System Overview

(CRAY T90™ Series)

HGM-002-C

Cray Research Proprietary

System Overview
-002-C-

Cray Research, Inc.

Any shipment to a country outside of the United States requires a letter of assurance from Cray Research, Inc.

Requests for copies of Cray Research, Inc. publications should be directed to:

CRAY RESEARCH, INC.
Customer Service Logistics
1100 Lowater Road
P. O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Comments about this publication should be directed to:

CRAY RESEARCH, INC.
Service Publications and Training
890 Industrial Blvd.
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

CRAY T90 Series System Overview

HGM-002-C
Last Modified: October 1996

Figures

Tables

Record of Revision

May 1995. Original printing.

Revision A, June 1996. Quadrant and module stack number changes were made to the CRAY T932 chassis map and a separate CRAY T916 chassis map was added.

Revision B, August 1995. Quadrant and module stack number changes were made to the CRAY T916 and CRAY T932 chassis maps.

Revision C, October 1996. Added CRAY T90 series enhancement information: CPE1, CM03, IO02, Scalable I/O, and GigaRing channel, along with various other new and changed pieces of information.

Purpose of this Document

This document introduces new technologies and components used in CRAY T90 series computer systems. It provides a basic summary of many subjects, most of which are described in more detail in other documents that are part of the CRAY T90 series hardware and maintenance documentation set. This document is used in the overview portion of the hardware training course.

You should read through this overview document to obtain a general understanding of the CRAY T90 series systems before you read the detailed documents in the set. Note that the CRAY T90 series hardware and maintenance documentation set is also available online; refer to the “Online and CD-ROM Documentation” section near the end of this document for more information.

CRAY T90 Series Systems

The CRAY T90 series computers are powerful, large-scale, general purpose parallel-vector supercomputer systems. They deliver high processing rates by using multiple CPUs, a large common memory, and a fast 2.27-ns/440-MHz clock speed.

There are three models of CRAY T90 series supercomputers: the CRAY T94, CRAY T916, and CRAY T932 systems. The numbers that follow the numeral 9 in the model number identify the maximum number of CPUs in each system.

CRAY T90 series computers use two types of chassis. The CRAY T94 system uses a small chassis; the CRAY T916 and CRAY T932 systems use a larger chassis. Both types of chassis house various configurations of logic and memory modules, as well as power supplies and other components. Both chassis contain tanks in which these components are submersed in a dielectric coolant. Figure 3 shows the small chassis and Figure 4 shows the large chassis. The large chassis contains a physical boundary that divides the tank into halves.

CRAY T90 series systems can easily be scaled up or down by physically adding or removing central processor and memory modules. Even though the CRAY T916 and CRAY T932 systems use the same basic chassis, the CRAY T916 system cannot be upgraded to a CRAY T932 system because the two systems use different system interconnect boards (SIBs), which cannot be exchanged on-site.

CRAY T90 series computer systems that use Cray Floating-point CPUs have two modes of operation: C90 compatibility mode and Triton mode (TRI bit in the exchange package). Systems that use IEEE floating-point CPUs (CPE1) run only in Triton mode; the TRI mode bit is removed from these systems. The A, B, and shared B registers in CRAY T90 series systems were expanded to 64 bits, and several new instructions were added to take advantage of other architectural differences. Some software programs written for earlier Cray Research systems may need to be recompiled before they can run in Triton mode. CRAY T90 series systems with IOS-E subsystems use the IO01 module. The IO02 module is used with GigaRing channel subsystems. More details are provided later in this document.

CRAY T94 System

The CRAY T94 computer system (illustrated in Figure 1) consists of a variety of standard and optional equipment. Cray Research can configure each computer system to meet customer needs and requirements.

A standard CRAY T94 computer system consists of the following components:

- Mainframe chassis (1 to 4 CPUs)
- 64 or 128 Mwords of memory
- I/O subsystems and maintenance platforms:
  - (with IO01)

    Combined IOS and SSD (600-, 700-, or 800-series chassis). Most IOS/SSD systems are shipped either the 600- or the 800-series chassis. The support system includes a support system chassis, MWS, OWS, and laser printer.
  - (with IO02)

    Scalable I/O: Single or multicabinet SIO cabinet (PC-10B). Support for up to 8 GigaRing high-speed I/O interconnects. I/O nodes that provide GigaRing connections for network, disk, and tape systems. The support system is a system workstation (SWS); a laser printer is optional.
- HEU-T90 heat exchanger unit
- HVDC high-voltage DC device (40 kW)
- Optional UPS and additional battery cabinets
- Optional remote support spares kit
  - Domestic sites: Telebit NetBlazer PN2, 28.8-Kbit/s modem, and cables
  - International sites: Telebit NetBlazer PN2 and cables
- Disk drives and other peripheral equipment and network interfaces

The components in the CRAY T94 systems, except disk drives and peripherals, are described in more detail later in this document.

The number of disk drives and other peripheral equipment used with the CRAY T94 computer system depends on individual customer needs. The disk drive configuration for a CRAY T94 computer system usually includes DE-60 and DE-100 disk enclosures or ND-12 and ND-14 ND Series Network Disk Array systems.

Figure 1. CRAY T94 IO01 Computer System

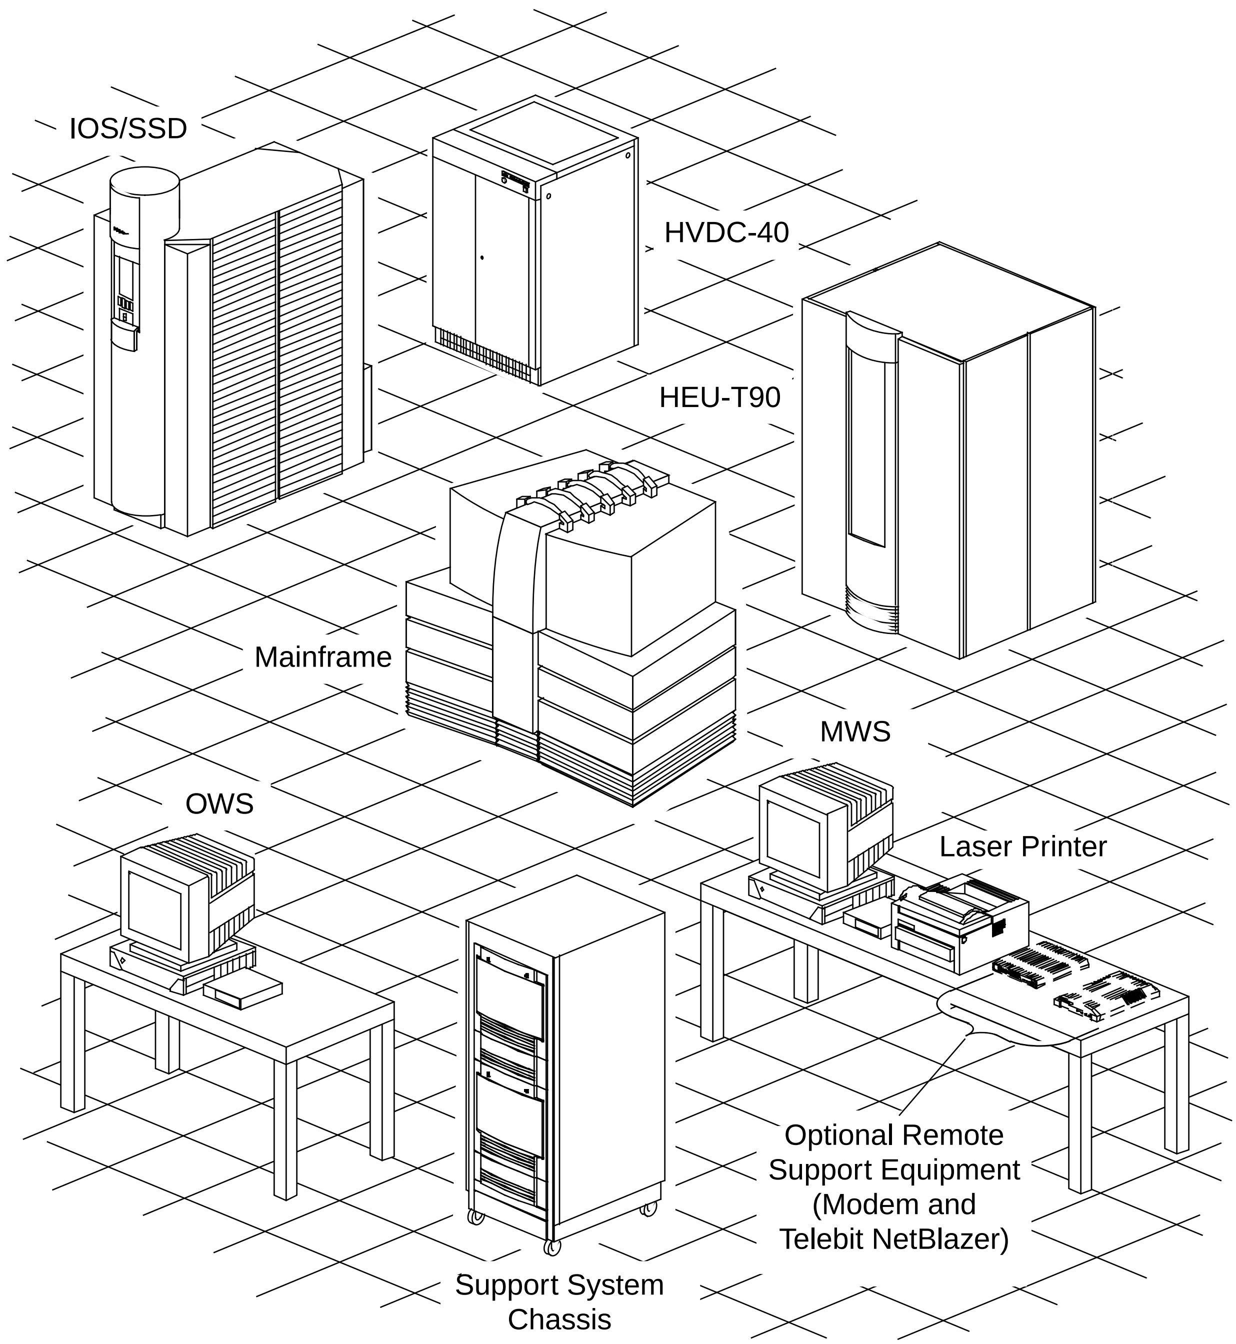

Figure 2. CRAY T94 IO02 Computer System

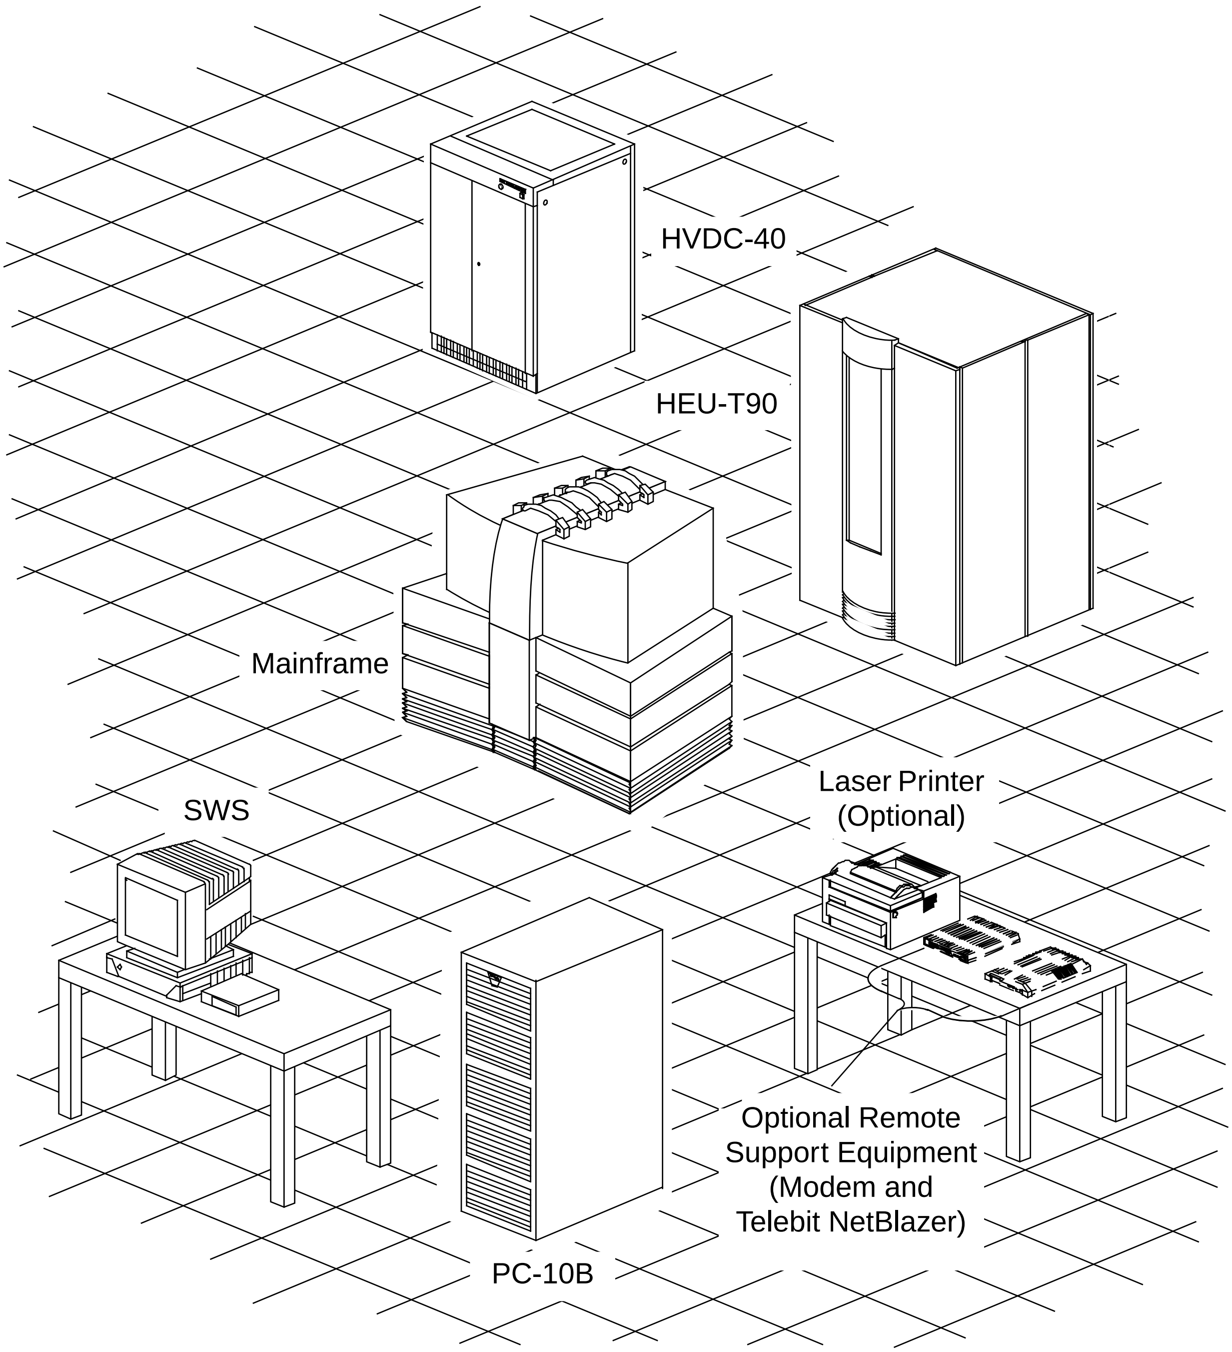

CRAY T916 System

A standard CRAY T916 computer system consists of the following components:

- Mainframe chassis (4 to 16 CPUs)
- 256 or 512 Mwords of memory
- I/O subsystems and maintenance platforms:
  - (with IO01)

    One or two IOS-Es (600-, 700-, or 800-series chassis)
    Most systems will ship with a 700-series chassis. (Each 700-series IOS/SSD-E chassis requires a dedicated HEU.) The support system includes a support system chassis, MWS, OWS, and laser printer.
  - (with IO02)

    Scalable I/O: Single or multicabinet SIO cabinet (PC-10B). With scalable I/O subsystems, the CRAY T916 mainframe uses one or two IO02 modules. Scalable I/O supports up to 16 GigaRing high-speed I/O interconnects (using two IO02 modules). It also provides I/O nodes for system, network, disk, and tape connections. A system workstation (SWS) is used to maintain and monitor the system; a laser printer is optional.
- HEU-T90 heat exchanger unit
- HVDC high-voltage DC device (160 kW)
- Optional UPS and additional battery cabinets
- Optional remote support spares kit
  - Domestic sites: Telebit NetBlazer PN2, 28.8-Kbit/s modem, and cables
  - International sites: Telebit NetBlazer PN2 and cables
- Disk drives and other peripheral equipment and network interfaces

In the CRAY T916 system, the system modules occupy one half of the chassis (refer to Figure 4). The other half remains empty, with no power or cooling connections.

