

Building and Measuring a High Performance Network Architecture

William T.C. Kramer, Timothy Toole, Chuck Fisher, Jon M. Dugan, David Wheeler, William R. Wing, William Nickless, Gregory Goddard, Steven Corbato, E. Paul Love, Paul Daspit, Hal Edwards, Linden Mercer, David Koester, Basil Decina, Eli Dart, Paul Reisinger, Riki Kurihara, Matthew J. Zekauskas, Eric Plesset, Julie Wulf, Doug Luce, James Rogers, Rex Duncan, Jeffery Mauth

and the

SCinet 2000 Team

The submitted manuscript has been authored by a contractor of the U.S. Government under contract No. DE-AC03-76SF00098. Accordingly, the U.S. Government retains a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

This work was supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Abstract

Once a year, the SC conferences present a unique opportunity to create and build one of the most complex and highest performance networks in the world. At SC2000, large-scale and complex local and wide area networking connections were demonstrated, including large-scale distributed applications running on different architectures. This project was designed to use the unique opportunity presented at SC2000 to create a testbed network environment and then use that network to demonstrate and evaluate high performance computational and communication applications. This testbed was designed to incorporate many interoperable systems and services and was designed for measurement from the very beginning. The end results were key insights into how to use novel, high performance networking technologies and to accumulate measurements that will give insights into the networks of the future.

Introduction and Background

SC2000 was the 12th annual SC conference on High Performance Communications and Computing. The conference series was previously known as Supercomputing. It is jointly sponsored by the IEEE and the ACM and is one of their larger activities. The conference attracts about 5,000 attendees from all areas of high performance computing, including all the major supercomputer facilities and vendors, the major network organizations and suppliers. More than 150 exhibitors demonstrate their latest accomplishments. This includes about 70 research exhibits with all the U.S. national laboratories and many major facilities from around the world demonstrating their research activities and actually trying new research ideas using the infrastructure of the conference. The conference also includes a stream of approximately 20 tutorials, 140 papers, invited talks and a unique program to introduce high school teachers from around the country to high performance computing. The details of the conference can be found at <http://www.sc2000.org>. This year the conference was held in Dallas, Texas, at the Dallas Convention Center (DCC).

SCinet Goals for SC2000

The goals of the SC2000 network project, known as SCinet 2000, were aggressive and expansions of previous work. They were:

- Provide a stable, high quality network infrastructure for SC2000 activities – This goal reflects for need to support the basic infrastructure of the conference, the attendees' needs and the extensive education program.
- Provide reliable, high performance network for exhibitors – This goal recognizes the fact that the conference provides a showcase for the research and development activities of both the industry and research exhibitors.
- Provide experimental opportunities and demonstrations for the latest high performance networking technology – This goal represents the fact that the opportunity to co-locate so much technology, expertise, and so many varied applications seldom occurs, and a role of SCinet is to demonstrate the use of this technology in novel ways.
- Support and facilitate applications that make use of high performance networks – This goal reflects the fact that it is one thing to build a large network; it is another to make it usable, and still another to actually use it. SCinet works to accomplish all three objectives.

SCinet Design

The overall network design consisted of four major networks that were designed to operate independently, but with significant overlap between them. Each network is explained below in more detail. The SC2000 network was intentionally designed with complexity to explore the issues typically encountered in real-world networking, such as interoperation between different network domains, using different network routers and technology, networks that interface with multiple wide area peers, carrying different protocols over the same network. A logical diagram of the network is shown below in Figure 1. The four levels of network are all interconnected, but can operate independently of each other.

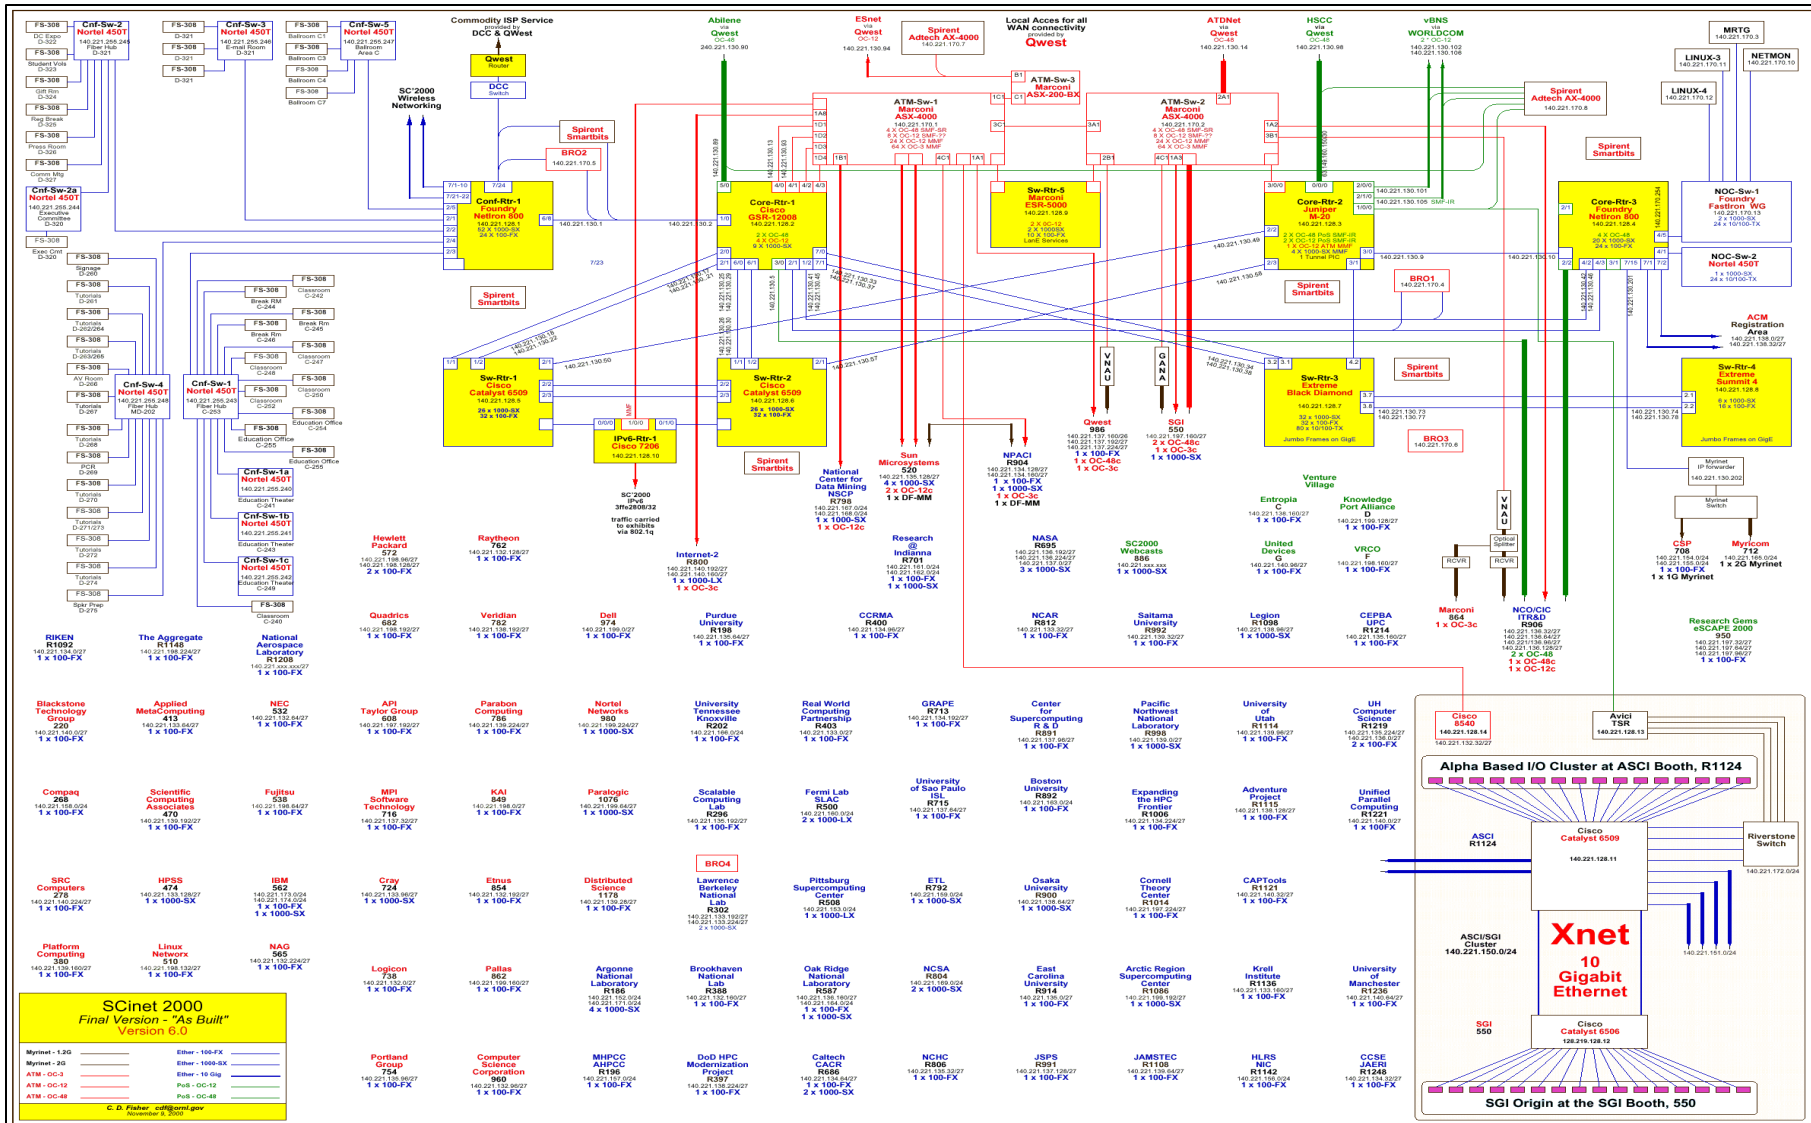


Figure 1 – The Logical Diagram of SCinet

In essence, SCinet is a self-contained ISP that peers with all the major research and government networks.

Commodity Network

At the first level, several days before the show started, there was a commodity Internet network to connect offices, the Education Program, and the email facilities. This network was expanded to include all the meeting rooms and lecture areas, including areas that webcast sessions, totaling more than 40 locations and over 300 drops. The network spanned about 200,000 sf over three floors of the DCC. Most of the network drops were connected using existing Cat-5 cables installed in the DCC to switches at 100 gigabits per second. These switches connected to the commodity router via multimode fiber. One connection for the email services was made at 1 Gbps using multimode fiber.

The DCC had an external 12 Mbs link provided by Qwestlink. For most conferences, this data rate is more than enough to support multiple events at any one time. Connections within the DCC were accumulated at an optical switch that was connected to a Cisco router managed by Qwest. The traffic then flowed over the Qwest backbone. Since the commodity network had to be up before the full SCinet network, and had to operate until the conference closed, it was decided that the best commodity service would be provided using the DCC external connections.

The commodity network connected to a single SCinet router, provided by Foundry Networks, a NetIron 800 router. This router, denoted Conf-Rtr-1, peered with the DCC Cisco router via BGP. There were three routers involved in the commodity BGP peering: Logically, the BGP peering was between the SCINet Foundry and the Qwestlink Juniper because the DCCC Cisco didn't have enough memory for all the tables. The DCCC Cisco only had a couple of static routes, including one pointing to 140.221.128.0/17 (SCINet) and default pointing to Qwestlink. Figure 2 shows how they were physically connected.

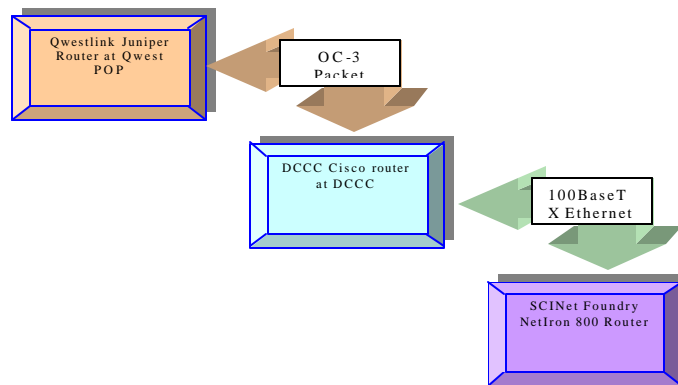


Figure 2 – Border Gateway Router Connections

Wireless Network

The second major network was an 802.11b 11 Mbps wireless network that spanned the entire conference area. Using Cisco Systems access points and a combination of Cisco Aironet and Lucent cards, SCinet created a large 11 Mbps wireless network throughout the conference space. An area of more than 200,000 sf on two floors of the DCC had continuous wireless service.

Wireless connectivity was designed to accommodate all 5,000 attendees to the conference, although only a subset actually had wireless cards with them. The wireless network was based on the 802.11b network standard, which allowed interoperability with a variety of network interface cards (Aironet, Lucent and other cards). IP addressing was provided using DHCP servers that covered the conference areas for both the wireless and the commodity networks.

