ETA SYSTEMS

# ETA10 System Overview: ETA System V

ETA10 Computer System

PUB-1232

Prepared by: ETA Systems, Incorporated
Technical Communication Dept.
1450 Energy Park Drive
St. Paul, MN 55108

# Table of Contents

# Preface

This document gives an overview of the ETA10 supercomputer with the ETA System V operating system, and a description of other products and services that are part of the overall hardware and software system.

It describes the makeup of ETA System V, and how enhancements added to the standard AT&T UNIX System V Release 3.0 operating system achieve the following:

- FORTRAN program performance
- Access to the ETA10's supercomputing capabilities
- Access to shared files
- Connectivity with other computing systems

This overview also tells how both the hardware and software work together for system performance. Information is presented in order of increasing detail for readers with different levels of interest.

For a quick entry into the system, the *Introduction to the ETA10 with ETA System V* highlights the advantages and performance characteristics of the system.

The sections on *Accessing and Using the ETA10* and *High Performance Processing* detail system features that give the user access to the ETA10's capabilities and outstanding program performance.

The sections titled *System Architecture*, *Enhancements to the UNIX System V Operating System*, *System Operation and Administration*, and *ETA10 Product Family* describe the aspects of the system named in the section titles in greater technical detail.

*The ETA10 Support Envelope* section describes customer support, training, and documentation.

The discussion assumes that you are familiar with the UNIX operating system from which the ETA System V operating system is derived.

For more information, contact:

Marketing Department
ETA Systems, Incorporated
1450 Energy Park Drive
St. Paul, MN 55108
Phone: (612) 642-3408

Or, contact your nearest Control Data Corporation sales office.

# ETA10 System Overview: **ETA System V**

# Introduction to the ETA10 with ETA System V

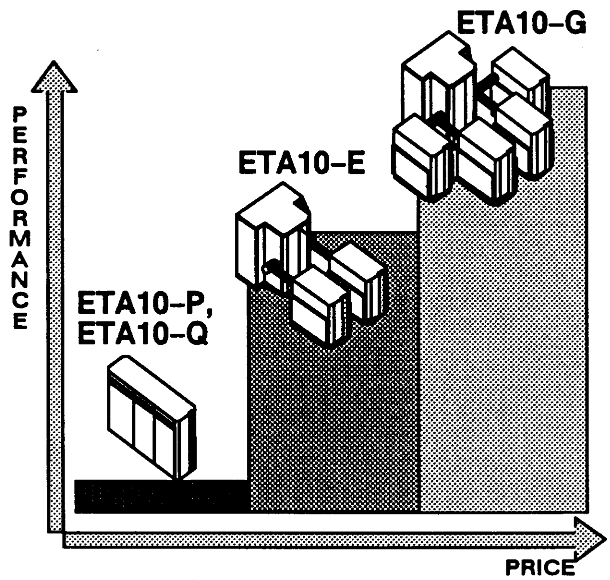


With more than 40 configurations, the ETA10 supercomputer family makes up the largest product line in the industry.

The performance range is also the broadest, at 27:1. A fully-configured ETA10-G with eight processors delivers 27 times the performance of a single-processor ETA10-P at the entry-level end of the range.

The ETA System V operating system, based on the industry-standard UNIX operating system with enhancements, runs on every model of the ETA10.

Every ETA10 has the same advanced hardware architecture based on the ETA10's central processor, a "supercomputer-on-a-board."

## ETA System V: The Industry-Standard Operating System Extended

The ETA System V operating system extends the familiar AT&T UNIX System V Release 3.0 operating system so that FORTRAN programs and other standard applications are easily ported to run on the ETA10 supercomputer. Organizations running a UNIX operating system can harness the ETA10's power without having to learn either a new operating system or non-standard network interfaces and without having to reprogram applications. Application compatibility is ensured now and in the future because ETA System V passes all 5,500 tests of AT&T's System V Verification Suite (SVVS), and future releases will continue to conform to industry standards. Familiar Berkeley extensions and other industry-standard facilities provide easy connectivity and file sharing between the ETA10 and other systems.

The ETA10 hardware and ETA System V support the TCP/IP protocol suite. Users access the ETA10 through an Ethernet local area network (LAN). Sun, Apollo, and CYBER 910 workstations are all verified to operate with the ETA10. Customers can connect almost any other type of TCP/IP-compatible terminal server, personal computer, minicomputer, mainframe, or supercomputer to the ETA10 through the LAN.

## Supercomputer on a Board

The key to the ETA10's power and its compatibility across the product family is the design of its central processor. Advanced manufacturing techniques put the entire central processor onto one 44-layer board about four times the size of this page. Just a few years ago, a central processor with comparable performance filled a room with cabinets and used many times the power of today's ETA10 central processors.

ETA's "supercomputer-on-a-board" is the fundamental building block of every model of the ETA10. Because the CMOS chips on each board are so fast, and generate so little heat, the central processor board achieves supercomputer speeds with the low-cost air-cooling used in the ETA10-P and Q models. With liquid nitrogen cooling, processing speed doubles in the higher-performance, super-cooled models E and G. Additional performance increases come from faster clock speeds and added CPUs.

The ETA10-P is a true supercomputer priced at the department level. Peak speed: 375 million floating point operations per second. It uses the same compact, powerful CPU as all other members of the ETA10 family.

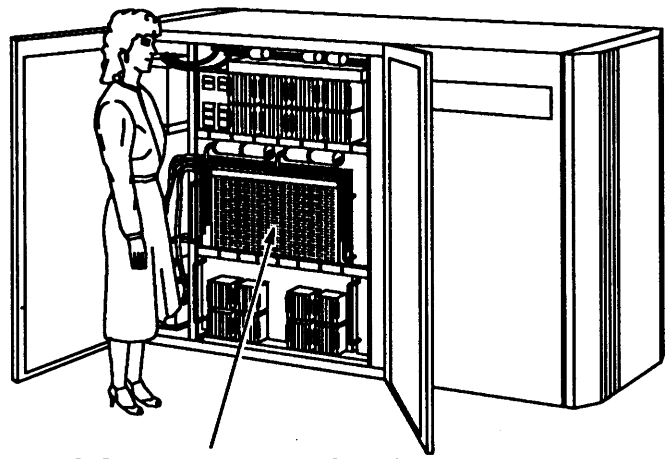


## FORTRAN 77 Compiler, Automatic Vectorizer, C Compiler

For FORTRAN performance on the ETA10 supercomputer, ETA System V includes an American National Standards Institute (ANSI) standard FORTRAN 77 compiler (ANSI X3.9-1978), so programs developed using standard FORTRAN 77 on other machines can be recompiled and run on the ETA10. The ETA VAST-2 preprocessor, an automatic vectorizer with unique interactive features, transforms FORTRAN scalar code into vector code that can be more quickly executed on the ETA10's vector-processing hardware. And, of course, ETA System V includes a portable C compiler with an optimizer.

## Designed for Top Performance

The ETA10 system achieves fast throughput because of its balanced architecture: very fast central processors, high-bandwidth memories, and a variety of I/O channel processors handling I/O functions that would otherwise slow the central processor. The hardware and the software work together to perform calculations faster on much larger problems than could ever be tackled before.

Each CPU is equipped with four million words (32 Mbytes) of contention-free memory and a central processor that includes 256 64-bit general purpose registers and scalar and vector processing units.

The system's large virtual memory, hierarchical physical memory, and fast I/O through fiber optic data pipes maintain the flow of data to keep the processors working at optimal capacity.

A large shared memory, which is where ETA System V puts the buffer cache, keeps data available to central processors much faster than it could be obtained from disk. A communication buffer that is part of the memory architecture synchronizes communications and can serve as shared memory for small data transfers.

The ETA10 provides a virtual memory address space of $2^{45}$ bytes, addressable to the bit. It supports large and small page sizes, making it possible for the system to accommodate large or small programs or data structures without explicit memory management.

## Large Family of Products

The ETA10 supercomputer family has the best price/performance range in the industry, and the largest product line, with over 40 configurations. The performance range is also the broadest, 27:1. An eight-processor, super-cooled model G at the top of the line has a peak performance of more than 10,000 million floating point operations per second (MFLOPS), running 27 times faster than the single-processor air-cooled model P at the low-cost end of the line, with a peak performance of 375 MFLOPS. Every ETA10, even the entry-level ETA10–P priced at less than $1 million (U.S.) is a true supercomputer with proven performance on industry benchmarks.

Because every ETA10 has the same central processor board design and can run the ETA System V operating system, customers easily move applications from a small ETA10 system to a larger one when their computing needs require more power. The upgrade potential in the ETA10 family protects the customer's investment in both software and hardware. Customers may add more central processing units, more shared memory, or more I/O units — up to the maximum configuration for each model. When computing needs grow beyond the maximum capacity for the initial configuration, customers can upgrade to a faster clock cycle, or to a larger number of CPUs, or they can move from an air-cooled to a super-cooled system.