CRAY T932 System

A standard CRAY T932 computer system consists of the following components:

- Mainframe chassis (8 to 32 CPUs)
- 512 Mwords or 1 Gword of memory
- I/O subsystems and maintenance platforms:
    - (with IO01)

      One to four IOS-Es (600-, 700-, or 800-series chassis
      Most systems will ship with one to two 700-series chassis. Each 700-series IOS/SSD-E chassis requires a dedicated HEU.) The support system includes a support system chassis, MWS, OWS, and laser printer.
    - (with IO02)

      Scalable I/O: Single or multicabinet SIO cabinet (PC-10B). With scalable I/O subsystems, the CRAY T932 mainframe must use two IO02 modules (to support boundary scan), but it may include a maximum of four IO02 modules. Scalable I/O supports up to 32 GigaRing high-speed I/O interconnects (using four IO02 modules). It also provides I/O nodes for system, network, disk, and tape connections. The support system is a system workstation (SWS); a laser printer is optional.
- Two HEU-T90 heat exchanger units
- Two HVDC high-voltage DC devices (160 kW)
- Optional UPS and additional battery cabinets
- Optional remote support spares kit
    - Domestic sites: Telebit NetBlazer PN2, 28.8-Kbit/s modem, and cables
    - International sites: Telebit NetBlazer PN2 and cables
- Disk drives and other peripheral equipment

In the CRAY T932 system, the two halves of the chassis have separate power and cooling connections. This enables personnel to power off one half of the chassis and drain the dielectric coolant, while the other half continues to operate.

Figure 3. CRAY T94 Mainframe Chassis

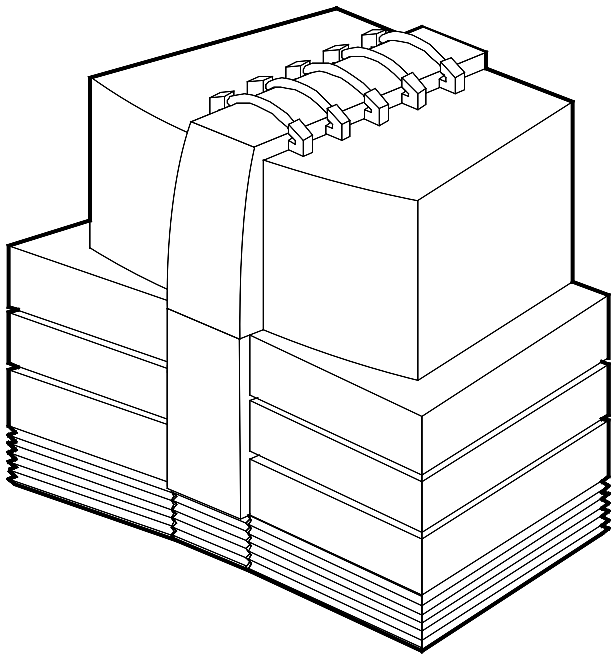

Figure 4. CRAY T916 and CRAY T932 Mainframe Chassis

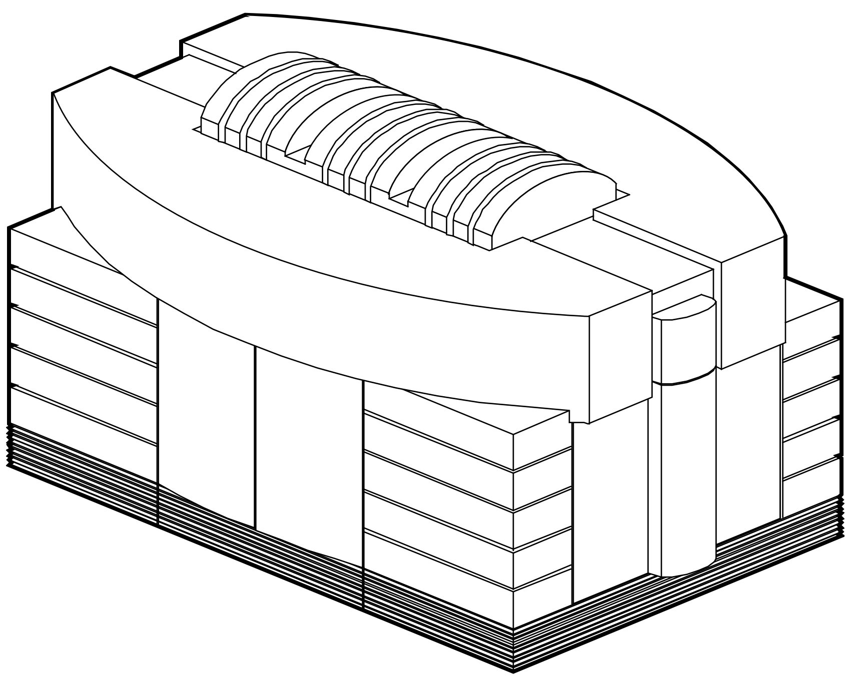

Chassis Map

Figure 5, Figure 6, and Figure 7 show chassis and module maps for CRAY T90 series mainframes that use IO02 modules.

Figure 5. CRAY T94 IO02 Chassis and Module Map

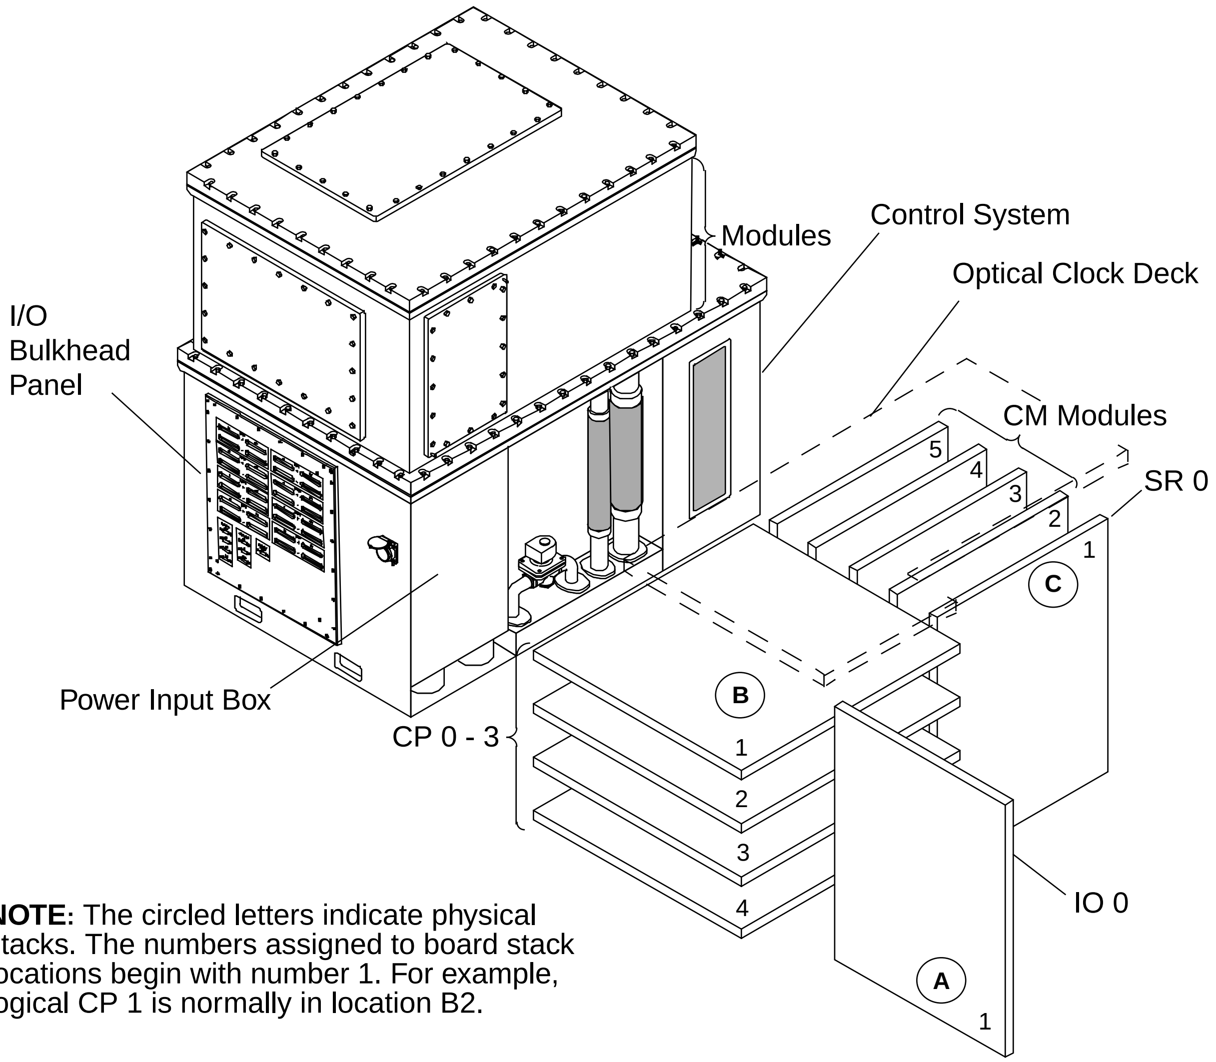

NOTE: The circled letters indicate physical stacks. The numbers assigned to board stack locations begin with number 1. For example, logical CP 1 is normally in location B2.

Physical Stack	Module Type	Module Number
A1	I/O	IO 0 (to CPs 0, 1, 2, 3)
B1 – B4	CPU	CPs 0 – 3 (CP 0 is normally on the top of the stack.)
C1	Shared	SR 0
C2 – C5	Common memory	CMs 1 – 4

Figure 6. CRAY T916 IO02 Chassis and Module Map

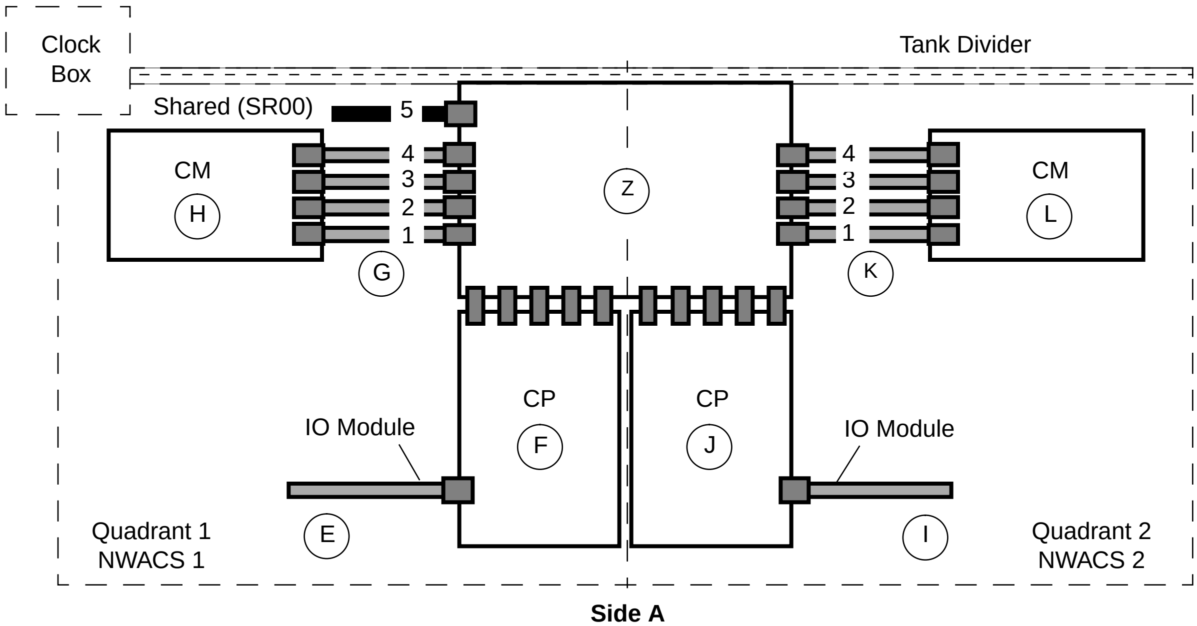

Physical Stack	Module Type	Module Number
E	I/O	IO 1 (to CPs 12, 13, 14, 15)
F	CPU	CPs 10 – 17
G	Networks and Shared	NWs 4 – 7 and SR 0
H	Memory	CMs 10 – 17 (Sections 2, 3, 4, 5)
I	I/O	IO 0 (to CPs 2, 3, 4, 5)
J	CPU	CPs 0 – 7
K	Network	NWs 0 – 3
L	Memory	CMs 0 – 7 (Sections 0, 1, 6, 7)
Z	System Interconnect Boards	SJs 1 – 8

Figure 7. CRAY T932 IO02 Chassis and Module Map

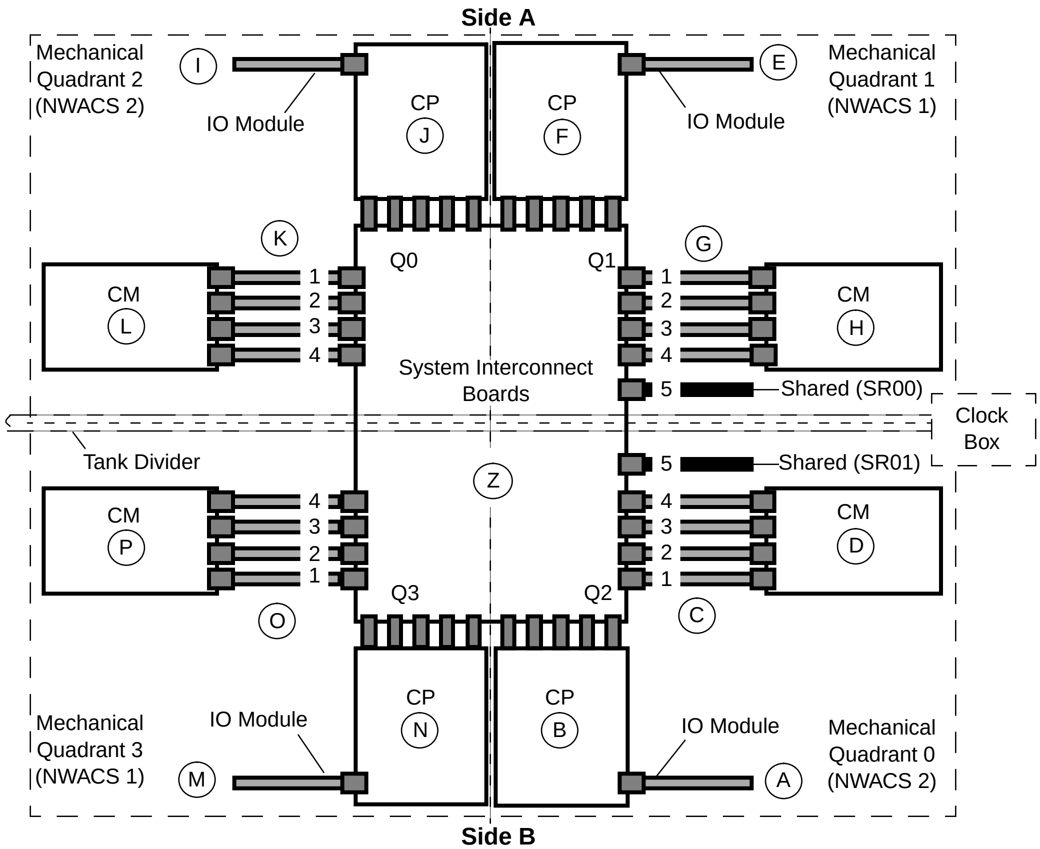

Physical Stack	Module Type	Module Number
A	I/O	IO 2 (to CPs 22, 23, 24, 25)
B	CPU	CPs 20 – 27
C	Networks and Shared	NWs 10 – 13 and SR 1
D	Memory	CMs 20 – 27 (Sections 2, 3)
E	I/O	IO 1 (to CPs 12, 13, 14, 15)
F	CPU	CPs 10 – 17
G	Networks and Shared	NWs 4 – 7 and SR 0
H	Memory	CMs 10 – 17 (Sections 4, 5)
I	I/O	IO 0 (to CPs 2, 3, 4, 5)
J	CPU	CPs 0 – 7
K	Network	NWs 0 – 3
L	Memory	CMs 0 – 7 (Sections 6, 7)
M	I/O	IO 3 (to CPs 32, 33, 34, 35)
N	CPU	CPs 30 – 37
O	Network	NWs 14 – 17
P	Memory	CMs 30 – 37 (Sections 0, 1)
Z	System Interconnect Boards	SIs 1 – 8

Module Assembly Components

Figure 8 illustrates the central processor (CP) module components. This illustration shows the basic components that are part of all mainframe modules. The sizes of various components differ among modules.

