APR 29 1997

SANDIA REPORT

SAND97-0748 • UC-705

Unlimited Release

Printed April 1997

Sandia's Network for Supercomputer '96: Linking Supercomputers in a Wide Area Asynchronous Transfer Mode (ATM) Network

RECEIVED

MAY 08 1997

OSTI

Thomas J. Pratt, Luis G. Martinez, Michael O. Vahle, Thomas V. Archuleta

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; distribution is unlimited.

Sandia National Laboratories

MASTER

SF2900Q(8-81)

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Sandia's Network for Supercomputer '96: Linking Supercomputers in a Wide Area Asynchronous Transfer Mode (ATM) Network

Thomas J. Pratt, Luis G. Martinez, and Michael O. Vahle
Advanced Network Integration

Thomas V. Archuleta
Telecommunications Operations Department II
Sandia National Laboratories
P. O. Box 5800
Albuquerque, NM 87185-0806

Abstract

The advanced networking department at Sandia National Laboratories has used the annual Supercomputing conference sponsored by the IEEE and ACM for the past several years as a forum to demonstrate and focus communication and networking developments. At *Supercomputing 96*, for the first time, Sandia National Laboratories, Los Alamos National Laboratory, and Lawrence Livermore National Laboratory combined their *Supercomputing 96* activities within a single research booth under the ASCI banner. Sandia provided the network design and coordinated the networking activities within the booth. At *Supercomputing 96*, Sandia elected: to demonstrate wide area network connected Massively Parallel Processors, to demonstrate the functionality and capability of Sandia's new edge architecture, to demonstrate inter-continental collaboration tools, and to demonstrate ATM video capabilities. This paper documents those accomplishments, discusses the details of their implementation, and describes how these demonstrations support Sandia's overall strategies in ATM networking.

CONTENTS

LIST OF FIGURES.....	3
LIST OF TABLES.....	3
1 INTRODUCTION.....	4
2 <i>SUPERCOMPUTING 96</i> NETWORKS.....	5
3 NETWORK DESIGN	6
4 WIDE AREA NETWORK	7
5 BUILDING THE CONNECTIVITY	9
6 CONGESTION CONTROL.....	9
7 TESTING THE NETWORK.....	11
8 TESTING THE APPLICATION.....	12
9 RESULTS OF THE EXPERIMENTS AT <i>SUPERCOMPUTING 96</i>	13
10 THE G-WAAT INTERNATIONAL CONNECTION.....	14
11 AN UNEXPECTED ATM VIDEO DEMONSTRATION.....	14
12 LESSONS LEARNED.....	15
13 CONCLUSION.....	15
14 ACKNOWLEDGMENTS.....	16
15 REFERENCES.....	17

List of Figures

FIGURE 1: ASCI BOOTH LAYOUT	5
FIGURE 2: ASCI BOOTH NETWORK PHYSICAL PLANT LAYOUT	8
FIGURE 3: NETWORK TO SUPPORT WAN SUPERCOMPUTING	9
FIGURE 4: LOGICAL MESH INTERCONNECT	10
FIGURE 5: NETWORK TEST SETUP	11

List of Tables

TABLE 1: SUPERCOMPUTING 96 NETWORKING REQUIREMENTS	6
--	---

1 Introduction

The advanced networking department at Sandia National Laboratories has used the annual Supercomputing conference sponsored by the IEEE and ACM for the past several years as a forum to demonstrate and focus communication and networking developments. The string of participation began in Minneapolis at the 1992 conference with a demonstration of the prototype Switched Multi-megabit Data (SMDS) and Synchronous Optical Network (SONET) technology that Sandia was intending to use in its consolidation of their supercomputers [1]. As a direct result of this participation, the National Information Infrastructure Testbed (NIIT) was born. At 1993 conference, in Portland, Sandia focused on the interoperability of emerging ATM technology and its efficacy in providing high quality video and multimedia capability [3]. This conference resulted in an early pilot of an interconnection of the three DOE Defense Program (DP) National Laboratories over a capable wide area network. The need for this type of network has continued to expand and current efforts in this arena include the DOE Laboratories Secure Network project. In Washington, the following year, the three Labs once again were connected over Sandia's extension of its production networks to the conference's trade show floor [7]. At the 1995 conference in San Diego, Sandia demonstrated, in collaboration with Oak Ridge National Laboratory (ORNL), Intel, and Giganet, a three node 622 megabit per second ATM interconnected Paragon network on the conference floor [8].