## Mass Production, Easy Maintenance

The compactness and speed of the ETA10's central processor derives from 240 very high density, 20K gate array CMOS chips doing the work of tens of thousands of less-advanced chips, and from the densely-laminated central processor board, which replaces more than a mile of wiring. The central processor board and other components of the ETA10 are manufactured in an automated, volume-production environment — putting it a generation beyond hand-built supercomputers. Advanced production and quality-control techniques result in shorted delivery lead times and new levels of reliability.

Each chip on the central processor and major system interface boards has its own Built-In Evaluation and Self Test (B.E.S.T.) logic, which identifies and pinpoints failures to specific chips and specific wires on boards and between boards. The ETA10 is the only supercomputer with this capability, which speeds up initial system checkout and makes testing and problem diagnosis easier once the system is out in the field.

When customers report problems, ETA Systems' Remote Systems Support center performs remote troubleshooting for software and hardware problems over a high-speed telecommunication link to a service unit at the customer's site.

The ETA10's design incorporates a large number of field-replaceable, plug-in units, making maintenance quick and simple and reducing support costs, while allowing most upgrades to be field-installed.

## Documentation and Training

ETA Systems' reference manuals and guides provide all the information about system hardware and software needed by programmers and users. Other manuals serve operators, administrators, and system support staff.

As part of customer support, ETA offers courses designed and taught by ETA instructors for programmers, analysts, operators, and system administrators. The courses are modular in design, so that training can be tailored to accommodate a customer's particular needs.

**Highlights of the ETA10 Supercomputer with ETA System V**

| | |
|---|---|
| AT&T UNIX Release 3.0 operating system enhanced | Standard operating system, user and application portablility |
| FORTRAN 77 with ETA VAST-2 automatic vectorizer | A FORTRAN engine |
| Supports Ethernet with TCP/IP and NFS | Connectivity and file sharing |
| 27:1 performance range, common architecture | Investment protection |
| Supercomputer central processor on a board | Speed, reliability, low operating cost |
| 32-bit math (in addition to 64-bit) | Double speed and data capacity |
| Advanced vector hardware | Very fast vector operations |
| Large virtual memory | Handles very large problems |
| Balanced architecture, large bandwidth | High throughput |

# Accessing and Using the ETA10

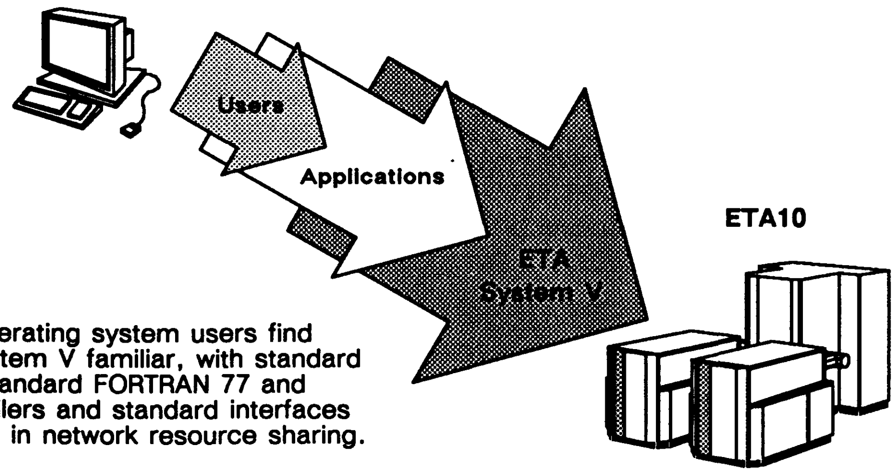


UNIX operating system users find ETA System V familiar, with standard tools, standard FORTRAN 77 and C compilers and standard interfaces for ease in network resource sharing.

## Users and Standard ETA System V

Users interact with the ETA10 through the ETA System V operating system, which is based on the AT&T UNIX System V Release 3.0 operating system. Users can count on familiar interfaces to a familiar operating system because both the operating system and the system calls conform to AT&T's System V Interface Definition (SVID), and pass all tests of the System V Verification Suite (SVVS). ETA System V will continue to conform to future industry standards as they are defined, including the Portable Operating System for Computer Environments (POSIX) standard being developed by the Institute of Electrical and Electronic Engineers (IEEE).

Users already familiar with the UNIX operating system can access the ETA10's supercomputing resources for running FORTRAN and other standard programs without having to learn a new operating system. Software applications based on any UNIX operating system run right away on the ETA10 as long as they are machine-independent and compatible with the SVID.

## Users and the FORTRAN and C Compilers

FORTRAN programs compatible with ANSI standard FORTRAN 77 also are easily ported to the ETA10. To enable users to take advantage of the ETA10's processing speed for their FORTRAN applications, ETA System V provides the FORTRAN 77 compiler. FORTRAN 77 comes with run-time libraries, and with single, double, and half-precision math libraries. If full precision isn't necessary, or if applications were created in a 32-bit environment, programs can be run at half-precision, with twice the processing speed of 64-bit full precision. If even greater than 64-bit accuracy is needed, precision to 128 bits is another option. FORTRAN 77, together with the ETA VAST–2 preprocessor, automatically vectorizes scalar code to make use of the fast vector processing capabilities of the ETA10.

Of course, ETA System V provides a standard portable C compiler, standard C libraries and a portable code optimizer.

## User Access to the ETA10

Users access ETA System V and the ETA10 through industry-standard networking facilities familiar to experienced UNIX users. The user commonly logs in to the ETA10 from a workstation on an Ethernet local area network (LAN). Once logged in, users have access to the ETA10's resources. They can also connect to other systems on the LAN and transfer files back and forth between other systems on local or remote networks. Berkeley *"r"* commands allow access to other UNIX systems, while *ftp* and *telnet* commands allow access to any TCP/IP-compatible system.

A possible network configuration for the ETA10 is shown below. In this example, workstations are connected directly to the Ethernet local area network (LAN), which in turn connects to the I/O Unit (IOU) of the ETA10. Almost any other type of computing system or network can be connected to the network by means of bridges, gateways, and servers. Authorized users connected to any of these systems then can use the ETA10's supercomputing power for all or part of their work.

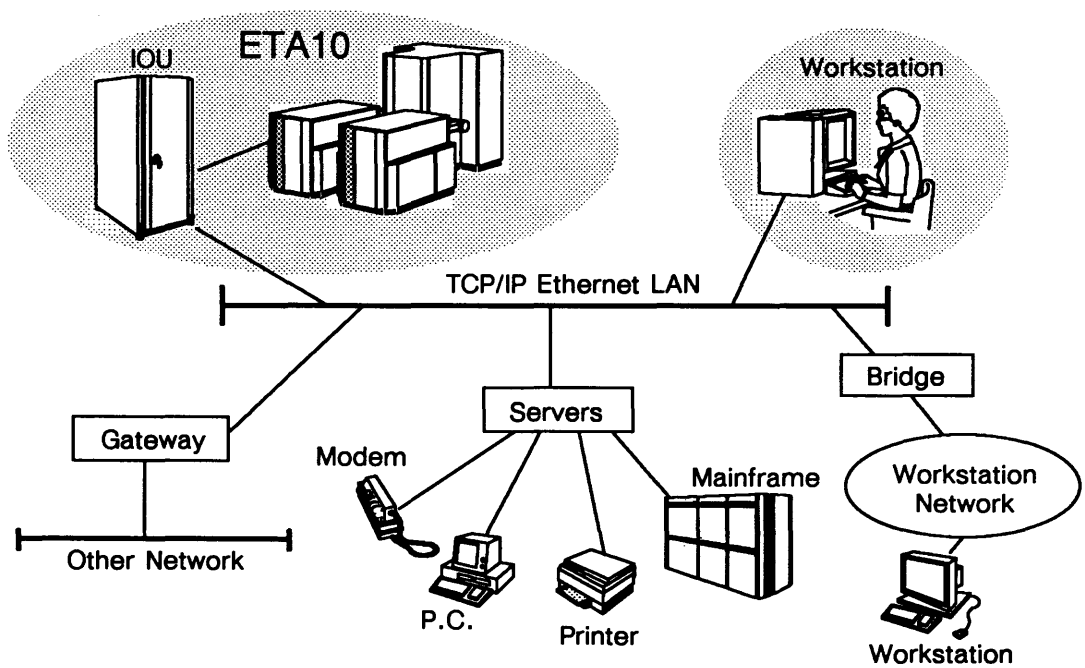


## Users Transparently Share Files Across the Network

Network File System (NFS) enables users to share files with other systems connected to the ETA10 through networks. Users are able to access files on other systems running NFS as if the files were stored on their own systems.

# High Performance Processing

A combination of hardware and software features work together to make the ETA10 a high-performance FORTRAN engine. The ETA10's performance is especially important when users need solutions to large or computation-intensive problems, because accuracy doesn't have to be sacrificed in order to get results within a reasonable time.