The wireless network used 34 Cisco AP340 base stations. There were 14 base stations positioned on the ceiling of the exhibit area, and one base station on the ceiling of each meeting room and selected other locations.

Figure 3 shows the logical diagram of the Commodity network, with most of the wireless base stations identified.

SCinet Production Network

This network was the heart of the entire project and had to provide service to the entire 150+ exhibitors, as well as connecting to the major external wide area networks. The network supported IPv4, IPv6, ATM, and Packet over SONET connections, Myrinet and multicast and Webcasting.

Using Qwest dark fiber in Dallas and Qwest SONET, ATM, and IP backbones nationwide, the wide area network featured multiple OC-48c (2.5 Gbps) and OC-12c (622 Mbps) connections as well as other connections. In addition to commodity Internet access, WAN connection links to SCinet 2000 included:

- ESnet¹ – provides highly capable leading-edge network services that support DOE's missions. ESnet emphasizes advanced network and distributed computing capabilities needed for forefront scientific research and other DOE programs.
- Abilene/Internet 2² – Internet2 is a consortium being led by over 180 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies to accelerate the creation of networking technology.
- HSCC³ – the High Speed Connectivity Consortium is a collaboration of universities, industry and other organizations to create a nationwide multi-gigabit network providing one to two orders of magnitude higher bandwidth than currently commercially available

¹ See <http://www.es.net> for more details

² See <http://www.internet2.edu> for more details

³ See <http://www.hsc.net> for more details

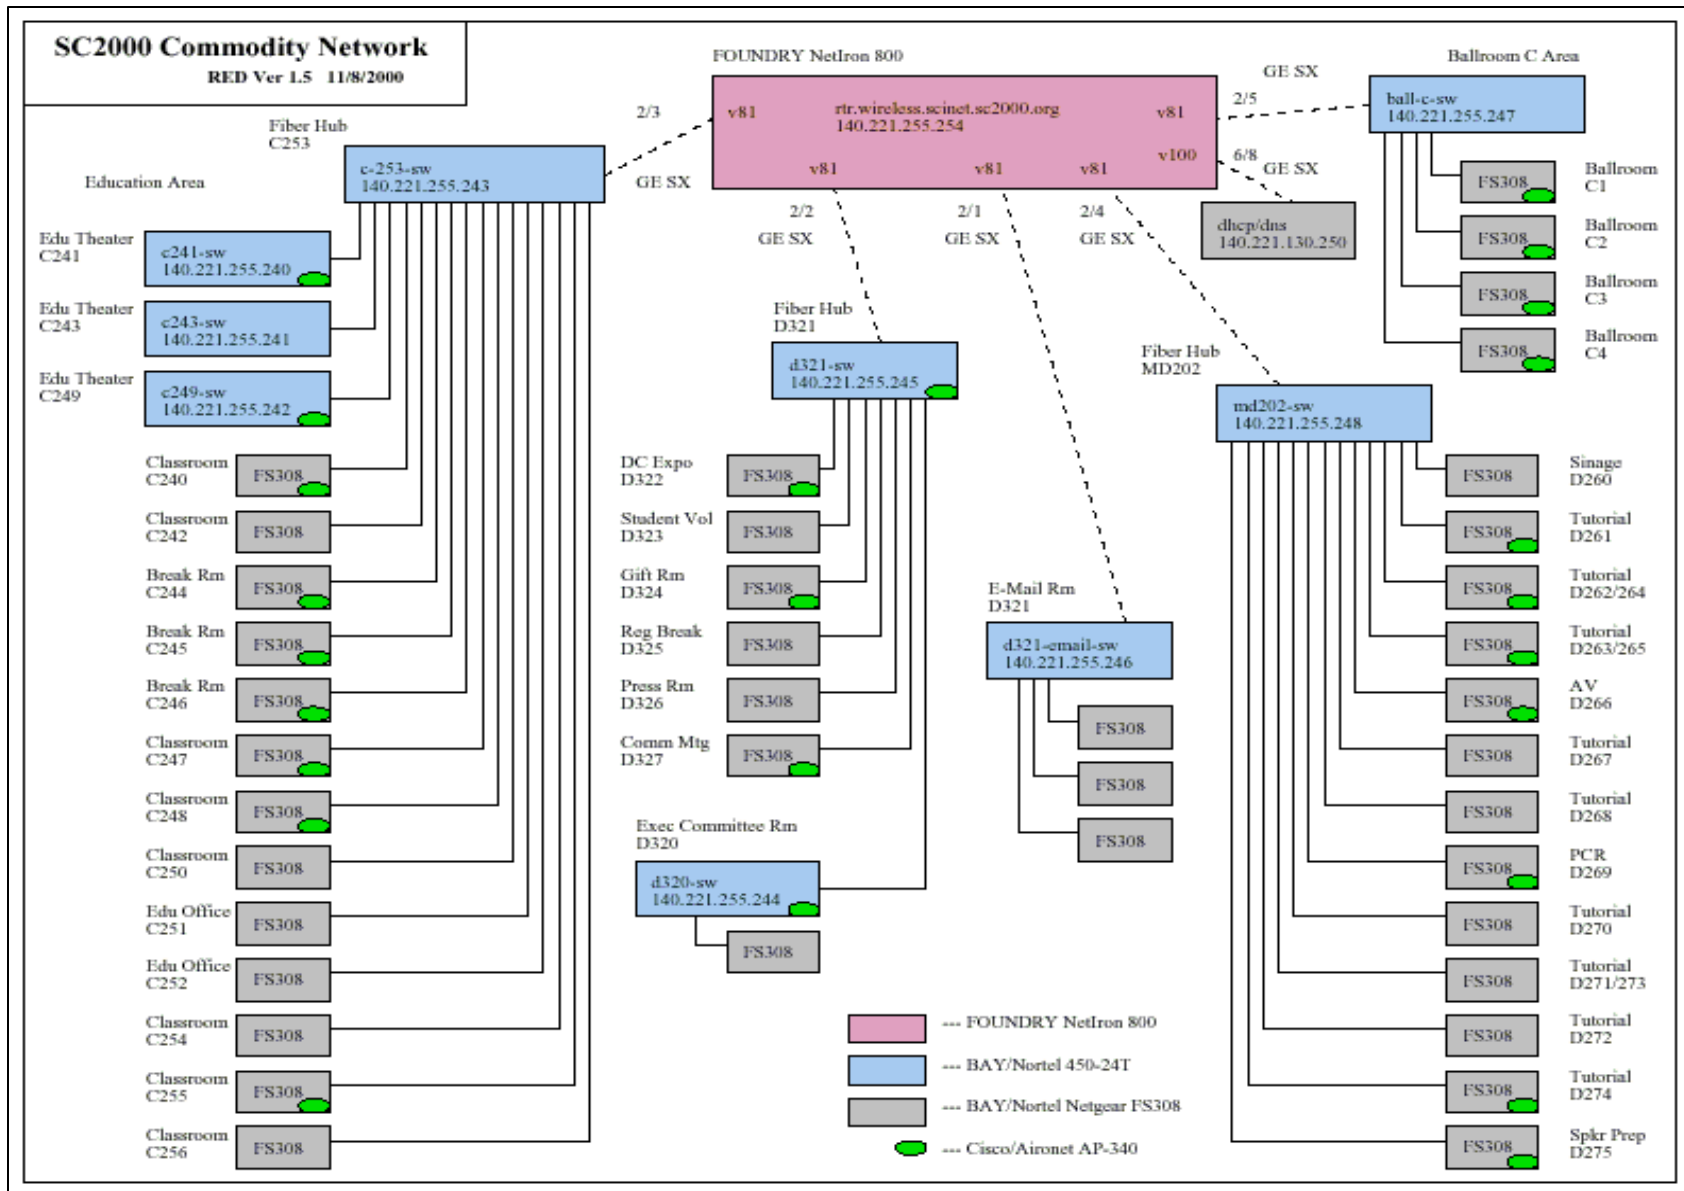


Figure 3 – Final Commodity Network Diagram

- ATDnet⁴ – The Advanced Technology Demonstration Network (ATDnet) is a high performance networking testbed in the Washington, DC area, intended to be representative of possible future Metropolitan Area Networks. Established by DARPA, ATDnet has a primary goal to serve as an experimental platform for diverse network research and demonstration initiatives. Emphasis is on early deployment of emerging Asynchronous Transfer Mode (ATM) and Synchronous Optical Network (SONET) technologies.
- VBNS+⁵ – a nationwide network that supports high-performance, high-bandwidth applications. Originating in 1995 as the vBNS, vBNS+ is the product of a five-year cooperative agreement between MCI Worldcom and the National Science Foundation.

Other major national networks were available because of peering relationships. One such network was

- NTON⁶ – a 2500 km 10-20 Gb/s Wavelength Division Multiplexed network deployed using in-place commercial fiber between San Diego, CA and Seattle, WA. NTON links government, research and private sector labs and provides the ability to interface with most of the broadband research networks in the U.S. NTON provides direct access to many of the major universities on the West Coast at data rates up to, and potentially beyond, 2.5 Gb/s. For SC2000, applications and demonstrations using NTON were routed through the OC-48 packet over SONET (POS) NTON-HSCC peering point at Los Angeles.

The total connectivity between SC2000 and the outside world was 8.4 Gigabits per second.

SCinet Router Design

The first layer of the production network consisted of three core routers and two ATM switches:

- **Core-Rtr-1:** A Cisco GSR 12000 with the following interfaces: two OC-48 ATM connections, four OC-12 ATM connections and nine Gigabit Ethernets. This router connected directly to the Abilene network and the Fore ATM switch.
- **Core-Rtr-2:** A Juniper M20 with the following interfaces: two OC-48 Packet over SONET connections, two OC-12 Packet over SONET connections, one OC-12 ATM connection, four Gigabit Ethernets connections and one tunneled PIC. This router connected to the HSCC and vBNS networks directly.
- **Core-Rtr-3:** Foundry Net Iron 800 with the following interfaces: four OC-48 ATM, 20 Gigabit (1,000 Mbps) Ethernet connections and 24 Fast (100 Mbps) Ethernet connections. This router connected to the Juniper M20.
- **ATM-Sw-1 and 2:** Two Marconi ASX-4000 ATM switches. They each had the following interfaces: four OC-48 ATM single mode fiber, eight OC-12 ATM single mode fiber, 24 OC-12 multimode fiber and 64 OC-3 multimode fiber. ATM-Sw-1 connected to ESnet, and ATM-Sw-2 connected to ATDnet.

Below the core level, a number of other routers provided specific connections to exhibit booths and switching and interface functions. A large number of interfaces had to be provided because of the many different types of media interconnections being used. Two Cisco Catalyst 6509 routers provided twenty-six Gigabit Ethernet interfaces and 34 Fast Ethernet connections. An Extreme Networks Black Diamond provided thirty-two Gigabit Ethernet, thirty-two Fast Ethernet and eighty 10/100 Mbps Ethernet interfaces. A second Extreme provided sixteen Gigabit Ethernets and sixteen Fast Ethernet connections. Both Extreme devices supported Ethernet connections that required Jumbo Frames. A Foundry FastIron WG provided one Gigabit Ethernet and twenty-four Fast Ethernet connections. A Marconi ESR-5000 provided two OC-12, two Gigabit Ethernets and 10 Fast Ethernet connections. A Nortel 450 provided one Gigabit Ethernet and twenty-four 10/100 Mbps Ethernet connections. Finally, a separate Cisco 7206 router provided support for the IP version 6 interfaces. Myrinet (both 1.2 Gbps and 2.0 Gbps) connections were provided between exhibit booths, and also forwarded IP packets to the Marconi router.

The Network Operations Center (NOC) was developed from scratch just for this event. This year, in addition to the traditional functions of supporting the network equipment and providing Help Desk and work areas for the network staff, the NOC had a variety of displays and information. All the equipment in the NOC was supported by two different Uninterruptible Power Systems (UPS), one from APC and another from Best.

Network addressing was handled in a number of ways. DHCP provided addressing for the wireless network and most aspects of the commodity network. For the SCinet Production network, a number of address ranges had to be provided for the different levels of the network. Many of the research experiments required permanent address assignments. Each exhibitor received a set of addresses that they controlled. Address assignment was managed using custom software developed by SCinet staff. This software also populated the DNS database.

⁴ See <http://www.atd.net> for more details

⁵ See <http://www.vbns.net> for more details

⁶ See <http://www.ntonc.net> for more details

Figure 4 shows the entire logical map of the network.

