GUIDE TO THE NERSC MACHINE ROOM
Lawrence Livermore National Laboratory, Livermore, CA
April 1990

The National Energy Research Supercomputer Center (NERSC) provides large-scale computational support to the DOE Energy Research community. NERSC applies advanced computer technology and computing techniques to the problems of magnetic fusion energy research, basic energy science research, high energy and nuclear physics research, health and environmental research, applied mathematical sciences, and Superconducting Super Collider (SSC) research.

Committed to the belief that pooling resources furthers progress, the NERSC encourages communication among American energy researchers and with European and Japanese energy researchers. The Center facilitates the sharing of information; codes, data, manpower, and computer power.

High-speed digital computers first worked on magnetic confinement fusion problems in the 1950's. Demand for computer time jumped twenty-fold in the late 60's due to the new possibility of simulating plasma behavior with computer codes. (Plasma is the gaseous fuel of the experimental reactors.)

The coast-to-coast NERSC network was designed to be a useful tool, dependable and responsive to the many needs of individual energy researchers. Research priorities, anticipated demand, and various orders of need were taken into account before the U.S. Energy Research and Development Administration established the facility for unclassified research in 1974.

Currently NERSC provides supercomputing services to about 4500 scientists and researchers at nearly 150 institutions throughout the country, including 26 government research laboratories, 92 universities, 11 private laboratories and 19 industrial sites. Its services are also available at several research centers in Europe and Japan, where foreign scientists collaborate with U.S. researchers. By means of two nationwide/international communications networks (MFENET, ESNET), our users may access the following supercomputers: a CRAY-1S, a CRAY XMP, and three CRAY-2's at NERSC as well as all of the VAX machines in the MFENET.

The computing power is concentrated in the NERSC "machine room." In this Guide, we discuss the computing facilities in four categories: Large Computers, Data Storage System, Communications System, and our local User Service Center. In addition we discuss the important support services of: machine-room operations, archival records, safety, and prime power, cooling, and heating. Maps of the machine room and MFENET are attached at the end.

LARGE COMPUTERS

1. CRAY-2 COMPUTER SYSTEMS (3)

- Serial 1 installed June 1985, serial 18 installed August 1988, serial 101 installed April 1990. The latter is the 30th and last CRAY-2 manufactured.
- Performs scientific calculations for energy researchers throughout the United States.
- Manufactured by Cray Research, Inc.
- First one manufactured was installed here in June 1985.
- Each central processing unit (CPU) performs up to a theoretical peak speed of 500,000,000 arithmetic operations per second. Serial 1 and serial 18 each have four CPU's; serial 101 has eight CPU's and it is the only CRAY-2 manufactured with eight CPU's.
- Each system has a 4.1 nanosecond cycle time. There are 1,000,000,000 nanoseconds in one second and it takes one nanosecond for electricity to travel across one foot of wire.
- Serial 1 has a memory capacity of 512,000,000 characters. Disk storage capacity is 27,200,000,000 characters at a cost of $115 per 1,000,000 characters.
- Serial 18 has a memory capacity of 1,024,000,000 characters. Disk storage capacity is 44,800,000,000 characters at a cost of $119 per 1,000,000 characters.
- Serial 101 has a memory capacity of 1,024,000,000 characters. Disk storage capacity is 24,800,000,000 characters at a cost of $119 per 1,000,000 characters.
- Each system is cooled by 250 gallons of circulating flourinert; 150 gallons in the computer and 100 gallons in the reservoir and pipes. Cost of the fluid is $200 per gallon.
- Each system uses 180 KW of electricity.
- Cost: $15,000,000 for serial 1; $17,500,000 for serial 18; $19,000,000 for serial 101.

2. CRAY 1S Computer System

- Serial 33 installed in October 1981.
- Manufactured by Cray Research, Inc.
- Fifty-five were manufactured and installed world wide.

- Performs up to 160,000,000 arithmetic operations per second.
- Has a 12 nanosecond cycle time.
- Has a memory capacity of 16,000,000 characters.
- Has a disk storage capacity of 7,800,000,000 characters at a cost of $212 per 1,000,000 characters.
- Cost: $10,000,000.

3. Cray X-MP/22 Computer System, Installed November 1984

- Performs scientific calculation for MFE and energy researchers throughout the U.S.
- Manufactured by Cray Research, Inc.
- Each processor (there are two) performs up to 210,000,000 arithmetic operations per second.
- Has a 9.5 nanosecond cycle time.
- Has a memory capacity of 16,000,000 characters. Disk storage capacity is 13,800,000,000 characters at a cost of $100 per 1,000,000 characters.
- Cost: $10,500,000.

DATA STORAGE SYSTEM

4. Common File System - CFS

- IBM 4381, Model P-21
  One processor
  Memory Capacity 16,000,000 characters
- IBM 4381, Model P-14
  Two Processors
  Memory Capacity 16,000,000 characters
- Disk capacity 180,000,000,000 characters,
  utilizing IBM and Storage Technology, Corp. Disk
- Cost of IBM mainframes $1,000,000.
- Cost of disk storage $2,000,000 or $11 per 1,000,000 characters.