## A Balance of Scalar and Vector Processing Capabilities

The ETA10 balances fast scalar processing with very fast vector processing. This balance speeds up throughput because typical computation-intensive programs have significant amounts of scalar work to do along with those parts of the work that can be vectorized. The ETA VAST–2 preprocessor included with the FORTRAN 77 compiler automatically vectorizes scalar FORTRAN code to make it able to take advantage of the vector processing units in the central processors. The FORTRAN 77 compiler also optimizes FORTRAN code for scalar processing.

To support the processing capabilities of the scalar and vector processors, the memory architecture of the system is designed to accomodate both large and small programs, to avoid I/O bottlenecks, and to feed the processors data as fast as they can process it.

The ETA10 is designed so that scalar and vector operations can be performed in parallel with each other and while I/O is going on. Vector processing is what gives the computation speeds needed for solving large-scale problems. After the startup time that it takes to get the first result, the vector processor will output two (64-bit) results per clock cycle.

If programs can be run at half-precision (or 32-bits), twice as many data items fit in the same memory space. Each vector pipeline can then process four 32-bit data values per clock cycle, and the two pipelines together can process a total of eight values and obtain four results per cycle.

If the code contains certain chain-type operations, called linked triads, in which two vector and one scalar operation or one vector and two scalar operations are performed in sequence, and when half-precision can be used, performance increases to eight operations per cycle.

## Pipelines as Computational Assembly Lines

Processing on the ETA10 takes place in pipelines. Just as a complex operation such as assembling a car can be broken down into steps to be performed one after another on an assembly line, mathematical operations such as ADD or MULTIPLY can also be broken down into a series of steps and each step can be accomplished within one segment of a pipeline. Operands are accepted at one end of the pipe, and a result is available at the other end after the entire operation is complete.

The pipeline shown below is separated into five segments, each of which performs one step or suboperation in a computation. (The number of segments in actual pipelines varies.) On the ETA10, the sub-operation in each segment takes one clock cycle to finish. When the first segment completes its processing, the results are sent to segment 2, and segment 1 is now free to begin processing a new set of operands.

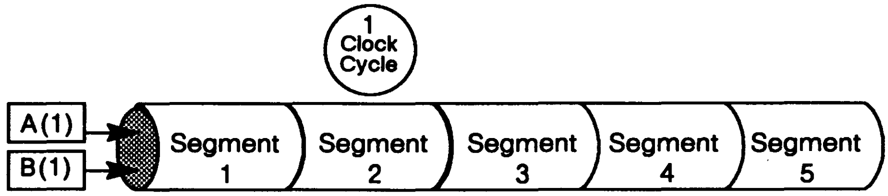


During the next and subsequent clock cycles, another pair of operands enters the pipeline, and results of the operands that entered in earlier clock cycles move to the next segment. Because this sample pipeline has 5 segments, the first pair of operands would move through the pipeline for 5 clock cycles before the first of the final results would be obtained.

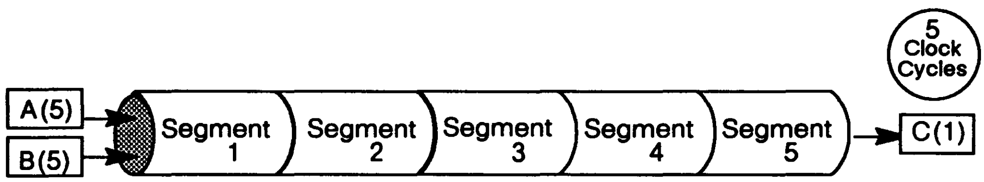


After the first of the final results is obtained, the subsequent final results stream out of the pipeline at the rate of one per clock cycle.

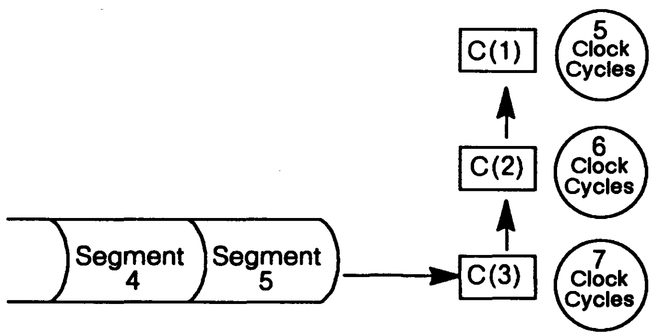


## Comparing Scalar and Vector Processing

To visualize the difference between scalar and vector processing, consider adding arrays A and B, each of length N:

C(I) = A(I) + B(I) where I equals 1, 2, ..., N

A scalar processor obtains the results by executing a series of instructions contained within a loop, as illustrated in the following figure.

```
LOAD A(I)
LOAD B(I)
ADD A,B
STORE C(I)
INCREMENT I
TEST I
BRANCH
```

(Many Clock Cycles) = 1 RESULT

A(1) and B(1) are loaded from memory into registers; A(1) is added to B(1); the result C(1) is stored back in memory; I is incremented by one; I is tested against N and, if it is not greater than N, control branches to the top of the loop. The process continues with A(2) + B(2), and so on, until all the operand pairs have been processed. The time required to complete this loop is the cost of each load, add, store, and branch multiplied by the size of the array. The scalar processor needs many clock cycles to obtain a single result. Working on the same problem, the vector processor, with its two vector pipelines working in parallel, achieves two results in one clock cycle. The vector statement shown below achieves the same result as the list of instructions in the scalar example.

```
C(1:N) = A(1:N) + B(1:N)
```

(One Clock Cycle) = 2 RESULTS

## Vectorization

Vectorization is the translation of scalar code into code that can be processed by vector processors. ETA System V provides an automatic vectorizer, ETA VAST-2, as a preprocessor to the FORTRAN 77 compiler. Both tools are in one utility called *ftn77*.

During the first pass of the code through *ftn77*, ETA VAST-2 automatically transforms vectorizable operations. It also recognizes code that could be vectorized if more information were available. In an output file, it provides diagnostic messages indicating where more information can identify whether or not it is safe to vectorize a certain portion of the code. The programmer can then insert directives to increase the total amount of code vectorized during the second pass. The following flowchart shows what happens during the first pass of compilation.

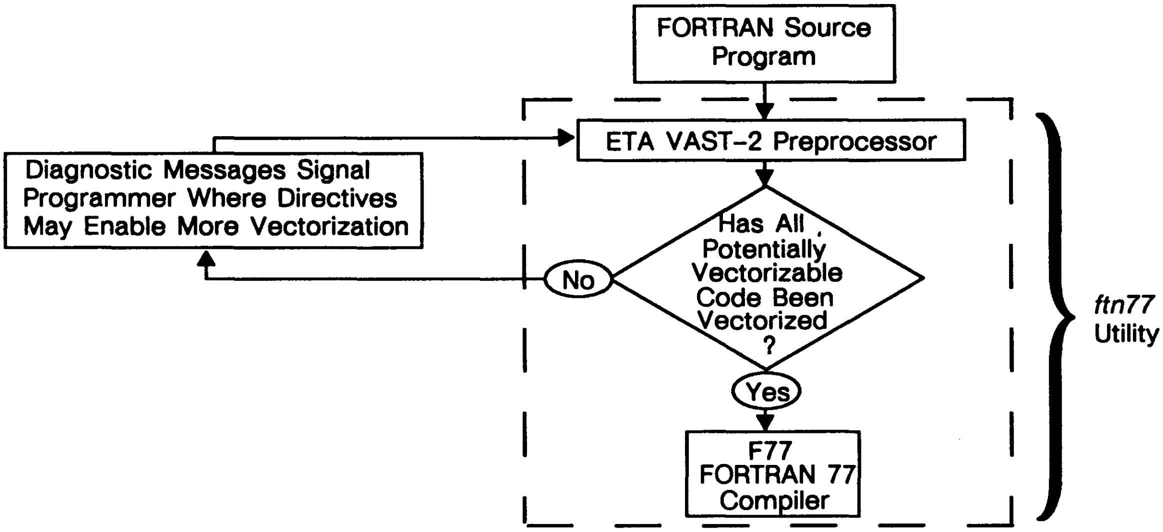


## Memory-to-Memory Architecture Speeds Vector Processing

The memory-to-memory ȧrchitecture of the vector processor speeds up computations. Unlike systems that use vector registers, the ETA10 supplies data to its vector pipelines directly from memory. Data streams through an input buffer directly to the vector pipeline where computations are performed. When results are obtained, data streams directly into memory again through the output buffer.