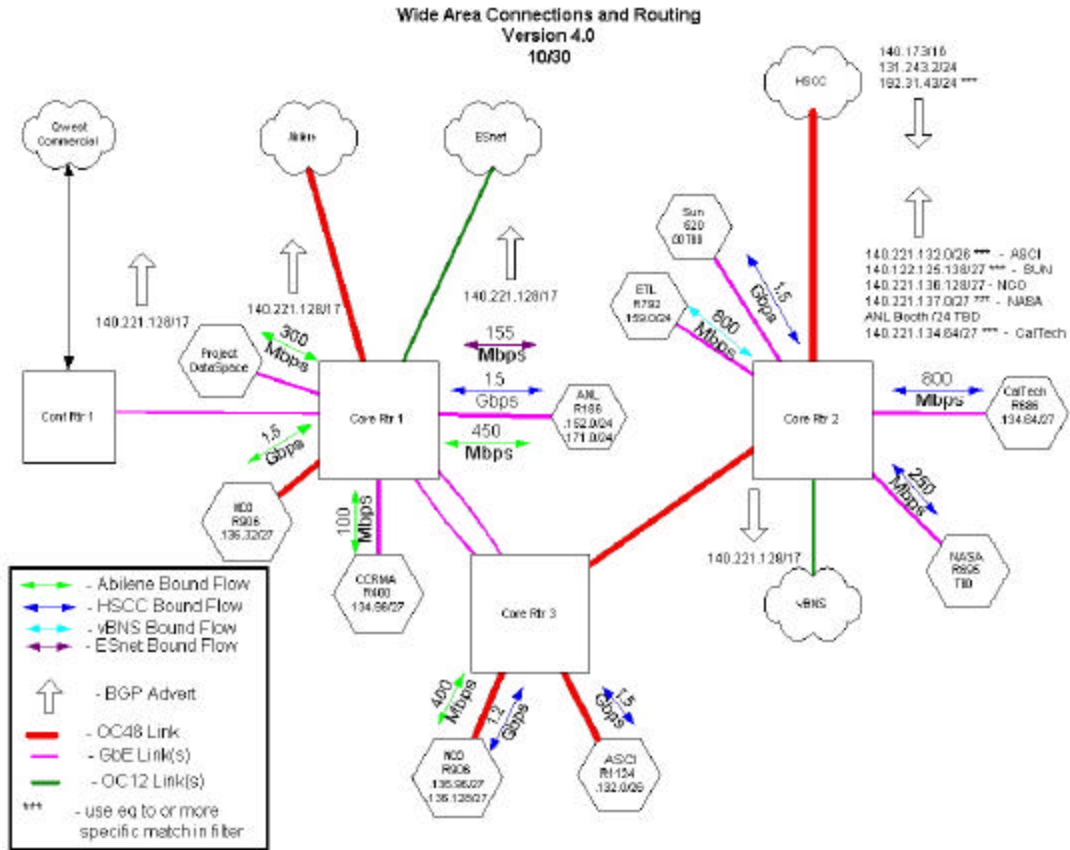


Figure 4 – Wide Area Connectivity Diagram

Experimental Network (Xnet)

The first three levels of networking have to provide relatively stable service, appropriate to the level of aggressive use of technology, but must be careful to provide redundancy and to use technology that is likely to provide reliable service. Vendors sometimes are reticent about showcasing bleeding-edge hardware in SCinet if it only were a production network. Thus, the fourth network incorporated into SCinet is an experimental network, labeled *Xnet*, which provides the solution to this dilemma. The goal of Xnet is to showcase pre-production network techniques, technology, or protocols that have (or will have) strong impact on high-performance networking, computing, and storage.

Xnet demonstrates possibilities, not production-quality products. It provides a context which is, by definition, bleeding-edge, pre-standard, and in which fragility goes with the territory. It provides vendors an opportunity to showcase network equipment or capabilities that typically does not exist outside the development lab. Xnet is the leading edge, technology-development showcase segment of SCinet and therefore is another research component for the network experiment.

At SC2000, the Xnet network was a point-to-point network arranged between the ASCII and SGI booths, provisioned using Cisco's pre-production 10 Gigabit Ethernet interfaces for their 6500 series switching routers. When forced to choose from the different optical interfaces they are working on (short haul serial, long haul serial, and parallel), Cisco selected the parallel interface to showcase at SC2000. This interface short-circuits the full serialization process by intercepting the 4 parallel XAUI streams and running them out directly as parallel data streams on optical ribbon cable. This ribbon cable has a reputation for being difficult to work with, so SCinet actually installed six separate spools of ribbon cable (the network required that four of them actually work). The goal of the demonstration was to show a 20-CPU storage cluster in the SGI booth (which was hooked to the switch through 20 separate Gig-E interfaces) feeding data through a pair of the 10-Gigabit Ethernet cards to a 20-processor compute cluster in the ASCII booth (again interfaced with 20 Gig-E links) which was processing the data and rendering images.

Fiber Infrastructure

A major part of creating the SCinet network is installing a completely fiber-based infrastructure in the exhibit and other selected areas. This year, 82.5 miles of fiber optics were run in the ceiling and throughout the exhibit areas. The fiber consisted of fifteen 170-meter, 24-pair multimode fiber spools, three 170-meter, 24-pair single mode fiber spools, nine specialized fiber spools and over 140 100-meter, two fiber patch spools.

The fiber infrastructure was a star-and-hub arrangement. The large 24-pair spools ran to different areas of the conference area. From the spools, a star of two fiber spools ran to each termination point. The other end of all the large spools ran back to the NOC. From the NOC, a 1,400-foot 24-pair fiber run was made to the demarcation point for connection to the Qwest dark fiber.

The wide area network connectivity was used both for individual applications and demonstrations as well as general usage. Qwest provided an individual fiber pair for each of these connections for the duration of the show. The fiber connections actually run in a ring to a major point of presence in Dallas, where many of the major network carriers exist. Qwest, which supports five of the networks, then patched that fiber to existing network terminations within the POP.

Measurement, Monitoring and Evaluation

The network is designed with measurement and monitoring technology incorporated from the very beginning. Several methods were used to monitor and measure the networks. Specific applications and events were monitored throughout the week.

- Spirent Systems [Smart Bits](#) and Adtech technology to monitor and measure aspects of SCinet. SmartBits is the industry standard for network performance analysis for 10/100/Gigabit Ethernet, ATM, Packet over SONET, Frame Relay, xDSL, Cable Modem, IP QoS, VoIP, Routing, MulticastIP, and TCP/IP.
- The Internet-2 “Weathermap”⁷ technology was used to monitor wide area flows.
- Further measurement was made by the Cisco Netflow software package.
- The “Bro”⁸ package from LBNL was used to monitor network traffic for intrusion.
- The SCinet team also created custom software to measure other aspects of the network such as the wireless usage.

Spirent, Adtech and Bro used optical splitters to tap into the actual network connections at various points in the network.

SC2000 Network Applications and the “Bandwidth Challenge”

In order to encourage the demonstration of bandwidth-intensive applications on this unique, once-a-year network, high-performance, bandwidth-intensive application demonstrations were developed. Twelve of these were evaluated in a formal judging called the “SC2000 Network Bandwidth Challenge.” These and others are applications that both stress the capabilities of the network and deliver innovative application value. A list of the applications is provided in Appendix A and details can be found at <http://www.sc2000.org/scinet> or http://www-fp.mcs.anl.gov/sc2000_netchallenge. A few thumbnail sketches of the applications are listed here to provide some idea of the use of the network.

- *Visapul*⁹ – *Using High-Speed WANs and Network Data Caches to Enable Remote and Distributed Visualization* – A prototype remote visualization application and framework for terascale data sets.
- *QOS* – *Enabled Audio Teleportation* – A real time demonstration using CD quality sound that shows Quality of Service to mark packets for expedited forwarding across intentionally congested network links.
- *A Data Management Infrastructure for Climate Modeling Research* – Demonstrated an infrastructure for secure, high performance transfer and replication management for large data sets.
- *ATDnet* – Greater than Gigabit per second applications between the SC2000 Exhibition floor and the Advanced Technology Demonstration Network (ATDnet) high performance networking testbed in the Washington, DC area.

Applications included one-way and two-way 1.5 Gbps uncompressed, progressive HDTV sensor and display streams and extension of GSN over a wide area network. The Qwest OC-48c service connected from one of the SC-2000 ATM switches to a Qwest location in Washington, DC. From there Verizon provided

⁷ The Weathermap technology is part of the Indiana University Network Administration Suite, which is a collection of programs developed at Indiana University for the maintenance and management of campus networks as well as the Abilene, TransPAC, and STAR TAP networks.

⁸ V. Paxson, “Bro: A System for Detecting Network Intruders in Real-Time,” *Computer Networks*, 31(23-24), pp. 2435-2463, 14 Dec. 1999. (<http://www.aciri.org/vern/papers/bro-CN99.html>).

⁹ See <http://www-vis.lbl.gov/projects/visapul/visapul-dpss.html> for more details

connectivity via an extension of a portion of the ATDnet optical network to the Qwest location.

The ATDnet ATM network provided a minimum of OC-48c ATM connectivity to all of the ATDnet sites. The ATDnet agencies (NRL, NSA, DISA, DARPA, DIA and NASA) partnered with Verizon, Qwest, SGI and Marconi to demonstrate several greater than Gigabit per second application “firsts” between Washington, DC and SC2000. One of the applications was the first long distance demonstration of uncompressed Progressive HDTV video conferencing. This two-way interaction over the ATDnet’s ATM and optical network layers was accomplished via a Qwest OC-48c connection from Dallas, Texas to Washington, DC and Verizon links in Washington, DC in full Progressive HD quality without processing or compression latency. The live video from cameras in the SC-2000 Qwest booth at NRL in Washington, DC was digitized, and the 1.485 Gigabits per second digital video stream (SMPTE 292M) was adapted to ATM using Tektronix Video Network Adapter Units (the network bandwidth is over 1.65 Gigabits per second in each direction). Uncompressed one-way streaming of Progressive HDTV was also demonstrated, including computer visualization, live video from NRL and NSA, and recorded ABC Network material. The computer visualization applications were fully interactive from SC2000 using SGI Teleffect software running on a network-connected O2 in the NCO/ITR&D booth. The primary application was NRL’s “mother of all databases, MOADB” which provides geospatial access to over 500 Gigabytes of still and motion imagery and other data types. This content was rendered in the 720-by-1280, 60 Hz progressive HD format by an SGI Onyx IR3 at NRL, and the digital video stream was again adapted to ATM. This Progressive HD video was displayed in the Qwest, NCO, and Marconi booths.

The ATDnet OC-48c ATM connection was also utilized to demonstrate the first extension of a Gigabyte Systems Network (GSN) high performance computer interface outside of the computer room and across a wide area network. This computer-to-computer network connection was accomplished using a GSN to ATM adapter (interim Gigabyte ATM Network Adapter, iGANA) developed by NRL and tested at SC2000 for the first time over a long distance. The iGANA tests were very successful, with repeatable data transfer rates of over 146 Megabytes per second (1.17 Gigabits per second) sustained for over 45 seconds at a time. (There did not appear to be a limit to sustaining this rate. The duration was simply determined by the 6.7 Gigabyte size of the test and line rate limitations of the iGANA.) These tests were accomplished using the Scheduled Transfer (ST) protocol and “gsnsttest” (similar to ttcp). The ST protocol in the SGI host demands very little processing power, and the CPU utilization related to these transfers was under 5 percent (<https://www.atd.net/sc2000/results/>). The interim GANA functionality encapsulates the GSN protocol and carries a great deal of overhead (this will be eliminated in the final GANA) so the one-way ATM network bandwidth was nearly 2 Gigabits per second during these tests. Because only 2.4 Gigabits per second was available to the ATDnet, these tests could not be conducted at the same time as the HDTV demonstrations, and most of this testing was conducted outside of exhibit hours Monday, Tuesday and Wednesday. Most of this time was consumed with finding and making all of the adjustments needed to window and buffer sizes to overcome the usual long fat pipe issues. Due to the indirect route over which the DC-to-Dallas OC-48c was provisioned, the round trip latency was nearly 40 milliseconds.

More information regarding the ATDnet demonstrations at SC-2000 is available at <https://www.atd.net/sc2000/>. Due to late determination of the specifics of the ATDnet demonstrations, they were not entered into the Network Challenge competition.

Results and Observations

The SCinet testbed was able to demonstrate a variety of new capabilities and insights. These broke down into several areas.

The Timeline

The activities for SCinet from first arrival at the DCC to completing teardown were 11 days, 12 hours and 30 seconds. The complete history is in Appendix C. Some major points were that 82 miles of fiber optic cable was installed in less than 51 hours for a rate of 1.4 miles per hour. The NOC was built and equipment installed and operating within 60 hours of arrival. The first high bandwidth external connectivity occurred in just under 80 hours, and the first high bandwidth user application – a videoconference at 30 frames a minute – was in just over 5 days.

Overall Bandwidth and Usage

A comparison of the connection information from SC2000 vs. SC99 shows a dramatic increase in high bandwidth connectivity (Table 1). If this trend continues, SCinet 2001 should be very interesting.

Type of Connection	Number of Connection	
	SC99	SC2000
OC-48c ATM	2	6
OC-48 PoS	1	5
OC-12c ATM	5	13
OC-12 PoS	0	2
OC-3c ATM	11	7
1,000 Mbps-LX	0	5
1,000 Mbps-SX	29	67
100 Mbps-FX	46	79
10 Mbps-FL	22	0

Table 1 – Connection List

The total bandwidth associated with the SCinet Production network routers is 118.31 Gbps in SC2000 (compared with 46.3 in SC 99.) These values do not include Xnet – which added 8 Gbps, commodity and wireless networks. Nor does this amount indicate the Myrinet connections (3 Gbps) nor the dark fiber connections provided for the exhibit floor. Adding these together brings the total internal bandwidth to over 130 Gbps.

