCBN, testing the network's components on the Lab's ICN, and ensuring the smooth operation of the other WCBN nodes.³³

As Los Alamos became the lead development site for networking efforts inside the DOE complex, the Lab continued its onsite supercomputer-networking development, which brought Los Alamos into the world of network standards setting. By the mid-1980s, the Los Alamos HSPI, the custom network interface the Lab had employed since the late 1970s, was reaching the end of its lifespan, as its 60Mbits/s throughput could not accommodate the anticipated performance increases of the next decade. The rapid increase in traffic across the ICN, and the introduction of new supercomputers with much faster I/O capabilities, such as the Cray Y/MP, meant that Los Alamos would soon require interface speeds far beyond what the networking industry offered in the late 1980s. However, developing interfaces in-house for each device the Lab purchased was expensive, and meant the Lab had to provide its own support. C Division had little choice but to create custom interfaces when it developed the HSPI in the 1970s and early 1980s, but vendors and standards-setting bodies associated with networking had evolved considerably between that time and the late 1980s. C Division proposed a standard interface for the supercomputing industry, which would provide an alternative to proprietary vendor interconnects.³⁴

6. The “Lunatic Fringe” and the HIPPI Standard

In 1987, a Los Alamos delegation, led by C Division networking expert Don Tolmie, approached the American National Standards Institute (ANSI) Networking Task Group with a proposal for the High-Performance Parallel Interface (HIPPI), the first gigabit network interface. Although ANSI initially labelled the Los Alamos group as “the lunatic fringe,” because of the extreme performance it was proposing, ANSI soon

³³ Ibid.

³⁴ Tolmie, interview by Nicholas Lewis, 2014; Don E. Tolmie and John K. Renwick, “HIPPI – Simplicity Yields Success,” *IEEE Network – The Magazine of Computer Communications* (November, 1992), 1-2, 4.

formed a new task group (the first ever in the supercomputer industry) to pursue the Lab's interface. Developed at Los Alamos, the HIPPI became an ANSI standard in 1991, and an ISO (International Organization for Standardization) standard soon thereafter, marking the first standard networking interface for supercomputers and their peripherals. Much as the Lab's collaborative development of the first supercomputer with IBM signified the maturation of digital scientific computing, the standardization of the HIPPI represented a similar maturation of supercomputer networking in the 1980s. Since the 1970s, the number of institutions utilizing supercomputers, including industry and centers of scientific research, had increased significantly, as the computers themselves became more compact and reliable, and as the software tools for using them improved. Older, advanced scientific computing centers, such as Los Alamos, had endured the challenges of early computing, which included the need to write locally much of the software used on their machines. This barrier largely precluded the deployment of supercomputers where the users lacked the ability or willingness to write their own application software.³⁵

By the mid-1970s, institutions like NCAR, the National Center for Atmospheric Research, began to demand from vendors the inclusion of basic software packages with the supercomputers they bought, as was the case with NCAR's purchase of the first commercially available Cray-1 in 1977. Advancing supercomputing hardware and software, and the growing number and range of users, fed into a rapidly maturing scientific computing environment by the 1980s, including the maturation of ancillary technologies, such as high-speed networking. With the Los Alamos HIPPI standard in place, and with a growing range of customers to provide pressure on vendors to use the standard, supercomputers and related devices began to ship with the HIPPI interface already built-in, reducing the cost and difficulty of supercomputer networking. In the 1990s, the HIPPI, and the high-speed networking interconnects it spawned and inspired, provided part of the vital networking infrastructure that underpinned the development of cluster computing.³⁶

As Los Alamos participated in the development and evolution of networking beyond the Lab, its local ICN continued to evolve in the late 1980s, transitioning away from custom protocols and systems, and toward the new range of industry standard solutions entering the market. While distributed processors and external network connections provided impetus for the Lab to explore X.25, the arrival of even smaller computing platforms and an increasing reliance upon connections to external networks spurred changes to the Los Alamos ICN in the second-half of the 1980s, along with preparations for even more extensive changes in the coming decade. By 1985, Sun Microsystems scientific workstations and IBM-compatible personal computers were rapidly entering use throughout the Lab as lower-cost alternatives to

³⁵ Tolmie, interview by Nicholas Lewis, 2014; Morrison, interview by Nicholas Lewis, 2014; Don E. Tolmie, "What's Happening with Supercomputer Networks" (Los Alamos, NM: Los Alamos National Laboratory, LA-UR-91-2704, 1991), 1-3.