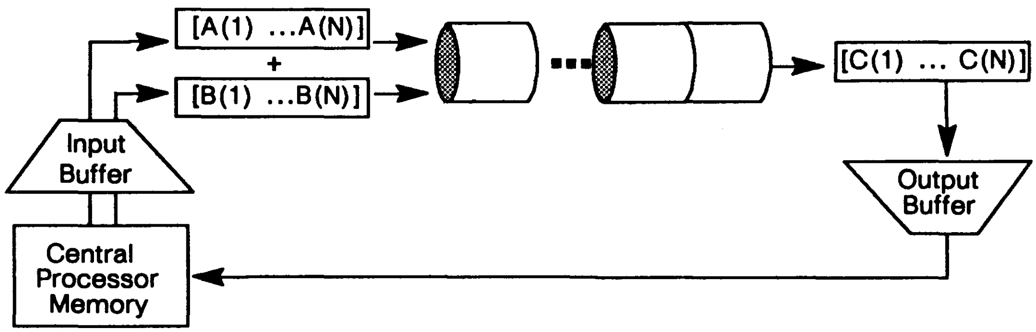


Since the startup time is fixed, independent of the length of the vector to be processed, the time it takes per operand decreases when there are longer vectors. The ETA10 hardware design minimizes the startup time for vector instructions by overlapping vector startup with other operations, so that vector processing is effective for short vectors as well as long ones. This is accomplished by buffering data needed for a subsequent vector operation along with the data needed for the current vector operation while the current operation is still being performed.

A vector shortstop capability shortens startup time for short vectors. If the next vector instruction uses the previous result as an operand, and if that result is in a shortstop buffer, it is immediately available to be fed back into the pipeline.

# System Architecture

The ETA10 supercomputer system is made up of a combination of hardware and software components that provide the balance of scalar and vector processing power and the fast access to data that are needed for high performance.

The main work of the system is done in CPUs, each of which has its own memory, and a processor with scalar and vector pipelines. The memory structure of the system is designed to keep enough data supplied to the central processors for computations to take place at top speed.

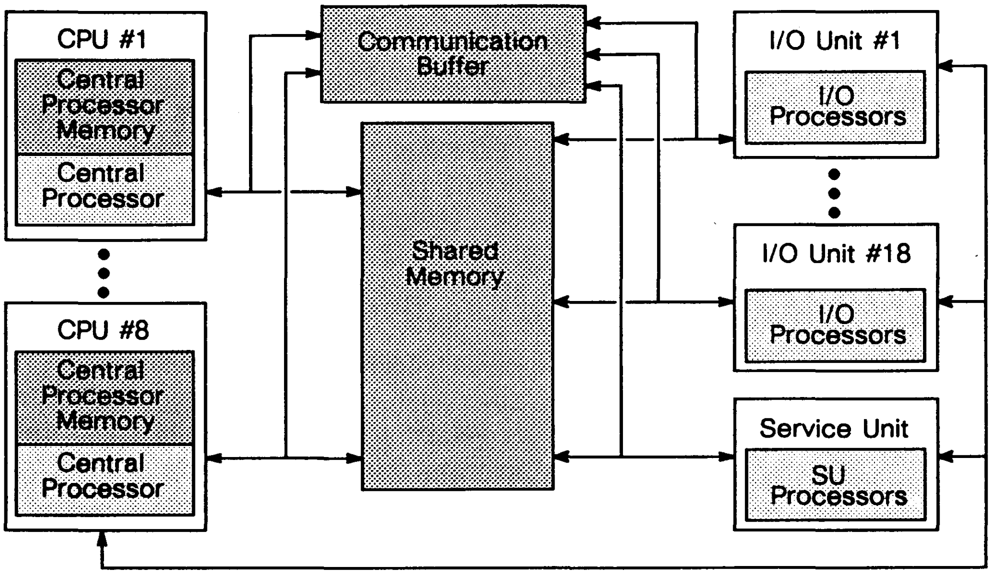


## Components of the System Architecture

As shown in the system diagram above, between one and eight central processing units (CPUs) connect to the shared memory, to the communication buffer, and to the service unit. Between one and eighteen input/output (I/O) units may be connected to an ETA10 depending on the model and on the number and type of peripherals connected to the system.

### Central Processing Unit

The CPU is where scalar and vector processing are done. Each central processor has its own central processor memory, which consists of four million 8-byte words (32 Mbytes). There is no contention from other CPUs for the local memory within each CPU.

Processing takes place in a scalar unit and a vector unit. Both scalar and vector operations take place simultaneously, in parallel with each other and with I/O.

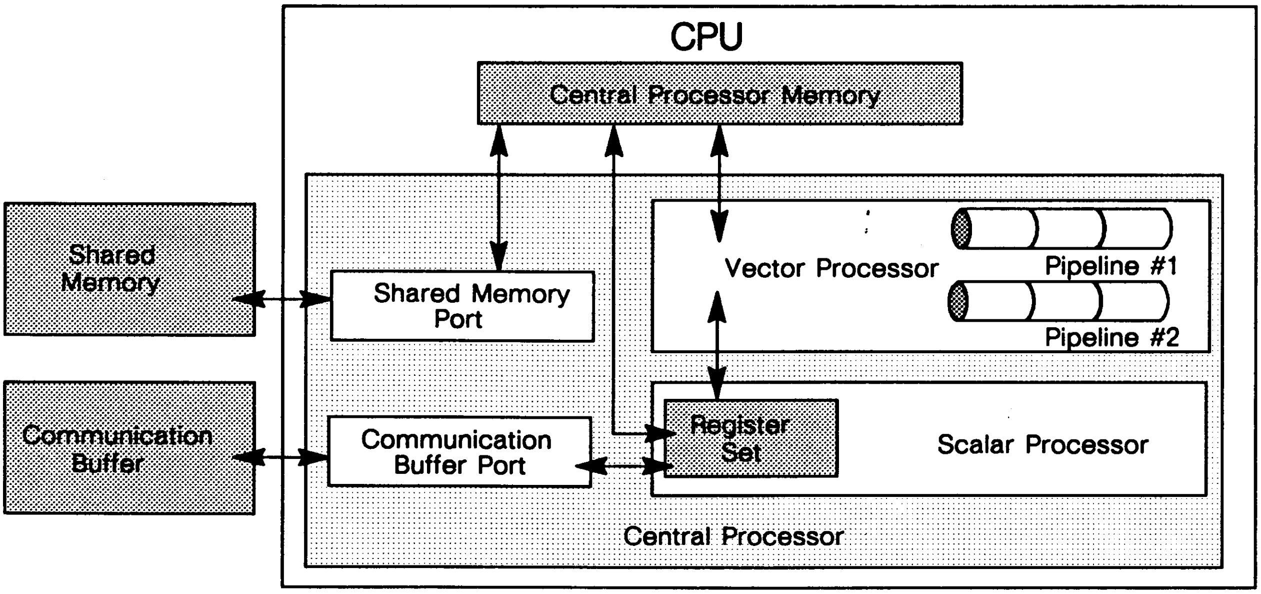


The illustration on this page shows a central processor memory and a central processor within a CPU. Within the central processor is a scalar processor, a vector processor, 256 64-bit registers, and ports to shared memory and to the communication buffer.

The kernel of the ETA System V operating system executes in the central processor. Processes running in the central processor communicate with the kernel by means of standard UNIX system calls.

### Shared Memory

The CPUs also access a large shared memory. The size of the shared memory is selectable from 64 to 2048 Mbytes.

Shared memory is primarily used as a buffer cache for files that are being read and written by processes executing on the ETA10. Data transfer between shared memory and central processor memory is very fast because of the high bandwidth between these two memories. The large shared memory allows users to run large and complex models faster than is possible with central processor memory alone. It also enhances system efficiency by reducing the amount of data transfer between memory and disk.

### Communication Buffer

The communication buffer has either 1/2 or 1 million words of memory (4 or 8 Mbytes). The communication buffer provides a means to synchronize I/O and pass messages among CPUs and I/O processors. Small amounts of data can be cached in the communication buffer for extremely fast transfer to the registers in a central processor.

### I/O Unit

I/O units connect the ETA10 to a broad range of high-speed data storage and network devices. Data is transferred between the I/O units and the shared memory through fiber optic data pipes. Each I/O unit will hold up to five channel processors, each of which connects to specific types of storage or network devices. All the channel processors share a single data pipe connection to shared memory. I/O processors within channel processors in the I/O units control data movement and message passing between the system and its peripherals, freeing the central processor from many low-level I/O functions.