Summarizing the commodity network usage, the maximum rate into the DCC was 10.9 kb/s (0.0%), and the average rate in was 8.0 kb/s (0.0%). The maximum rate outbound was 70.5 Mb/s (70.5%) while the average outbound rate was 730.3 kb/s (0.7%).

The total external connectivity to the Conference is summarized in Table 2

Network	Type	Maximum Speed
Abilene	OC-48 ATM	2.5 Gbps
ATDnet	OC-48c ATM	2.5 Gbps
HSCC	OC-48	1.5 Gbps
Esnet	OC 12 ATM	622 Mbps
vBNS	OC 12 ATM	622 Mbps
vBNS	OC 12 POS	622 Mbps
Commodity	ATM	12 Mbps
Total		8.477 Gbps

Table 2 – External Connection List

Several interesting points can be made about this configuration. First, as shown in Figure 4, a number of major peering points were set up. Most of the peering traffic, and indeed 9 of the 12 bandwidth challenges, used HSCC to route to the NTON and other networks. HSCC actually routes traffic over the Qwest backbone network, which is an OC-48. The backbone traffic accounts for approximately 500 Mbps, peaking at times to almost 1 Gbps. In order not to impact the backbone traffic for the large number of Qwest clients, SCinet agreed to limit traffic over the HSCC link to 1.5 Gbps. This limit was implemented through careful scheduling of the demonstration applications that used HSCC, self-throttling these applications, and monitoring the traffic in great detail. This limit turned out to be the major performance limitation for some of the applications.

A second interesting issue arose because both the HSCC network traffic and the commodity network traffic flowed over the Qwest backbone. While routing out of the conference could be specified, it was not possible to separate the return, acknowledgment packets, so SCinet designated the commodity network as the official network.

Network Usage

Individual bandwidth measurements show that at least one application achieved over 3.2 Gbps on a sustained basis, transferring HDTV data streams with real time control of images between Dallas and Washington DC. The Visapult application transferred 1.56 Gbps on a 5 second average, and 1.76 Gbps on a 0.1 second sample that directly monitors the application by associating sockets and IP address for traffic analysis. The 5-second sample was measured by using SMNP polling of the routers involved, while the 0.1 sample rate came from the Adtech measurement devices. This application reported a 1.48 Gbps sustained rate during the demonstration run on November 8. . The application transferred 266 Gigabytes during the one-hour demonstration period. Figure 5 shows the Adtech plot of the Visapult application performance over the hour demonstration period. This graph shows several interesting aspects, including a time when there was little traffic. This was when the application was resetting to add another server, which boosted peak and sustained performance. Overall, this application was judged the “Fastest and Fattest” for reaching the highest measured speed and transferring the most data. According to the authors, this application could have reached close to

the 2.5 Gbps level if they were allowed to use the entire HSCC link. A second application, *A Data Management Infrastructure for Climate Modeling Research*, also sustained performance over the hour of more than 1 Gbps, connecting several wide area sites.




Figure 5 – Bandwidth Demonstration of Visapult Application

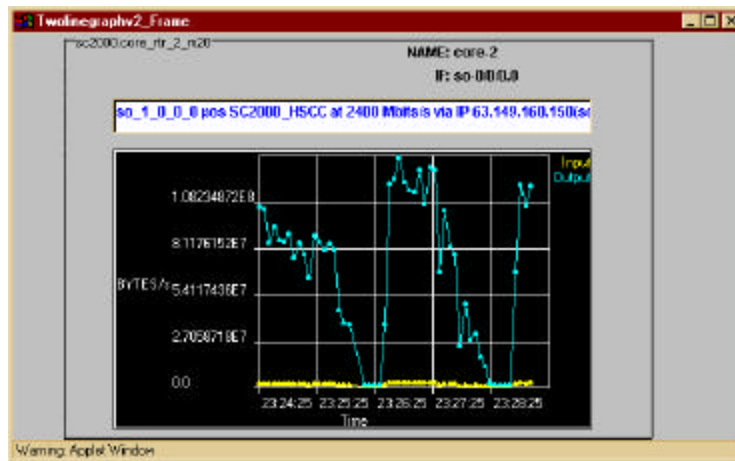


Figure 6 – Example plot of the 5 second SNMP polling for Bandwidth Tracking during Visapult Application

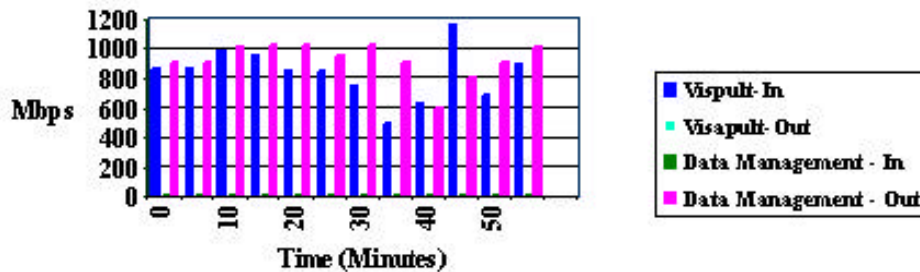


Figure 7 – Tracking during Visapult and Data Management Applications

There are three network performance measures: within the application, monitoring packets associated with the distributed parts of the application, and monitoring the router. All three of these measures show agreement.

SCinet also experimented with how much bandwidth the entire network could support in and out of the conference at the same time. Figure 8 shows a snapshot of the external network usage from Friday, November 9. Interestingly, while a number of bandwidth intensive applications were running at the time, Visapult and the other HSCC bound applications were not running. Still, the high water mark bandwidth usage was observed at 4.92 Gbps out of the maximum 8.477 Gbps – or a sustained 58% utilization. The network “weather map” below shows the last measurement of the network that totals 4.4 Gbps.

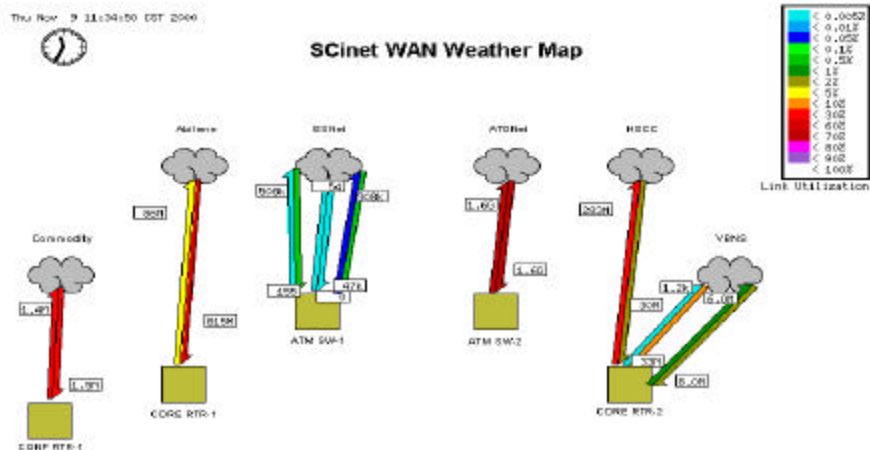


Figure 8 – The Last SCinet Weather Map snapshot - Aggregate bandwidth of 4.4 Gbps

It may be that the high number of connections and the aggregate bandwidth provided and used indicated the accelerating pace of network technology and usage. It is also an indication that the SC conference series is succeeding in its efforts to expand the high performance networking activities of the conference. Clearly, there are several applications that showed their ability to use a significant share of the high bandwidth provided for meaningful applications. This is over three times the usage from a year earlier.

Wireless Network

The wireless network was used and well received by a significant number of attendees. The donation of wireless cards to every teacher in the education program guaranteed 120 wireless clients. Measurements show that up to 300 clients were simultaneously using the wireless network. Figure 9 shows the wireless network usage for a single access point. Figure 10 shows the number of clients on the network.

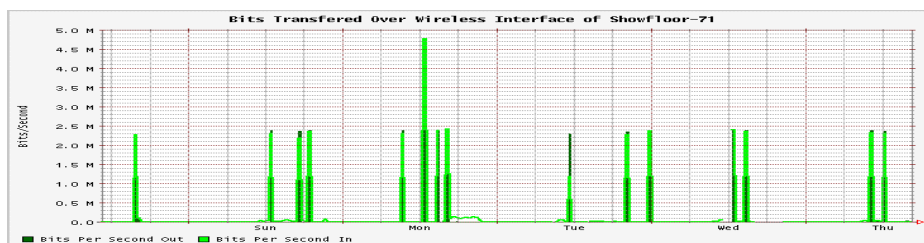


Figure 9 – Bandwidth Usage for the Wireless Network for one Access Point

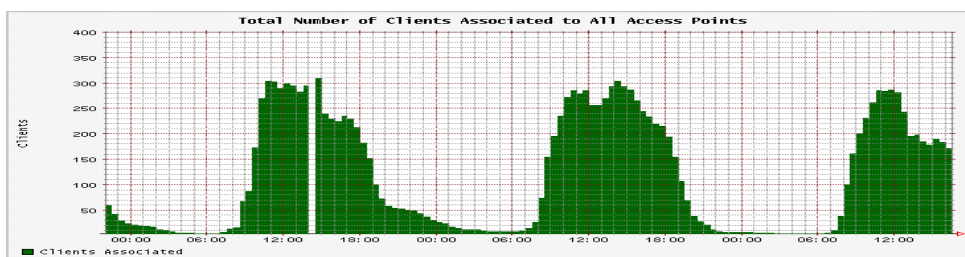


Figure 10 – Clients associated with the Wireless Network

The biggest issue with the wireless network was identifying and resolving interference with the access points. There were at least two sources of interference. The first was non-SCinet access points and the second was other 2.4 GHz equipment such as wireless video extenders. Three methods were used to identify sources of interference. Broadcasted SSIDs (i.e., "NPACI rox") sometimes provided a pointer to the responsible organization for a competing access point. A laptop system with utility software provided by Lucent was used to look at noise levels for the different channels. SCinet staff wandered around the areas to find the source of interference using this laptop. The Cisco utilities also showed SSID mismatches that gave hints of rogue Access Points.

The SCinet team wrote software for automating the installation of a large-scale wireless infrastructure. These automated configuration tools make it easy to configure a large number of APs and guarantee consistent configuration for good performance. The staff was able to provide support for Lucent bronze cards by setting the APs for “basic” 1 Mbps speed and “yes” for the 2, 5.5, and 11 Mbps speeds. This was in addition to the 2 Mbps client support button in the AP configuration.

Using *cricket* software to monitor the APs was helpful by showing when an AP stopped working so the reasons could be investigated.

Initially the access points deployed in the show floor were on channels that were not optimized. The channels were reallocated to better use channels 1, 6 and 11, which gave the maximum separation of frequencies. This provided the best channels to minimize overlapping radio transmission in spite of the physical overlap of coverage. To further decrease interference effects, the power of some of the APs was lowered from 30 mW to 5 mW. This resulted in better performance than at 30 mW, but the wireless link was sporadic in at least one area (the Network Operations Center) at this power level. The APs’ power was raised back up to 15 mW. The improvement of wireless performance when the access point frequencies are varied within the frequency channel range (adjacent access points are kept at least two channels apart) has a significant impact.

An experiment was tried to shut off a few APs to see if all were needed. The number of operating APs in the exhibit areas was decreased by 25%, from 12 in to 9, but the performance of the network went down. Cisco support suggested a maximum user count of 30 clients on an AP, which appeared to be correct.

At times it was useful to force a re-association of clients when performance on a particular AP decreased. An example of this was when a large number of clients started up in one area shared by a small number of APs. The vast majority (>80%) of the clients were assigned to one AP. The algorithm used is unknown, but this behavior may be due to insufficient randomness in the search algorithms used in client NICs. The overloaded AP demonstrated more errors and lower performance. Gradually, some of the clients migrated to the underused AP, yet performance remained uneven. The transition was relatively slow (over a period of hours). Experiments were done to adjust the power level of some APs to force reassociation, which worked well by hand, but is not a scalable operation. Next year, SCinet plans to write software to automatically monitor and control the load-using power adjustment.

A small number of APs periodically dropped off the net. They would sometimes come back and sometimes not. This did not appear to correlate to the AP software version. The one unknown is the serial number of all the APs: there may have been some older APs in our deployment which might have had problems. There are some known problems with older AP hardware having difficulty changing connection speeds down from 11 Mbps.

Xnet

Initial tests showed the system was working, but was showing such a high error rate that the actual throughput was unusable. The issues were purely fiber related, or at best fiber interface related. After fiber swapping, the error rate came down to a fully usable level, and the system did actually deliver the promised level of performance. The 10 GE interfaces performed well and did not require swapping or replacement.