Figure 8. CP Module Assembly Components

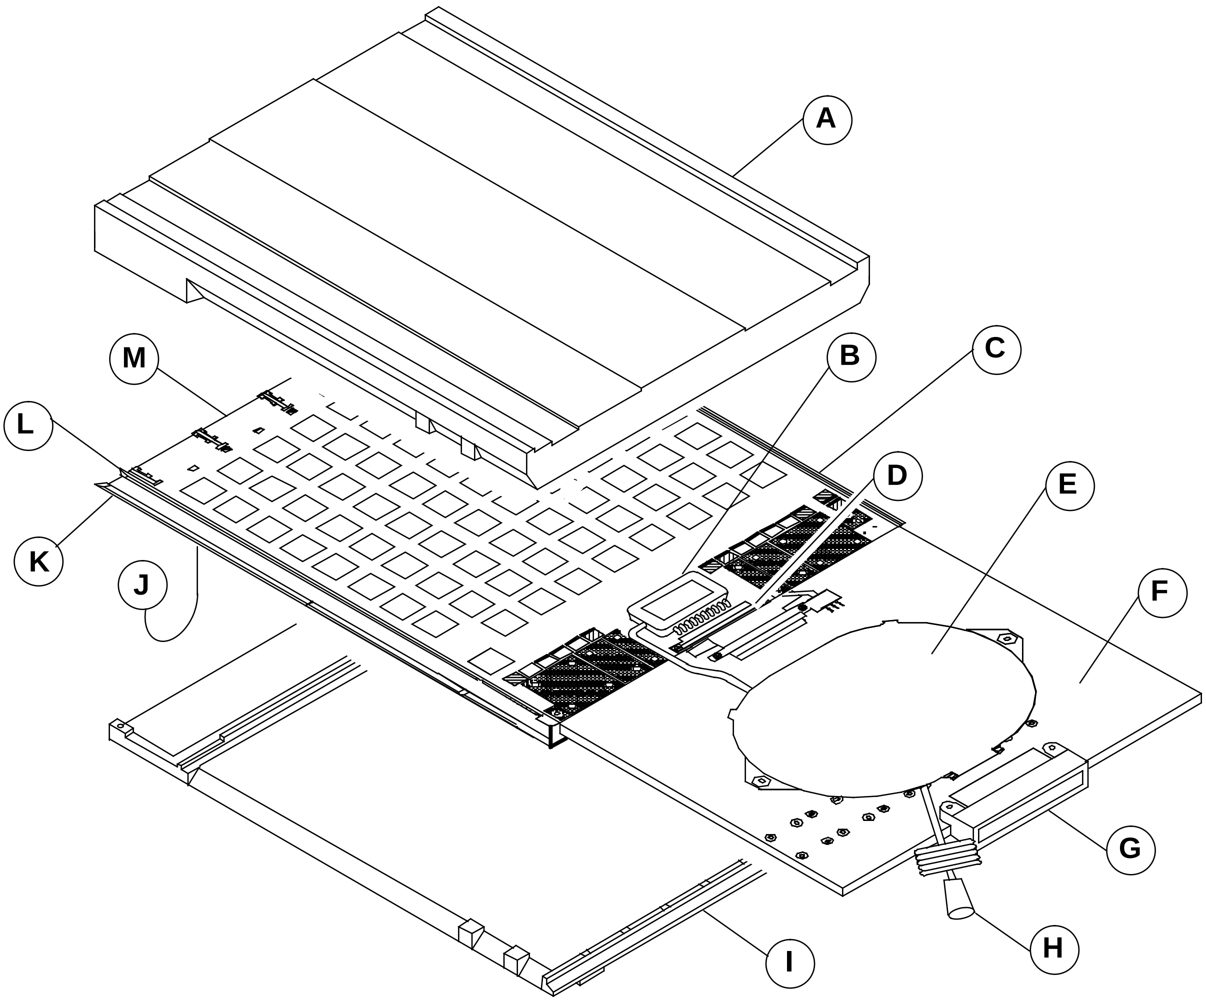

A	Flow Block, Board 1	**H**	Fiber-optic Coupler
B	Optical Receiver	**I**	Flow Block, Board 2
C	PC Board Edge Shim	**J**	PC Logic Board 2
D	Maintenance Connector Flex Assembly	**K**	Outer Rail
E	Fiber-optic Spool Assembly	**L**	Inner Rail
F	Voltage Regulator Board Assembly	**M**	PC Logic Board 1
G	Maintenance Connector		

Module Types

A CRAY T90 series system includes the following six types of logic modules:

- Central processor modules (CP02 and CPE1)
- Central memory modules (CM02 and CM03)
- I/O modules (IO01 and IO02)
- Shared module (SR)
- Network module (NW)
- Boundary scan module (BS, used with IO01, not IO02)

Each CRAY T90 series module consists of one printed circuit (PC) board. Unlike CRAY C90 style modules, the CRAY T90 series PC boards are not mounted on a coldplate. Because the modules are submersed in dielectric coolant, module coldplates are not needed. The dielectric coolant flows around the PC boards.

Central Processor (CP) Modules

The CP02 and CPE1 modules each contain one CPU. The CRAY T90 series system can contain from 1 to 32 CPUs, depending on the chassis type. CP modules are arranged in a stack of up to four modules in a CRAY T94 chassis. CP modules are arranged in stacks of up to eight modules in CRAY T916 and CRAY T932 chassis. The CRAY T916 chassis contains two physical stacks, and the CRAY T932 chassis contains four stacks.

CRAY T90 series CPUs have double vector and floating-point functional units. There are eight physical ports connected to the memory interface. Each port can move 1 word per clock period when reading data and 1 halfword per clock period when writing data. The vector pipelines use four logical read ports and two logical write ports. A CPU can transfer 8 data words per clock period to and from memory. The double pipeline enables two elements or operand pairs to access two functional units in 1 clock period. With CP02 modules, the CPU is instruction compatible with the CRAY C90 series instruction set and can run in C90 mode.

Figure 12 shows a block diagram of the CPU.

CPE1: 64-bit IEEE Floating-point CPU Module

The new CPE1 module supports the 64-bit IEEE floating-point arithmetic standard. The new CPE1 module changes functional unit options. The CPE1 CPU provides many enhancements to the CP02 CPU; several new instructions have been added to support the new arithmetic functional units. CPE1 modules are the same size as CP02 modules and are plug compatible in any CRAY T90 series mainframe chassis. CRAY T90 series mainframes can use either CP02 or CPE1 modules or both types of modules in any combination (OS support of this capability is expected in early 1997).

Use of the CPE1 modules requires UNICOS version 9.1 or later and version 2.0 of the Cray programming environments. Compilers that support CRAY T90 series systems with IEEE floating-point format include CF90 Fortran version 2.0, Cray standard C version 5.0, and Cray C++ Compiling System version 2.0. The CF77 programming environment used in previous systems does not function in IEEE floating-point systems.

Customers can gain the following benefits by using the IEEE format:

- An industry standard format for data, which makes data easier to use on other computer systems.
- Greater precision with floating-point numbers by adding 4 bits to the coefficient field.
- The addition of new divide and square root functional units.
- The capability to set floating-point data rounding modes and specify numeric interrupt modes.
- The capability of representing infinity and non-numeric numbers by using special values, such as NaN (not a number), signed zero, and infinity.

This enhancement adds several new instructions that compare and convert numbers in both floating-point and integer formats. The CPE1 module does not support a C90 compatibility mode or an original CRAY T90 series Cray floating-point mode. The CPE01 module does not support the IEEE floating-point standard for 32-bit formats or for denormalized numbers.

Several mode, flag, and status bits have been removed from the exchange package. Other bits have been added in the exchange package and status register 0 to support IEEE floating-point operations. Status registers 1 through 7 (SR1 through SR7) have not changed.

Central Memory Modules (CM02 and CM03)

The CM modules contain central memory that is common to all CP modules. Like the CP modules, the memory modules are arranged in stacks. The CRAY T94 chassis contains one stack of up to four CM modules. The CRAY T916 chassis contains up to two stacks of eight modules, and the CRAY T932 chassis contains up to four stacks of eight modules. Central memory is organized by sections, subsections, and banks. However, CRAY T94 systems do not have subsections.

Each CM module consists of three laminated printed circuit boards. A power distribution board is laminated between two logic boards. The logic boards are populated with memory stacks and logic on one surface and logic options on the other surface. Each CM module has 16 memory stacks (Figure 9). Each memory stack consists of two printed circuit boards (ST02 A and ST02 B) that contain forty 4-Mbyte SRAM chips. The stack has two columns of 20 chips each. Nineteen chips in each column store data. The bottom chip in each column functions as the spare chip for the column.

Figure 9. Memory Stack

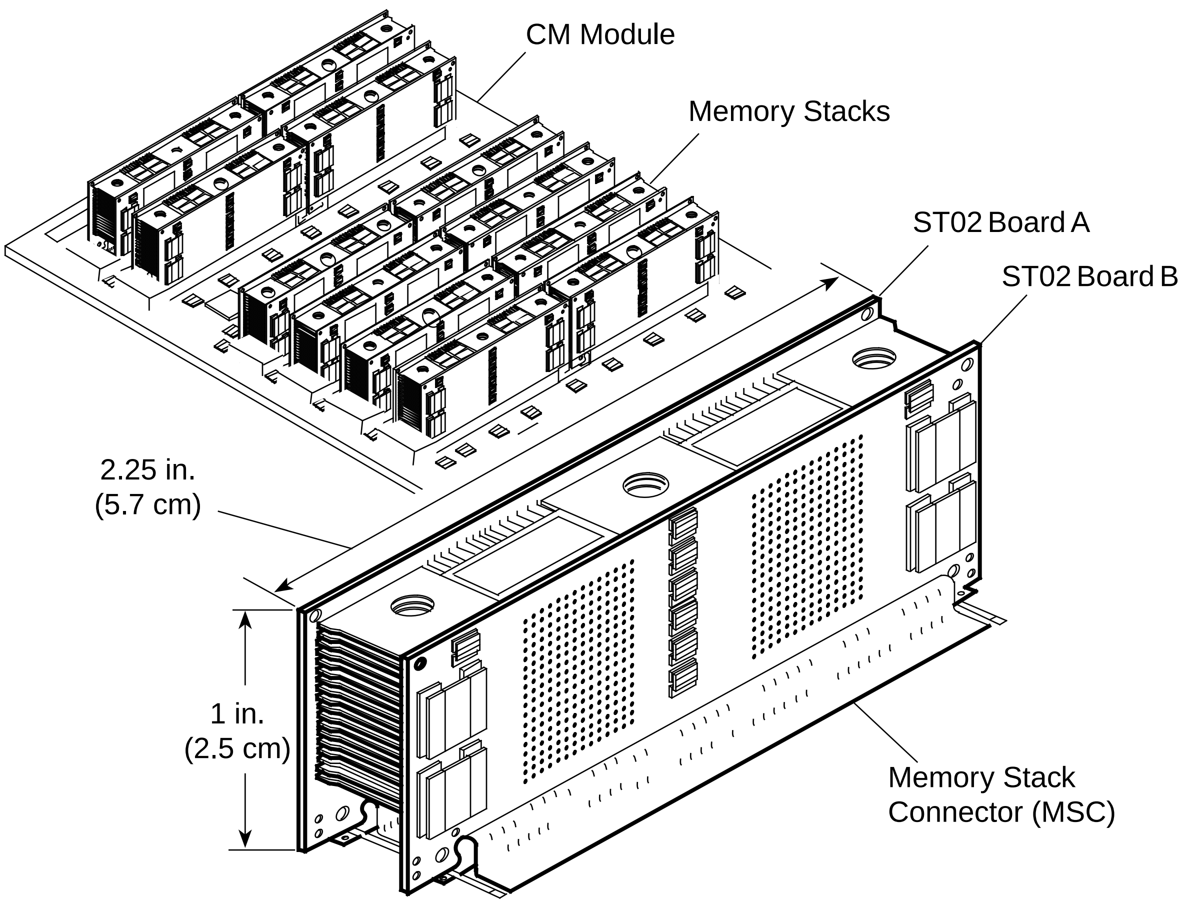

When fully populated with 4-Mbit chips, a CM module contains 32 Mwords of memory. Memory stack connectors (MSCs) connect memory stacks to the memory module. The MSC is an electronic zero insertion force (EZIF) connector that requires the application of electric current to insert or remove the memory stack from the module. The electric current heats a spring in the connector, which forces the connector open.

In CRAY T916 and CRAY T932 systems, each module stack connects to four network modules. The network modules enable any CPU to access any area of common memory. In CRAY T94 systems, a module stack connects directly to the CP modules; therefore, the final level of memory interconnecting and routing takes place on the memory modules.

CM02 modules use a 2-bit byte error-correction scheme and 2-bit wide memory chips so that whole chip failures can be corrected. Failures on different chips can be detected; multiple failures might not be detected. CM03 modules use the same error-correction scheme as CM02 modules, but they use 4-bit wide memory chips. On CM03 modules, failure of a whole memory chip can be detected but not corrected.

CM03 Memory Module

The CM03 memory module is designed as an upgrade to or replacement in CRAY T90 series systems that use the CM02 memory module. The CM03 module is built with faster 2-stage, 4-Mbit synchronous SRAM (SSRAM) chips.

The CM03 module offers the same storage capacity as a CM02 module. However, a CM03 memory system provides twice the number of memory banks as a CM02 system and a bank busy time of 4 cycles versus 7 cycles. The additional memory banks and reduced bank busy time of a CM03 system results in an overall bandwidth increase of 3.5 times that of a CM02 memory system. CM03 modules use the same error-correction scheme as CM02 modules, but they use 4-bit wide memory chips. On CM03 modules, failure of a whole memory chip can be detected but not corrected.

CM03 memory offers the greatest enhancement in CRAY T90 series systems that are over half populated with CPUs and that run highly vectorized, highly memory-intensive applications. CRAY T90 series systems with fewer CPUs and those running applications that do not contain high vector content are not likely to experience any performance improvements. Use of CM03 memory modules is supported by UNICOS version 9.0 or later.

I/O Modules (IO01 and IO02)

The IO01and IO02 modules provide an interface between CP modules and external channels. In CRAY T916 and CRAY T932 systems, the IO module connects directly to four of the eight modules in a CP stack. Table 1 lists the maximum number of IO modules and channels that each CRAY T90 series system supports. The term *LOSPX* refers to an error logger channel, a maintenance channel, a support channel, or an MCU channel on an IO01 module.

Table 1. Maximum IO01 Module and Channel Configurations

Model	IO01 Modules	LOSPX Channels	LOSP Channels	HISP Channels	VHISP Channels
CRAY T94 System	1	4	8	8	4
CRAY T916 System	2	8	16	16	8
CRAY T932 System	4	16	32	32	16

Table 2. Maximum IO02 Module and Channel Configurations

Model	IO02 Modules	GigaRing Channels
CRAY T94 System	1	8
CRAY T916 System	2	16
CRAY T932 System	4	32

IO01 Module

Each IO01 module is divided into four quadrants.   Each quadrant is logically connected to 1 of the 4 CPUs that physically connect to the IO module. Each quadrant on the module can handle transfers for the following channels:

- Two LOSP channels
- Two HISP channels
- One VHISP channel
- One LOSPX channel

The location of LOSPX channels on an IO module are as follows: maintenance channel on quadrant 0; support channel on quadrant 1; error logger channel on quadrant 2; and MCU channel on quadrant 3.

LOSP, HISP, VHISP, and LOSPX channels connect to the IO module from 1 of 16 “fuzz-button” connectors. Refer to “Fuzz-button and Pogo-pin Connectors” on page 49 for more information and an illustration.

All I/O channels can operate simultaneously, sharing memory bandwidth with individual CPUs. However, only one data channel from each quadrant (LOSP, HISP, and VHISP) can operate simultaneously because each quadrant has only one path to a CPU.

All CPUs can perform data transfers on all channels. Any CPU within the system can request the services of any LOSP or VHISP channel. Control of HISP channels is performed by the I/O system. Because a direct connection between the requesting CPU and receiving IO module may not exist, any I/O command is always initially sent to the SR module. The SR module can either access the IO module or it can send the command to the remote SR module in a CRAY T932 system. The SR module sends the command to an I/O channel through the CP module that lies between the SR module and the IO module where the channel resides.

IO02 Module: Scalable IO and GigaRing Channel Support

The new IO02 module provides a CRAY T90 series interface to scalable I/O (SIO) systems and supports up to 8 GigaRing channels. Boundary scan operations are now performed by the IO02 module. The IO02 module can be used with either CP02 or CPE1 modules.