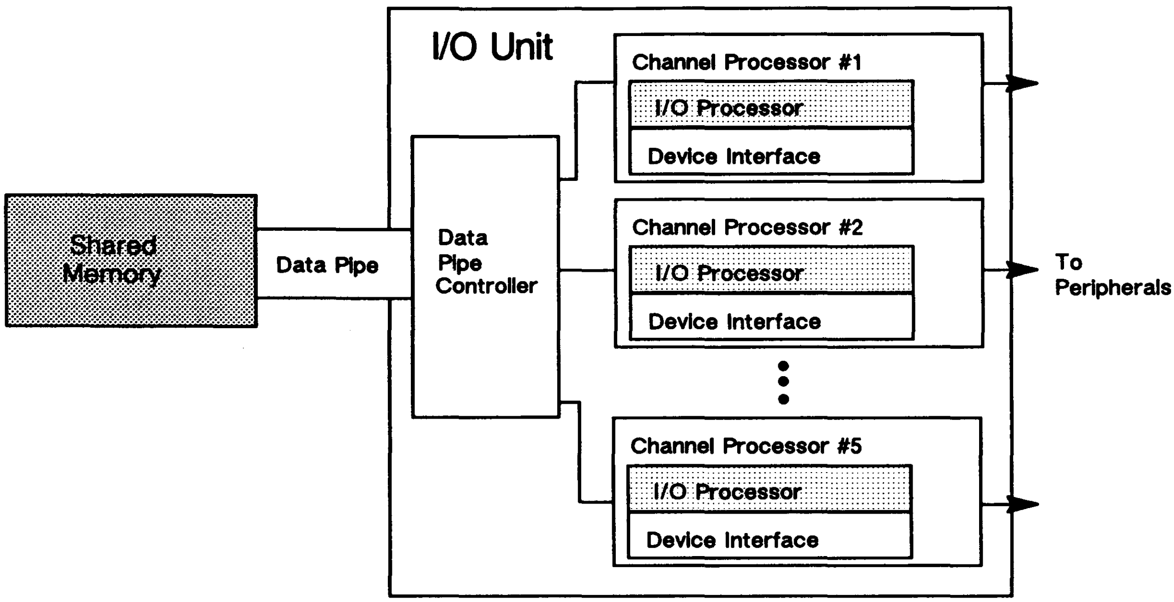


### Service Unit

The service unit consists of one or more operations and/or maintenance consoles. By using the Built-In Evaluation and Self Test (B.E.S.T.) logic on each CPU chip, and other diagnostic utilities, engineers at the service unit can troubleshoot hardware problems. The service unit also contains a high-speed modem that enables the engineers in the ETA Systems' Remote System Support (RSS) group to access and troubleshoot the system, at the customer's request.

## Hierarchical Memory Structure

The ETA10 has a hierarchical memory structure. The path that I/O takes, across the hierarchical memory, is between central processor memory, shared memory, and the peripherals. The data transfer rate between peripherals and shared memory is high; the channels from peripheral to I/O unit are multiplexed into a faster path between I/O unit and shared memory. CPUs can use data faster than this and, because the bandwidth between the shared memory and the central processor memory is higher, shared memory is used as a large, very high-speed peripheral for the CPUs. Files being used by executing processes in the CPU are stored in shared memory. Data cached in the shared memory can be supplied to the central processor memory much faster than from disk.

The fastest memory of all for the central processor to access is the central processor memory. The rate of data transfer between central processor memory and the CPU matches the CPU's processing capability: 512 bits (64 bytes) per clock cycle, more than the bandwidth needed to run the vector processor at peak speed. To ensure accuracy, single error correction/double error detection (SECDED) is performed on every 32 bits.

### Large Virtual Address Space

Programs running in the central processors can make use of a large virtual address space of 35 trillion bytes for their code and data. Data in virtual memory is bit-addressable. The virtual addressing mechanism is supported both by hardware and by virtual memory functions in ETA System V.

Underlying the large virtual address space are the three types of memories as they are shown in the following illustration, from quickest access and smallest on the left to the slowest access and largest on the right: the CPU's central processor memory, shared memory, and disk storage.

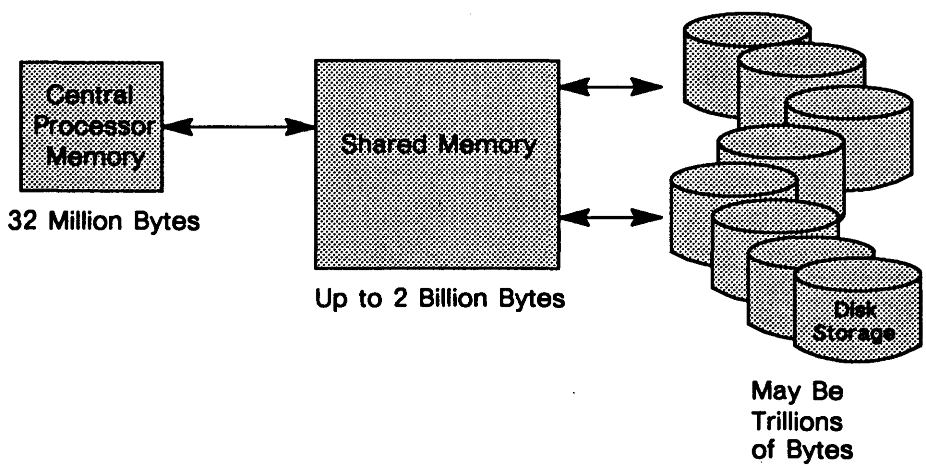

# Enhancements to the UNIX System V Operating System

While implementing standard functions from the AT&T UNIX System V operating system, ETA System V contains significant enhancements that extend its supercomputer capabilities. Some enhancements fine-tune the kernel to match the capabilities of the ETA10 hardware. Other enhancements extend the basic UNIX software system for FORTRAN support. A third type of enhancement extends the network capability of the ETA10 and provides transparent file sharing. These enhancements are described in this section.

## Enhancements in the ETA System V Kernel

### Large Block Size

To reduce the overhead associated with multiple I/O transfers and to support high performance disks, ETA System V has a large block size (16 Kbytes) for data being read from disk.

### Shared Memory Buffer Cache

ETA System V has changed the traditional UNIX procedure of putting the buffer cache in central processor memory. To free up the central processor memory and to make it possible to provide more space for executing processes, ETA System V puts the buffer cache in shared memory. Shared memory buffer caching provides the central processor with very fast access to data and reduces the need to read and write to disk.

The following illustration shows the path over which data are transferred from disk, through the IOU to the buffer cache in shared memory, and then transferred from shared memory to the central processing memory when needed. The size of the arrows indicates the comparative speeds of the data transfers.

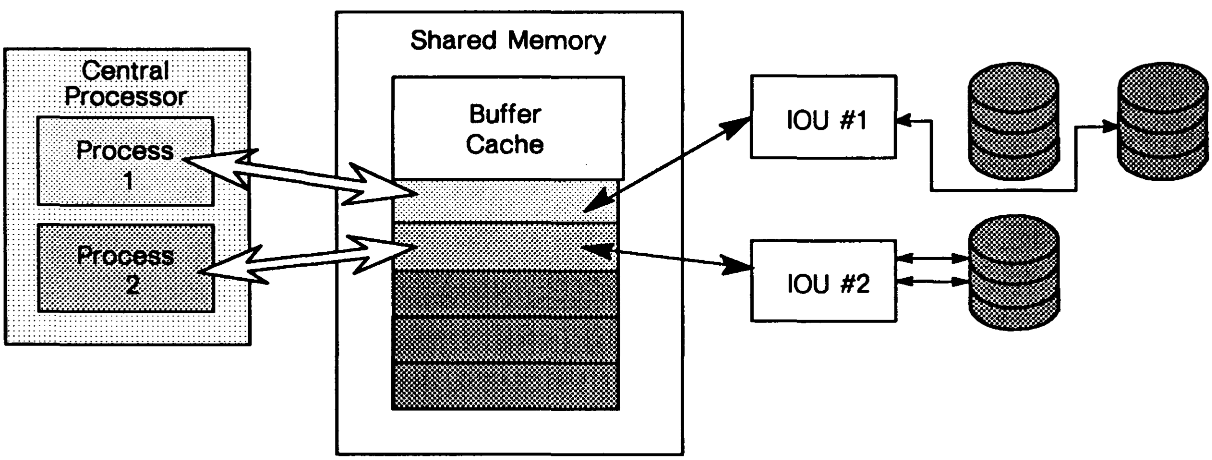

## Asynchronous I/O

ETA Systems has extended ETA System V with new system calls to support asynchronous as well as synchronous I/O. Asynchronous I/O is a requirement where high performance is essential.

When a standard synchronous UNIX *read*(2) or *write*(2) system call is made to the kernel of a UNIX operating system, control returns to the user only when the call is complete. The new system calls take advantage of asynchronous I/O capabilities built into the ETA10 hardware. A caller does not have to wait for the I/O operation to complete before getting control back.