The demonstration implementation used “striping” at the application layer, and none of the drivers or protocol stacks had been optimized. Eight 1 Gbps streams between a pair of 10-GE ports was consistently demonstrated. Due to fiber limitations, it was not possible to utilize all four 10-GE boards. The maximum that would have been transferred was 8 Gbps due to the number of GE feeder ports.

Routing

The routing for SCinet was complex. Part of the complexity was due to the number of WAN connections being used, and the fact that different exhibits were using different paths to accomplish their demonstrations. Another reason was the desire of the testbed to try to mimic real wide routing that exposed issues that often crop up in subtle manners. Running I-BGP on all the distribution switches was the proper thing to do, as it made for more effective traffic flow to the appropriate edge router for the high performance external links. Dynamic routing was used, in that traffic could switch to alternate links if there was a complete failure on link. Adaptive routing was not used.

Getting commodity IP transit service from Qwestlink, while at the same time using the Qwestlink HSCC product as a high performance external link, made external routing considerably more complicated than it might be otherwise.

It was particularly interesting how the layer 3 switches performed under load. These observations are particularly interesting since they were made under real load in a real networking environment. Real world experience suggests there are many subtle issues that can only be evaluated in complex networking, operating with a diverse set of applications.

Support for IPv6 still has room for improvement. Cisco devices worked, but with some issues that need to be resolved. Juniper, Foundry, Extreme and other equipment did not support IPv6 sufficiently for usage at the show.

Network Monitoring

To truly monitor the network at high resolution, as was done for the network bandwidth challenge, more precision than SNMP polling was very valuable. The one-tenth (0.1) of a second resolution on the Adtech devices showed significant detail; surges as much as 13% higher than the SNMP polling at 1 second were clearly shown for some applications.

While there were no major intrusion attempts, Bro demonstrated the effectiveness of such monitoring. Bro detected a number of things that under normal circumstances would be less than desirable system management practices. An example is that some exhibitors logged directly into root accounts from remote locations. Seeing how Bro was implemented with optical splitters and kernel reassembly of monitored data proved that splitters are a very effective way to deliver the data to be studied.

Network Operations

There were many pieces of equipment that could accept wide voltage ranges. There were only a couple that required a specific voltage. The same holds true with the plugs. Most were typical 15A plugs. A few were non-standard, but they were exceptions rather than the rule. The APC Symmetra unit ran at 85% load for a 16kVA system. The 18kVA Best Axxium Pro ran at 43% load.

Device	Amperage (Amps)	Voltage (Volts)	UPS Percent Utilized
First Power to NOC Racks			8%
Foundry NetIron on-line			6%
Cisco GSR, Foundry NI	8	237.4	15%
Fore ASX40000 #1	13.7		28%
Fore ASX 4000 #2	20		42%
Juniper M20	21.3		46%
Extreme Black Diamond	28		58%
Extreme Summit 4	29.3		60%
Marconi ESR 5000	20.7		63%
Cisco 7507	32		66%
SPIRENT Equipment	28.7		76%
PCs, Laptops, Sun	40		79%
Cisco 6509 #1	42		83%
Cisco 6509 #2	44		86%

Table 3 – Power usage by different network devices

Conclusions

The effort and timeliness of SCinet 2000 is shown in the Table 4.

Total miles of fiber installed	82.2 hours
Time from first lift to outside connectivity	59.7 hours
Miles of fiber per hour	1.4 fiber miles per hour
Time to first OC-48 connection (Abilene)	80.2 hours
Total theoretical peak external bandwidth	8.477 Gbps – self limited, 9.477 Gbps actual interface
Estimated theoretical peak show floor bandwidth	More than 130 Gbps
Wireless coverage area	Entire show area 207,338 square feet
Total effort of volunteers	11.27 people years
Value of volunteer efforts at \$200,000 per year	\$2,225,000
Estimated value of donated equipment and services for the testbed	Greater than \$25,000,000

Table 4 – SCinet Summary

The degree to which real applications are able to take advantage of very high network bandwidth is impressive. The Berkeley Visapult application and the Argonne climate modeling application both sustained over a gigabit per second in the wide area for an extended period of time. Yet, it is still the case that not all developers appreciate the fact that the

WAN is not a LAN, since there were several applications demonstrated that did well in the local network (without hops and delays) and did very poorly with wide area demonstrations. The stunning performance of the remote, interactive, digital video demonstrations from NRL carried over HSCC achieved sustained rates between 3.2 and 3.4 Gbps. This application might have been the most impressive of the show.

SCinet accomplished all of its objectives. It provided a very functional and stable network for basic show functions. The wireless network was a huge success but at the same time pointed to key insights that will need to be addressed in the future before a true large-scale, production-quality implementation is accomplished within such an open environment.

The intentionally complex network design yielded valuable information and experiences for both the vendors and the network engineers that will be put to good use. The growth in the network capacity and demand indicates that acceleration is taking place in high performance computing that will continue to drive the need for such unique testbed activity.

Finally, no matter what the application portfolio or the amount of equipment, this scale project could not be successful in this time period if not for the expertise and commitment of the SCinet volunteer staff. Clearly the most limited and essential factor to continuing the ever-increasing network usage is the people doing it.

Acknowledgements

This year, the team consisted of people from many organizations. The effort was more than 11 people-years of effort spread out over approximately 75 people who contributed to building, running and measuring the network. A number of companies and organizations loaned equipment and/or services to this effort. These groups include:

Aaronsen Group, Argonne National Laboratory (DOE), Army Research Laboratory (DOD), Avici Systems, Caltech, Corp of Engineers Waterways Experimental Station (DOD), Cisco Systems, the Dallas Convention Center, the Dallas Convention and Visitor's Bureau, Extreme Networks, Foundry, GST Telecom, Internet-2 (NSF), Juniper Networks, Lawrence Berkeley National Laboratory (DOE), Lawrence Livermore National Laboratory (DOE), Marconi, MCI, MITRE Corporation, National Center for Supercomputing Applications (NSF), Northeast Regional Data Center/University of Florida, Nichols Research/CSC, Nortel Networks, Oak Ridge National Laboratory (DOE), Oregon State University, Pacific Northwest National Laboratory (DOE), Qwest Communications, Sandia National Laboratory (DOE), Spirent Communications, Texas A&M University, University Corporation for Advanced Internet Development, University of Tennessee/Knoxville, the very high performance Backbone Network Service+ – vBNS+ (NSF).

Appendix A

SC2000 Network Challenge Entries

Visapult – Using High-Speed WANs and Network Data Caches to Enable Remote and Distributed Visualization – W. Bethel, J. Shalf, S. Lau, D. Gunter, J. Lee, B. Tierney, V. Beckner, J. Brandt, D. Evensky, H. Chen, G. Pavel, J. Olsen, B.H. Bodtker

World Wide Metacomputing – M. Mueller, S. Sanielevici, A. Breckenridge, S. Sekiguchi, J. Brooke, F.-P. Lin, T. Imamura

Development of a Telescience Portal – M. Hadida, T. Hutton, M. Martone, A. Gupta, R. Moore, S. Peltier, S. Khetani, M. Wong, A. Lawrence, M. Ellisman, S. Mallen, J. Haynes, F. Berman, B. Fink, M.-H. Su, C. Kesselman, M. Sany, R. Wolski, A. Shamir, C. Bajaj

QoS Enabled Audio Teleportation – C. Chafe, S. Shalunov, B. Teitelbaum, M. Groger, R. Roberts, S. Wilson, D. Chisolm, R. Leistikow, G. Scavone

Project DataSpace – R. Grossman, E. Creel, M. Mazzucco, S. Connelly, A. Turinsky, H. Sivakumar, S. Wahlston, B. Hollebeek, P. Proropapas, R. Williams, R. Irwin, D. Rocke, T. Arons, Y. Guo, S. Hedvall, P. Milne, G. Williams, G. Becker, J. Hubshman, W. Martinez

Reservoir Simulation and History Matching – Grid Based Computing and Interactive Dataset Exploration – J. Saltz, T. Kurc, U. Catalyurik, M. Wheeler, S. Bryant, M. Peszynska, A. Sussman

Gigabyte per Second File Transfer in a Clustered Computing Environment – T. Pratt, J. Naegle, L. Martinez, M. Barnaby

Gigabit/sec High Definition TV over IP – C. Perkins, L. Ghari, A. Mankin, T. Gibbons, D. Richardson, G. Concher

High Resolution Visualization Playback on Tiled Displays – M. Papka, R. Stevens

Scalable High-Resolution Wide Area Collaboration over the Access Grid – L. Childers, T. Disz, B. Olson, R. Stevens

Bandwidth Thirsty Particle Physics Event Collection Analysis and Visualization Using Object Databases and the Globus Grid Middleware – J. Bunn, H. Newman, J. Patton, K. Holtman

A Data Management Infrastructure for Climate Modeling Research – A. Chervenak, C. Kesselman, I. Foster, S. Tuecke, W. Allcock, B. Drach, D. Williams, A. Sim, A. Shoshani

Appendix B
SCinet Team Members

Bill Kramer, Conference Vice-chair, UC Berkeley/Lawrence Berkeley National Laboratory/NERSC
Tim Toole, Deputy Chair, Sandia National Laboratories
Eli Dart, Network Security Chair, Sandia National Laboratories
Jon Dugan, Wireless Chair, National Center for Supercomputing Applications/University of Illinois
William "Bill" Wing, Experimental Network (Xnet) Chair, Oak Ridge National Laboratory
Rex Duncan, Committee Networking Chair, Oak Ridge National Laboratory
Chuck Fisher, Production Chair, Oak Ridge National Laboratory
Greg Goddard, Network Monitoring, University of Florida
Ian Foster, Application Evangelist Chair, Argonne National Laboratory
Paul Daspit, On-site Challenge Coordinator, Nortel Networks
Doug Luce, Information Management / Customer Support Chair, Aaronsen Group
Jeff Mauth, Physical Infrastructure Chair, Pacific Northwest National Laboratory
Martin Swany, Network Management/Monitoring Chair, University of Tennessee, Knoxville

Steve Corbato, Internet2 UCAID
David Wheeler, National Center for Supercomputing Applications
Zaid Albanna, MCI
Greg Almes, Internet2
Warren Birch, Army Research Laboratory
Bryan Bodker
Roberta Bourcher, Lawrence Berkeley National Laboratory
David Crowe, Oregon State University
Julie Wulf, Argonne National Laboratory
Patrick Dorn, National Center for Supercomputing Applications
Adam Duke, Florida State University
Johnny Pak, Cisco
Larry Dunn, Cisco
Hal Edwards, Nortel Networks
Stacy Eubanks, DCC
Eric Plesset, Spirent
Joseph Perches, Spirent
Riki Kurihara, Spirent
Basil Decina, Naval Research laboratory
Paul Reisinger, Marconi
Thomas Hutton, University of California at Dan Diego
Kevin Walsh, San Diego Supercomputer Center
Linden Mercer, Naval Research Laboratory
Jason Hasse, Cisco
Roland Gonzalez, Juniper Networks
John Jamison, Juniper Networks
Matthew J Zekauskas, Internet2,
Steve Jones, CEWES
Wesley K. Kaplow, Qwest
Ed Kempe, Dallas Visitor's Bureau
Tom Kile, Army Research Laboratory
Dave Koester, MITRE Corporation
Bill Lennon, Lawrence Livermore National Laboratory
E. Paul Love, Internet2
George Miller, MCI
Bill Nickless, Argonne National Laboratory
Kevin Oberman, Lawrence Berkeley National Laboratory
James Patton, Caltech
Jim Rogers, CSC/Nichols
Jim Ross, Sandia National Laboratories
Ralph McEldowney, Wright Patterson Air Force Base
Glen Smith, Qwest
Robert Spenser, Qwest