In all cases the significant contributions of Sandia's technical partners made the results possible and added significantly to the accomplishments.

Some common themes and benefits at each of the conferences have been:

- partnering with industry to gain early access to new technology,
- focusing current projects and activities by preparing challenging demonstrations,
- engendering new and evolving partnerships with industry, academia, and the other government labs and agencies,
- discovering and establishing new partnering opportunities,
- highlighting the synergy that results from the tight coupling of networking and communication technologies and organizations,
- providing a stage to professionally interact with colleagues and associates from other organizations in order to challenge and validate our current thinking.

For *Supercomputing 96* in Pittsburgh Sandia built on the success of *Supercomputing 95* along with the preceding ATM efforts to build a wide area network between the show floor and large supercomputing platforms at Sandia in New Mexico, ORNL in Tennessee, and the Pittsburgh Supercomputer Center. A material science application was run across this configuration. The result won a High Performance Computing Challenge gold medal.

This paper documents those accomplishments, discusses the details of their implementation, and describes how these demonstrations support Sandia's overall strategies in ATM networking [4,5,6]. Additionally, it describes the construction of a network to support the DOE DP National Laboratories at the conference.

2 SUPERCOMPUTING 96 Networks

At *SuperComputing 96*, for the first time in conference history, the three DP Laboratories put together a single integrated research booth. The three Laboratories, Sandia National Laboratories, Los Alamos National Laboratory, and Lawrence Livermore National Laboratory teamed under the Advance Strategic Computing Initiative (ASCI) rubric. While all the laboratories contributed equipment and personnel to the effort, Sandia was selected to lead the effort to provide the communication network needed to support the booth demonstrations.

Figure 1: ASCI Booth Layout

3 Network Design

The network design goal was to create a flexible environment that allowed the on site networking personnel to meet initial requirements while accommodating late changes.

The initial requirements collected from the Labs were:

Los Alamos National Laboratories (LANL)	Lawrence Livermore National Laboratory (LLNL)	Sandia National Laboratories (SNL)
24 IP addresses all Ethernet	10 IP address 6 Ethernet and up to 4 FDDI	10 IP addresses 9 Ethernet and 2 ATM
Internet Access Via SCINET ¹	Internet Access Via SCINET	Internet access via SCINET
2 machines must be located on the ASCI wall		Access to Pittsburgh Supercomputer Center and G-Waat network via SCINET

Table 1: Supercomputing 96 Networking Requirements

Considering these initial requirements it was decided that the primary service offering would mirror Sandia's latest edge communication architecture. This architecture provided high density switched Ethernet, 10Base-T, commodity ports to the users. It also contained the ability to integrate FDDI and ATM physical ports. A router was included in the design to provide the users within the booth isolation from problems that have tended to crop up in the past with the show network (SCINET). The connection to SCINET would provide access to FDDI, ATM and Ethernet. The FDDI connection was the primary data connection. The Ethernet connection was used to provide a backup connection should the FDDI network break. The ATM connection connected the ATM interfaces within the booth to the show network. FDDI was chosen as the commodity connection due to the nature of the show's historical ATM services. The show's ATM network is typically populated by demonstrations that are early development efforts that have the ability and the need to pass large data sets around the show floor. This experimentation tends to make the show's ATM network less reliable than its FDDI network.