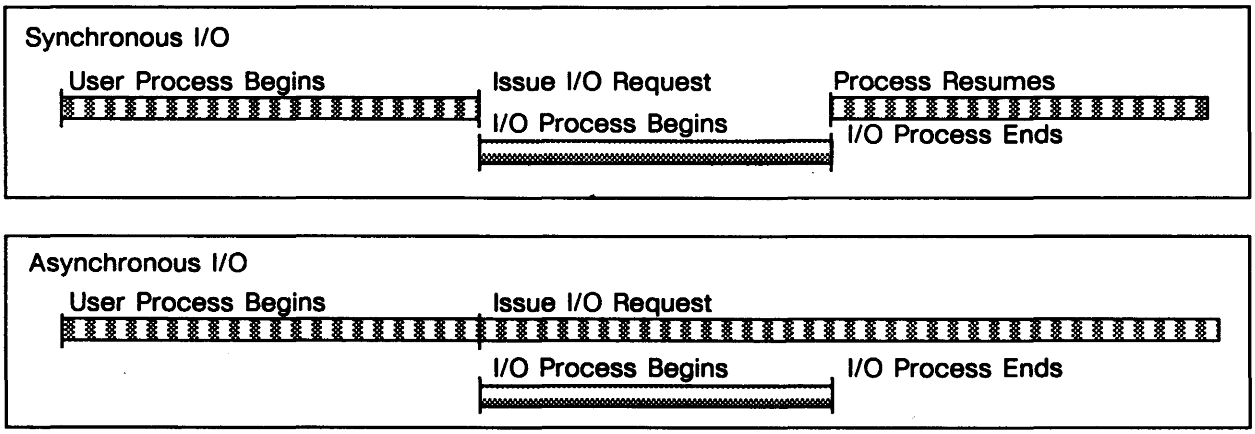


## Virtual Memory/Large and Small Pages

As mentioned in the system architecture discussion, the ETA System V software works with the virtual memory capabilities of the hardware to support a virtual address space of $2^{45}$ or about 35 trillion bytes.

The ETA System V operating system supports two small page sizes of 16 and 64 Kbytes (2 and 8 Kwords) and a large page size of 512 Kbytes (64 Kwords), which are also supported by the ETA10 hardware. Large page sizes allow large FORTRAN and other types of programs to execute at high speed by reducing the number of times the system has to access shared memory and secondary memory. Users can organize programs to run efficiently by making best use of large and small pages and can compile and link programs in such a way as to control how virtual memory space is used and how pages are allocated.

## Distributed I/O Intelligence

The ETA system design delegates many low level I/O functions to the I/O units, freeing CPUs from most I/O management responsibilities so they can spend more time executing user tasks. This distributed intelligence is built upon a real-time, multitasking kernel in each I/O processor. This kernel provides support for the processes that handle the low level I/O functions, and those that handle communications between the I/O unit and peripheral devices and between the I/O unit and the central processors.

The message-passing scheme both within the ETA System V kernel and between the I/O processors and the central processors is based on the AT&T UNIX Release 3.0 operating system feature, STREAMS. STREAMS allows data to be passed in streaming mode for increased throughput.

STREAMS also provides a standard mechanism that programmers can use when developing networking applications and individual device drivers.

## Enhancements for a FORTRAN-to-Kernel Interface

To adapt ETA System V for high performance FORTRAN, ETA Systems has customized a FORTRAN run-time library to serve as the interface between ETA System's FORTRAN 77 compiler and the C language-based ETA System V kernel. The compiler supports the UNIX Common Object File Format (COFF) standard.

## Enhancements for Networking

ETA System V software and the ETA10 hardware provide support for Ethernet, with network communications based on the Department of Defense (DoD) standard TCP/IP protocol suite. Communications facilities include our implementation of Berkeley sockets based on AT&T's STREAMS, and AT&T's Transport Level Interface (TLI). In addition, ETA Systems has extended the ETA10's networking capabilities with the following utilities and file-sharing facilities.

### Berkeley and DoD Networking Utilities

ETA System V provides two sets of networking utilities: the Berkeley "*r*" commands for communicating with UNIX hosts, and the DoD *ftp* and *telnet* utilities for communicating with hosts running any type of operating system.

### Transparent File Access

With Network File System (NFS), the ETA10 can share data with other computers transparently across the network. Authorized users can use files from remote directories as if they were local.

NFS uses remote procedure calls (RPC) and external data representation (XDR) procedures. ETA System V also includes Yellow Pages (YP), a centralized database that provides password, group, network, and host information, a service that eases the job of administering networked machines.

## Summary of Networking

Besides the features and protocols already described, ETA System V provides support for these networking protocols: Address Resolution Protocol (ARP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

The following figure shows how the networking protocols and applications fit together for network functionality by indicating where the protocols and applications would fit into the seven levels in the ISO/OSI networking model. This reference model was prepared by the International Standards Organization for Open Systems Interconnection as a basis for international standardization of networking protocols.

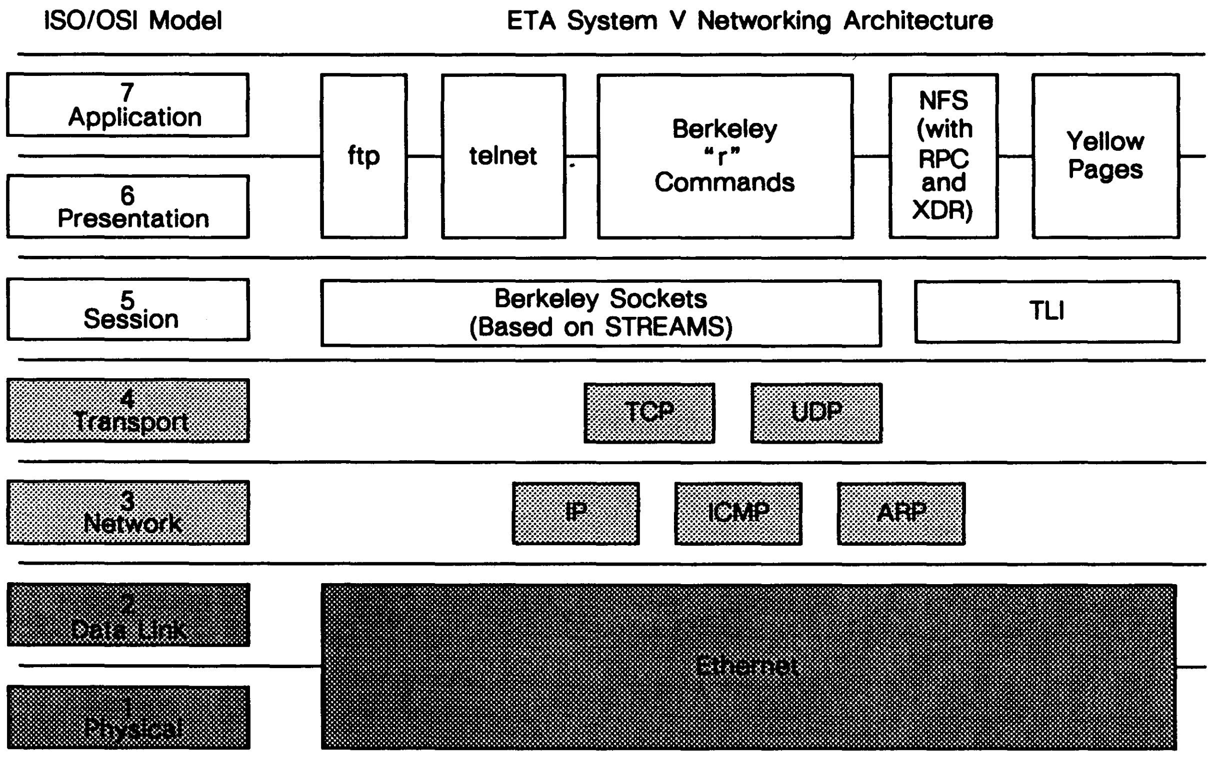

# System Operation and Administration

System administration for ETA System V on the ETA10 generally conforms to standard UNIX system administration procedures, with some additions that may be required to support the supercomputing environment of the ETA10. The decision about which procedures need to be performed by which personnel at each site must be tailored to the size and goals of your operation and to your organization's management system.

System monitoring is generally performed through the ETA10 service unit by ETA System's system support personnel. Certain operations functions, such as booting the system, are performed at the service unit console by customer personnel at the customer's site. Other operations functions may be performed by customer personnel through any computing system that has access to the ETA10. These functions may be assigned to more than one person.

Traditionally, UNIX sites assign system administration and operation tasks to system administrators. UNIX sites with small systems may have a single person acting as system administrator and performing other duties as well.

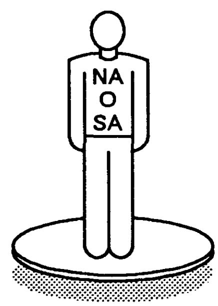


ETA10-P or Q

KEY:
NA = Network Administrator
O = Operator
SA = System Administrator

On the other hand, a traditional supercomputer installation where the supercomputer is run as a back-end usually requires a full-time system administrator supported by operators and, perhaps, a full-time network administrator. Operators generally perform computer-room tasks such as mounting tapes. This breakdown of tasks is used at some sites where the larger ETA10-E or ETA10-G multiprocessor super-cooled units are installed.