Existing CRAY T90 series systems can be upgraded in the field from IO01 to IO02 modules. This process requires I/O bulkhead, I/O rack, I/O harness, and maintenance cable changes. A CRAY T90 series IOS-E site must also remove the IOS-E and replace it with an SIO system; the MWS and OWS must also be removed and replaced with a single SWS. There are no plans to support the use of partitioning to run an IOS-E based logical system and a GigaRing logical system in the same chassis.

The IO02 module also contains a Cray Research proprietary supervisory channel interface that provides access for the maintenance channel, the error logger channel, the support channel, and the boundary scan system. The supervisory channel provides a general low-speed connection used for maintaining and monitoring the mainframe. An Ethernet connection from the system workstation (SWS) through a multipurpose node (MPN) in the SIO chassis to a mainframe I/O bulkhead attaches the workstation to the supervisory channel.

The supervisory channel runs on an industry standard SBus platform that makes it adaptable to a variety of environments such as multipurpose nodes (MPNs) and SBus-based workstations (only MPN is supported at the time of this printing).

A GigaRing channel consists of a pair of unidirectional, counter-rotating rings. The rings contain nodes that allow clients to connect to the GigaRing channel. The GigaRing channel implements a packet-based communication protocol. GigaRing node chips implement the logical-layer protocol of the GigaRing channel and support the I/O protocol.

Eight 514-pin I/O connectors reside on the module: four on each side; Figure 10 shows the connectors on IO02 board 1. Six of these connectors route GigaRing channel traffic to the eight GigaRing node chips. (Each GigaRing node chip uses 2/3 of an I/O connector.) The remaining two I/O connectors are used for the supervisory channel and boundary scan testing.

Figure 10. IO02 Board 1 with GigaRing Node Chips

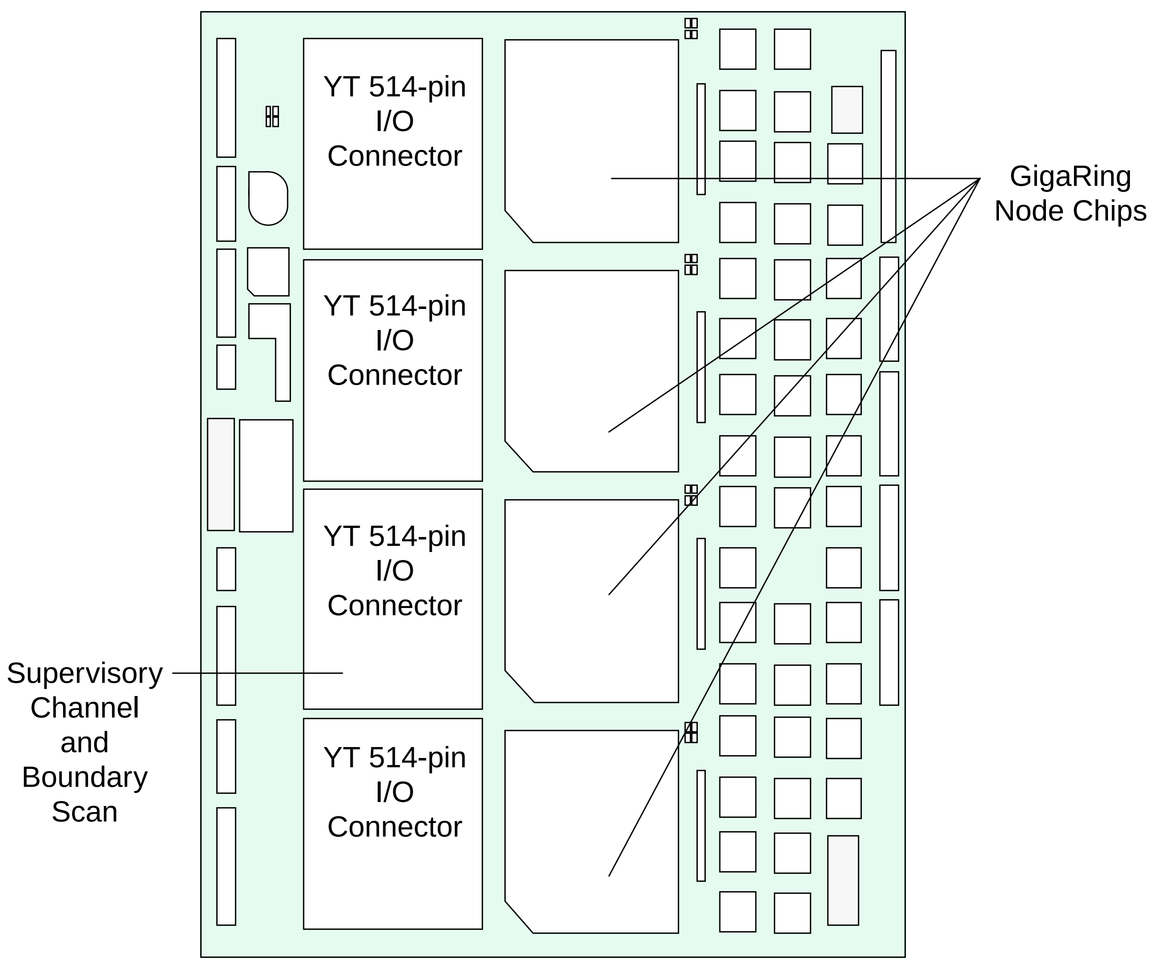

Shared (SR) Module

Each shared module contains 18 clusters of shared registers. Each cluster contains the following shared resources: 16 SB registers, 16 ST registers, and 64 semaphore registers. CRAY T94 and CRAY T916 systems each contain one shared module; CRAY T932 systems contain two shared modules (36 clusters of shared registers). A shared module contains all of the shared register logic, some of the I/O control logic, and provides data paths for all maintenance and configuration information.

The SR module functions as a central point for I/O command and status routing. Because a direct connection does not always exist between the CPU making a request and the IO module receiving it, all I/O commands are always sent first to the SR module. The SR module forwards the request information to the IO module that services the requested channel. When the I/O interrupt occurs, it is sent to the SR module. The SR module determines which CPU will handle the interrupt.

The SR module acts as the point of control for configuration of the sanity tree and error logger. Status information is returned to the SR module, and the SR module forwards it to the requesting CPU.

Network (NW) Module

The network module provides an interface between CM and CP modules. There are four network modules per stack of CM modules. Each network module is capable of steering data from eight sources (all the CP modules in one stack) to eight destinations (all the physically connected CM modules). The network module is also responsible for subsection arbitration. CRAY T94 systems do not contain NW modules.

Boundary Scan (BS) Module (IO01 systems)

The BS module on CRAY T90 series IO01 systems performs the following functions:

- Serves as an interface for all boundary scan operations
- Acts as a substitute maintenance interface (in place of an IO module) in the CPU tester only (only with BS01 module)

CRAY T94 and CRAY T916 systems use one BS module, and CRAY T932 systems use two BS modules, one in each half of the mainframe. In a CRAY T932 system, each BS module connects to 44 modules (88 in the system). Each BS module has connectors to support 48 system modules.

The boundary scan module is not needed in IO02 systems because boundary scan operations are performed by the IO02 module. CRAY T932 systems require two IO02 modules to perform boundary scan operations, one in each half of the mainframe.

System Interconnect Board (SIB)

A CRAY T90 series system has no wire mat. The system interconnect board (SIB) replaces the wire mat that was used in previous Cray Research computer systems. All module interconnections are accomplished with an SIB and two types of EZIF module connectors.

The SIB passes data and control signals between modules that connect to it. CRAY T916 and CRAY T932 computer systems contain eight SIBs; CRAY T94 systems do not contain SIBs.

CRAY T916 and CRAY T932 systems use two different SIB module types: SJ and SI modules. CRAY T916 systems use the SJ module type, which is approximately half the size of an SI module. CRAY T932 systems use the SI module type. The SI is a single module that is present on both sides of the tank divider in order to connect to module stacks on both sides of the system.

Each SIB module consists of two printed circuit boards bonded together. SIBs have no logic components and no interconnections between the two printed circuit boards. The SIB is not a field replaceable unit.

Optical Clock Module

The optical clock module generates and sends an optical clock signal to all modules in the mainframe chassis. The optical clock module consists of two printed circuit boards: an oscillator board and a laser board, as shown in Figure 11. The oscillator board contains the crystal oscillators and integrated circuits. The laser board contains the optical transmitter modules (OTMs), which house the lasers. The optical clock module has its own power supply.

In CRAY T94 systems, the optical clock module mounts on a clock deck located on top of the card cage (Figure 11). The clock deck is a field replaceable unit (FRU) if the coupler fails; however, if the optical clock module fails, then the module becomes the FRU instead of the clock deck. In CRAY T916 and CRAY T932 systems, the optical clock module is located between memory module stack D and memory module stack H and is labeled *Clock*.

Clock Signal Path

Figure 11 shows the path of the optical clock signal. The oscillators create a differential electronic clock (ECL level) signal to drive the OTMs. Four oscillators on the optical clock module provide three different clock frequency levels: 430, 440, and 450 MHz (slow, normal, and fast). The normal clock frequency runs at 440 MHz or 2.27 nanoseconds. Two normal frequency oscillators are available for redundancy purposes.

The OTMs convert the differential electronic clock signal received from the oscillators into an optical signal. The OTMs contain the lasers that send the optical signals off the optical clock module through fiber-optic couplings. Fiber-optic waveguides transfer the optical signals to the optical receiver on mainframe modules. Specifications for OTM lasers used in CRAY T94 systems differ from those used in CRAY T916 and CRAY T932 systems.

Couplers split clock signals coming from an OTM and send them to modules in the mainframe chassis. There are two levels of couplers on each optical clock module. A 2-by-2 coupler is used in CRAY T94 systems; it has two inputs (one from each OTM) and two outputs. This coupler sends the optical clock signal to two 1-by-8 couplers. Each 1-by-8 coupler splits the signal and sends it to a maximum of 8 different system modules. The optical clock signal is sent to 11 logic modules in a CRAY T94 system.

The optical receiver on each module converts the optical signal back to an electrical signal and sends it to a clock distribution option on the module. The clock distribution option then distributes the electrical clock signal to all options on the module.

Figure 11. Optical Clock Signal Path

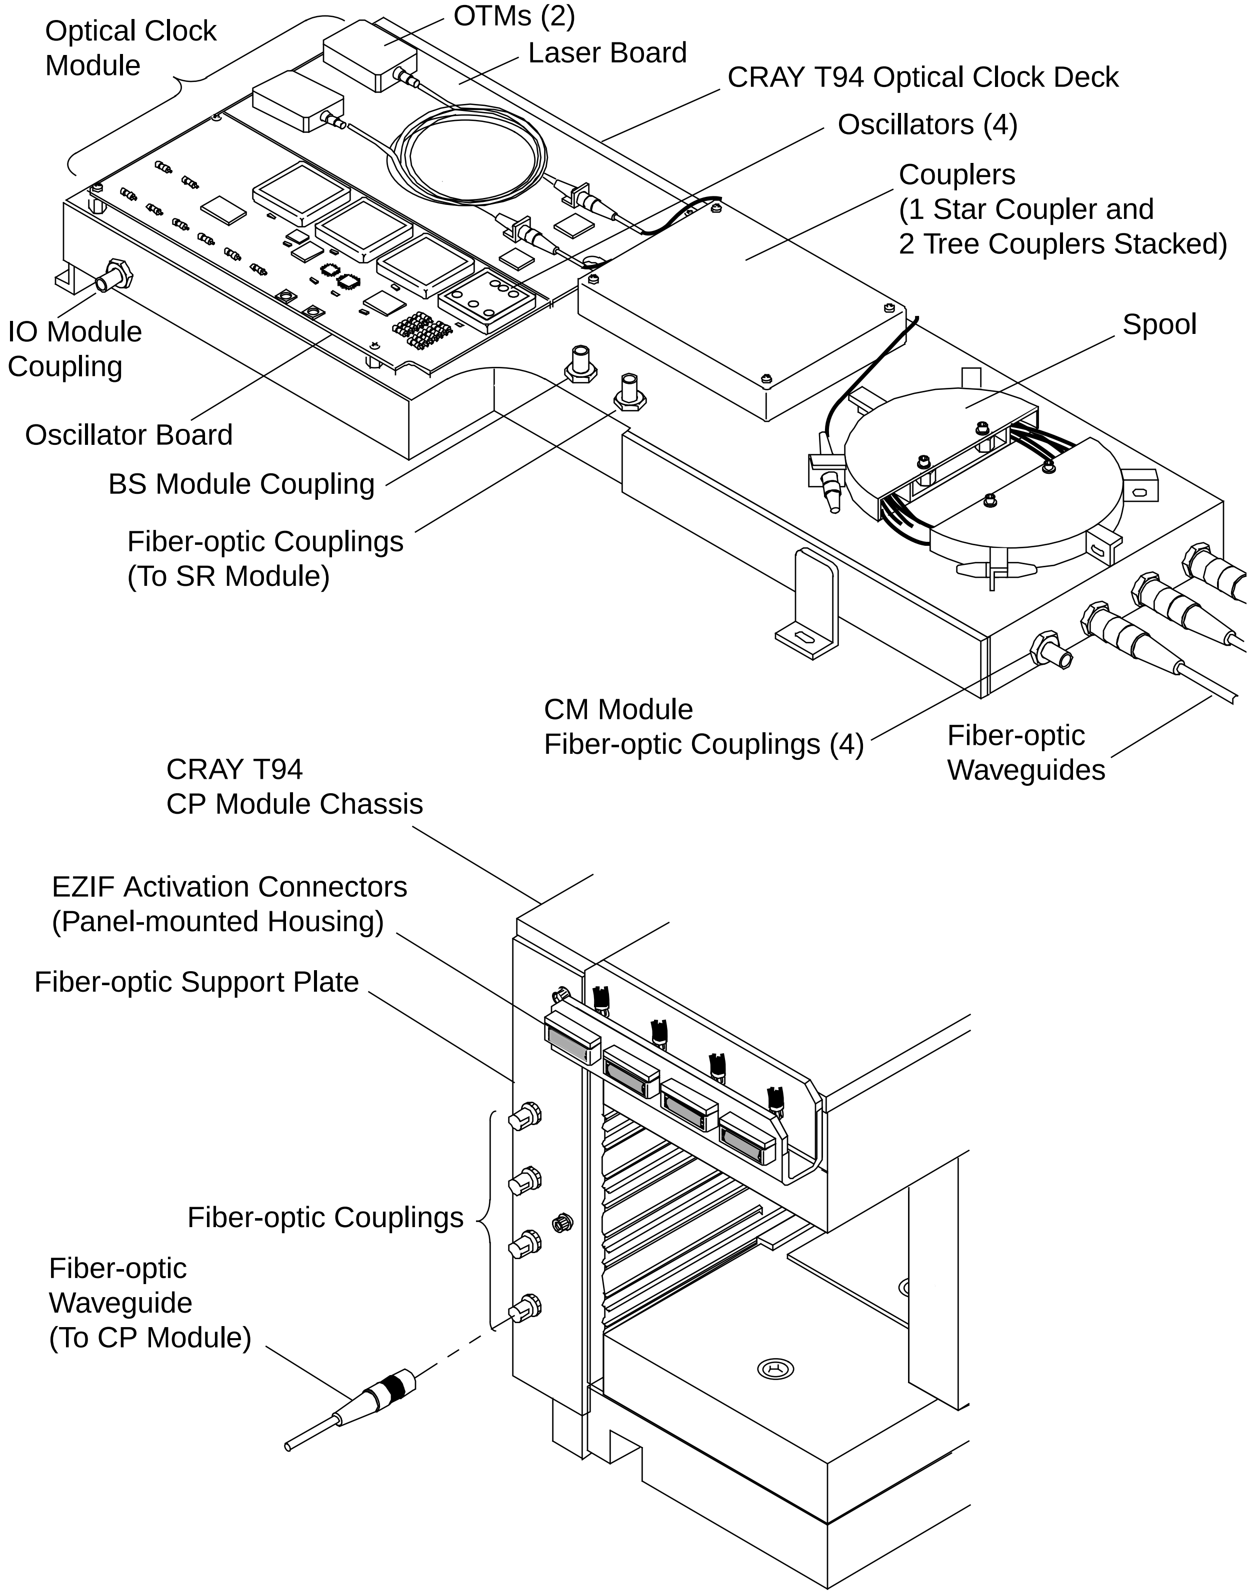

PC Board Technology

All CRAY T90 series modules are built with 22-, 52-, or 54-layer printed circuit boards (PCBs). NW and CP boards have 54 layers. The BS module has one 22-layer board. Boards used in other modules in the system have 52 layers.