A Fore Systems switch was selected for the booth because a joint SNL and Fore Systems demonstration was planned for the show. By using the Fore Switch we ensured easy interconnectivity with our partner.

¹ SCINET is the high performance network built each year by volunteers for the supercomputing conference

As the conference date neared, as expected, some changes to the preliminary network design were required due to the changes in the requirements of the exhibitors. The primary change was the drop of all the FDDI requirements by the LLNL and LANL exhibitors. The other change was a requirement for a 10Base-2, ThinNet, connection for the SNL exhibitors. This requirement was satisfied by the addition of a 10Base-2 Ethernet bridge. During the conference, SNL exhibitors requested an additional two ATM connections to support a video demonstration. The request could only partially be met through the use of an extra fiber pair that was installed as a contingency against any late additional requirements. The lack of a second fiber meant that the video demo could only operate in one direction at a time.

The combined booth was divided into 5 sections. These sections were defined as LLNL, SNL, LANL, ASCI Wall and the Networking Area (see Figure 1). The individual laboratories controlled their particular section of the booth. The ASCI Wall section was shared by the three labs to show ASCI demonstrations and videos. The Networking Area section was used to house the networking equipment and network monitoring equipment as well as a network available printer. Figure 2 shows the cable and fiber run used to support each area of the booth. In each of the laboratories section a cable was provided that contained 12 10Base-T connections. The SNL section had two multimode optical fiber pairs to support the ATM requirements. In addition to providing the required cabling, an extra multimode optical fiber pair was provided to accommodate any late interconnection requirements. The ASCI Wall section was supported by four individual 10Base-T cables. LANL also ran an additional five individual 10Base-T cables into the LANL section as an additional contingency.

4 Wide Area Network

At *Supercomputing 96* SNL, ORNL, and PSC teamed to create a wide area supercomputing demonstration where the largest massively parallel supercomputers available at each site were connected together to solve a single problem. The ground work for this demonstration was laid at *Supercomputing 95* when a small scale version of the demonstration was done on the exhibit floor. That demo used smaller, portable versions of the Intel Paragon supercomputers [8]. The supercomputers involved in this year's demonstration were scattered across the United States, in New Mexico, Tennessee, and Pennsylvania, and consisted of Sandia's 140 gigaflop 1842 processor nodes, Paragon named Acoma, Oak Ridge's 150 gigaflop, 1157 processor nodes, Paragon named XPS150, Oak Ridge's 35 gigaflop, 548 node Paragon named XPS35, and the PSC's Cray T3D. DOE Energy Science network, ESNET, provided the wide area networking infrastructure. The material science code that was run was a first principles Copper-Nickel alloy code².

² To get more information on the application see the web page: <http://www.ccs.ornl.gov/GC/materials/MShome.html>

Figure 2: ASCI Booth Network Physical Plant Layout

The ATM network (see Figure 3) that the demonstration used was built using permanent virtual circuits (PVCs). Six PVCs were used to create a mesh (see Figure 4) to interconnect the machines. The lowest speed section within the network was a DS3 (44.763 MBps) link provided by ESNET. The highest speed links in the network were the 622 megabits per second OC-12 local connections to the Paragons. The Intel Paragons in the demonstration were connected to the ATM network via Giganet's ATM Protocol Acceleration Engine [8]. The T3D was connected via a Hippi channel to an SGI workstation that then attached to the ATM network. The message passing protocol

running across the machines was a version of Parallel Virtual Machine (PVM). PVM used native ATM adaptation Layer 5 (AAL5) as its transport protocol.