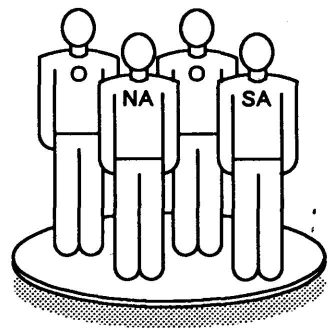


ETA10-E or G

The traditional UNIX single-administrator configuration is still appropriate for many installations, usually those with the smaller air-cooled models, ETA10–P and ETA10–Q. However, because all models may be accessed interactively without any front-end computing system between the user and the supercomputer, even sites with the larger models can be administered the traditional UNIX way.

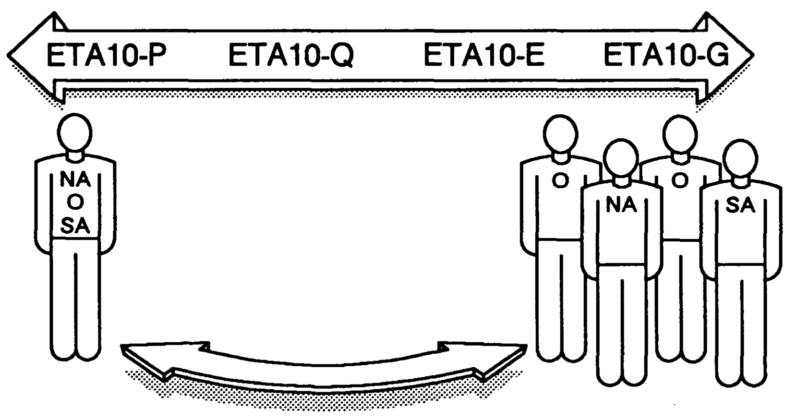

# ETA10 Product Family

The ETA10 product family consists of four supercomputer models differing primarily in the clock cycle time and cooling of their central processors. The ETA10-P and ETA10-Q models are single cabinet systems, configured with either one or two air-cooled processors.

Performance and super-cooled processors distinguish the ETA10-E and ETA10-G models — their processors are housed in cryostats containing liquid nitrogen. The E model has up to eight processors. At the top of the product line is the ETA10-G, which also may be configured with as many as eight processors. An eight-processor ETA10-G has 27 times the peak MFLOPS performance of a single processor ETA10-P.

The following table summarizes basic characteristics of the four ETA10 models.

| | ETA10 P | ETA10 Q | ETA10 E | ETA10 G |
|---|---|---|---|---|
| Peak Performance (MFLOPS) | 750 | 947 | 6,858 | 10,286 |
| Processors | 1 to 2 | 1 to 2 | 1 to 8 | 2 to 8 |
| Shared Memory (megabytes) | 64 to 512 | 64 to 512 | 256 to 2000 | 512 to 2000 |
| Total Central Processor Memory (megabytes) | 32 to 64 | 32 to 64 | 32 to 256 | 64 to 256 |
| Cycle Time (nanoseconds) | 24.0 | 19.0 | 10.5 | 7.0 |
| Cooling Medium | air | air | liquid nitrogen | liquid nitrogen |
| Maximum Number of I/O Units | 3 | 4 | 9 | 18 |

## Software Compatibility on Every ETA10

The continuity of architecture across the ETA10 models means that the same software is usable on every model in the product family — no software retooling is required when customers upgrade. Customer applications are totally portable from the entry level ETA10-P through to the ETA10-G. When customers move up the ETA10 product line, they take their applications with them.

## Cooling Options for Speed

In the ETA10 family there are two methods of cooling central processors, air-cooling and super-cooling. The P and Q models are air-cooled: their central processors and memory are cooled with ambient air. Air-cooling makes these models cheaper to buy, install, and operate. In the E and G models, the central processors are super-cooled to double their speed. The super-cooled central processors are immersed in liquid nitrogen and operate at 77 degrees Kelvin/–320 degrees Farenheit/–196 degrees Celsius.

## Special ETA Technology

ETA Systems has developed state-of-the-art circuit board and chip technologies to build very fast, very reliable central processors. The central processor itself is a 0.25 inch (.64 cm) thick 44-layer printed circuit board, measuring 16.5 by 22.5 inches (42 by 57 cm). A single board replaces over one and one-half miles of what would be manually installed wiring in other systems. The processor is fast and reliable because of the high level of circuit integration and resulting short signal paths. Processor components such as adders and registers are built from the ETA set of special 20K gate array chips. Central processor and major interface boards populated with these CMOS chips have two advantages: they function with low operating costs due to low CMOS power requirements, and they last longer because CMOS chips generate relatively little heat.

## System Upgrades Are Available

Two types of upgrade are available. Customers upgrade either by moving up to a model with a faster clock cycle, or by remaining at the same cycle speed and adding central processors, I/O units, or expanding shared memory. Most upgrades are field-installable.

## Redundancy Options

An ETA10 set up with redundant components quickly resumes productive operations when hardware problems occur. Additional central processing units, service unit server nodes, I/O units, and second units of shared memory, communication buffer, and memory interfaces can upgrade a non-redundant ETA10 system to a redundant configuration. Redundancy options are all field-installable.

# The ETA10 Support Envelope

ETA System's support envelope is a set of three components that work together to support the customer's ETA10 system and particular site needs.

The three components – customer support, training, and documentation – are designed so each complements and overlaps the others:

Customer support activity begins with site planning and includes coordination of the delivery, installation, and start-up of the system. Support for ETA10 operations continues on a day-to-day basis.

Training can be tailored to individual customer needs so personnel take only the classes they need and come up to speed quickly.

Documentation is available for site planning, supports all training classes, and includes a full set of user reference and how-to manuals.

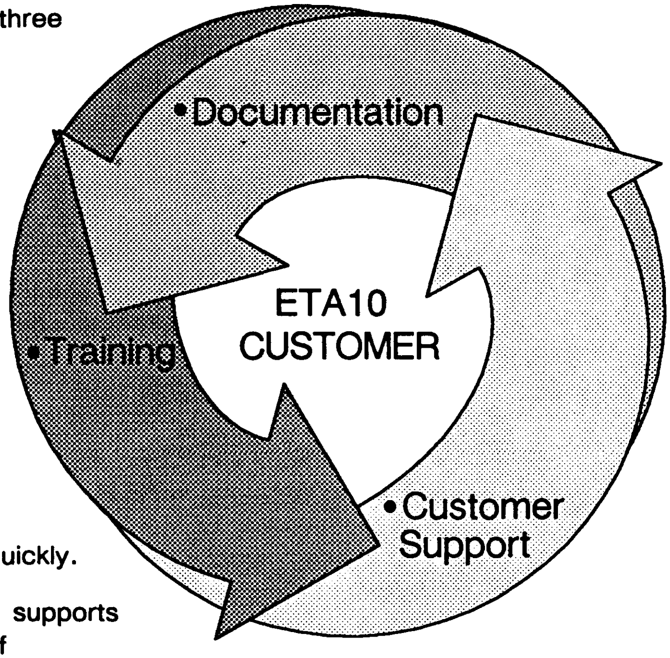


## Support Envelope: Customer Support

To assure customer satisfaction, each ETA10 system is supported by:

- a multilevel support organization
- 24 hour-a-day problem reporting/resolution capability through the Remote System Support (RSS) center
- Control Data's worldwide support resources

### Multilevel Support for the ETA10

The multilevel support organization provides ETA10 customer support through:

- *the account manager*
  A customer's continuing support resource is their account manager. An account manager is assigned when the contract is finalized and remains active throughout the life of the account.
- *remote support*
  The Remote System Support center is the foundation of ETA10 support, providing 24 hour a day dial-in hardware and software troubleshooting and problem-resolution services.
- *local or on-site analyst and customer engineer support*
  Support is available at the customer's facility through on-site systems and support analysts, or through local on-call analysts and customer engineers

- *the technical and management problem-resolution process*
  This process brings in the additional technical resources and management personnel required to resolve the problem.

### Remote System Support Center

The RSS center offers full hardware and software support to ETA10 sites, and assists on-site support personnel as needed.

RSS is staffed with system support specialists available 24 hours a day, 7 days a week. When authorized by the customer, RSS connects to the data link opened into the ETA10's service unit and performs remote troubleshooting functions — diagnostic execution, status checks, and problem isolation. If required, RSS dispatches customer engineers to sites and supervises their activities over a telephone connection. RSS has access to internal ETA10 hardware/software engineering databases, and uses them as an additional support resource as needed.

### On-Site Support Personnel

Super-cooled ETA10 systems typically have an analyst-in-charge (AIC) on site who directs all service and support activities. The AIC carries out traditional analyst functions such as software maintenance, dump analysis, network administration and assistance, and software update installation. Additional tasks include limited hardware maintenance, resolution of software problems, and working with RSS to diagnose hardware/software problems. Customers with air-cooled systems may contract to have either on-call or on-site analyst support.

Customer engineers (CE) are called to a site by the AIC, or are sent by RSS, and work under their guidance and direction. CEs perform preventive and corrective maintenance on ETA10 hardware as well as on a variety of peripheral devices.