Appendix C
SCinet Time Line

Event	Date	Time from Start
Arrival	10/30/00 9:00 AM	
First Fiber Lifted	10/30/00 4:00 PM	7.00
First Newspaper article	11/1/00 8:00 AM	
First Light to DCC	11/1/00 7:01 PM	58.02
First Power to NOC Racks	11/1/00 7:48 PM	58.80
First Light Extended to Electronics (GSR)	11/1/00 8:43 PM	59.72
Foundry NetIron on-line	11/1/00 8:56 PM	59.93
Cisco GSR, Foundry NI on-line	11/2/00 12:17 PM	75.28
Fore ASX40000 #1 on-line	11/2/00 1:32 PM	76.53
Fore ASX 4000 #2 on-line	11/2/00 1:35 PM	76.58
Juniper M20 on-line	11/2/00 3:20 PM	78.33
Extreme Black Diamond on-line	11/2/00 4:10 PM	79.17
Abilene Circuit up to Dallas POP	11/2/00 4:10 PM	79.17
Extreme Summit 4 on-line	11/2/00 4:15 PM	79.25
Marconi ESR 5000 on-line	11/2/00 4:20 PM	79.33
Cisco 7507 on-line	11/2/00 4:22 PM	79.37
SPIRENT gear (collectively) on-line	11/2/00 5:09 PM	80.15
Abilene Circuit Completed - First OC 48 for SC2000	11/2/00 5:14 PM	80.23
Abilene Peering up	11/2/00 5:25 PM	80.42
Cisco 6509 #1 on line	11/2/00 5:44 PM	80.73
Cisco 6509 # 2 on-line	11/2/00 5:45 PM	80.75
vBNS OC12 Packet over Sonnet up	11/2/00 6:40 PM	81.67
Best Power UPS on-line: All power changes complete	11/3/00 1:40 PM	100.67
HSCC OC-48 POS up	11/3/00 2:30 PM	101.50
Wireless APs installed in Education area	11/3/00 5:00 PM	104.00
Address data purified and in the DB	11/4/00 7:15 AM	118.25
GPS Working for network monitoring	11/4/00 12:00 PM	123.00
Second vBNS OC-12 POS up, ATM Juniper, Marconi Up	11/4/00 12:45 PM	123.75
Completed Help Desk Software	11/4/00 12:59 PM	123.98
Began accepting drop requests	11/4/00 1:00 PM	124.00
First Video Conference DCC to NSF at 30 frames a second	11/4/00 2:00 PM	125.00
ESnet up	11/5/00 1:35 PM	148.58
Bro 3 tap	11/5/00 3:00 PM	150.00
First TV report on local ABC affiliate	11/5/00 5:30 PM	152.50
260 Wireless Clients	11/6/00 11:00 AM	170.00
Bandwidth challenge 1.56 Gbps 1 second sample peak	11/7/00 10:15 PM	205.25
SCinet Production Network Shut down	11/9/00 4:00 PM	247.00
SCinet completely torn down and shipped	11/10/00 6:30 PM	273.50
Total Time from set up to tear down		11 Days, 12 Hours, 30 Seconds

Appendix D

Lessons Learned for SC 2001 and other conferences

SCinet Access and Badging Issues

The badging of SCinet needs to be redesigned. Currently there are three indicators of SCinet usage, almost all of which overlap to a significant degree. They are:

- SCinet Staff sticker label – given out before the conference registration starts, it provides access to the NOC and allows complete access to all areas of the conference. It also allows transporting equipment in and out of areas.
- SCinet registration – provided to 35 people. This is not a registration to the technical program and does not provide a proceedings not attendee gift.
- SCinet ribbon – allows access to the NOC and allows complete access to all areas of the conference. It also allows transporting equipment in and out of areas.

Problems that were related to this scheme included the fact that a large number of the people really do not need complete SCinet ribbon access but that is the only level available.

An improved access plan would include four tiers.

1. SCinet Key Personnel – these are the workers year round who also spend a number of weeks at the conference site. The number should be flexible but budgeted to 40. These workers should get the registration goodies. These people would also have complete access to the entire conference, either with "blinkies" or some other indicator
2. SCinet Ribbon – a ribbon acknowledging significant contributions to SCinet, but a not a key personnel. The ribbons are honorary
3. SCinet Support engineers – these people are field engineer and vendor personnel that just need access to the NOC for installing, maintaining and deinstalling the equipment. Many of these people are involved for only one or two days, or are on call for problems.
4. NOC Access – many exhibitors need some access to the NOC to do network experiments, help with setup and work with SCinet staff. They do not need the complete access allowed by the traditional SCinet ribbon.

SC/SCinet should consider purchasing one of the laminated ID card systems. This way it would be possible to generate picture IDs for all SCINET staff, increasing security over the sticker system. The badge backgrounds can be modified year to year, so the investment can be amortized over several years. Consider magnetic strips on the IDs that can be used with door locks to control access to the NOC. Some vendor staff said SCinet should never let sales people into the NOC

Wide Area Connectivity

As in the past, arranging for wide area connectivity was the most difficult issue for SCinet 2000, not from a technical point, but from the need to find the right people to work with and the right companies. The "last mile" problem continues to exist. This is a job for an experienced volunteer within SCinet. There were also issues after the agreement with Qwest to provide direct access to the DCC, such as designing a flow map of all the external networks. It was in this process that discovered the HSCC link that was thought would be a full 2.5 Mbps really shared a 2.5 Mbps, OC-48 link with all the Qwest backbone service. Therefore it had to be limited to 1.5 Mbps. Issues were also discovered that resulted from using Qwest as the commodity service provider and the major carrier for HSCC.

The lesson is that there should be an explicit position on SCinet chartered to manage the provisioning of the connection out of the conference centers and coordinating with the other national networks.

Network Operations Area (NOC) Suggestions

Increase distance from the rack fronts to the glass to a minimum of 6 feet (2 meters) to allow better traffic flow during patching and physical connection debugging. There is no anticipated adverse affect on the viewing angle for exhibitors and attendees, simply a loss of square feet from the NOC. Distance from rear (mirrored) wall to rack backs was sufficient. I suggest a standard of approximately 48 inches. This works out well when considering an overall width of the platform at a standard 4 meters.

Never request equipment from Dublin, Ireland and expect it to get past Customs. Never let Sales people into the NOC (coming from a vendor).

The NOC services network should be placed so as to be as stable as possible. There were a couple of times where, because of issues surrounding the bandwidth challenge, the main NOC network was not reachable. This meant that access to the web server, database server and Bro boxes was cut off for that time. One thought is to not connect the NOC service network to a core router (see point 3, below) and instead connect it to the switch router that looks like it

will be the most stable. The corollary to this is to try to have one and only one connection between the NOC services network and whatever switch router it's connected to. That way, if a router turns out to be unable to perform to the required level, it is possible to move the connection to the NOC network easily to wherever is needed. The idea here is to design the portion of the infrastructure that is critical to be as easy as possible to move to a more stable location should the circumstances require it.

Revise the audiovisual area significantly. The projectors require a focal distance of approximately 12-16 feet. Accommodate this by revising the A/V room. Incorporate 8 x 8 ft displays (suggest two) into the exterior framework of the booth. The booth is made out of tinker toys. Surely this is a workable problem. Incorporating the displays into the framework will eliminate those annoying bars at 1 meter increments. As an aside, there needs to be increased attention given to what is intended to be displayed before well before the conference. SCinet needs someone to review A/V requirements for the NOC – since some of the plans were not coordinated with the physical infrastructure and it was not possible to use the 8 x 8 screens due to inadequate projection space.

Keep the conference room, but expand it in size and use clear glass. This way, it is possible to provide a conference room that will accommodate private conversations and provide a quiet work environment for use as needed. With glass walls, it is possible to tell at a glance if the SCinet chair/deputy is “in office.”

Keep the white foam core panels with the logos. These made a great impression. Do not put SC logos or year indicators on the panels so they can be reused.

Power needs to be simplified. There were many pieces of equipment that could accept wide voltage ranges. There were only a couple that required a specific voltage. It holds true with the plugs. Most were typical 15A plugs. A few were non-standard, but they are exceptions rather than the rule. I suggest keeping the APC relationship intact and pressing them for more of the Symmetra units. That unit ran at 85% load for a 16 kVA system. The 18 kVA Best Axxium Pro ran at 43% load. Estimate 16 kVA as the minimum and hook up only the critical power supplies in the critical routers and leave PCs, projectors, and all sorts of other garbage on the house power. I have an as-built schedule of power services used, and some rough approximations of power actually drawn by individual devices. To simplify things in future years, individual rack power distribution units that provide some set number of 20A 110V outputs and a couple configurable outputs for non-standard plug/voltage situations should be manufactured in advance. The pigtailed on each of these should plug into a standard outlet that attaches directly to the UPS. Thus it will be possible to bring in the hard wired main to the UPS, plug in the pigtailed, and have all the power needed, regardless of the last minutes changes.

Increase table space for NOC inhabitants. There was unauthorized use of the food service area for laptops. Keep the cooler. Keep the sofa, which was used by several staff after “all-nighters.”

Reduce the Help Desk footprint. It was excessive. NOC space should be optimized for people bringing laptops rather than providing displays and keyboards.

Consider PDAs or a Palm application so that drop teams out on the floor (and patch teams too) can request their next assignment electronically when they complete a drop rather than fooling with a stack of paper assignments that may have aged. This arrangement will also increase the FIFO nature of our service level. I got several complaints about non-sequential completion of drop requests. A Palm interface to the database would make this very easy. And since there is a year to write it, the schedule is not a big problem.

Routing

Layer 3 configurations should be done in advance. In the past this hasn't taken too awfully long – depending on the router it can take 10 minutes to 2 hours or so. However, if it is put in before beginning booth drops, it will make the booth drops go more smoothly.

Staging

Pre-assigning VLANs and IP addresses as early as possible is very helpful.

Trying to bring up WAN circuits at a staging 2 weeks before the show allowed the carriers to find the inevitable problems. This was extremely valuable.

Wireless

Getting the wireless network up quickly and routing to the outside (even over the existing commodity net) was a huge help, allowing both the NOC team to use the net and getting vendors on quickly. It also allowed vendors a separate path to the Help Desk system. In essence, Wireless networking is ready to become a permanent, fully supported feature of the SC conference stream if SC can assure the equipment availability. It will eventually be considered just part of the

commodity network infrastructure. As such SC needs to invest in the equipment to provide that service – just as there have been ongoing investments in commodity networking, rather than relying on vendor loans and/or donations.

As noted above, the wireless issues were mostly due to unexpected interference with other devices. SCinet should be delegated the control of all RF devices at conferences and manage the implementation of those devices in a way that minimizes overlap and interference. This is mostly an education and coordination issue.

Services

SCinet should have backup DNS and DHCP servers that are not located in the same place physically, connected to the same router, or using the same power as their primary counterparts. Preferably, one of each is on a UPS. Also, the network should be able to survive being power cycled and come back to full functionality unattended. It is probably most efficient for SCinet to actually buy and configure these core servers rather than relying on equipment loans. Low-cost PCs running FreeBSD, Linux, or some other free operating system would suffice.

Help Desk

During the Gala Monday evening, the Help Desk was closed, but in the future it the Help Desk may be open whenever the show floor is open (especially when it's first open) as well as during the vendor setup period – perhaps even half an hour before the show floor opens to handle problems vendors find when they come in each morning. Same for NOC staffing. Making sure that that both the Help Desk and the NOC have somebody there while the exhibit floor and/or sessions are taking place is important, especially 30 minutes before events such as a keynote. This would allow for the inevitable “xxx is broken” just before some really visible presentation/event.

Some sort of sheet listing times of coverage is needed in the NOC and Help Desk, and folks could be assigned. The only two major network problems occurred after major SCinet efforts, on Tuesday morning due to a UPS failure for the DNS, and on Friday morning due to problems internal to the Qwest commodity network. Unfortunately, a late night before and the fact that the exhibit floor was closed meant there was no one at the NOC to handle these problems first thing.

Web-based trouble ticketing is on the cusp of being awesome, especially with wireless access available to SCinet and attendees. Several vendors are quite willing to directly interact with the Help Desk database to shepherd their problems or add new tickets entirely. The NOC team was also very receptive to interacting directly with the database.

There are several areas that still need to be addressed such as quickly assigning problems to responsible NOC team members and distributing the responsibility for making address/path assignments to more than one person.

Monitoring

Supporting vendors view it as a success in many respects: at minimum, meeting and working with some of the best networking minds in the world, and providing tools that can facilitate development on U.S. Government sponsored research projects in the SCinet participants' real life!

The Adtech participation was a real success for both Spirent and SC2000, and I am personally very pleased that Spirent was able to help demonstrate performance measurements for the HPC Challenge that previously were not possible. I am not as pleased that the SmartBits GPS demonstration didn't go off as well (due to problems from SCinet and SmartBits both), but confident that this year's lessons will go a long way towards an even more robust measurement and monitoring role in the demonstration network and other areas in the future.

The first of them from my perspective would be:

- The monitoring team should be designated quickly, and get together much earlier and act as a subcommittee.
- Measurement and monitoring methodologies should be agreed to much earlier, along with sufficient thought so as to minimize last-minute major changes to resource-intensive requirements such as the Adtech demonstration.
- More attention to the deliverables such as the displays
 - a. Maybe find a flat panel display manufacturer to contribute 3 x 3 ft screens to create a 12 x 12 ft presentation.
 - b. Maybe the wall display is set up at the front of the show floor just before you enter the exhibits area.
 - c. Maybe this is tied into the show guide somehow..(this might be kind of radical and maybe I'm totally off base, but since the whole show is about high performance networks, this would be the most visible way for the high performance agency representatives to get emotionally involved immediately when they walk in!)

- The Spirent team will be getting organized and ready much earlier and be prepared to co-lead the effort in a more coordinated fashion.
- If the Weather Map is used again, have the display superimposed over a representation of SCinet to make it more meaningful to observers on the floor.

Bandwidth Challenge After Action Report

If the Bandwidth Challenge is to remain a centerpiece of future SC conferences, the experiences gained from SC2000 can provide a number of “lessons learned.”