5. Storage Technology Corp. 4400 Automated Cartridge Systems (ACS), (3)

   . Each silo has a data storage capacity of 1,200,000,000,000 characters stored on 6,000 magnetic tape cartridges. Each cartridge holds 200,000,000 characters.

   . Access time is 15 seconds or less to mount a cartridge, 30 seconds (average) to find the desired data.

   . Each silo has four cartridge read/write drives attached.

   . Cost is $.31 per 1,000,000 characters.

6. Automated Tape Library (ATL) (CalComp 7110) - 2 Systems

   . Provides fairly cheap storage of data that are used less often than once a week.

   . Manufactured by Braegen, Inc.

   . The four tape units attached to the side of the ATL system 1 are ordinary tape units. The heart of the ATL is a robot that runs on a track down the middle of the "tunnel," fetches the desired tape, and then reaches through a window in the side of the ATL to mount the tape on the tape reader. There are four of these tape units attached to the ATL system 2.

   . Access time is about 2 minutes, including the time needed to find the desired data on the 2400-foot tape reel.

   . Total capacity of the two systems is 811,560,000,000 characters; stored on 4263 reels in system 2 and 2500 reels in system 2; each reel holds 120,000,000 characters.

   . Cost: $.76 per 1,000,000 characters.

   . These systems will be retired by June 1990 when data in them has been moved to the Automated Cartridge Systems.

7. Shared Disk System ("Disk Farm")

   . Provides high-speed storage for frequently used data. Also used to pass data between the five CRAY computers connected to it.

   . Access time is about 1/20th of a second.

   . Total storage capacity is 19,200,000,000 characters.

   . Cost: $905,000 or $47 per 1,000,000 characters.

USER SERVICE CENTER

8. VAX 8650

. Manufactured by Digital Equipment Corporation.

. Allows development and testing of computer programs prior to distribution to similar sites in the MFENET.

. Cost $400,000.

. Has a memory capacity of 32,000,000 character and disk storage capacity of 4,200,000,000 characters.

COMMUNICATIONS SYSTEMS

9. Data Communications Network - MFENET

. Connects

-- Five large computers in the machine room at MFECC; a CRAY-1, CRAY-XMP, three CRAY2's, and a DEC-VAX.

-- Thirty-three DEC-VAX computers and ten VAX-Clusters located at twenty-nine sites.

-- Over one hundred DEC-PDP 11 computers located at thirty geographical sites, over sixty of which control data flow; the remainder provide user services such as terminal access, printing and graphics capabilities.

. Via

-- Three, dual, 56,000-bit-per-second satellite links.

-- One, dual, 112,000-bit-per-second satellite links.

-- Two 56,000-bit-per-second DDS telephone lines.

-- Forty-five 4,800/9,600-bit-per-second telephone lines, and

-- Seventy-five short-distance cable connections.

(Note: 8 bits = 1 character.)

10. High Speed Data Communications Network - ESNET

    . Brought on-line during Jan., Feb., Mar., of 1990.

    . Replaces older network with 27 T1 circuits, over 25 times faster than 56,000 bit-per-second links.

    . Capable of carrying multiple communications protocols simultaneously.

    . Connects over 25 major Energy Research funded institutions.

    . Interconnects with 'back bone' networks of NASA, NSF, and DCA.

    . Interconnects with over 10 regional networks.

    . Has international connections to Europe and Japan.

11. Network Control Center

    . Used to monitor the operation of ESnet.

    . Problems may be diagnosed and corrected remotely from the Control Center.

    . Data about the performance of ESnet is collected every 16 minutes and saved.

12. Computer Standard Interface Channels (CSIC)

    . Provide high-speed communications links between computers at the NERSC.

    . Equipment designed and built by NERSC personnel.

    . Data may be sent and received simultaneously on up to eight channel.

    . Data transfer rate may be up to 2,500,000 characters per second.

13. Satellite Communications

    . Data is sent to a "stationary" (synchronous orbit) communications.

    . satellite 22,000 miles over the equator and then re-transmitted to the receiver.

    . Dual communications links using separate satellites are used to ensure reliability. Dual earth stations include 7-meter (23-foot) diameter dish antennas.

. Data transfer rate is 7,000 characters per second.

14. Local Terminal Concentrator

. Collects characters, as typed into keyboard terminal, into text lines and forwards them to any of the computers on MFENET. Process is reversed for sending output to terminal.

. Simultaneously connects up to 256 terminals to the CRAY-1, the CRAY XMP and three CRAY-2's as well as the VAX's on MFENET.

. Terminals may be connected to system by:

-- Direct wire.

-- Commercial telephone dial-up lines.

-- Via TYMNET, a commercial data-communications network.

. Terminal connection speed may be set by the user from 10 to 960 characters per second.

15. Hyperchannel Network

. Provides high speed communication links between computers in the NERSC Machine room at 50,000,000 bits-per-second.

. Equipment is built by Network Systems Corporation.

ARCHIVAL RECORDS

16. Microfiche Recorder

. Records on film the graphic or printed results of computations for future reference.

. Manufactured by Dicomed Corporation.

. Each 4 x 6 inch microfiche film sheet contains up to 250 image of a photograph reduced page of computer output. Page are recorded at the rate of 95 pages per minute.