Each 52-layer printed circuit board consists of the following three individually fabricated and tested boards:

- One 8-layer printed circuit board (BDA) with the following features:
  - Two power-ground pairs
  - Two surface layers/ground pairs
- Two 22-layer printed circuit boards (BD1 and BD2) with the following features:
  - Five vertical-horizontal signal pairs
  - Five power-ground pairs
  - Two surface layers/ground pairs

The primary purpose of the 8-layer PCB is power distribution and ‘Z’ directional interconnects. This board serves as the core of the 3-board assembly. The two 22-layer PCBs serve as the logic signal transmission boards, and they also provide some power/ground distribution. The cell arrays, connectors, and capacitors mount to these 22-layer printed circuit boards.

CPU Block Diagram: Differences

The following bulleted lists describe new CPU features. Refer to the CPU block diagram shown in Figure 12.

Registers

- The A, B, and shared B registers expanded to 64 bits.
  Shared B and T registers expanded to 16 registers per cluster.
  Communication semaphores expanded to 64 per cluster.

- The number of clusters was expanded to 18 for CRAY T94 and CRAY T916 systems. CRAY T932 systems have a total of 36 available clusters.

- Several new instructions use the 64-bit A registers for logical and shift operations to reduce loading on the S registers. Any instruction preceded by a 005400 instruction is run in Triton mode.

Scalar Cache

- The scalar cache feature was added for A/S register load operations. New instructions were added for cache data use and control. Cache is organized into 8 pages, with 8 lines per page and 16 words per line, for a total cache size of 1,024 words.

- The performance monitor can now gather information about cache use.

Functional Units

- New 64-bit iota, compress, and expand functional units and a new 40-bit integer multiply functional unit were added. The bit-matrix multiply functional unit is now standard on all CPUs. Instructions were added for the new functional units.

- The address adder functional unit was expanded to 64 bits wide, and the address multiplier was expanded to 48 bits wide. In CPE1 modules, the address multiplier was dropped; the floating-point multiplier performs 128-bit multiplication.

- The scalar registers can now access the bit matrix multiply functional unit.

- The second vector logical functional unit was separated from the floating-point multiply functional unit, to eliminate a hold issue condition.
- Vector mask capabilities were expanded for use in the new iota functional unit. Internal logic was modified to reduce issue conflicts.
- The capability to perform two gather operations and one scatter operation at the same time was added. A new double gather instruction was added.

Exchange Package

- The exchange package expanded to 32 words and changed format. Refer to Figure 13. Five exit addresses were added, which allows jobs to exchange to other jobs or the operating system to have multiple entry points, thereby reducing system overhead. The operating system does not currently support the capability to exchange to other jobs.
- The exchange address is no longer lost on a deadstart exchange.

Memory Control

- A new memory management scheme uses logical address translation (LAT) tables rather than base and limit addresses. Each LAT table is based on a 40-bit logical address space and a 38-bit physical space. LAT tables enable direct access to a maximum of 256 Gwords of memory. The memory system supports up to 35 address lines.
- A chunked chaining of common memory read operations was created for use with vector registers. The entire vector is no longer needed for memory read operations. Chaining proceeds when eight contiguous memory operands are available instead of requiring the entire vector from memory before processing can begin.
- A read followed by a write to the same A*k* address will not necessarily occur in order in memory.

Figure 12. CPU Block Diagram

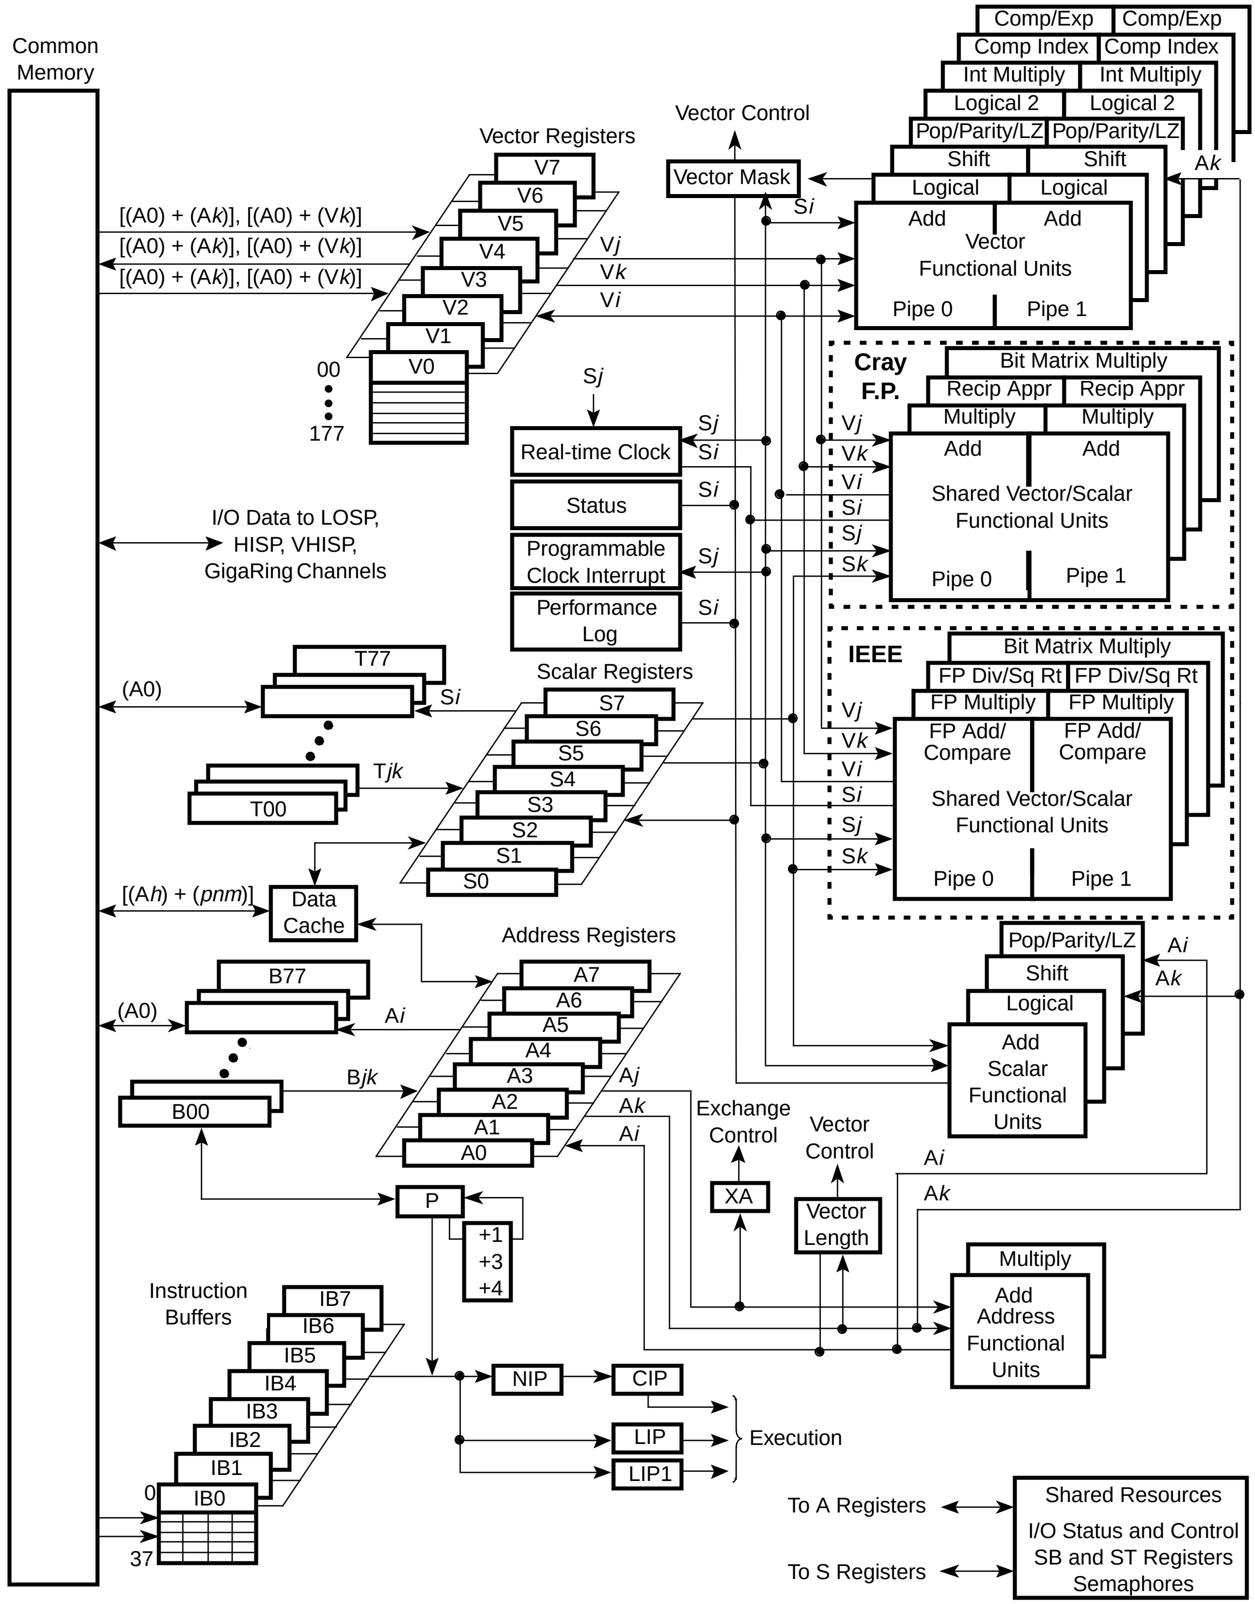

Memory Control Scheme

The memory control scheme of the CRAY T90 series system differs from the memory scheme of earlier systems. In earlier systems, Cray Research used a memory system referred to as a *parallel access system*. In this type of memory, all CPUs were involved in every memory reference.

The CRAY T90 series system uses a *network access system*. In this type of memory, the requesting CPU arbitrates the access to its own ports only. Once the arbitration is complete, the reference leaves the CPU, and no other CPU ever gets involved in the arbitration process. The network module and the memory module must resolve access conflicts to the memory bank.

The network and memory modules work in a *handshaking* fashion; they contain first-in-first-out (FIFO) buffers and arbitrators. If a FIFO buffer for an arbitrator is full because of an unresolved conflict, that buffer will not send a Resume signal to its supplier. The supplier then starts filling up its FIFOs. This process continues all the way back to the CPU. When the CPU FIFOs are full, the system stops issue. The buffer sizes minimize the number of hold issues caused by memory backup.

Because of memory arbitration, when a CPU makes a request to the memory system for a read operand or instruction, the CPU does not know when the read data will be returned. CPUs add a *tag* to the memory request. The tag contains the requesting CPU number and an identifier for the register to be loaded.

Each memory bank holds the associated tag while the requested word is being referenced. The tag is returned with the read data. The tag is used to route the data to the proper destination. This scheme also works for I/O traffic by replacing a register identifier with an I/O identifier.

Figure 13. CRAY T90 Series Exchange Package

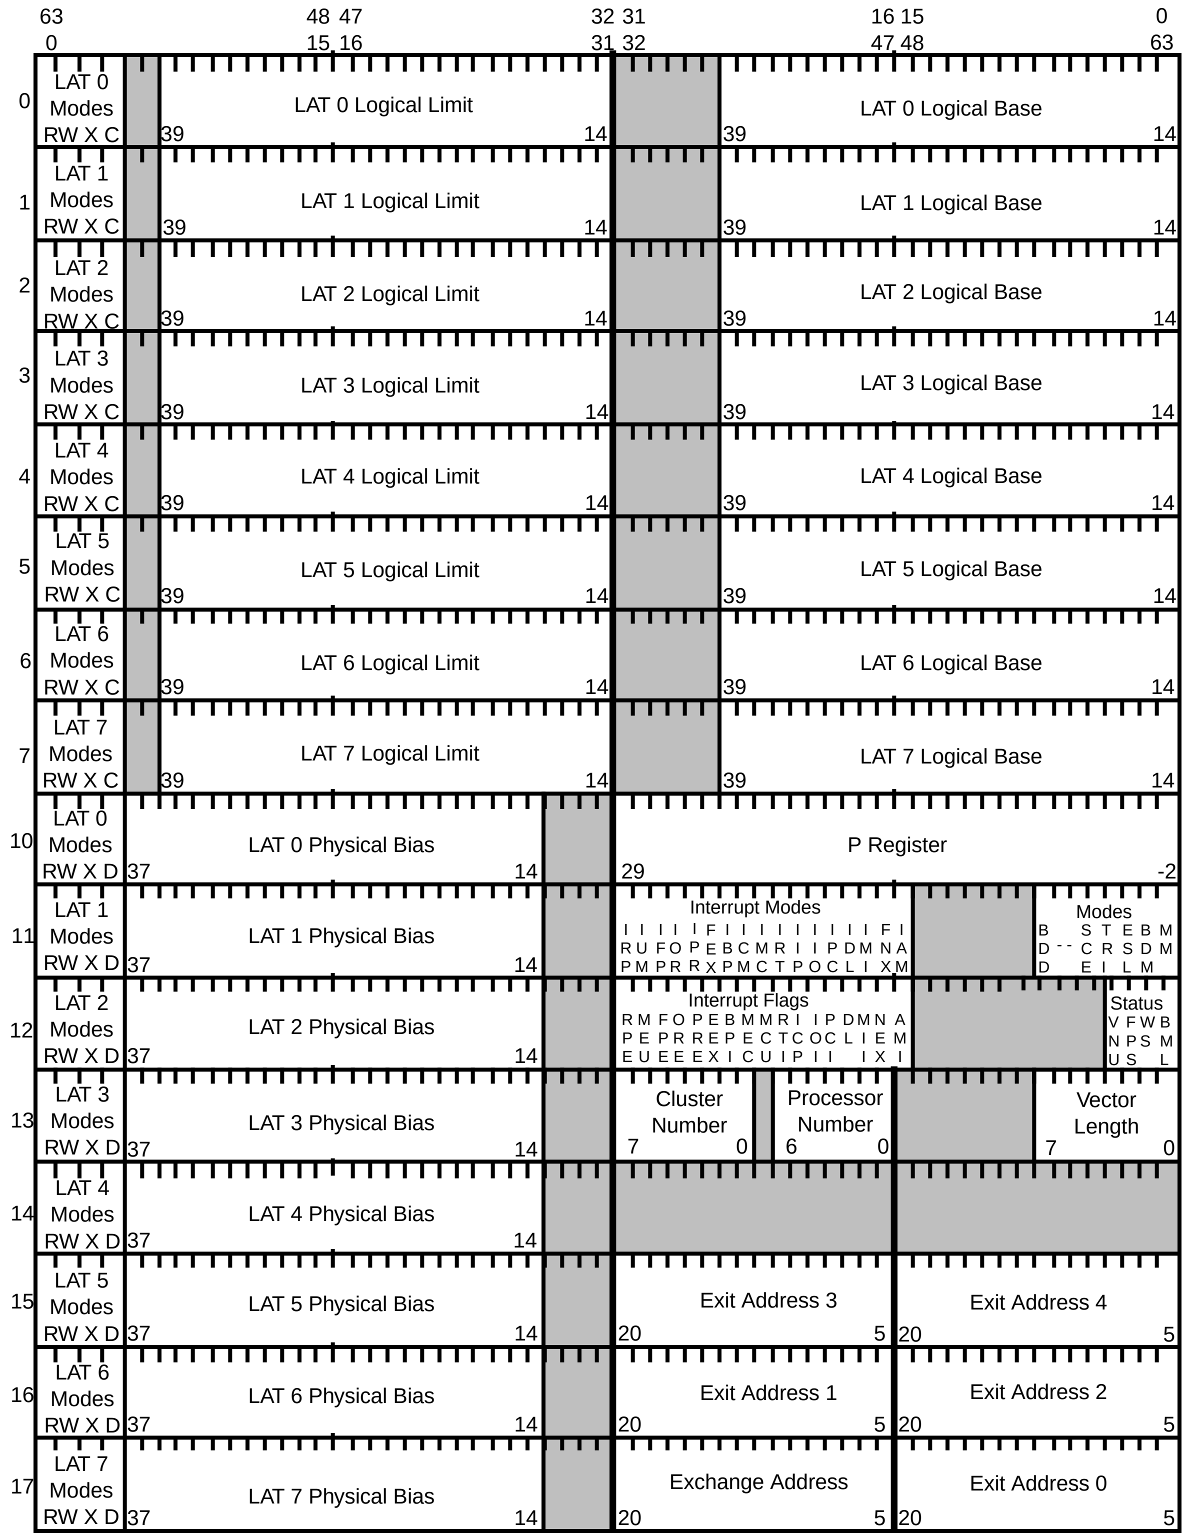

Words 20 – 27: A Registers 0 –7 Words 30 – 37: S Registers 0 – 7

Logical Address Translation (LAT) Table

The CRAY T90 series computer system uses a system of memory management called logical address translation (LAT) tables. The LAT tables replace the traditional base address registers and limit address registers, which were limited to 32-bit addresses and were not flexible enough. LAT tables enable use of multiple address spaces.