- Contest categories – Establish different categories ... suggest one category for high bandwidth challenges (throughput and peak) and another category for applications that compete for more efficient bandwidth and innovative use of bandwidth.
- Planning – Integrate Bandwidth Challenge planning with other aspects of SCinet planning early in the planning cycle to include A/V displays and monitoring, bandwidth scheduling and time-sharing, contestant performance baseline measures. The Bandwidth Challenge committee should be augmented by a technical subcommittee.
- Contestant performance baseline – Insure that performance baselines have been established for SCinet interconnects to high bandwidth contestants. Throughput measures from the contestant’s booth through the SCinet switches and routers should be verified and documented.
- Scheduling access to time-shared bandwidth – At SC2000, this almost became a full-time job. Suggest SCinet be responsible for scheduling only during exhibit hours and whatever time is needed to conduct the Bandwidth Challenge. Contestants themselves should be permitted to schedule up to 90-minute time blocks at other times via a web site.
- Monitoring and Displays – Display Bandwidth Challenge activities in real time using large format displays and monitors positioned at strategic locations throughout the show floor. Identify which application(s) is/are running. Show an overall bandwidth weather map for all WAN interfaces and selected high-speed interfaces on the show floor.

Strategic Issues for the Steering Committee and Sponsoring Organizations

SCinet 2000 and past SCinet efforts demonstrate conclusively that a very high performance, complex networks can be created and run effectively at the conference and there are valuable applications that use very substantial amounts of the bandwidth provided. Indeed, some experiments run at the show could not have been done otherwise. This level of networking and usage is essential to achieving the goal of SC to make high performance networking an equal partner with high performance computing at the conference. SC2000 achieved other improvements as well – for example, it doubled the number of IAC members from networking companies and it set the foundation for what could be a very productive long-term relationship with Qwest. It also engaged new partners who made substantial contributions, such as Internet2, SBC DataCom and ATDnet. There was also an increase in participation among industry exhibitors from networking companies. In this sense SC2000 was a rousing success from the networking perspective.

As with any milestone, now is the perfect time for the next step. Without being critical of the other aspects of the conference, multiple people observed that, while the technical program was excellent and made a great contribution to the success of SC2000, it was much more focused toward computation than networking. There was only one paper session out of 23 devoted to networking (or 3 out of 63 papers); no networking tutorials; not one of the four State-of-the-Field was on networking topics; only one half of one of the nine Masterworks presentations had a networking theme, and one of the nine panel sessions was a networking topic. Admittedly, this count is somewhat harsh, since some other topics, like MPI and grids, have network components. But these mostly deal with networking as an underlying infrastructure rather than as an explicit topic.

If SC conferences are to sustain the momentum created by SCinet, Escape, webcasting, the Bandwidth Challenge and the new partners, there must be more balance throughout all aspects of the conference. This is particularly important for the networking vendor community and the large research networks. There must be enough technical networking content to attract their clients and stakeholders as well as using their features to attract the traditional computational oriented attendees.

The steering committee should set an explicit goal of having at least 40% of all the conference activities devoted to network themes – mostly to networking over the wide area. The establishment of an award program at the level of the

Gordon Bell prizes is a step in the right direction. Whether this goal should be set for SC2001 or whether the goal should be progressively left to the steering committee and future conference committees is open to discussion, but if the conference is not close to balance by SC2004, it would be a missed opportunity for SC.

Another strategic issue is the importance that networking capability should play in site selection and contract negotiations for future conference locations. There are three aspects to site selection decisions that should be adjusted. The first is that a very important criterion that should be added is the locations being considered have excellent networking infrastructure – more so in the local area as within the actual conference site. This includes having usable dark fiber between the conference site and a major “fiber hotel” where many network providers have a point of presence. It is well known that the “last mile” is at least as difficult as all the other miles in any network implementation. SCinet has to deal with the last mile problem every year and spends tremendous effort solving it. Future site selection must take into account the existence of the last mile – or its cost to implement it if it does not exist.

The second aspect is that certain cities have a lot of network cross connects in nearby locations for the major carriers. It is much easier for a national carrier to run a patch between floors of a fiber hotel than it is to install an entirely new circuit over a long distance. These circuits are often priced by bandwidth, and some national networks cannot afford the highest bandwidth runs for just a short time. Thus, site selection should be kept to only the locations that are major hubs for the major carriers and at locations that have cooperative local companies.

The third aspect is most conference facilities now recognize providing Internet services is a money making activity for them, and have some level of networking services they sell to their standard conferences. This results in conflict that should be resolved up front with the conference site rather than well after the fact, as was the case with SC2000. From the conference center point of view, SC doing their own networking is akin to SC coming and saying the conference needs special food so it will cook our own. Clearly a conference site expects a lot of revenue from catering, and this would be a problem. By the middle of this decade, conference sites may be getting about that much revenue from networking (\$750 per day per network IP address?). SC must continue to control and provide its own networking services, since clearly the growth in demand and capability far exceed anything a normal conference site can provide (despite their perceptions).

The resulting recommendations for the steering committee are:

1. The selection committee should include a networking expert that spans multiple years. This person could be a volunteer or a paid consultant. The responsibilities would be to create the criteria used in evaluating sites relative to their ability to provide the networking needs of the conference in the out years and to participate in that evaluation and subsequent contract discussions.
2. Be skeptical that conference sites that have a good networking infrastructure in the conference site for standard conferences are good for SC. It may be the case that this is more of a hindrance than a benefit, particularly if the site is not provisioned for quick and cheap expansion.
3. Consider repeating sites more often because there is good networking infrastructure in place. A 10-year return cycle is of little benefit since the technology installed will clearly be out of date. By returning to a site often, SC would know the infrastructure and be able to accomplish incremental improvement. Also there is more likelihood of consistency in the local support staff for networking .

The final point is that the base contract with the conference site must provide the ability for SC to do its own networking (run wires and fibers, bring in connections, control the area of the show, etc.). This should be negotiated early and up front. Just the same as there needs to be long term conference management expertise involved with negotiating these contracts, there should be long term network expertise involved so that the conference. Without these strategic steps in site selection and arrangements, it is likely the SC conference series will face an ever increasing cost of doing networking on the scale needed to make the conference succeed. A proposed agreement/specification for this is below, along with the documents of the agreement eventually negotiated at DCC..

Appendix E
Suggested IEEE/SCinet agreements and requirements for future conferences
(in a generic format)

AGREEMENT ON NETWORKING FOR CS 'XY CONFERENCE
BETWEEN
THE IEEE/ACM AND XYZ CONVENTION CENTER (CC)

Background information on SCinet

Since 1993, the SC conference series has set up "SCinet", a sophisticated high speed gigabyte network infrastructure that links the high performance computers of exhibitors and research exhibitors. SCinet's goals are to provide experimental opportunities and demonstrations for the latest high performance networking technology and to support and facilitate applications that make use of high performance networks. When completed, it is the fastest network in the world.

SC sets up its own fiber optic network on the exhibit floor which links exhibits to each other, to wide area networks and to the Internet. Preference shall be given to convention centers with existing fiber optic infrastructure, which shall be used to distribute this show floor network to other locations in the convention center. (It should be noted that SC provides its own optics and electronics, it only require access to existing passive infrastructure of fiber and/or copper wiring)

The show floor network exists in large part to showcase agency networks (e.g., DARPA, NASA, DOE, and DOD) and the research they support. Thus, in any particular year, the network bandwidth required into the wide area tends to be slightly greater than the bandwidth of the current fastest agency network; almost, but not quite the sum of all the agency networks being brought to the show floor. For example, in 2000, DAPRA Supernet, HSCC and Internet-2 were each using OC48 backbones, Esnet and vBNS were at OC12 and several others used lower speeds. The off floor aggregate network bandwidth was approximately an OC 192. The external bandwidth brought into the center by these organizations totaled 8.4 Gbps. The current year, 2001, if DARPA's Supernet is provisioned at OC192, ESnet and NASA were each provisioned at OC48, and there were several other OC12s. The off-show floor bandwidth would be greater than OC192 plus 2xOC48.

Note that this bandwidth is typically donated to SC for the duration of the show by IXC (Inter-eXchange Carriers) such as Qwest and WORLDCOM. We believe it would be illegal (section 271) for it to be donated to an intermediate for-profit company acting as a CLEC (Competitive Local Exchange Carrier) such as a local ISP. Furthermore, it would not be possible to have this donated to any for-profit organization providing network services to the convention center.

SCinet Management

The SC2000 SCinet team is a volunteer group led each year by a volunteer Chair. It is composed of volunteers are drawn from US national laboratories and networking companies who are selected on the basis of their technical expertise. They are simply the best in the world in terms of their networking experience and knowledge.

The SCinet team spends more than two years designing and planning the installation of the network. The design of the network must be completed nine to ten months prior to the SC meeting to allow enough time for implementation of the infrastructure and coordination with exhibitors. The team holds bi-weekly teleconference calls to discuss the latest developments. The team is divided into subcommittees responsible for specific parts of the network operation: such as experimental, production, physical infrastructure, wireless, network management/monitoring, information management/customer support, and committee networking.

The SCinet on-site operation is an intensive four-week to build, deploy and operate what is often the world's fastest network. SCinet is installed on the Sunday or Monday prior to the actual exhibitor move-in. The network takes a minimum of a week to install, and starts operation with the opening of the exhibit hall move in of exhibitors, typically the following Thursday morning. SCinet provides reliable, high performance network services starting Saturday morning for the education program and is fully operational – as is all the exhibits – the following Monday evening for the Gala Opening.

Infrastructure Requirements

For SC2000, the SCinet team installed more than 82 miles of fiber and over 150 fiber drops, and supported thousands net devices, including many wireless/mobile devices. It supported multiple applications using more than 1 Gbps of measures bandwidth sustained over the WAN. In the local area the highest performing application was over 8 Gbps on one connection. SCinet connected 100Base FX; 1000BaseSX; ATM-OC3c; ATM OC12c ATM OC-48, Dark single and multimode fiber connections; OC12c; Myrinet; 2 Experimental networks; SONET OC 48 and many IEEE 802.11b wireless access points. There are multiple Wide Area connections into different routers, some using advanced WDM technology.

SCinet usually operates by finding the nearest high speed access point, often a major network interconnection facility ("fiber hotel"), a National Science Foundation site, a National Laboratory or a university center that offers OC 192. SC requires direct access to DREN (the Department of Defense network), Supernet (another DOD network), ESnet (the Energy Sciences Network), Internet2 (the National Testbed Network), and vBNS (the National Science Foundation network), which are high speed requirements an Internet Service Provider (ISP) cannot provide.. In essence, SCinet becomes its own ISP. OC 192 or above is outside the usual spectrum of Internet services and is a highly specialized, cutting edge technology.

To operate SCinet, SC purchases fiber optic cable and connectivity equipment and provides some support for its volunteer team, with net cost to the conference of approximately \$40,000. However, the majority of expense of renting and purchasing equipment is defrayed by contributions and loans of equipment from major computer companies such as CISCO and Sun Microsystems and from the national laboratories. In the past, SCinet has received the equivalent of \$25 million in donated computer and networking equipment. Contributors to SCinet also send their top engineers to help with the installation. SCinet is a highly visible operation for these companies and they want to ensure its success.

More information can be found in the attached report on SC 2000, documenting the entire network creation and usage.

Convention Center Requirements

To successful operate SCinet at SC 'XY, we have the following requirements – to be provided at no cost to SC or the IEEE/ACM:

- 1) SCinet shall have complete and unlimited access to all aspects of installing, operating and maintaining network infrastructure in the exhibit halls, the meeting rooms, and the common areas assigned to the conference. CC, its agents, contractors or others will place no restrictions on the ability of SCinet to install cabling, fiber, wireless and other networking technology, whether it is overhead or under the floor, using existing conduits and pathways - or creating temporary paths that do not damage CC. This also includes use of any fiber infrastructure running to meeting rooms.
- 2) SCinet shall have complete and unlimited access to all aspects of installing, operating and maintaining a route of fiber from the exhibit hall to the location where telecommunications enters CC. CC will make at least one straightforward pathway accessible to SCinet for this, well before the conference move begins. SCinet shall be able to independently arrange external connectivity for SC2000, exclusive of any existing or planned CC networking.
- 3) SCinet shall have access to bundle of at least 24-pairs of single-mode dark fibers will be installed between the CC and a multi provider the local telecommunications hotel, preferably where QWest/MCI and GST are also present. Location shall be determined by SC. SC shall be able to use any unallocated dark fiber already running to the CC, assuming SC makes independent agreements with the owners of that fiber.
One end of the bundle will be terminated on a patch panel in the telecommunications room of the CC. The fibers (non-zero dispersion shifted, unspliced), their polish (UPC or better, and loss (0.25 db/km or less) will be suitable for DWDM, optical CDMA, or other leading-edge technologies. The other end of this bundle will be terminated at the telecom hotel in a patch panel that can be easily patched to other CLEC's, ILEC's, and IXC's. There will not be a charge for the use of this fiber. A charge to patch to a fiber to any for any CLEC, ICLE, or IXC in the telecom hotel will not to exceed \$500 per patch.
- 4) SCinet shall have access and use of any CC wireless networking infrastructure if such infrastructure exists and is compatible with SCinet needs. This includes conforming to appropriate standards for 11 or higher Mbps service. SCinet shall have the ability to implement any other wireless service independent of the CC service. If at any time

there is a conflict in the SC conference space between the CC wireless equipment and SCinet equipment that cannot be expeditiously resolved, the CC shall shut down the components causing the conflict for the duration of the SC show.