Figure 3: Network to Support WAN Supercomputing

5 Building the connectivity

Since its inception, ESNET has provided IP services to ORNL and SNL. During fiscal year 1996, ESNET was updating its backbones to ATM. ORNL's ATM connectivity was completed in August 1996. Sandia's connection to ESNET was upgraded to ATM in October 1996. PSC connection to ESNET was not put into place until the *Supercomputing 96* SCINET network was built in November. The network running between SNL and ORNL had a round-trip latency of 106 milliseconds. About half the latency in the link was attributed to the speed of light. The rest of the network's latency was attributed to the network equipment electronic processing times. The wide area network consisted of a link from New Mexico, with drops at SNL and LANL, to ESNET's California terminal, a link from the California terminal to the Sprint's Chicago Terminal, and from Chicago two links, one to ORNL in Tennessee and one to PSC in Pennsylvania. The route through the West Coast added about 2000 miles to the network distance between the sites.

6 Congestion Control

Because the application traffic would compete for bandwidth with the ESNET production backbone traffic, a method was needed to ensure that the combination of demonstration

traffic and production traffic would not cause unacceptable network congestion. Originally, ESNET attempted to protect their backbone networks using a policing parameter on the Fore System switches. During testing of the network, it was discovered that the Fore switch failed to maintain the policing contract, thereby dropping traffic even when it met the established policing criteria. The result was the switch dropped application data even at a data rate less than what the policing setup allowed. This was an expected finding because Sandia had seen this same effect during earlier switch testing in Sandia's ATM Testbed. Through negotiation with ESNET, an agreement was reached that allowed the application to have full access to the network. The agreement relied on the ability of the application to control the amount of data it transmitted over a given period of time. The agreement called for the application to not exceed a rate of 20 megabits per second as measured over any 100 milliseconds of time. The 20 megabits per second amounted to about one half of the usable bandwidth on the slowest speed link on the network. The application was capable of limiting data transmission through its flow control methods. The application's flow control was based on the handshake the application was doing to achieve reliable data delivery. To meet the traffic levels that were agreed upon, the application would burst out four 32,768 byte data and then wait for a positive acknowledgment from the data receiver before sending another data packet.

Figure 4: Logical Mesh Interconnect

This application flow control is similar to a TCP sliding window implementation. While this method doesn't ensure that the network wouldn't get congested, it did ensure that the additional traffic that the application would insert onto the network wouldn't congest the network to the point that production services would suffer.

7 Testing the Network

To ensure that ESNET was able to successfully absorb and pass the agreed upon data rate and data bursts, a network throughput test suite was established. The test suite was designed to mimic the worst case traffic allowed by the negotiated agreement. Test traffic was sent into ESNET that allowed ESNET network technicians to monitor the effect on their network. ESNET was interested in the effects of the testing to verify that the agreed upon traffic level didn't cause ESNET's production traffic to become sluggish. The test allowed ESNET to prove that they could successfully pass the application data a large percentage of the time. The test was monitored to verify that our test traffic was not being dropped within the ESNET network. Figure 5 shows the testing setup.

Figure 5: Network Test Setup

Actually two tests were designed. Test 1 verified ESNET's ability to reliably deliver the promised data rate over a long period of time. A Hewlett Packard Broadband Test System was used as the data generator and data monitor. Additional data monitoring was provided by the diagnostic capabilities of Sandia's internal Globeview 2000³ core ATM switch. The test passed a large AAL5 data packet that was shaped by the HP test set to a 20 megabit data stream that was non-bursty in nature. It was during this testing that it was discovered that the policing parameter within the ESNET network was not operating correctly. The testing indicated that ESNET with the policing parameter turned on could only successfully deliver about 10 megabits per second of the shaped data through the network. By sharing this test data with ESNET, an agreement was reached with ESNET

³ A Globeview 2000 is a high performance ATM switch manufactured by Lucent Technologies

to remove the policing from the network. Once the policing was removed, the shaped 20 megabit per second traffic pattern was successfully passed through the network. This traffic pattern was sent into ESNET for long periods of time, from one hour to four hours, over several days. Only once during the testing period did the test produce dropped data. The amount of data lost during this time was less than one percent of the data transmitted. Another good indication from the testing was that ESNET didn't receive any reports of sluggish network response during this testing.