Each LAT has a 16-K block size and is based on a 40-bit logical address space and a 38-bit physical address space. This means that the maximum logical address is 1 teraword and that space is allocated and checked in 16-K blocks. Because of the 38-bit physical address space, the LAT implementation has a 256-Gword limit. However, the CRAY T90 series system supports 35 address bits for a total of 32 Gwords of memory.

The exchange package provides space for eight LAT entries. There are hardware limitations in the memory control logic that allow for a maximum of three types of LAT entries: read, write, and execute. There can be three entries with read and/or write permission and three entries with execute permission. Entries can be combined (for example, read and write) in the same LAT entry. If the user violates these limits, the LAT information could be lost or ignored. If a LAT table has no mode bits set, the rest of the entry is ignored.

The CRAY T90 series system does not return memory requests in a predetermined order. This means that the order of the LAT entries into the hardware cannot be determined from their order in the exchange package. Because of this, there is no way to determine which LAT entry was lost if mode limits are exceeded.

Refer to Table 3 for descriptions of the various LAT fields.

Table 3. LAT Fields

Field	Description
Logical base	The first logical address of this LAT
Logical limit	The last logical address +1 of this LAT
Physical bias	A logical-to-physical mapping constant (Physical_Bias = Physical_Base_Address - Logical_Base_Address)
Modes	Bits that define the use of this LAT (R)ead (W)rite (X)execute (C)achable, (D)irty. The Dirty bit is not a user set bit; It is set by hardware.

LAT Calculations

LAT calculations are based on three fields: logical base address, logical limit address, and physical bias. Unlike previous systems, there is no physical base address.

The physical bias field maps logical addresses to a physical space in memory. The operating system calculates the physical bias, over which the user has no control. The operating system calculates the physical bias by using the following equation:

Physical Bias = Physical Base Address - Logical Base Address

NOTE: Because there is no relation between physical space in memory and logical space, the physical bias could be a negative number.

Refer to Figure 14 for an example of the CRAY T90 series memory mapping, which is managed by the operating system. In this example, three *pages* or areas of memory are mapped.

Page 1 has a physical base address of 0 and a logical base of 30K. Using the formula to calculate physical bias yields a physical bias of -30K. If a requested address was 35K and you wanted to map that to a physical space in memory, you would add that address to the physical bias, and the results would be at address 5K in memory.

The other two pages can be calculated the same way as the first page. Note that when the physical bias is 0, the request address equals the physical address.

Figure 14. LAT Memory Mapping

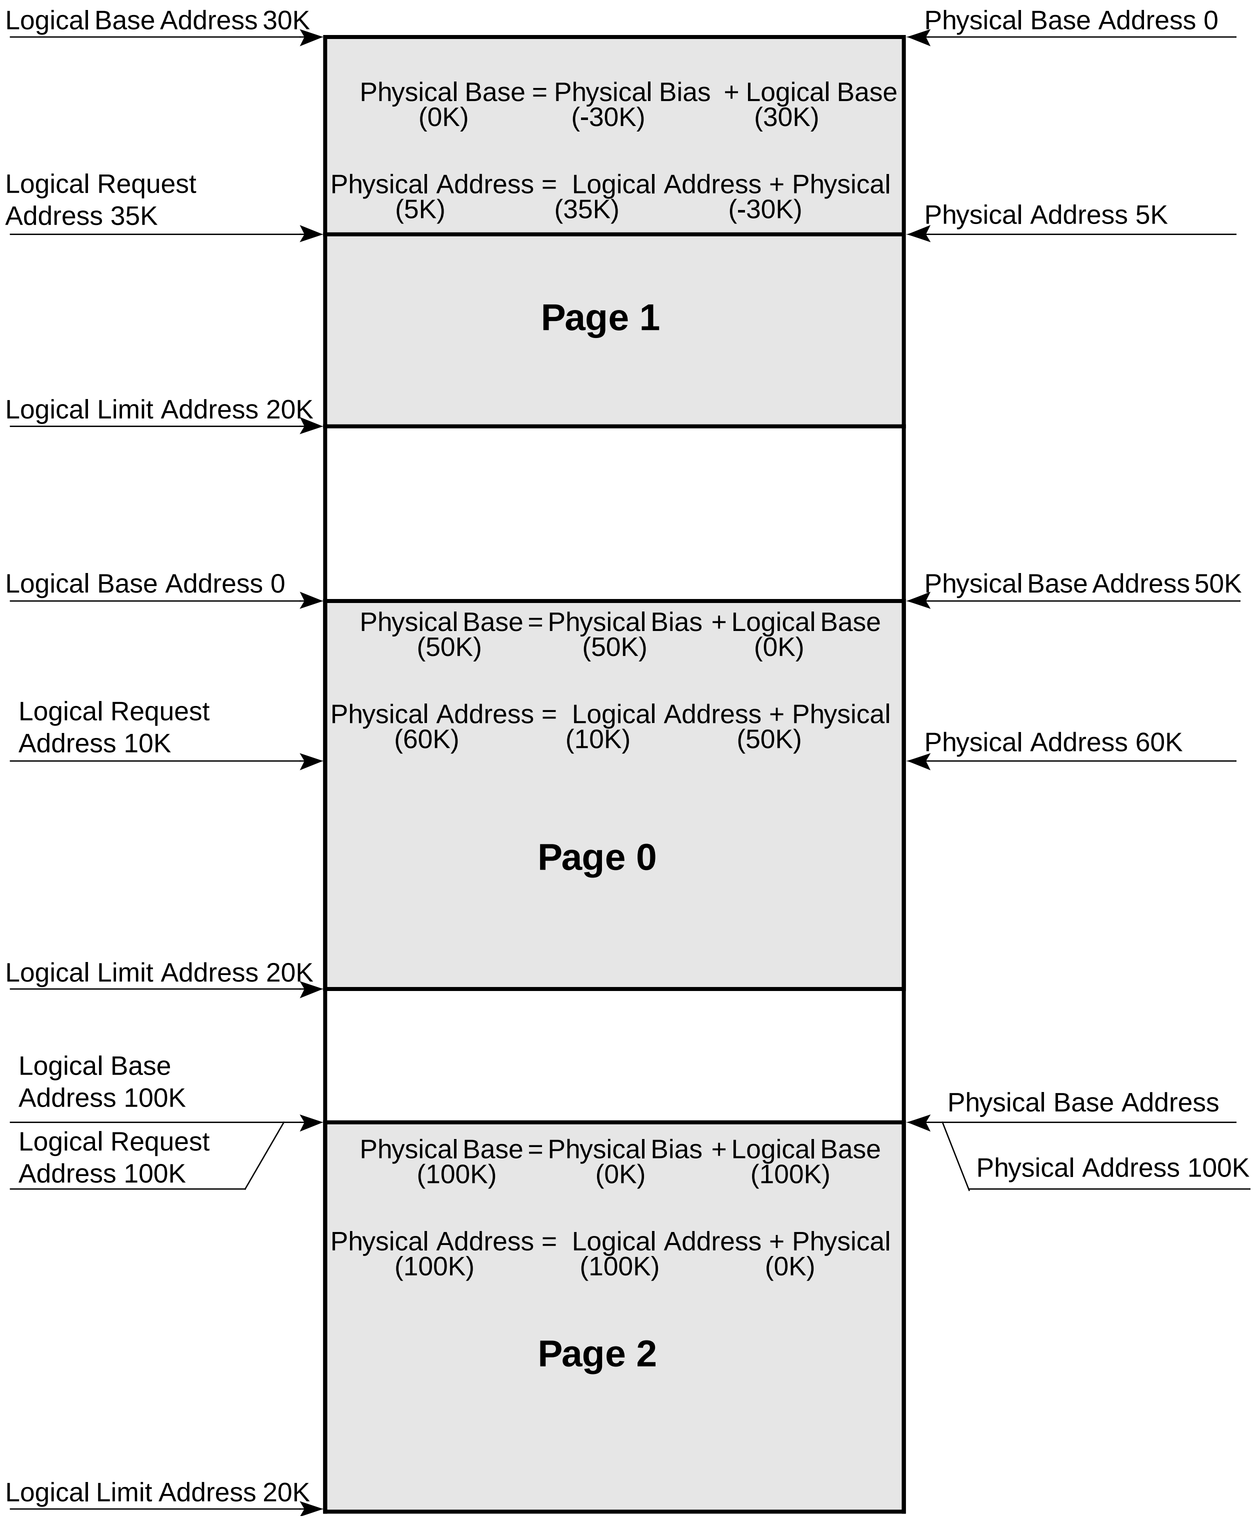

System Configuration

The System Configuration Environment (SCE) program is used to create and manage the logical configuration of the system. CRAY T90 series mainframes use logical configuration: the configuration is set by manipulating soft switch values rather than by physically rearranging the hardware. Logical configuration enables introduction or removal of modules within a system configuration without powering down any modules or the mainframe.

The SCE program establishes the numbers of memory sections, subsections, and banks in a particular system or system configuration. The SCE program establishes other settings such as logical CPU numbers and whether particular I/O channels are active. SCE is also used to degrade or partition memory.

By degrading memory, the system can continue operating, or maintenance can be performed in the presence of memory system failures. Systems can be degraded to a level in which one CPU is addressing 1 bank of memory in each of four memory sections (a total of two memory modules with only one-eighth of the memory on each of those modules in use).

By partitioning the system, parts of the I/O system can be configured to direct I/O traffic to any quarter or half of memory. CPUs must be configured consistently with I/O configuration. The other memory configuration settings also apply to any I/O traffic through a CPU.

All configuration/degrade selections apply symmetrically within a processor; for example, selecting 4-bank mode within a CPU means that all requests on all memory ports will use 4-bank addressing. It is not possible during normal operation to get 4-bank addressing through one port and 8-bank addressing from the remaining ports. Other configuration features are listed below.

- Software performs system configuration. Software can logically configure modules into or out of the system. A module could be physically absent from the system while the system operates normally (a possible reduction in performance and capabilities may result).

- The system can be logically partitioned, which allows different operating systems to be run concurrently.

- Separate power and cooling connections in CRAY T932 systems allow one half of the mainframe to be powered down while the other half continues to function.

- Mainframe major components (common memory, I/O, and clusters) can be configured into separate partitions. Each CPU can be assigned to a partition.

Sanity Code

Sanity code is a 6-bit pattern that controls data flow and interconnecting processes between modules. Sanity code performs three primary functions:

- Establishes paths for sending maintenance commands and data
- Enables module-to-module communication
- Establishes a return path for reporting errors

Sanity code ensures that a logical path exists to all configured portions of the system. Sanity codes control all signals that cross module boundaries. A module must receive sanity code through an interface or connector to another module before it can respond to signals or commands from the other module. A module must also return a sanity code to the interface or connector before a module can receive a return response. The interface between modules is either a single connector or a set of serial connectors [a connection between a CP and a network module goes through two connectors, one on each side of a system interconnect board (SIB)].

A module must receive sanity code from another module for 96 consecutive clock periods to establish a connection. Once the connection is established, the sanity code generator continuously rebroadcasts the 6-bit sanity code pattern throughout the tree every 6 clock periods. If a module loses sanity code from a particular port, it still may be receiving sanity code from another port. A module must continue to receive valid sanity code in order to remain in an active state in a sanity tree structure.

The absence of valid sanity code on all ports of a module causes that module to go into a master clear state; all outputs from that module are then ignored. If module A sends sanity code to module B, and module A loses its original sanity code, then module A also drops its sanity code connection to module B.

Sanity Trees

A sanity tree is a logical structure that provides a flexible and reliable control network. A sanity tree is a sanity code distribution network that is defined when the mainframe is configured.

The actual path or structure of the sanity tree is not rigid; its structure is configured by the operator. Once a tree structure is established, it remains set at the same structure until it is changed or reconfigured. Remember that configuration of the mainframe defines the order of the tree structure.

Boundary Scan

Cray Research implements a boundary scan feature on CRAY T90 series systems. Boundary scan testing can test any interconnection in the system for proper connection, except memory chip pins and I/O channel connections (between boundary scan output and input cells). This feature enables field personnel to test the interconnections on the module and the interconnections between modules.

The boundary scan feature is provided by a boundary scan module (BS02) in IO01 systems. (You cannot boundary scan test the actual BS02 module.) In IO02 systems, boundary scan logic is located on the IO02 module. IO02 systems do not use boundary scan modules.

Boundary scan has the capability to detect shorts between option interconnections, shorts between interconnections and power/ground grids, or open foils in the interconnections. Boundary scan cannot test memory chips or IO channels because there are no boundary scan capabilities built into those connections.

Boundary scan provides the following advantages over previous diagnostic and edge-connector test-vehicle techniques:

- Reduces the need to access and probe pins and test points
- Shortens test and checkout times for chips and modules
- Reduces times and costs of development and production
- Reduces mean time to repair (MTTR)
- Eliminates the need to understand a function in order to test it

Boundary scan testing enables most connections within the module and connections between the modules to be tested without diagnostic tests. Boundary scan can identify the module type, module serial number, and option type and revision of all logic on a module.

Boundary Scan Register

Each chip or option in a CRAY T90 series mainframe contains a chain of shift-register-based cells around the periphery of the option. These cells are the boundary scan register (BSR). There is one BSR location (or bit) for each input/output pad on the option. Sixteen additional bits store an option identification number.

Boundary Scan Testing

Boundary scan testing can be run whenever a failure occurs that shuts down the system, and it is run to verify the integrity of the system after a repair procedure has been completed. Boundary scan requires control of the system and, therefore, cannot run simultaneously with the operating system. However, CRAY T932 systems can be degraded, which enables boundary scan to run on one half of the system while the operating system runs on the other half of the system.

Boundary scan testing is performed by a set of Cray Research programs called boundary scan system test. The main programs that are part of boundary scan system test are `bsb`, `bscan`, `breport`, and `runbscan`.

Using the `runbscan` shell script is the recommended method of running boundary scan system test. The `runbscan` shell script utility automates the process by starting `bscan`, which runs a sequence of tests and uses `breport` to automatically generate a report in an easy-to-interpret format. The `bscan` program also invokes `bsb`, a program that builds the serial boundary scan interconnect path between modules and calculates expected scan data values. All boundary scan programs are described in detail in the *Boundary Scan System Test* document, publication number HDM-117-A.

Module-to-module Connectors

CRAY T90 series systems use the following types of module-to-module connectors (refer to Figure 15):

- OIMs D/D and D/L
- SIMs

Figure 15. Module-to-module Connector Types

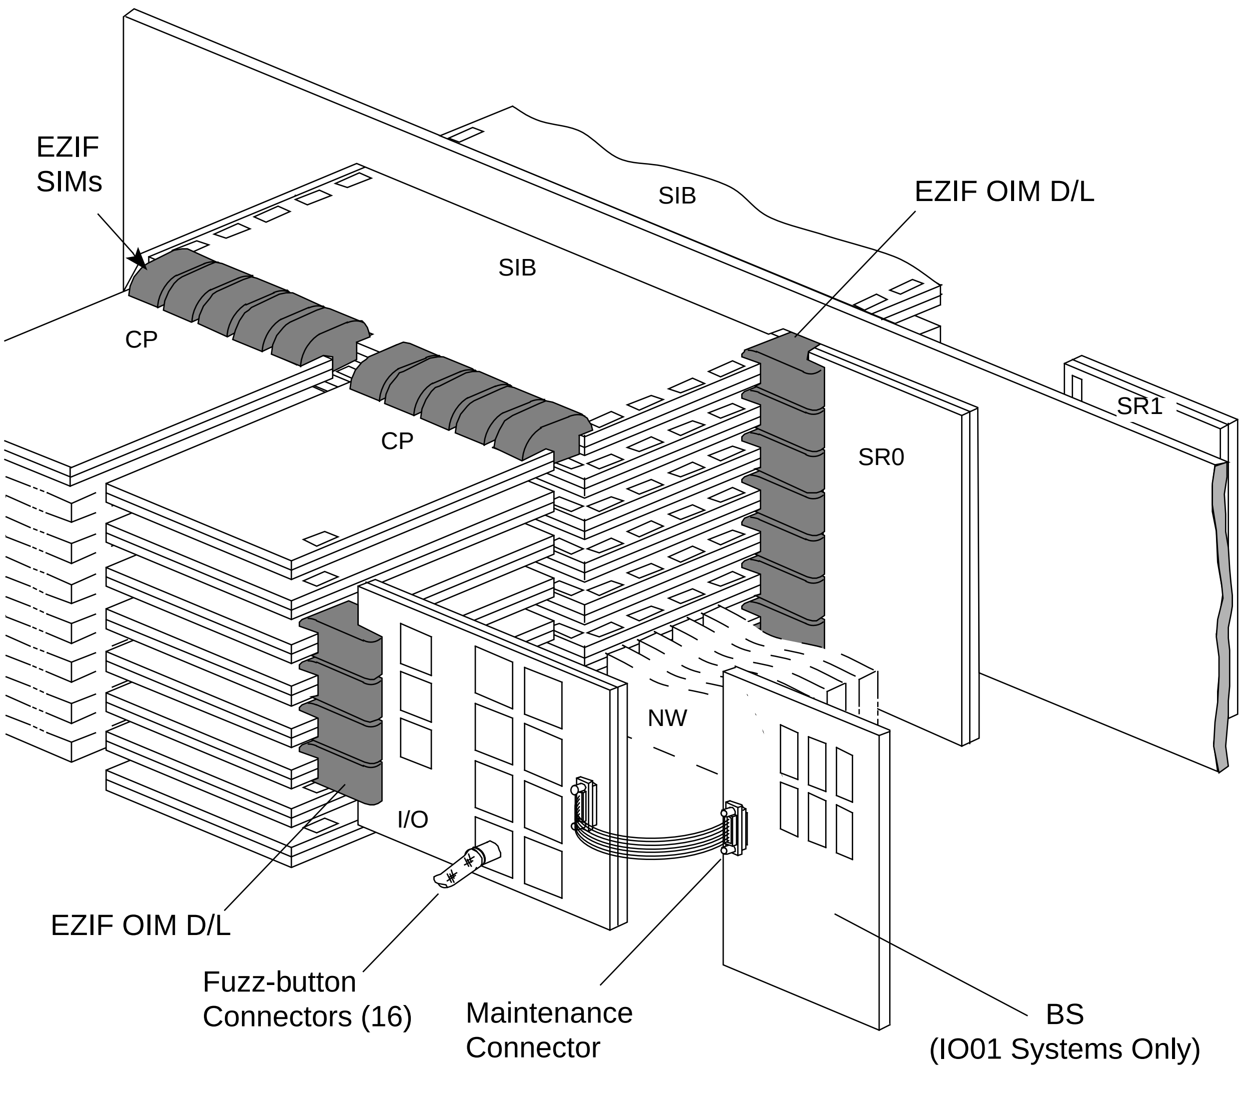

OIM Connectors

Orthogonal interconnect module (OIM) connectors connect modules that are linearly rotated 90 degrees to each other.