. The microfiche recorder can also produce images on 35 millimeter slides or film strips, in either black and white or color.

MACHINE-ROOM OPERATIONS

17. Operator Console Area

- Provides communications and monitoring facilities for computer operators to keep an eye on the four CRAY's, two ATL's, the shared disk system, the three ASC's, and the two communications systems.

- A significant part of the operator's job is diagnosing problems and deciding who to call for repair and management of archival data.

18. Fire Protection

- Smoke detectors are located in the ceiling and under the floor. Overhead sprinklers and under-floor fire-suppressant gas can be discharged either automatically or by an operator. Steel overhead fire doors immediately drop behind the lobby windows when smoke is detected, and the Lab's Fire Department is automatically notified.

19. Monitoring Panel

- The panel in the hall leading to the machine room displays information about the environmental and safety conditions in the machine room and sounds an alarm if anything is amiss, such high humidity, cooling system failure, or smoke.

PRIME POWER, COOLING,

- Power consumption -- The computers and other equipment in the machine room consume large quantities of electrical power; In October 1988 use was 1,318,296 kilowatt hours.

- Cooling -- Sophisticated computers can do very fast computations if their vital parts are very close together: this can be done only if there is a very high rate of cooling. In the CRAY-1 and the CRAY XMP cooling is accomplished by refrigerant cooling coils mounted directly in the structural framework of the computer assemblies. The CRAY 2's are fluorinert cooled. Each of the five CRAY's requires 60 tons (air equivalent of direct cooling. In addition the machine room and building are cooled by a 150-ton air-conditioning system, which has an extra 150 tons capacity in standby.

- The comfort heating for building 451 is provided by the heat generated in the computer room (primarily from the disks.)