Test 2 was designed to mimic the nature of the application data traffic pattern. This traffic pattern consisted of bursts of data being passed onto the network. The addition of an ADTECH 3000 ATM data generator was added to enhance the test ability to manipulate the data burst patterns. The HP test set and the Globeview 2000 were still used to monitor the success of the test. The goal of this test was to pass a burst of 131,072 bytes of data into ESNET and have the data successfully returned through the loopback. Early results showed that this test failed nearly one hundred percent of the time. When the data bursts were reduced to the size of 5000 bytes, then ESNET successfully delivered the data. ESNET technicians together with SPRINT were notified and they tried to find the congested network link. Sandia assisted the search by providing a known test load that allowed the technicians to quickly find the link where the congestion was taking place. The bursty nature of IP traffic combined with the way TCP responded to congestion made it difficult to find network congestion. Typically, the TCP/IP protocol will back away from congestion in a way that actually lowers the amount of traffic that passes over a link. Although, the actual traffic passed over a given period is reduced, the congested nature of the link will reappear as the TCP/IP back off algorithm runs its course. In short, the link was congested, but the network monitoring equipment registered that the data traffic was getting smaller. Adding a controlled background load to the IP traffic running across the network, allowed ESNET and Sprint technical staff to more quickly discover the congestion points in the network. Sprint took action once the congestion point was discovered by rerouting the test traffic into an uncongested link.

8 Testing the Application

A lot of effort was put into characterizing the communication capability of PVM and the communication requirements of the application. Testing at ORNL showed that the peak sustained communication rate that the current version of PVM could achieve was 20 megabytes per second. The material science application's mode of operation consisted of a compute cycle followed by a communication cycle. The application would compute for 20 seconds then exchange the computed data by communicating via PVM. In a local environment with an OC-12 link between two supercomputers, the communication cycle lasted 2 seconds with about 52 megabytes of data passing between the machines, 26 megabytes in each direction. The results of this testing indicated that the application running over the distributed network that consisted of a 20 megabits per second data circuit would result in a 20 second compute cycle, followed by a 10 second communication cycle. This indicated to the application programmer that doubling the problem size only cost an additional 36% more time on the distributed.

To ensure that the application ran successfully in the distributed environment, a large scale version, running on more than 500 parallel compute nodes, of the application needed to be successfully run on individual platforms at SNL, PSC, and ORNL. This was an important undertaking because all of the machines and all of the computing environments at each site varied greatly. PSC's hardware and operating system was supplied by Cray. Even the machines at ORNL and SNL, although both Intel Paragons, had different hardware for the compute nodes. The production operating system at each site was different. ORNL used a released version of Intel OSF operating system on both the compute nodes and the service nodes. At SNL an internally developed operating system named SUNMOS was used on the compute nodes while OSF was used on the service nodes. To further complicate the internal environments, SNL was replacing SUNMOS with a new internally developed operating system named PUMA. The requirement to run a large scale version of the application in the local standalone environment was not a problem for ORNL as they owned this particular application. They had made several large production runs of this application within their environment. Because of this fact, and as ORNL had two large Paragons on their computing floor, they were able to concentrate on running the application in a locally distributed environment. SNL's first standalone attempt required that the application be run in a special fully dedicated environment that was all OSF. As *Supercomputing 96* came closer, the application was ported to run on PUMA compute nodes making it easier to run the application in a less dedicated setup of the SNL supercomputer. Eventually, the application was successfully run on 512 standalone nodes within the SNL environment. Once successfully tested, an attempt to connect the application over the WAN distributed environment was begun. Several unsuccessful attempts were made to run a large scale version of the application across the combined SNL ORNL environment. Although unsuccessful, these tests were critical to get ready to run the application during *Supercomputing 96*. Prior to *Supercomputing 96*, only small versions of the application were successfully run in the PSC environment. The application at PSC was limited by the data transformation that was needed to get the data out of the T3D. The data passed from the T3D to an SGI workstation via a TCP/IP Hippi channel. The SGI transformed the TCP/IP datagram into an AAL 5 datagram. This limited the communication performance of the application to under 2 megabits per second.