There are two types of OIM connectors:

- OIM D/D - direct-direct OIM
- OIM D/L - direct-lateral OIM (also known as OIM direct-slide-by)

Modules that use the OIM D/D connector plug into each end of the connector. With OIM D/L connectors, one module plugs into one end of the connector, and the other module must slide through the entire row of connectors before seating in place.

SIM Connectors

Straight interconnect module (SIM) connectors connect modules that are oriented in the same plane. All SIM connectors are the direct-direct type, which is denoted SIM D/D.

T-rail Connector Assemblies

OIM and SIM connectors attach to T-rail assemblies as shown in Figure 16. A printed circuit board is mounted to each T-rail assembly. Each CM/CP T-rail assembly contains five OIM D/D EZIF connectors, which provide interconnect paths between the CM, CP, and SR modules. The IO/CP T-rail assembly contains four OIM D/L EZIF connectors, which provide interconnect paths between the IO and CP modules.

CRAY T94 systems contain four CM/CP T-rail assemblies and one IO/CP T-rail assembly. All T-rail assemblies in CRAY T94 systems are accessed and replaced through the CP module hatch. Each quadrant in CRAY T916 and CRAY T932 systems contains eight CM/NW, NW/SIB, and CP/SIB EZIF T-rail assemblies and one IO/CP T-rail assembly.

If a system is run with a module removed from a stack or if a stack is not fully populated, a flow block must be installed. The flow block maintains the even distribution of Fluorinert liquid throughout the system.

Figure 16. OIM T-rail Assemblies

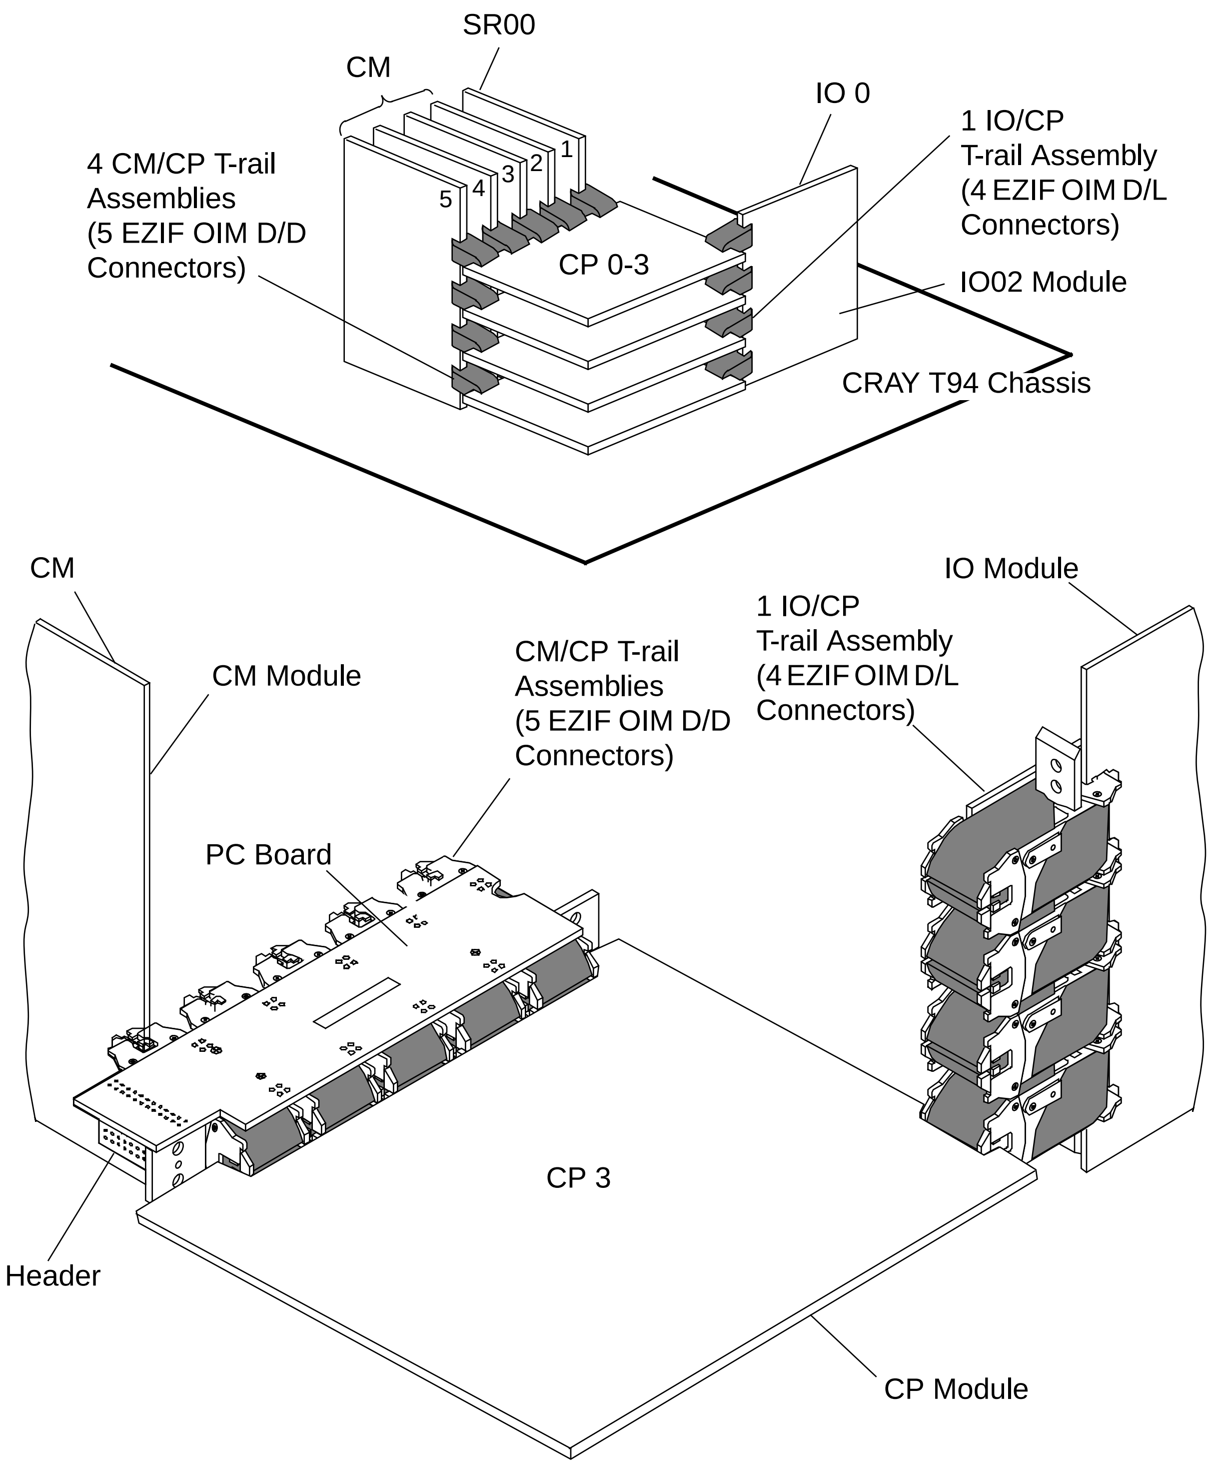

Flex Circuit

Flex circuits consist of three layers of copper that are encapsulated in layers of a flexible insulating film. The use of flex circuit technology in CRAY T90 series systems accommodates higher signal speeds and the orientation of modules on different planes. The pin and socket connector technology used in previous mainframe systems does not provide the electrical characteristics needed to attain the higher speeds of CRAY T90 series systems.

OIM/SIM Connector Operation

Each OIM and SIM connector has an embedded heater-actuator spring (HAS) assembly to open and close the flex circuit jaws of the connector. Figure 17 illustrates the HAS assembly.

An EZIF heater controller device opens and closes different OIM and SIM connectors. The controller applies 12-Vdc current to maintain a preset temperature to each connector on a T-rail assembly. The heater controller device also monitors the status of each T-rail connector to prevent damage to the heater elements in the HAS. The heater controller selects the primary or secondary heater element. The secondary heater element is used in case the primary heater element fails.

NOTE: The secondary heater can be used only once (to remove a module). Once the module has been removed, the faulty connector must be replaced.

When 12-Vdc current is applied to the HAS, the internal resistance causes the temperature of the selected heater element to rise gradually. Heat is transferred by conduction from the heater element to the actuator. The actuator begins to expand at a temperature of 70 °C (158 °F), which starts to force open the inner and outer springs in the HAS. The opening and closing of the HAS causes the OIM/SIM connector to open and close. The HAS will open fully after current has been applied for approximately 2 minutes, when the temperature of the actuator reaches 100 °C (212 °F).

The actuator applies a force to the springs only when opening the connector. When current to the heater is turned off, the heater and actuator cool, and the outer spring forces the HAS to close.

Figure 17. Heater-actuator Spring (HAS) Assembly

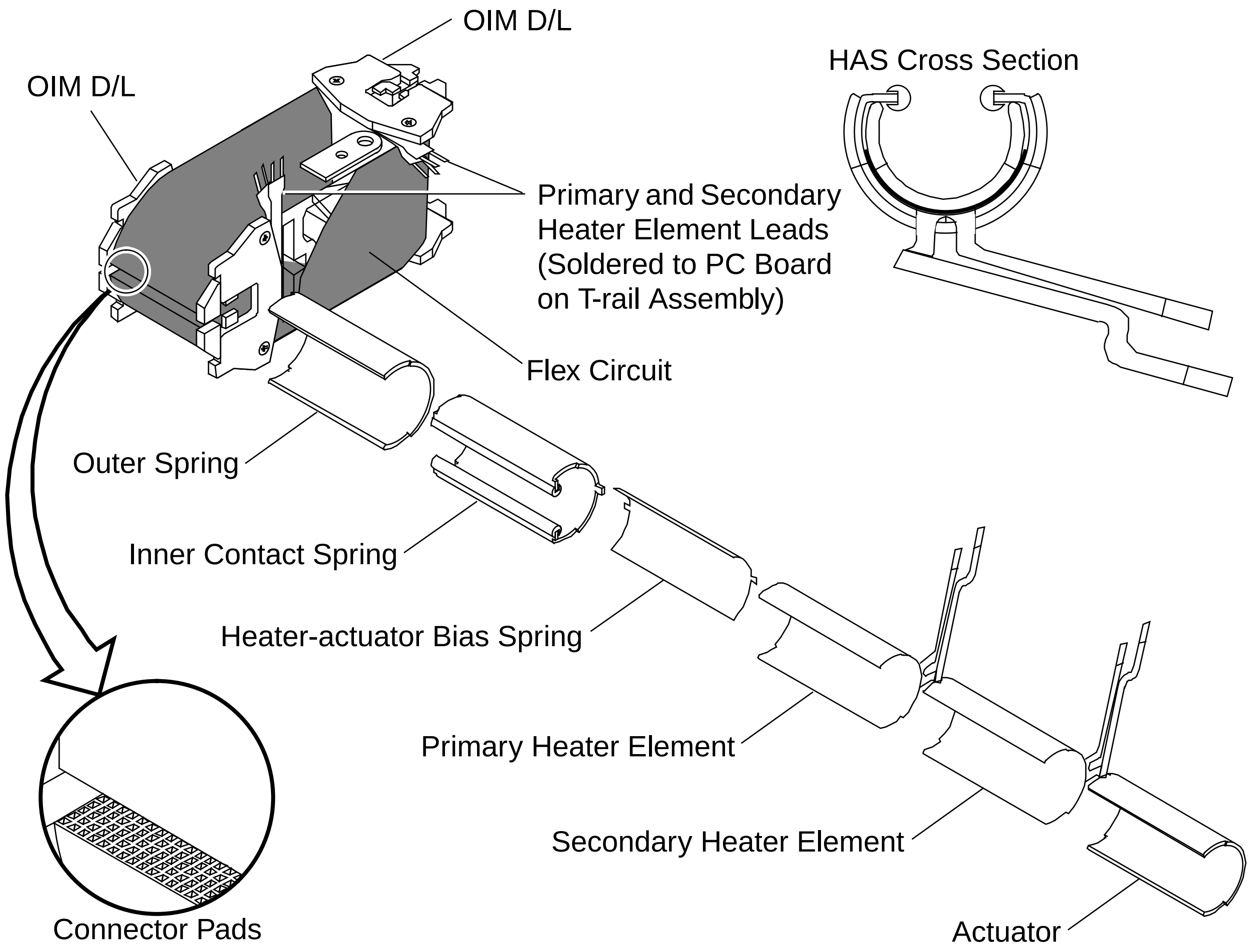

The external 12-Vdc temperature controller plugs into a heater-board assembly inside the chassis tank. Each heater board is labeled for the appropriate module type and location. From the heater board assembly, 12-Vdc wire harnesses connect the header on the T-rail assembly.

On CRAY T94 systems, the wire harness from the heater board assembly housing connects to a housing on the EZIF interconnect board (Figure 18). A harness from the EZIF interconnect board plugs into the header on the T-rail PC board. In CRAY T916 and CRAY T932 systems, the wire harness from the panel-mounted housing connects directly to the header on the T-rail PC board.