³⁶ Norman R. Morse, "C-Division Annual Review and Operating Plan, August 1990" (Los Alamos, NM: Los Alamos National Laboratory, LA-11945-MS, 1990), 6, 24, 60; Tolmie, "What's Happening with Supercomputer Networks," 1-3.

minicomputer-based distributed processors and computer-based terminals. C Division spent the latter-half of the 1980s exploring whether Ethernet-based LANs (Local Area Networks) of small computers could be utilized at Los Alamos and still maintain the integrity of the larger ICN.³⁷

Scientific workstations and PCs initially operated as stand-alone units, while C Division experimented with secure means of connecting them with the ICN's custom architecture. PCs soon entered use as terminals on the Terminal module of the ICN, but only in the unclassified partitions. Between 1985 and 1987, the Lab connected workstations and personal computers to the ICN by creating small LANs whose traffic went through the VAX gateways used to interface with ARPANET. Routing traffic through the ARPANET gateways was a temporary workaround solution, while the Lab investigated the routing and interface hardware needed to connect LANs more directly to the larger network. In 1989, a KCC team developed a more permanent solution, creating a new DEC MicroVAX-powered Keyboard Communications Concentrator that supported file traffic between workstations and the ICN. C Division likewise enabled TCP/IP compatibility on the XNET module for the distributed processor network, after DEC had added TCP/IP support to its VAX operating systems in response to DoD paid incentives.³⁸

Although Los Alamos had little interest in the DoD's TCP/IP protocol suite in the early 1980s, by the mid- to late-1980s, the suite had spread far beyond its original application on ARPANET. Los Alamos networking expert John Morrison later recalled the issues with TCP/IP, particularly its poor performance with large file transfers, its lack of built-in security features, and its error-detection deficiencies, which initially dissuaded the Lab from pursuing the protocol suite beyond developing gateways with ARPANET and MILNET. However, the snowball effect of TCP/IP's adoption among business, consumer, and government networks and computing systems expanded the market for TCP/IP-compatible systems and devices, and vendors followed the demand. As a consequence, in the late-1980s, C Division began to develop a strategy to convert selected modules of the ICN to TCP/IP, and migrate away from the custom ICN protocol suite, as the promise of practical vendor-supplied networking solutions outweighed the deficiencies of the DoD protocols.³⁹

In 1987, C Division began experimenting with building TCP/IP support into CTSS, the Livermore and Los Alamos co-developed operating system used on the Lab's Cray supercomputers. The Lab soon dropped the proposed modification of CTSS in favor of converting its supercomputers to the Cray-sourced version of Unix,

³⁷ Yvonne Martinez, et al., "Computing Division Two-Year Operational Plan, FY 1984-1985" (Los Alamos, NM: Los Alamos National Laboratory, LA-9978-MS, 1985); Morse, "C-Division Annual Review and Operating Plan," 34, 37, 49, 57.

³⁸ Abbate, *Inventing the Internet*, 140; Ewald, "Computing Division Two-Year Operational Plan, FY 1983-1984," 81-82; Romero, "Overview of the Los Alamos Integrated Computer Network," 9-10; Morse, "C-Division Annual Review and Operating Plan," 17, 51, 66, 73, 110, 113.

³⁹ Morrison, interview by Nicholas Lewis, 2014; Morse, "C-Division Annual Review and Operating Plan, January 1987," 79.

called UNICOS, which supported TCP/IP. C Division also tested an experimental Ethernet network, using TCP/IP to connect a DEC VAX to a Cray supercomputer through an IP router on the Secure partition. Ethernet's poor performance in the test gave further justification to pursue the HIPPI standard. The Lab eventually returned to Ethernet in the mid-1990s as a solution for message passing between the supercomputers and their storage systems, saving the faster HIPPI links for data transfer.⁴⁰

7. The ICN2

The Lab spent the late 1980s and early 1990s preparing for a gradual transition of the ICN away from custom protocols and devices as part of its "Network Modernization Project." Modernization was planned to occur in stages, proceeding from the least to the most critical modules and systems, to avoid unexpected performance or security issues. Once under way, the overall conversion was expected to take place over a decade or more, but DOE intervention necessitated a more abrupt break with the homegrown Los Alamos network. The resulting conversion to TCP/IP-compatible systems occurred as one part of a total restructuring of the ICN architecture. In 1991, the DOE Office of Safeguards and Security determined that the Lab's partitioning of the ICN into Open, Administrative, Secure/Unclassified, National Security, and Secure/Unclassified partitions via custom security controller hardware and software would not meet new DOE-wide security requirements. Safeguards and Security sought uniform network security policies and enforcement across the Department of Energy, and custom systems were difficult, if not impossible, to verify in accordance with standardized metrics. Faced with the new DOE requirement, Los Alamos planned and constructed the "ICN2" between 1991 and 1996, which involved creating a physical separation between the Secure and Open portions of the network. The Lab completed its transition to industry standard systems as part of the restructuring, replacing the custom gateways and other bespoke systems with vendor-supplied, TCP/IP-compatible components. Taking five years to complete, the transition occurred rapidly, considering the size and complexity of the Lab network, and the need to prevent disruptions of vital services. While not the gradual transition Los Alamos had anticipated, the ICN restructuring provided the compatibility with industry networking standards envisioned in the late 1980s.⁴¹

⁴⁰ Morse, "C-Division Annual Review and Operating Plan, January 1987," 73; Lynn Jones, interview by Nicholas Lewis, Los Alamos, NM, 8/6/2015; Raymond Miller, interview by Nicholas Lewis, Los Alamos, NM, 6-25-2014.