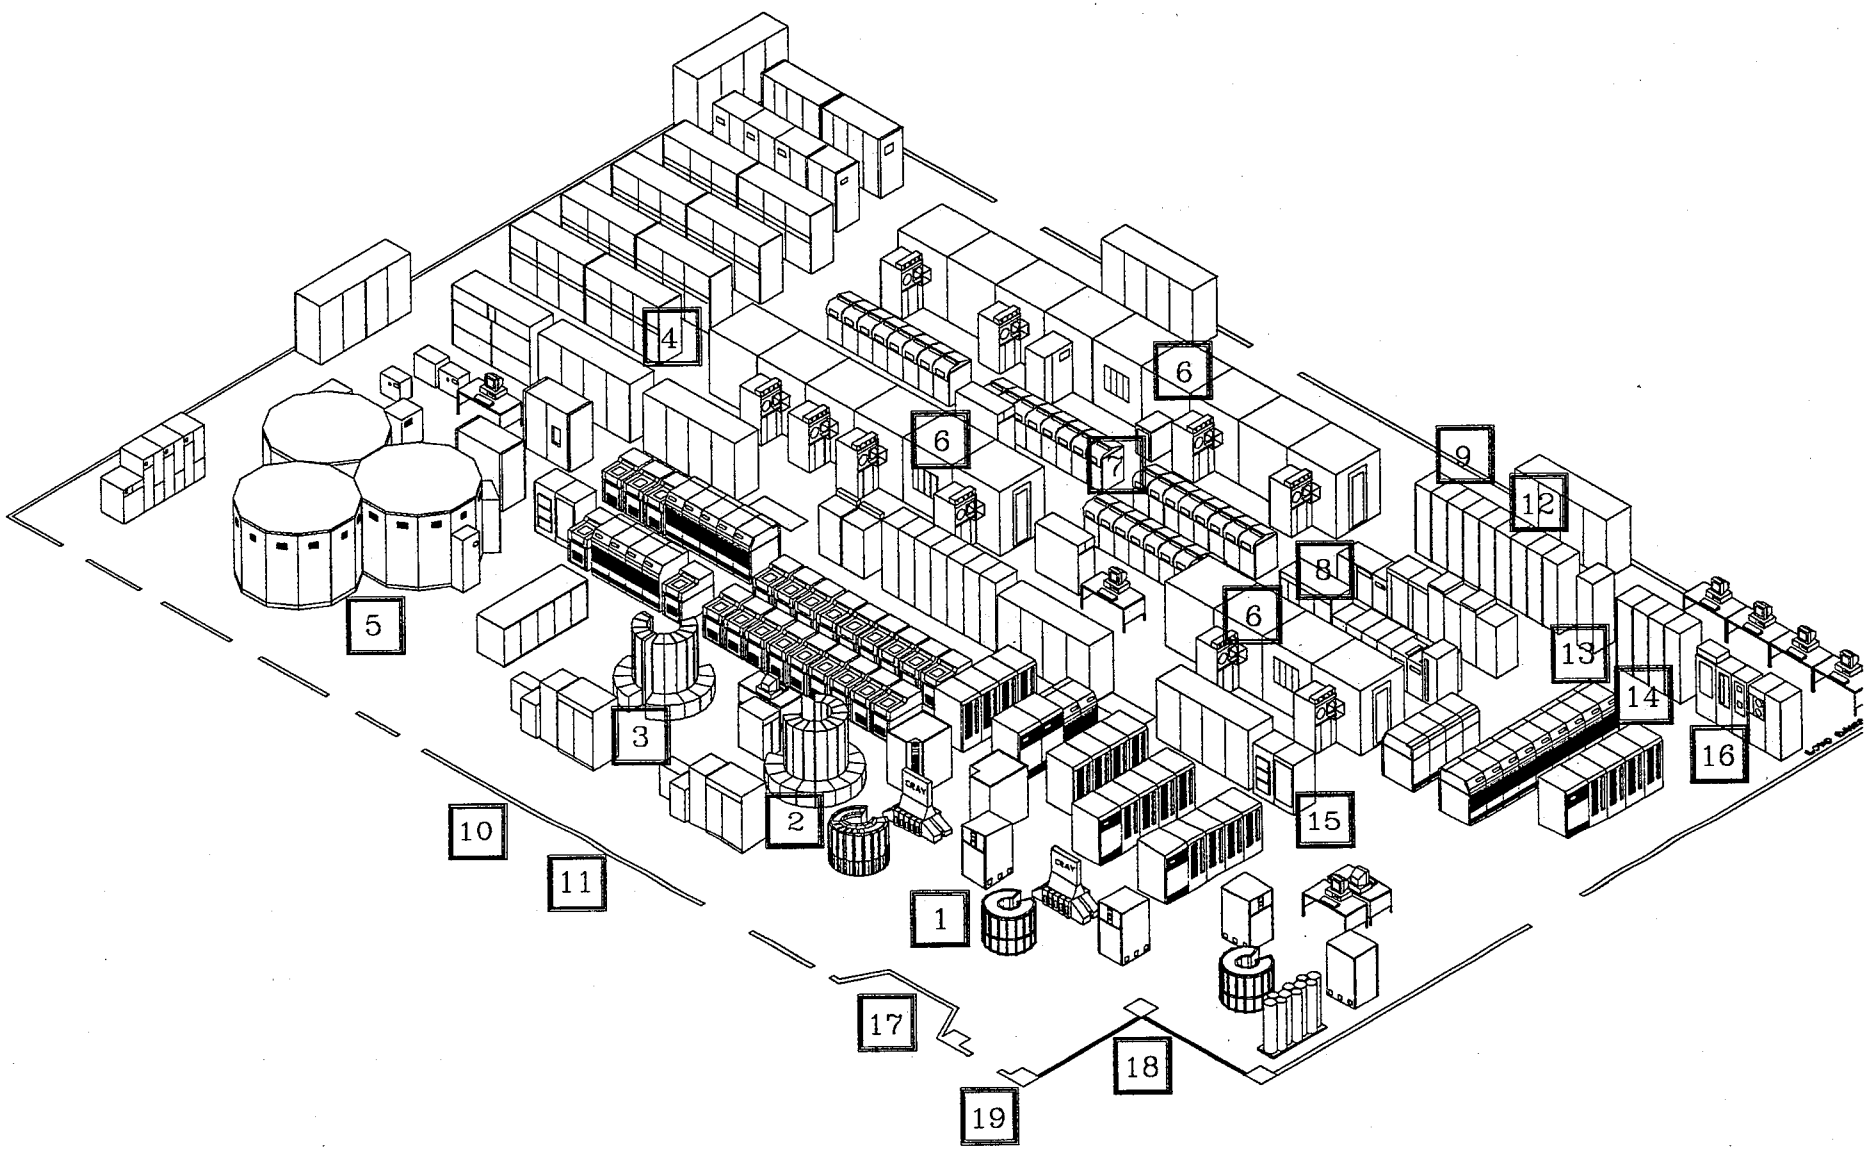
1
2
3
4
5
6
6
6
7
8
9
10
11
12
13
14
15
16
17
18
19