9 Results of the Experiments at *SUPERCOMPUTING 96*

SNL and ORNL ran a 1024 compute node version of the application. The application ran for two hours. During that period the application passed over four gigabytes of data. The sustained rate of the data over the entire period was 1 megabit per second. The highest sustained data rate was 40 megabits per second. Twenty megabits per second in each direction was sustained over a 10 second period. The application was started and ran for a short period of time during each day of the conference with similar peak traffic rates. The demonstration was awarded a High Performance Computing Challenge Gold Medal for Concurrency. ESNET production traffic was not impacted during the demonstration. Because the PSC branch of the proposed network performance was so poor, it was

decided to limit the PSC interaction. An eight node version of the application was run between ORNL and PSC just to demonstrate that PSC was connected to the distributed environment.

10 The G-WAAT International Connection

At *Supercomputing 96* Sandia participated with the University of Stuttgart in Germany and the PSC. The demonstration used the Global Wide Area Applications Testbed (G-WAAT) backbone to show a G7 application. That application was a part of the Global Interoperability for Broadband Network (GIBN) program. This application emphasized collaboration tools for interactive design. A video connection between the sites was used to provide human interactivity. Design data was also passed between the sites. The design data could then be operated on interactively by all participants. A high capacity telecommunication line was planned for the application. Unfortunately, this high capacity line never was realized so the exhibitors chose to run the demonstration using the conference's Internet. As expected, this led to some good performance periods and to some poor performance periods. During off peak hours, the video connection was capable of providing about 4 frames a second that seemed to be acceptable to the users. During most of the other time the video was only capable of supplying four frames in seven seconds, slightly better than 1/2 frame a second while at the worst of times the video was not usable by the demonstration. The exhibitor's reaction was that as a cheap solution the results were better than they had expected. The visualization tools that were at the heart of the demonstration were capable of being controlled by either end. The bandwidth of that portion of the demonstration appeared to be adequate, providing a good interactive environment.

11 An Unexpected ATM Video Demonstration

An ATM video demonstration was done in cooperation with Fore Systems. Fore had developed a new video product and asked Sandia to help them demonstrate it at the conference. This is a typical example of how demonstrations often materialize at the conference. Since Sandia needs an ATM video distribution system [3,6], the demonstration provided a fortuitous opportunity to test a potential solution. The goal of the demonstration was to pass video traffic on the show floor between the Fore System booth and the SNL booth. Although this demonstration successfully passed video between the booths, a simplex transfer of the video was made necessary because of the lack of sufficient fiberoptic cable to support a duplex session. The video end equipment required four fiberoptic cables to run a duplex video. This requirement was not based on bandwidth but instead on the physical needs of FORE video equipment. After SC96 it was discovered that the video equipment could run in a duplex mode using a single fiber pair. The transmit and receiver unit can be daisy chained together. Because this demonstration materialized during the show, it was only partially successful. However, it was an opportunity to get an early look at a possible solution for Sandia's emerging real time video needs.

12 Lessons Learned

The construction of the network within the booth was complicated by the fact that we elected to let the show's contracted labor place the cables without any on site supervision. The drawings that were sent to the contractor did not contain detailed measurements. The result was that the network drops needed to be moved once the network designers arrived. In the future, a network designer needs to arrive on the first day of setup or the expense of making detailed drawing needs to be made. The addition of extra infrastructure was worthwhile; however, additional late requirements could not be satisfied because the booth setup required a static deployment of the network infrastructure. A booth layout that allows the network to expand as needed is a potential solution. A face-to-face meeting of the networking personnel from the separate laboratories before the conference began would also have been beneficial.