### Ease of Problem Reporting

Help is always available to ETA10 customers and site analysts via a telephone call to Control Data's Customer Service Support center. This call connects ETA10 customers and site analysts to the RSS center. RSS troubleshoots and isolates the problem, and when needed, provides corrective code or dispatches a customer engineer to the site. Additional help for software problems is available through Control Data's SOLVER system, a dial-up interactive program/problem reporting and tracking database.

## Support Envelope: ETA10 Documentation

The ETA10 documentation set provides manuals for system usage and system support. Usage manuals meant for programmers and applications users include guides for the operating system as well as for programming language compilers and preprocessors, program development utilities, and other software applications. Support manuals, including operations and administrative references and guides, meet the needs of managerial and other support staff of an ETA10 system.

Documentation and training are coordinated to ensure that customer needs are met. Typically, user documents are used in conjunction with classes. As an example, this overview provides a general, system-level description of the ETA10 for prospective or new ETA10 customers, as well as for anyone just interested in the ETA10 supercomputer, and it also serves as a reference for the *Introduction to the ETA10* training class.

### Documentation for ETA10 Customers

Customer documentation can be divided into support and usage documentation, as shown below. The manuals in this chart are used to suggest the types of manual in each division, and may not be actual titles. A complete list and description of available ETA10 documentation is contained in the *Pricing and Policy Communicator*, available from ETA Systems.

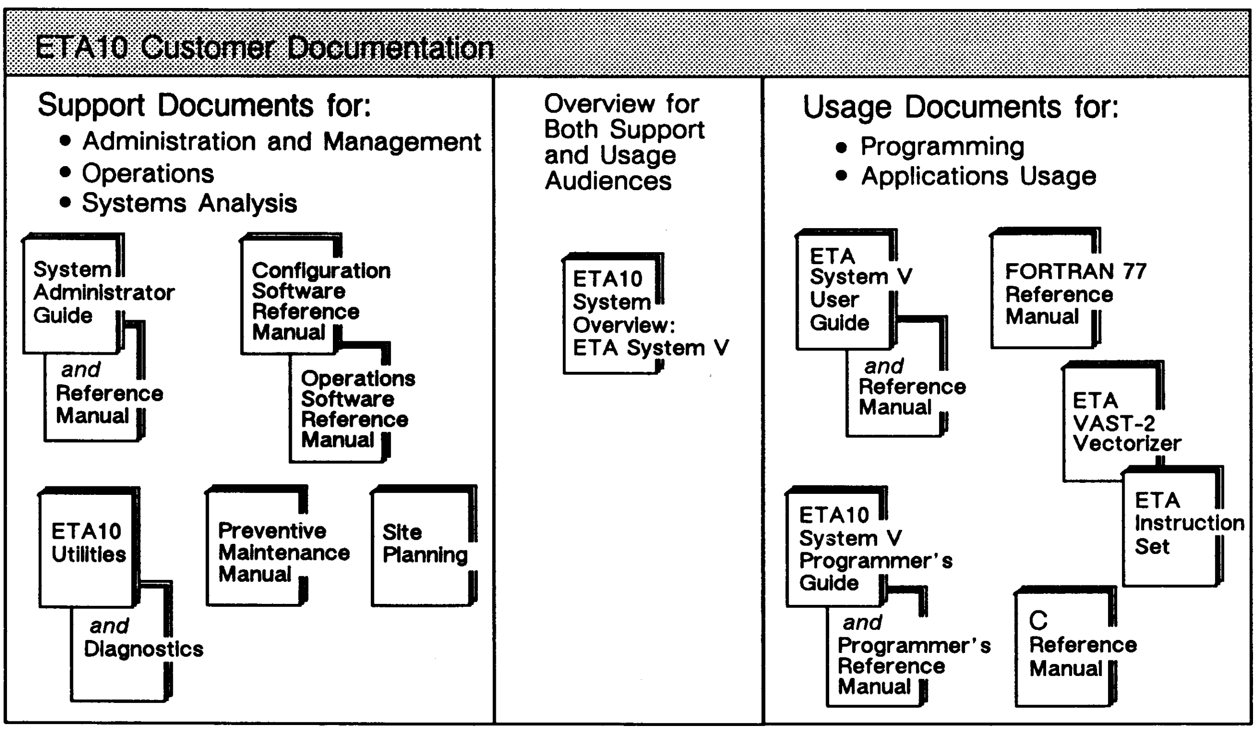

## Support Envelope: ETA10 Training

The ETA10 is supported by a wide range of formal training, with courses oriented to the needs of site staff and users. All classes are developed and taught by ETA instructors.

### What Is Available

The following diagram shows the core curriculum offered to site staff and users.

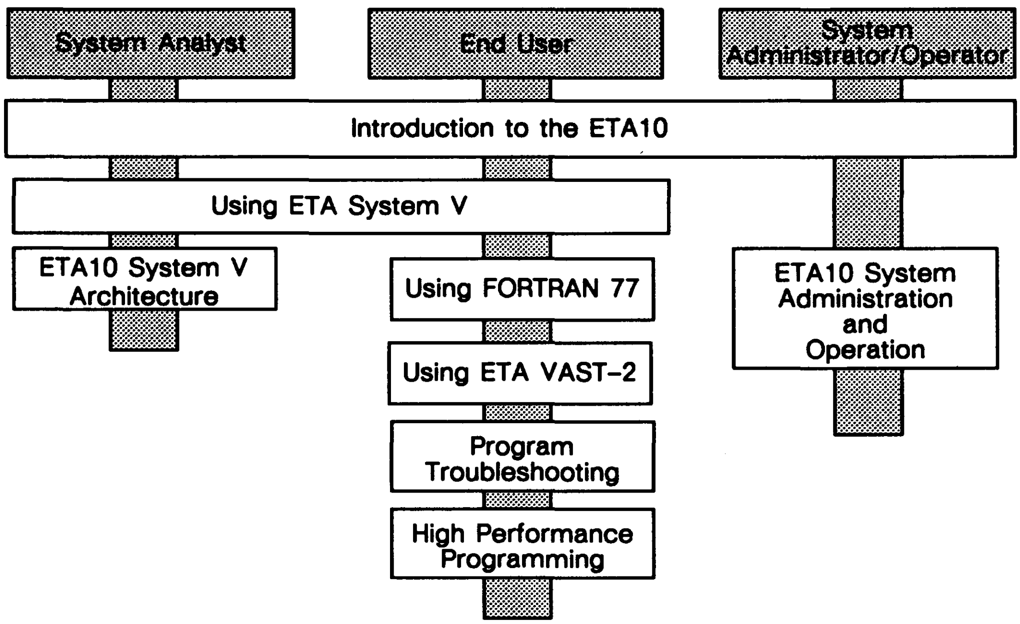


### How Training Is Integrated with Documentation

ETA10 user documentation and customer training classes have been integrated to serve the programming, operations, and administration activities of ETA10 users. Classes are based upon and use standard ETA10 documentation supplemented with special training materials. In this way, students use the same manuals in class that they will use in their work.

### Locations for Training

Courses are taught at customer sites, at U.S. and international facilities of Control Data Corporation, and at ETA Systems headquarters in St. Paul, Minnesota.

Classes are modular and can be customized to accommodate the individual needs of a customer's staff and users. Most classes are one or two days in length.

# Acronym Guide

| | |
|---|---|
| **ANSI** | American National Standards Institute |
| **ARP** | Address Resolution Protocol |
| **AT&T** | American Telephone and Telegraph company |
| **B.E.S.T.** | Built-In Evaluation and Self-Test (on-chip test logic) |
| **CDC** | Control Data Corporation |
| **COFF** | Common Object File Format |
| **CMOS** | Complementary Metal Oxide Semiconductor |
| **CPU** | Central Processing Unit |
| **DoD** | Department of Defense |
| **FORTRAN** | Formula Translation (programming language) |
| **ICMP** | Internet Control Message Protocol |
| **IEEE** | Institute of Electrical and Electronics Engineers |
| **I/O** | Input/Output |
| **IOU** | Input/Output Unit |
| **ISO/OSI** | International Standards Organization for Systems Interconnection |
| **LAN** | Local Area Network |
| **MFLOPS** | Millions of Floating Point Operations per Second |
| **NFS** | Network File System |
| **POSIX** | Portable Operating System for Computer Environments |
| **RPC** | Remote Procedure Call |
| **RSS** | Remote System Support |
| **SECDED** | Single-bit Error Correction Double-bit Error Detection |
| **SVID** | System V Interface Definition |
| **SVVS** | System V Verification Suite |
| **TCP/IP** | Transmission Control Protocol/Internet Protocol |
| **TLI** | Transport Level Interface |
| **UDP** | User Datagram Protocol |
| **VAST** | Vector and Array Syntax Translator |
| **XDR** | External Data Representation |

ETA SYSTEMS

A Control Data Company

**ETA Systems, Incorporated**
1450 Energy Park Drive
St. Paul, MN 55108