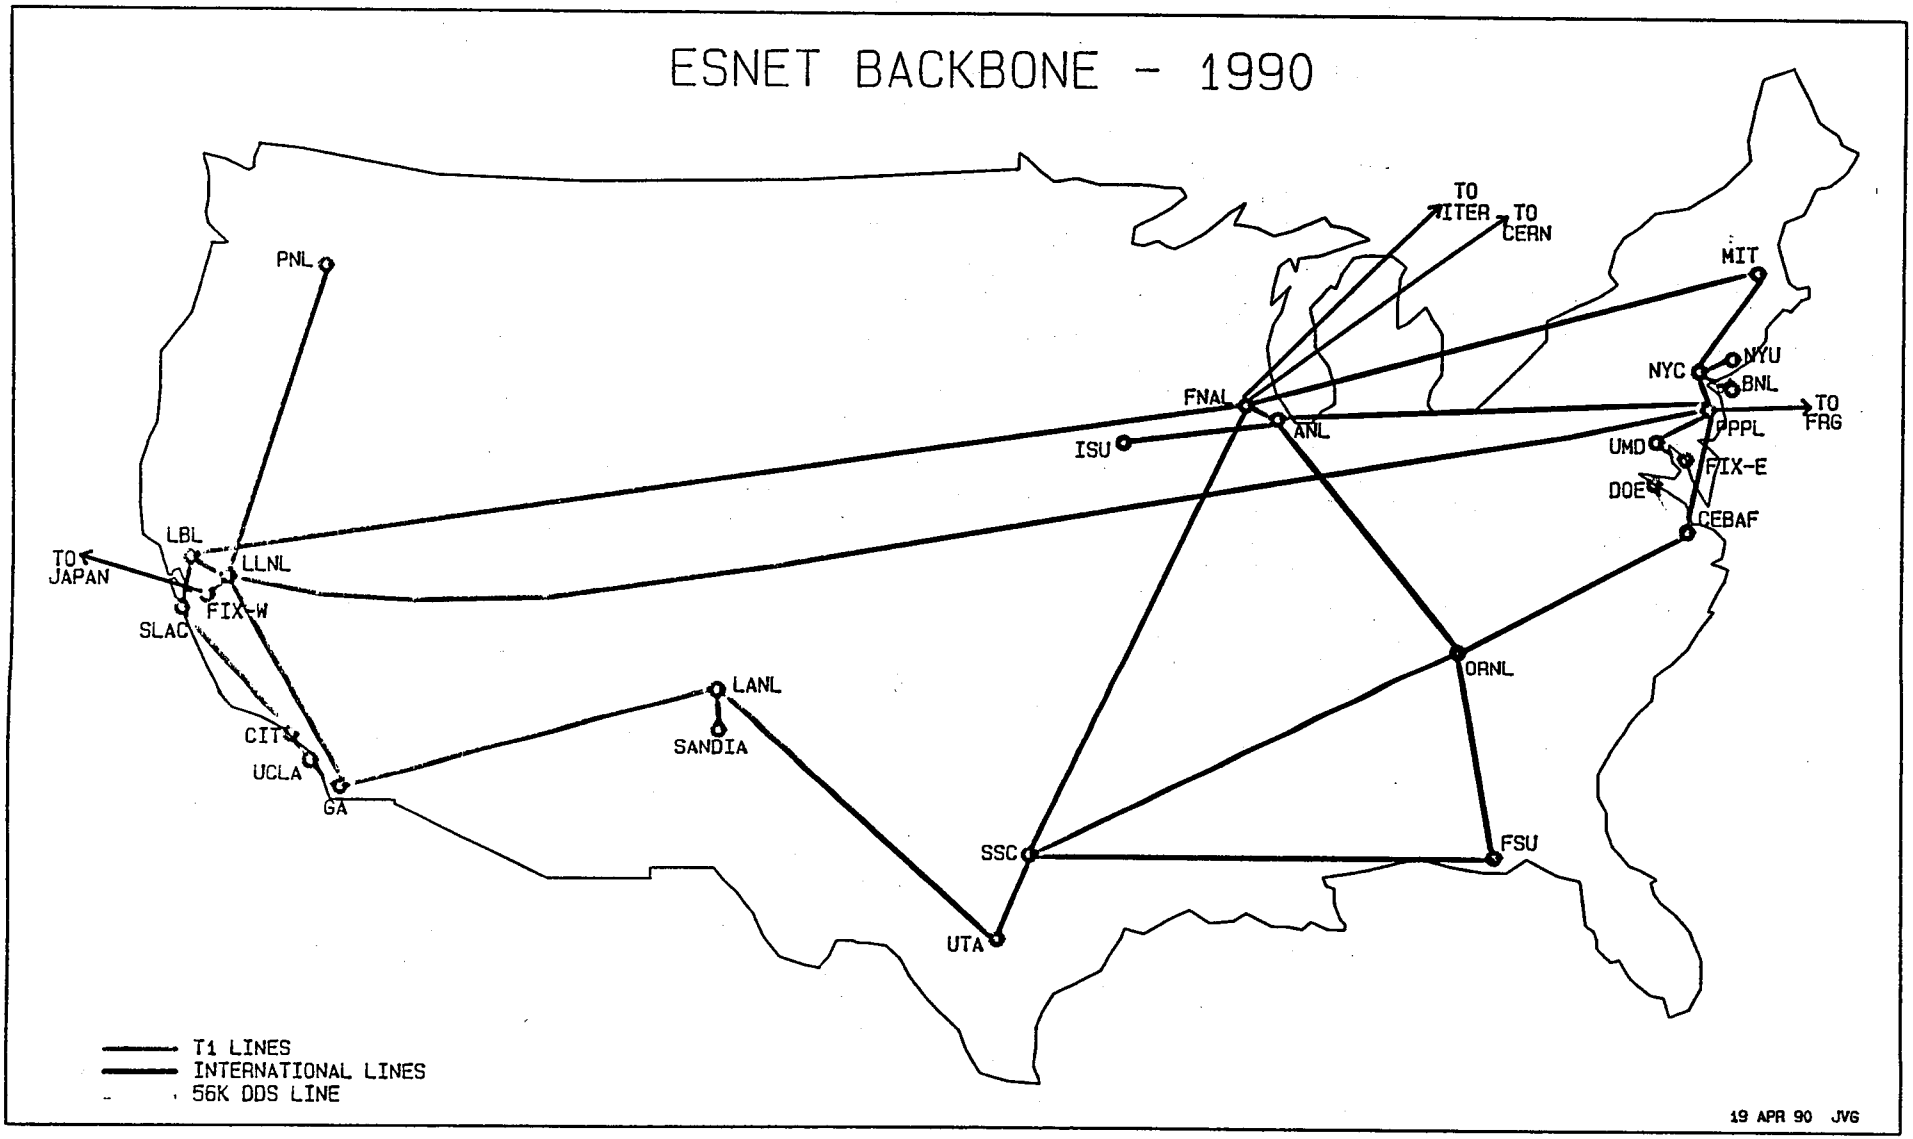
ESNET BACKBONE - 1990
TO ITER
TO CERN
PNL
MIT
NYC
NYU
BNL
FNAL
ANL
TO FRG
PPPL
ISU
UMD
PIX-E
DOE
CEBAF
LBL
TO JAPAN
LLNL
FIX-W
SLAC
ORNL
LANL
SANDIA
CIT
UCLA
GA
SSC
FSU
UTA
T1 LINES
INTERNATIONAL LINES
56K DDS LINE
19 APR 90 JVG