5) If the CC has networking infrastructure that is used by SC, it shall support/repair it if such repair is necessary for proper operation. If SC changes any configuration and/or equipment, SC will return it to the original state before departing

6) The CC staff, and/or appropriate contractors shall support SCinet planning by providing wiring and equipment information, the locations of all conduit, building diagrams (including infrastructure and conduit runs, and any other information requested by SC.

7) The CC shall make an equipment staging area to be made available to SCinet in the DCC starting in 1 October . This area, at least 1,000 square feet, will be used to stage and test equipment before move-in. The area should be secure, climate controlled and appropriate for computing equipment.

8) If the CC has available Uninterrupted Power available in the CC, SCinet would appreciate UPS services for 15 KVA of equipment at a cost below that of renting a UPS independently.

Time Schedule

Exhibit Space - 150,000 - 200,000 gross square footage

SCinet Staging Move-in: October 1 prior to the conference at 7:00 am

SCinet Move-in: Sunday, 8 days prior to the conference start at 7:00 a.m.

Exhibitor Move in: Thursday, 5 days prior to the conference start, at 7:00 am.

Event Dates: Monday through Thursday at 5:00 p.m.

Move-out: Thursday at 5:00pm to Saturday at 12:00 noon

Our Offer to the Center

In return for flexibility with the ISP, WAN, and fiber installation, SCinet offers our expertise to center staff:

Any CC staff member can participate in the bi-weekly planning meetings and work with our team on-site. The SCinet volunteers are experts in the field, willing to share their knowledge of the latest technology with CC staff. Center staff attended the SC conference prior to the schedule and come away with information that they could use in the future for other events. The SCinet volunteers have accumulated a great deal of expertise about laying fiber in convention center that we believe would be beneficial to your staff.

SC would also consider assisting in installing, configuring an deploying infrastructure and other technology (WAN connections, fiber and wireless, electronics, etc.) that is provided by the CC that will remain place in the CC on a case-by-case basis.

We view the successful operation of SCinet as opportunity to work together for technical and marketing advantages. SC can help build a high quality network that can be promoted by the CC to attract future business.

Agreed

Representative from IEEE/ACM

Representative from DCC

SCinet 'XY Chair

Appendix F
IEEE/SCinet agreements with DCC

Anne Marie Kelley, IEEE

28 January 2000

Mr. Ron Melton
Executive Vice President and Chief Financial Officer
Dallas Convention & Visitors Bureau
1201 Elm Street, Suite 2000
Dallas, Texas 75270

Dear Mr. Melton:

I write concerning issues of critical importance to the operation of SC2000, a technical and educational conference and exhibition. The SC2000 committee has been working with staff of the DCC on the implementation of our conference network. We ask your assistance to resolve questions concerning the operation and cost of the network.

Background information on SCinet

Since 1993, the SC conference series has set up "SCinet", a sophisticated high-speed gigabyte network infrastructure that links the high performance computers of exhibitors and research exhibitors. SCinet's goals are to provide experimental opportunities and demonstrations for the latest high performance networking technology and to support and facilitate applications that make use of high performance networks. When completed, it is the fastest network in the world.

SCinet Management

The SC2000 SCinet team is a volunteer group led by Bill Kramer of Lawrence Berkeley National Laboratory. Volunteers are drawn from US national laboratories and networking companies and are selected on the basis of their technical expertise. They are simply the best in the world in terms of their networking experience and knowledge.

The SCinet team spends more than a year designing and planning the installation of the network. The design of the network must be completed nine to ten months prior to the SC meeting to allow enough time for implementation of the infrastructure and coordination with exhibitors. The team holds bi-weekly teleconference calls to discuss the latest developments. The team is divided into subcommittees responsible for specific parts of the network operation: experimental, production, physical infrastructure, wireless, network management/monitoring, information management/customer support, and committee networking.

The SCinet on-site operation is an intensive two week commitment. SCinet is installed on the Sunday or Monday prior to the actual exhibitor move-in. The network takes a minimum of a week to install, and starts operation with the opening of the exhibit hall the

Infrastructure Requirements

For SC99, the SCinet team installed 40 miles of fiber and 136 fiber drops, and supported 2539 net devices. It transported over 55 Terabytes into the WAN. It connected 46 SCinet 100Base FX; 27 SCinet 1000BaseSX; 22 SCinet 10BaseFL; 11 SCinet ATM-OC3c; 10 Dark MM; 7 Dark SM; 4 SCinet ATM OC12c; 3 SCinet Myrinet; 2 Xnet 1000 Base SX; 2 SCinet ATM OC48c; 1 Xnet SONET OC 48 and 1 10BaseT. All were carried into the wide area on DWDM.

SCinet usually operates by finding the nearest high speed access point, usually a National Science Foundation site or a university center that offers OC 192. We need direct access to DREN (the Department of Defense network), ESNNet (the Energy Sciences Network), Internet2 (the National Testbed Network), and vBNS (the National Science Foundation network), which are high speed requirements an Internet Service Provider (ISP) cannot usually provide. SCinet links this point via local telephone service to the backdoor of the convention center. In essence, SCinet becomes its own ISP. We have avoided using an outside ISP because the risk is high – ISPs do not normally work with high speed connections and do not have the expertise to support SCinet’s requirements. OC 192 is outside the usual spectrum of internet services and is a highly specialized, cutting edge technology.

To operate SCinet, SC purchases fiber optic cable and connectivity equipment and provides some support for its volunteer team. However, the majority of expense of renting and purchasing equipment is defrayed by contributions from major computer companies such as CISCO and Sun Microsystems. In the past, SCinet has received about \$10 million in donated computer and networking equipment. Contributors to SCinet also send their top engineers to help with the installation. SCinet is a highly visible operation for these companies and they want to ensure its success.

Convention Center Issues

We have been told verbally that the DCC claims exclusive rights to run any fiber optic cable at a cost of \$2.00 per linear foot. SCinet has used volunteer labor in the past to lay fiber. Using the same requirements as SC99, this would be more than a \$400,000 expense. Given that SC2000 is using a larger exhibit hall than SC99, we estimate the cost of running fiber to be higher.

We have been told that we would be required to use the DCC’s ISP, and that the ISP will be able to support up to OC 192. We are concerned whether or not this service can actually be provided. There would be no reason for an ISP to establish the infrastructure needed for SCinet, including the connectivity and routers required. We also have not received any prices for ISP services.

There is no pricing yet established for the wireless services. The plan that has been shared with us is that the upgrade to 10 MSPS will be completed by the end of the first quarter, with the possibility of 100 MBPS wireless implemented by our arrival in November.

As an educational conference, SC has tried several experiments in the past using SCinet. In Dallas, we may want to partner with a telecommunications company to provide us with wireless connectivity between the center and a hotel. We need the flexibility to continue these educational activities.

Our Expectations for the DCC

We need to resolve these issues now because they affect our SCinet operation, both technically and financially. To successfully operate SCinet at SC2000; we have these expectations for the DCC:

1) SCinet shall have complete and unlimited access to all aspects of installing, operating and maintaining network infrastructure in the exhibit halls, the meeting rooms, and the common areas assigned to the conference. No restrictions will be placed by DCC or others on the ability of SCinet to install cabling, fiber, wireless and other networking technology, whether it is overhead or under the floor, using existing conduits and pathways - or creating temporary paths that do not damage DCC. This also includes use of any fiber infrastructure running to meeting rooms.

2) SCinet shall have complete and unlimited access to all aspects of installing, operating and maintaining a route of fiber from the exhibit hall to the location where telecommunications enters DCC. DCC will make at least one straightforward pathway accessible to SCinet for this, well before the conference move begins. SCinet shall be able to independently arrange external connectivity for SC2000, exclusive of any existing or planned DCC networking.

3) SCinet shall have access to bundle of at least 24-pairs of single-mode dark fibers will be installed between the DCC and a multi provider the Dallas Telecommunications hotel, preferably where QWest/MCI and GST are also present. One suggested location is 323 Bryan Street, Suite 2500 Dallas, TX 75201. (Currently access providers in that POP are MCI, TCG, and SWB. Upcoming carriers are ICG, Nextlink, and Time Warner).

One end of the bundle will be terminated on a patch panel in the telecommunications room of the DCC. The fibers (non-zero dispersion shifted, unspliced), their polish (UPC or better, and loss (0.25 db/km or less) will be suitable for DWDM, optical CDMA, or other leading-edge technologies. The other end of this bundle will be terminated at the telecom hotel in a patch panel that can be easily patched to other CLEC's, ILEC's, and IXC's. There will not be a charge for the use of this fiber. A charge to patch to a fiber to any for any CLEC, ICLE, or IXC in the telecom hotel will not to exceed \$500 per patch.

4) SCinet shall have access and use of DCC wireless networking infrastructure if the infrastructure is compatible with SCinet needs. This includes conforming to appropriate standards for 11 Mbps service. SCinet anticipates attendees will be bringing their own equipment as well as some using the DCC so the DCC infrastructure shall accommodate that. Cost per host card should be less than \$20 for the week of the conference. SCinet shall have the ability to implement any other wireless service independent of the DCC service.

5) DCC shall make an equipment staging area to be made available to SCinet in the DCC starting in 1 October 2000. This area, at least 500 square feet, will be used to stage and test equipment before move-in. The area should be secure, climate controlled and appropriate for computing equipment. Bill Kramer has already discussed this space request with Paula Tait.

6) If DCC has available Uninterrupted Power, SCinet would appreciate UPS services for 15 KVA of equipment at a cost below that of renting a UPS independently.

Our Offer to the Center

In return for flexibility with the ISP, WAN, and fiber installation, we offer our expertise to center staff:

Any DCC staff member can participate in the bi-weekly planning meetings and work with our team on-site. The SCinet volunteers are experts in the field, willing to share their knowledge of the latest technology with DCC staff. Center staff attended the SC99 conference in Portland and came away with information that they could use and a sample of a PCV bracket that holds fiber. The SCinet volunteers have accumulated a great deal of expertise about laying fiber in convention center that we believe would be beneficial to your staff.

We will assist with the review of proposals received for the ISP and wireless provider. Bill Kramer has already provided wording for both RFPs.

We will offer to negotiate leaving the infrastructure in place in the DCC.

We view the successful operation of SCinet as opportunity to work together for technical and marketing advantages. We can help build a high quality network that can be promoted by the DCC to attract future business.

We look forward to speaking with you about these issues.

Sincerely,

Anne Marie Kelly
Director, Volunteer Services

IEEE Computer Society

cc: Louis Turcotte, SC2000 Chair
Betsy Schermerhorn, SC2000 Vice Chair
Dennis Duke, Conference Vice Chair for Conference Showcase
Bill Kramer, SC2000 Conference Vice Chair for Information Architecture
William Wing, SC2000 Experimental Network Chair

Agreement Between IEEE and Dallas Convention and Visitor’s Bureau

The following provisions apply to the network related services, facilities and equipment that will be used by the SC2000 conference between Oct 16 and Nov 12, 2000. Specifically,

Usage:

DCVB Internet Services will allow SC 2000 use of all existing networking infrastructure owned by Internet Services within the Dallas Convention Center as needed for the conference, including fiber optic and copper cabling within the confines of the space allocated for IEEE. IEEE will be allowed to collocate equipment in the Internet Services MIS room, up to one full height 19” rack with prior approval and based upon availability at the time of approval. The list of collocated equipment must be pre-approved by technical and administrative staff of Internet Services prior to allocation of rack space. IEEE will be allowed to make connections as necessary to the ODS Infinite switch and any other network infrastructure hardware owned and operated by Internet Services. IEEE may utilize Internet Services 12 Mbps connection to Qwest Communications for the purposes of Internet connectivity on the following schedule:

Oct 16-19	2 Mbps	Staging
Oct 20-Oct 28	.1 Mbps (100 Kbps)	Emergency connections and monitoring while equipment is in DCC
Oct 29-Nov 3	2 Mbps	Setup
Nov 4 - Nov 10	12 Mbps	Conference
Nov 11	.1 Mbps (100 Kbps)	Final tear down

Relative to 1 month at a 12 Mbps level, this equates to 28% of the monthly aggregate usage.

Support:

Support will be provided for the existing infrastructure with the exception of any infrastructure (cabling, equipment, etc.) which has been modified by IEEE.

Wireless:

Internet Services will waive its exclusive right to wireless networking in the Dallas Convention Center during the IEEE conference for the areas of the DCC assigned to the IEEE. Further, IEEE may replace existing wireless equipment with their own for the duration of the show.

Limitations, Requirements and Exclusions

A list of all equipment owned or operated by Internet Services to be replaced for the duration of the show must be submitted for approval prior to replacement. It is understood that the all networks will be restored to their original operating condition, and that any modification made to the network will not impact Internet Services, its customers, or the Dallas Convention Center networks at any time. Approval is considered given if any DCC or Internet Services staff/consultants are aware of and verbally or in writing concur with plans and actions.

Cost

The all-inclusive cost for these services, facilities and equipment is \$20,000.

Agreed

Ron Melton. Dallas Convention and Visitors Bureau