Because the networking personnel were isolated from the main booth, the networking activities did not get the exposure that previous shows have provided. This also possibly resulted in some of the booth demonstrators going directly to SCINET with networking problems instead of coming to the booth's on site networking staff. This complicated the resolution of some networking problems and also complicated keeping the network in a known stable condition. The addition of IP address blocking within the Internet caused the original plan of using a corporate class C address to be unsuccessful and resulted in some last minute redesign.

For the ORNL, Sandia, and PSC WAN demonstration, it was decided early to do a demonstration that mirrored the goals of a production project. The efforts required to do the project would not have been reasonable for a one time demonstration. The effort to bring in PSC started too late in the project to adequately address all of the difficulties that the connection required. The interconnection of Sandia and ORNL a month and a half prior to the conference was just adequate to get the demonstration operational by the conference.

13 Conclusion

By all measures, the conference proved successful for Sandia. The conference provided a forum for Sandia to feature a wide variety of state-of-the-art networking and communication technologies and associated applications. The demonstrations benchmarked the current state of the SONET and ATM technologies, both of importance to many Sandia initiatives, and the evolving partnership with Oak Ridge National Laboratory. The success stories were the culmination of work accomplished by many people both within and outside of Sandia. In order to meet the challenging goals of the state-of-the-art networks, many teams were formed that crossed corporate and organizational boundaries. The conference also provided an opportunity to identify future goals and to plan joint activities. The teamwork amplified the accomplishments and achievements of all the participants. Similarly, the conference provided many Sandians an individual opportunity for professional growth, friendly competition, and professional

association. Still, on another level, the conference challenged its participants to take stock of their individual projects and to focus them for the demonstration. In all these ways, Sandia benefited from its participation in *Supercomputing 96*.

14 Acknowledgments

To put together the activities surrounding the Supercomputing conference takes a large number of talented and dedicated individuals. Without their efforts, Sandia could not have accomplished the demonstrations that were done at the conference. We would like to thank the following individuals for their efforts in making Sandia's networking efforts for *Supercomputing 96* a success.

SNL	-	Dennis Bateman, Arthurine Breckenridge, Joseph Brenkosh, Debbie Brunty, Pang Chen, Lee Ann Fisk, Steve Gossage, David Greenburg, Michael Hannah, Tan Chang "Richard" Hu, Cassandra Shaw, R. Michael Cahoon, Bruce Whittet, and Vicki Williams
LANL	-	Alice Chapman and Daryl Grunau
LLNL	-	Wayne Buteman and Mary Zosel
ORNL	-	Arthur "Buddy" Bland, Tom Dunigan, Lawrence MacIntyre, Philip Papadopoulos, Tim Sheehan, Bill Shelton, Malcolm Stocks and Bill Wing,
ESNET	-	Mike Collins, Jim Gagliardi, Jim Lieghton, and Rebecca Nissan
PSC	-	Jamshid Mahdavi
Intel	-	Kenneth Lord, Anthony Ralph, and Paul Work
Fore Systems	-	Russell J. Kuba and Frank Yeh