National MFE Network 1990

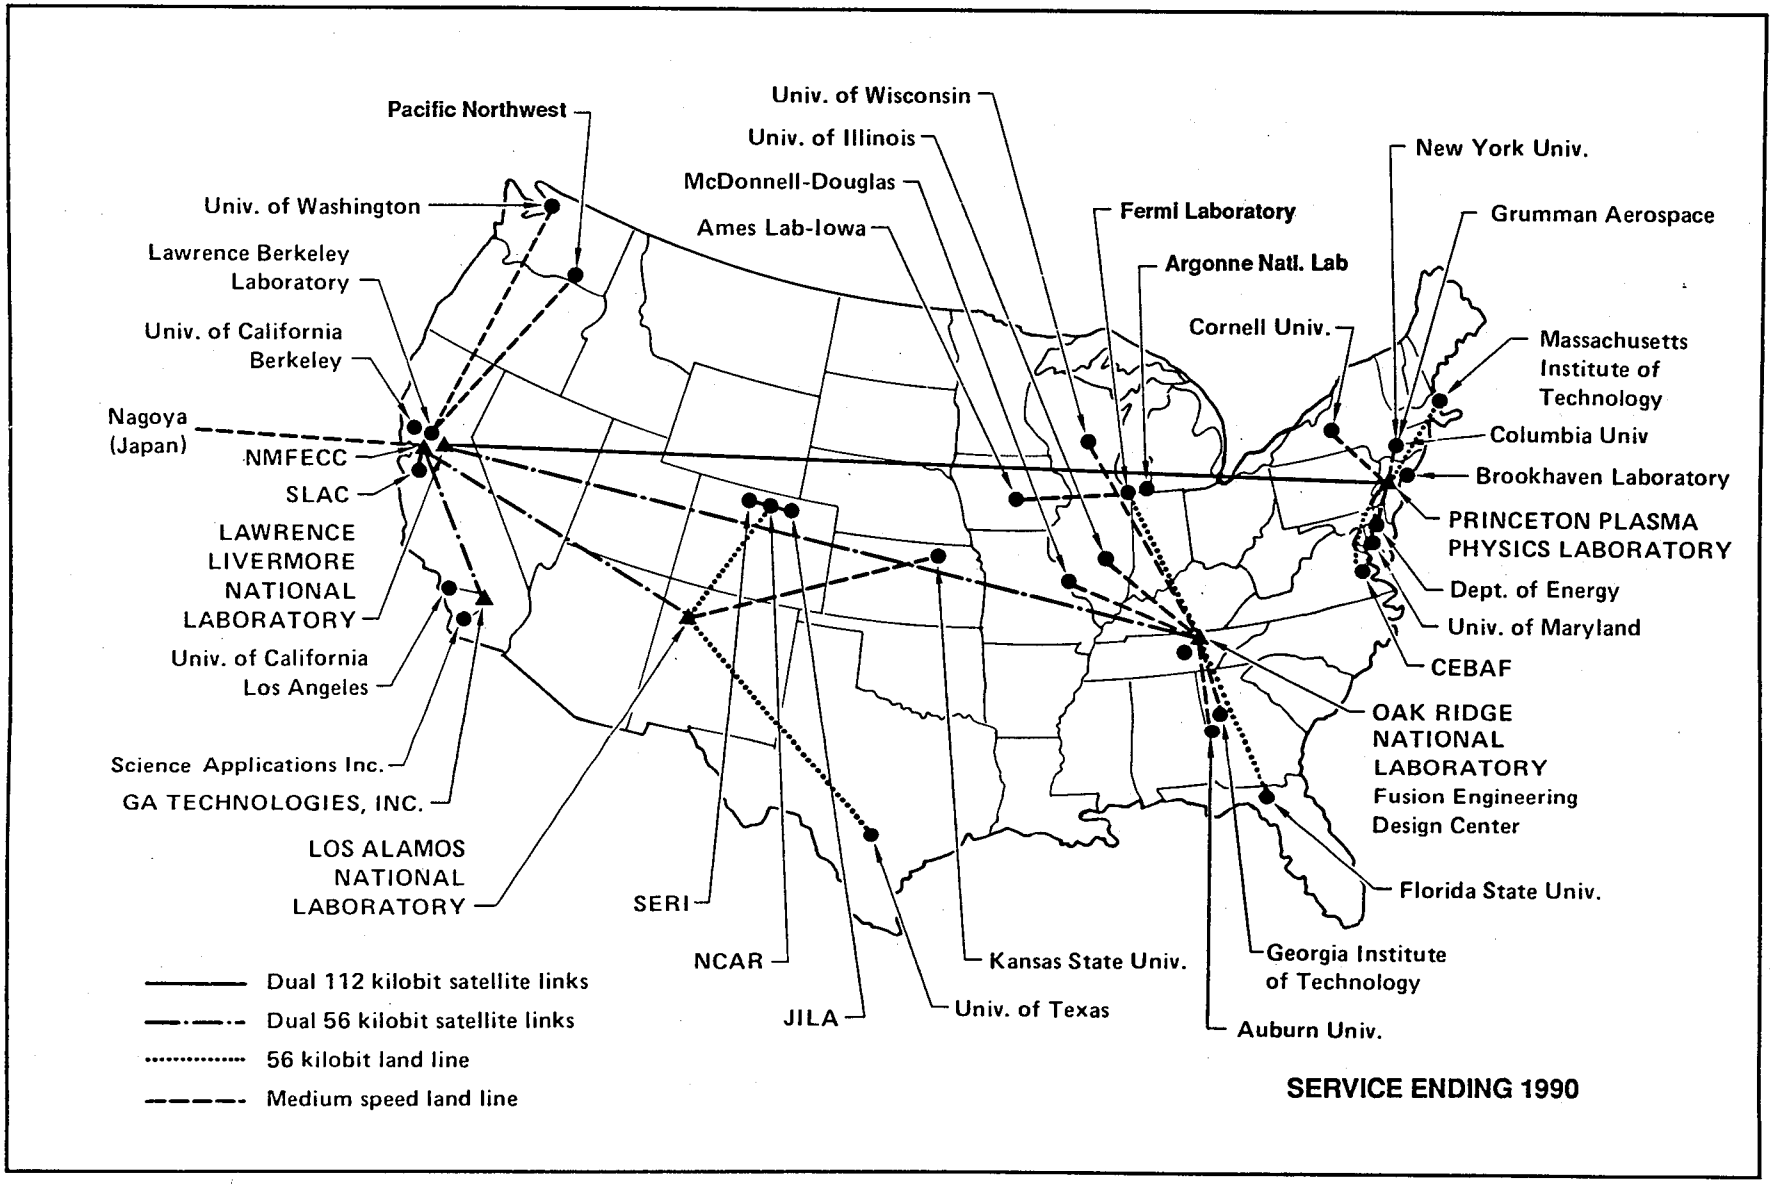

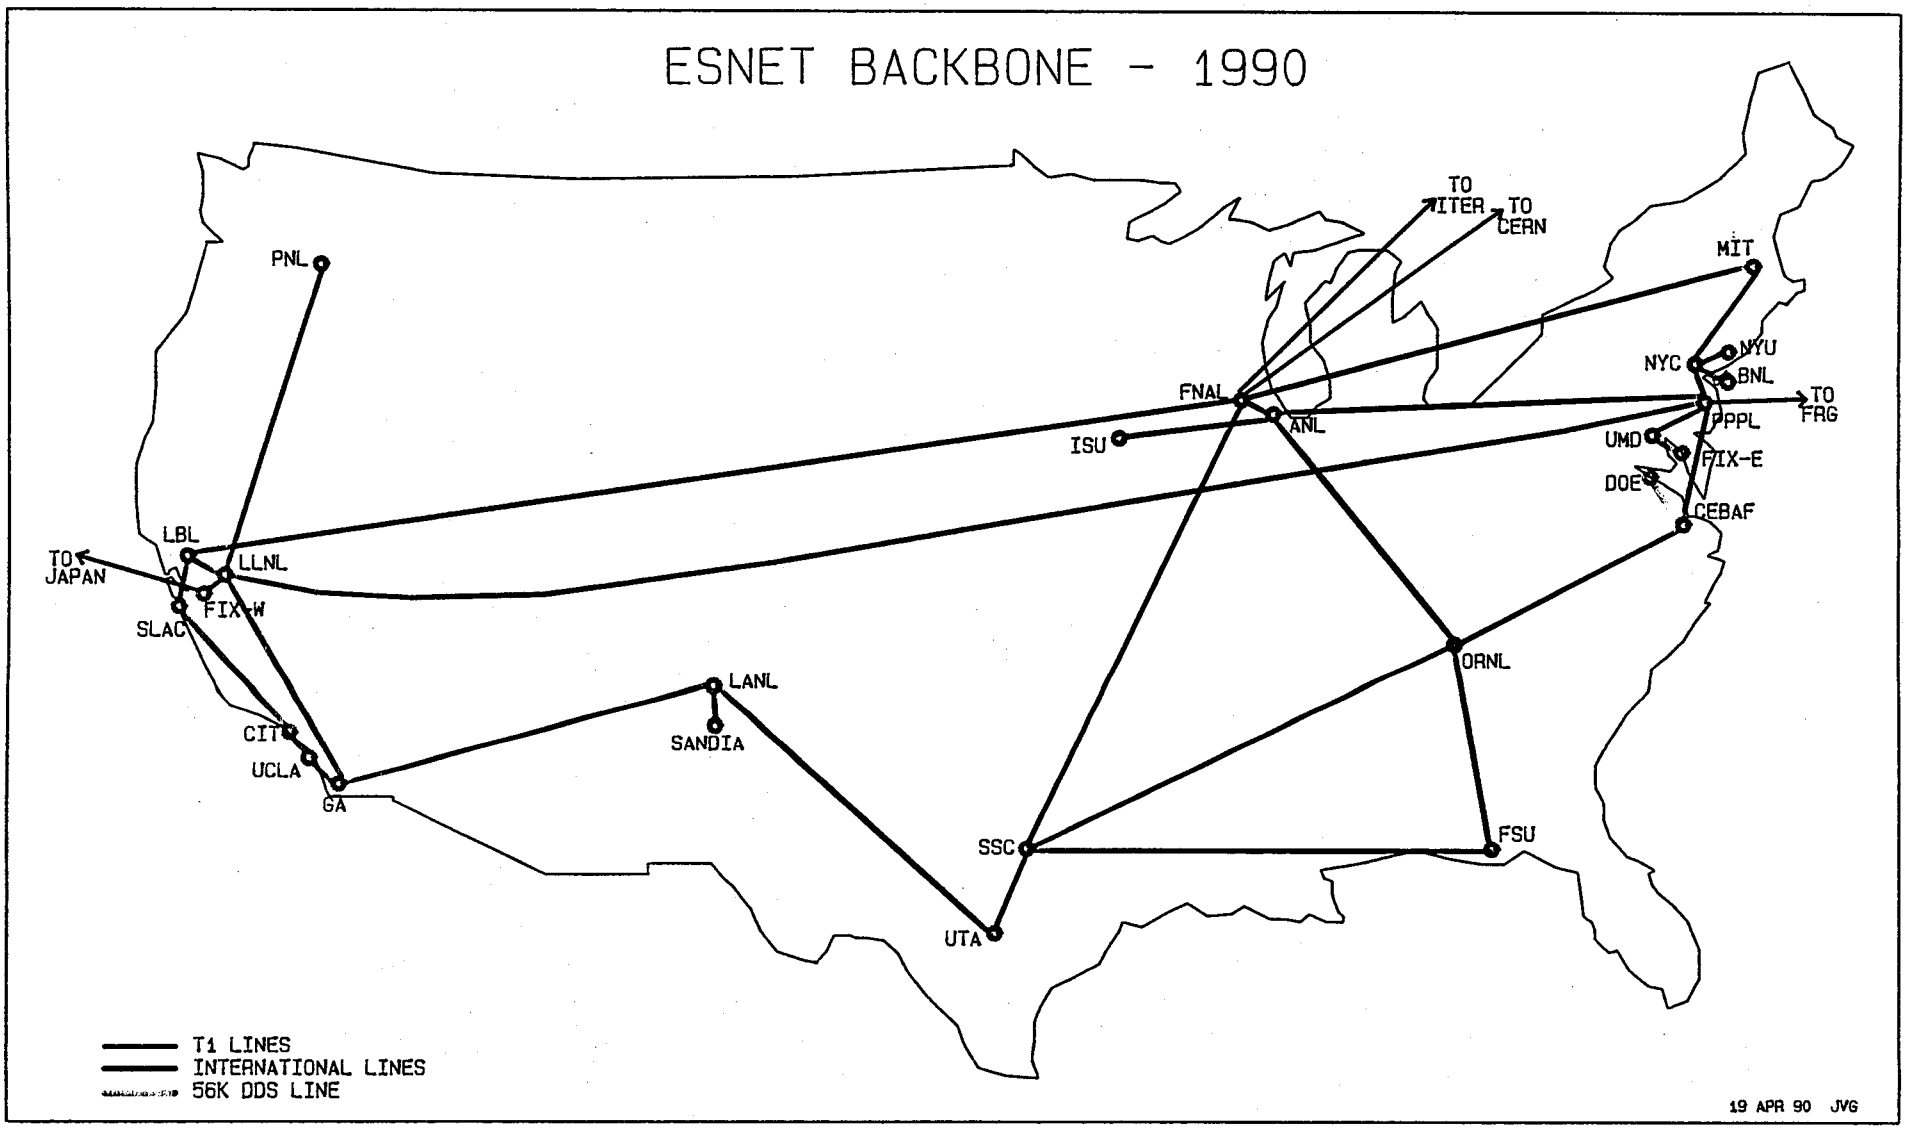
ESNET BACKBONE - 1990
TO ITER
TO CERN
MIT
PNL
NYC
NYU
BNL
FNAL
ANL
TO FRG
PPPL
ISU
UMD
FIX-E
DOE
CEBAF
LBL
TO JAPAN
LLNL
FIX-W
SLAC
ORNL
LANL
SANDIA
CIT
UCLA
GA
SSC
FSU
UTA
T1 LINES
INTERNATIONAL LINES
56K DDS LINE
19 APR 90 JVG