⁴¹ Morse, "C-Division Annual Review and Operating Plan, August 1990," 52; Hassan Dayem, "C-Division Annual Review and Operating Plan, May 1992" (Los Alamos, NM: Los Alamos National Laboratory, 1992), 54-55; Hassan Dayem, "Computing, Information, and Communications, 1995-1996" (Los Alamos, NM: Los Alamos National Laboratory, LALP-96-35, 1996), 44.

8. New Custom Solutions

Although the Los Alamos network was no longer comprised of mostly in-house systems, custom networking solutions never completely disappeared. Awaiting the development of its replacements, CFS remained in use until 2002, using a custom gateway to translate the ICN's TCP/IP protocols for the CFS controller. As it had done before, the Lab returned to custom networking when vendor solutions were either inadequate or missing from the market altogether. Before the term "cyber security" was coined, in the late 1990s Los Alamos developed its own high-speed security monitoring hardware to sift through Web-based traffic passing through the ICN. A purely hardware-based solution that split data traffic into multiple streams made it possible to screen the data for malicious packets without slowing network traffic. Beyond security, Los Alamos broke a world speed record in 1994 when it completed a 792Mbits/s data transfer via a Lab-developed HIPPI/SONET gateway to the San Diego Supercomputer Center, a distance of over 1,200 miles. This experiment helped to push the boundaries of remote-computing performance, allowing distant centers and users to share supercomputer resources. These examples are characteristic of the Lab's continuing ability and drive to pursue custom approaches in order to explore and push boundaries, and to stimulate the research, development, and application of new capabilities at and beyond Los Alamos.⁴²

9. Conclusion

The homegrown Los Alamos Integrated Computer Network of the 1970s and 80s represented a unique path of networking development, where the technical demands of a nuclear weapons facility required the pursuit of custom solutions to problems that few other computing centers had yet to encounter. Although unique, Los Alamos was not isolated from the broader networking community. Just as it fostered intellectual and financial ties with computer vendors, researchers, and other users, the Lab played an active role in outside networking development, lending its expertise to create and disseminate technical innovations, and incorporating new technologies as they offered benefits to ICN users. The Los Alamos network was not a linear development, changing its technical and philosophical orientations in relation to varying internal and external pressures, and sometimes adopting technically inferior solutions, such as TCP/IP, for economic or bureaucratic expedience. Examining the Laboratory at Los Alamos highlights how often-overlooked institutions experienced, and contributed to, computer networking differently than the better-known cases from academia,

⁴² Wally St. John and Dave DuBois, "Wide-Area Gigabit Networking: Los Alamos HIPPI-SONET Gateway," In *Supercomputing '95: Proceedings of the 1995 ACM/IEEE conference on Supercomputing*, 1995; Dayem, "Computing, Information, and Communications, 1995-1996," 45; Jones, interview by Nicholas Lewis, 2015; Grider, interview by Nicholas Lewis, 2015.

industry, and the US military, which serves to complicate and enrich the narrative of computer networking history.

Acknowledgements

Special thanks to the HPC History Project, including: Gary Grider, Carolyn Connor, Jeff Johnson, and Freddie Marshall. Additional thanks to the Los Alamos National Laboratory Archives. The HPC History Project gratefully acknowledges the support of the Advanced Simulation and Computing (ASC) Program (FOUS Program Managers: Paul Weber, acting; Jason Hick).

References

1. MacKenzie, D.: The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing. *Annals of the History of Computing*, vol. 13, no. 2, 179-201 (1991).
2. St. John, W., DuBois, W.: Wide-Area Gigabit Networking: Los Alamos HIPPI-SONET Gateway. In: *Supercomputing '95: Proceedings of the 1995 ACM/IEEE conference on Supercomputing* (1995)
3. Abbate, J.: *Inventing the Internet*. MIT Press, Cambridge, MA (1999)
4. Russell, A.L.: *Open Standards and the Digital Age: History, Ideology, and Networks*. Cambridge University Press, New York (2014)
5. DeNardis, L.: *The Global War for Internet Governance*. Yale University Press, New Haven (2014)
6. Bashe, C.J., Palmer, J.H., Johnson, L.R.: *IBM's Early Computers*. MIT Press, Cambridge, MA (1986)
7. Fitzpatrick, A.: *Igniting the Light Elements: The Los Alamos Thermonuclear Weapon Project, 1942-1952*. The George Washington University (1999)
8. Haigh, T., Priestley, M., Rope, C.: Los Alamos Bets on ENIAC: Nuclear Monte Carlo Simulations, 1947-1948. In: *IEEE Annals of the History of Computing* vol. 36, no. 3, 42-63 (July-Sept. 2014)
9. Michael, G., Ghausi, M.: An Interview with Sidney Fernbach, <http://www.computer-history.info/Page1.dir/pages/Fernbach.html>.
10. Collins, B., Devaney, M., Kitts, D.: Profiles in Mass Storage: A Tale of Two Systems. In: *Ninth IEEE Symposium on Mass Storage Systems, 1988. Storage Systems: Perspectives*, pp. 61-67. IEEE Press, New York (1988)
11. Tolmie, D.E., Renwick, J.K.: HIPPI – Simplicity Yields Success. In: *IEEE Network – The Magazine of Computer Communications*, pp. 28-32. (November, 1992)