15 References

- [1] Arthurine Breckenridge and Michael O. Vahle. An account of Sandia's research booth at *Supercomputing '92*: A collaborative effort in high performance computing and networking. Technical Report SAND 93-0224, Sandia National Laboratories, Albuquerque, New Mexico, March 1993.
- [2] Steven A. Gossage. Delivery of very high bandwidth with ATM switches and SONET. Technical Report SAND92-1295, Sandia National Laboratories, Albuquerque, New Mexico, October 1992.
- [3] Steven A. Gossage and Michael O. Vahle. Sandia's research network for *Supercomputing '93*: A demonstration of advanced technologies for building high performance networks. Technical Report SAND93-3807, Sandia National Laboratories, Albuquerque, New Mexico, December 1993.
- [4] Nicholas Testi, John H. Naegle, and Steven A. Gossage. ATM: Sandia user case study. Telecommunications, February 1994.
- [5] Nicholas Testi, John H. Naegle, and Steven A. Gossage. Combining ATM and SONET at the LAN. Business Communications Review, pages 47-50, January 1994.
- [6] Nicholas Testi, John H. Naegle, Steven A. Gossage, Michael O. Vahle, and Joseph H. Maestas. Building networks for the wide and local areas using asynchronous transfer mode switches and synchronous optical network technology. *IEEE Journal on Selected Areas In Communications*, 1995.
- [7] Michael O. Vahle, Steven A. Gossage, and Joseph P. Brenkosh. Sandia's network for *Supercomputing '94*: Linking the Los Alamos, Lawrence Livermore, and Sandia National Laboratories using Switched Multimegabit Data Service. Technical Report SAND 94-3096, Sandia National Laboratories, Albuquerque, New Mexico, December 1994.
- [8] Thomas J. Pratt, Michael O. Vahle, and Steven A. Gossage. Sandia's network for *Supercomputing '95*: Validating the Progress of Asynchronous Transfer Mode (ATM) Switching. Technical Report SAND 96-0820, Sandia National Laboratories, Albuquerque, New Mexico, April 1996.

Distribution

1	0103 R. J. Detry, 12100	1	0807 G. A. Yonek, 4417
1	0159 L. D. Bertholf, 4500	1	0807 I. C. Alexander, 4417
1	0318 A. Breckenridge, 9215	1	0807 M. A. Schaefer, 4417
1	0318 G. S. Davidson, 9215	1	0807 R. L. Adams, 4417
1	0321 W. J. Camp, 9204	1	0807 R. M. Cahoon, 4918
1	0431 S. G. Varnado, 6500	1	0807 S. D. Nelson, 4417
1	0449 M. R. Sjulín, 6512	1	0807 T. V. Archuleta, 4417
1	0449 T. D. Tarman, 6512	1	0820 P. Yarrington, 9232
1	0630 M. J. Eaton, 4010	1	0826 J. D. Zepper, 9111
1	0660 W. D. Swartz, 4619	1	0833 J. H. Biffle, 9103
1	0661 M. H. Pendley, 4612	1	0841 P. J. Hommert, 9100
1	0801 M. J. Murphy, 4900	1	0877 J. M. McGlaun, 5903
1	0803 J. F. Jones, Jr., 4600	1	1109 A. L. Hale, 9224
1	0806 C. D. Brown, 4621	1	1111 S. S. Dosanjh, 9221
1	0806 D. C. Jones, 4603	1	1361 W. F. Mason, 4022
1	0806 J. H. Naegle, 4616	1	9003 D. L. Crawford, 8900
1	0806 J. M. Eldridge, 4616	1	9011 C. L. Yang, 8920
1	0806 J. P. Brenkosh, 4616	1	9011 F. T. Bielecki, 8930
1	0806 L. F. Tolendino, 4616	1	9011 H. Y. Chen, 8910
1	0806 L. G. Martinez, 4616	1	9011 J. A. Friesen, 8920
1	0806 L. G. Pierson, 4616	1	9011 J. A. Hutchins, 8910
1	0806 L. Stans, 4616	1	9011 J. C. Berry, 8930
10	0806 M. O. Vahle, 4616 (10)	1	9011 J. M. Brandt, 8910
1	0806 S. A. Gossage, 4616	1	9011 P. W. Dean, 8910
1	0806 T. C. Hu, 4616	1	9011 R. E. Palmer, 8901
50	0806 T. J. Pratt, 4616 (50)	1	9012 J. Costa, 8920
1	0807 B. C. Whittet, 4417	1	9012 R. D. Gay, 8930
1	0807 D. Brunty, 4417	2	9018 Central Technical Files, 8940--2
1	0807 F. J. Castelluccio, 4417	5	0899 Technical Library, 4414 (5)
		1	0619 Review and Approval Desk, 12690 for DOE/OSTI