Figure 18. HAS Wire-harness Assembly

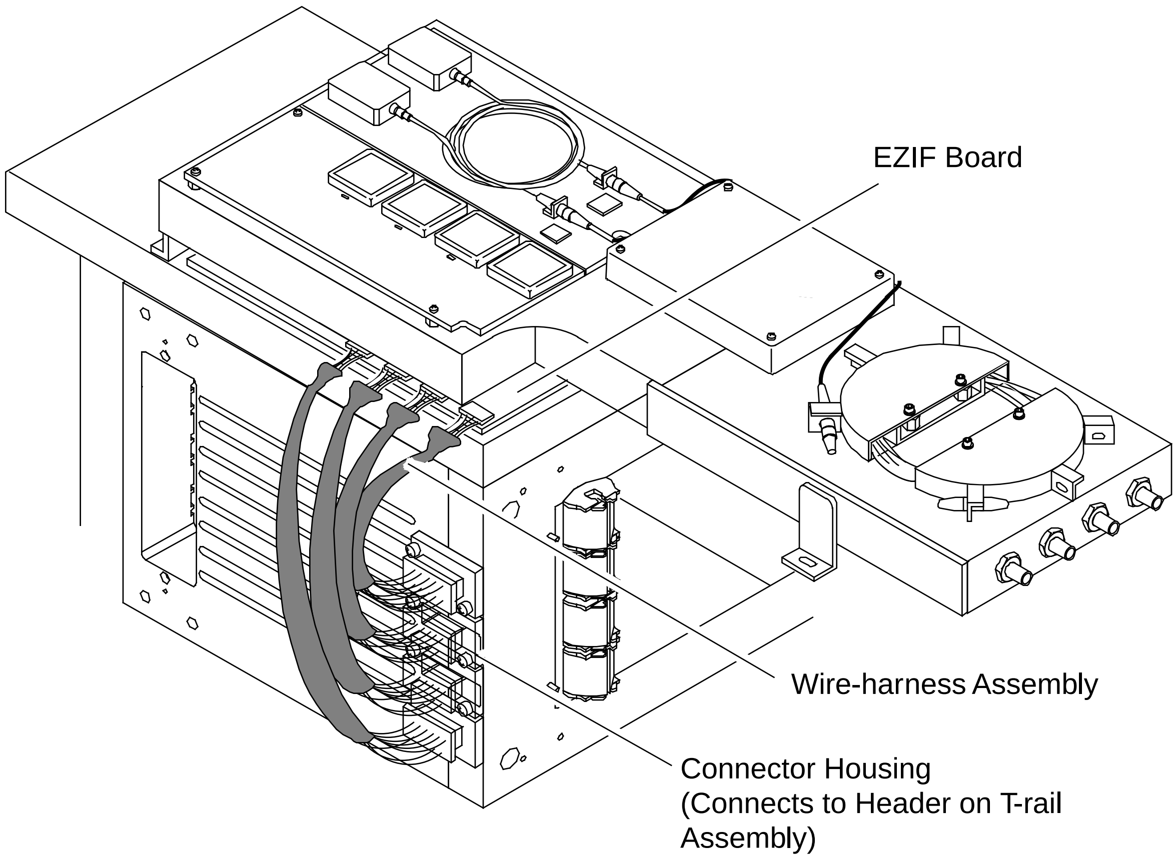

Fuzz-button and Pogo-pin Connectors

Data, address and control signals pass to or from an IO module through several I/O harness assemblies. Figure 19 illustrates an I/O harness assembly and fuzz-button connector. The IO02 module uses a new style of connector called a pogo pin. Each pogo pin is a round pin that uses a socket connector versus the pressure-fit pad connectors used with fuzz-button connectors.

Up to 16 connectors may be connected to an IO module; each connector has 514 pins. The I/O harness assembly divides from the connector into several micro-D connectors, which connect to the I/O bulkhead. The entire I/O harness assembly resides within the cooling tank.

Figure 19. I/O Harness Assembly

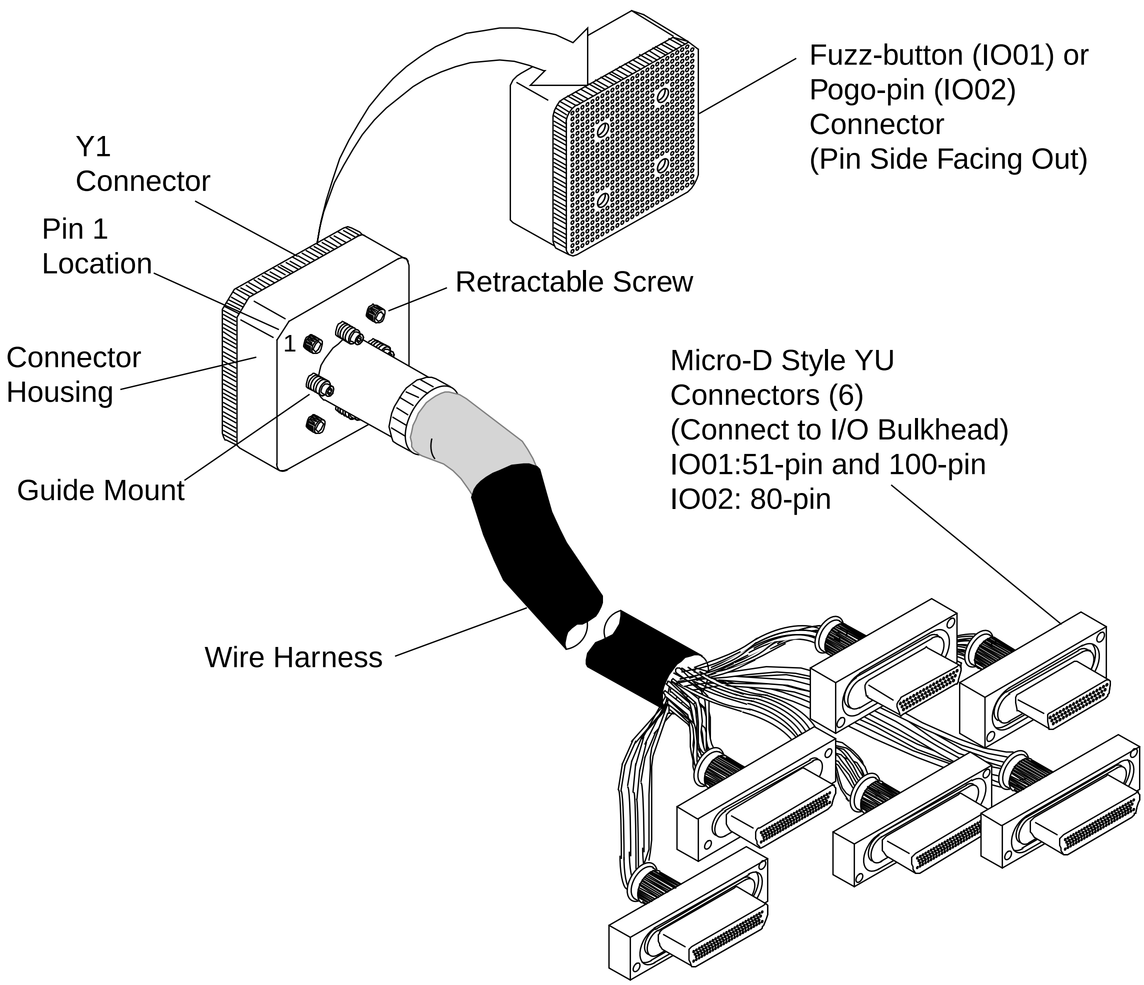

NOTE: All wires in the I/O harness assembly are shielded.

Maintenance Connector

Each SR, IO, CP, NW, and CM module has a maintenance connector. The maintenance connector is located on the power regulator of each module. The maintenance connector sends the following signals to a module and receives the following signals from a module:

- Boundary scan signals - TCLK, TM, TDI, TDO, and RCLCK
- Boundary scan cable ground
- OIM and SIM connector interlock signals
- Continuity line signals
- Test-point inputs for connectors and test-point probing (STCO only)
- Thermal diode to output module temperature status

NOTE: The OIM maintenance connector on the BS module is used only in a test vehicle; it is called a TV maintenance connector.

Thermal diodes embedded in options sense module temperature. The diodes enable an analog signal to be sent through the maintenance connector to a test bulkhead; these thermal diodes are not used in the field.

Maintenance Features

Cray Research has built many new maintenance features into CRAY T90 series systems to decrease mean time to repair (MTTR). The CRAY T90 series design includes concurrent maintenance and redundancy features to greatly reduce customer downtime.

The systems include many maintenance features, such as spare memory chips, the capability to down processors, the use of sanity codes to reconfigure a system, and a physical boundary on CRAY T932 systems that allows half of the mainframe to remain operational while the other half is maintained.

The following list describes important maintenance features:

Maintenance Channel

The maintenance channel enables the following functions: to define and control system configuration, initialize the system, enable and disable sections of the system, monitor system activity, run and control diagnostic tests, access test points, and master clear CPUs and memory. In CRAY T90 series systems that use GigaRing channel architecture, maintenance channel logic is accessed through the supervisory channel.

No Hardware Switches

A CRAY T90 series system does not have hardware switches. All aspects of power-up and power-down, configuration, and diagnostic maintenance and monitoring processes are performed through the use of the monitoring and control system, which is controlled through the MWS, MWS-T, or SWS.

Boundary Scan

The CRAY T90 series system also incorporates a boundary scan feature that enables the system to check the chip-to-chip and module-to-module interconnect paths and report which paths are failing. This feature also enables the system to read the chip type, revision level, and serial number of all the options.

System Interconnect Board

There is no longer a wire mat; it has been replaced with a printed circuit board called a system interconnect board (SIB). To reduce costs, only two types of SIBs are used: the SJ module for the CRAY T916 system and the SI module for the CRAY T932 system.

Spare Memory Chips

Common memory modules contain spare memory chips. In the event of a memory chip failure, memory can be repaired and normal memory operation continued, without physical maintenance.

P Register and Test-point Information

The P register address of each CPU can be read back to the OWS/MWS/SWS through the logic monitor.

Each CPU can be set to trap on a P register address, an issued instruction, or any test point.

System Configuration

System configuration information can be read from the mainframe. Modules can be logically removed to perform concurrent maintenance.

Diagnostics

Instruction buffer load and dump functions are available for diagnostics.

Exchange load and dump functions are available for diagnostics.

Numerous other test, diagnostic, and maintenance functions are available.

Power Supplies

Power supplies are arranged in an N+1 configuration to keep the system operational in the event of a power supply failure.

Partitioning

Partitioning involves dividing the mainframe into logical partitions. Partitioning is supported only on IO02 systems. CRAY T94 and CRAY T916 mainframes each have one physical partition, and CRAY T932 mainframes can have one or two physical partitions. Each partition can be divided into a maximum of four logical partitions. Partitioning allows multiple customer functions or maintenance functions to run in different areas of the mainframe at the same time. The SCE program controls all partitioning of the mainframe by setting soft switch values.

IO01 (IOS-E Based) System Equipment

This subsection describes equipment that is used with CRAY T90 series IO01 systems. A section describing equipment that is specific to IO02 systems follows this section.

Support System

The support system is the point of first-level access, control, and status of a CRAY T90 series computer system that uses the IO01 module (model E I/O support). Figure 20 shows the cable connections of support system chassis components.

The support system consists of the following components:

- Support system chassis
- MWS and OWS workstations
- HP LaserJet printer

NOTE: Remote support equipment is provided in optional spare parts kits.

Figure 20. Model-E (IO01) Support System Interconnects Block Diagram

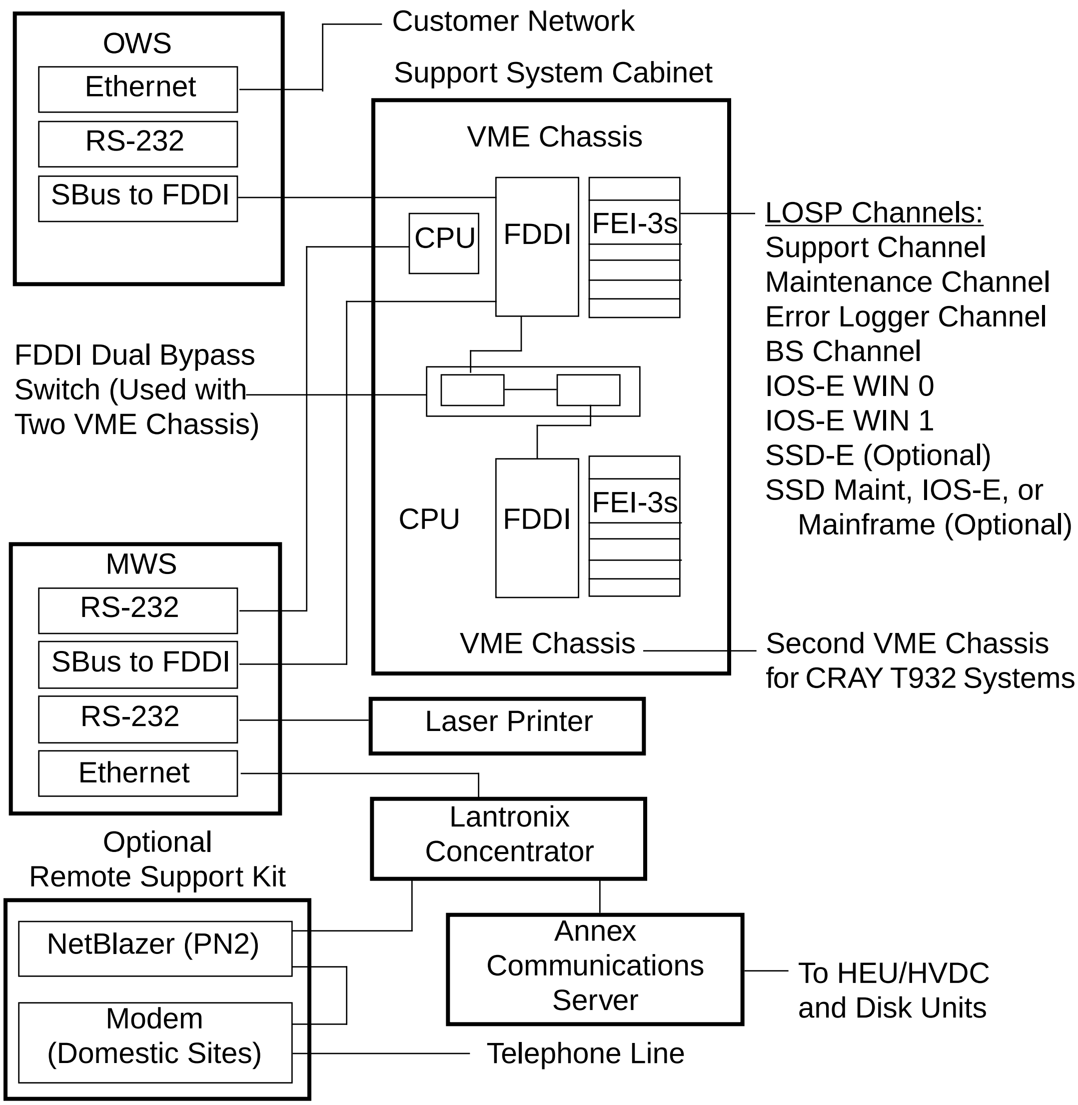

Support System Chassis

The support system cabinet is a cabinet enclosure that contains up to two 20-slot VME chassis, one stacked above the other. CRAY T94 and CRAY T916 systems use only the upper chassis (the lower chassis is empty). CRAY T932 systems use both the upper and lower VME chassis. Figure 21 illustrates the support system chassis.

Figure 21. Support System Chassis

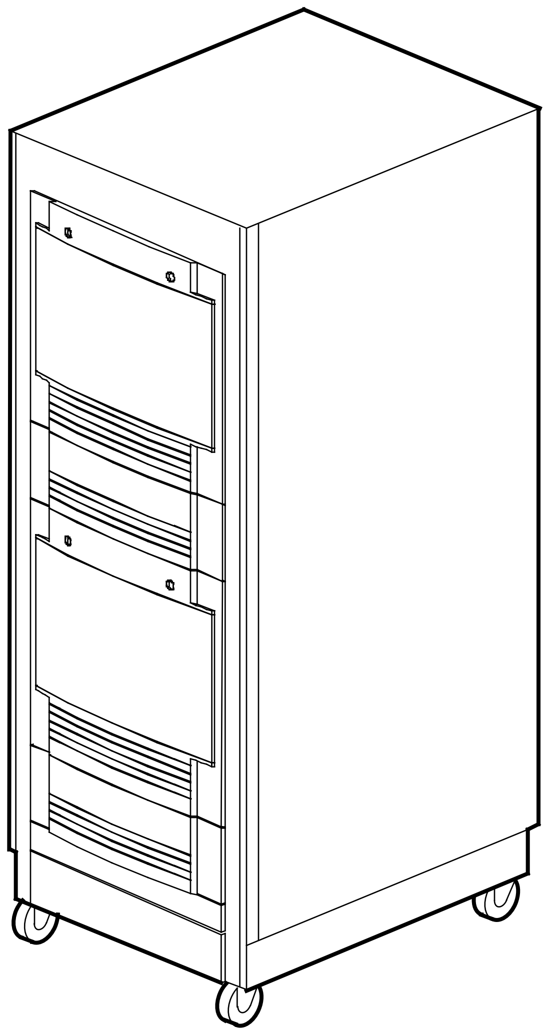

Each VME chassis is configured with a microprocessor-based CPU, up to eight FEI-3 two-board sets, and a fiber distributed data interface (FDDI) controller. A Motorola serial interface module provides eight RS-232-D serial channels on the CPU to provide maintenance connections to system disks and to the control system. A fan tray that contains three rotary fans provides airflow to cool the boards and power supplies in the VME chassis. Figure 22 shows the back view of the support system chassis.

The FDDI network is a 100-Mbit/s fiber-optic local area network that connects the MWS and OWS to the VMEbus interface in the support system cabinet. FDDI is a full-duplex network used as a high-performance interconnection between computers and peripheral equipment and as a high-speed backbone network for medium-performance local area networks (LANs).

A customer cannot connect his or her network to the FDDI ring. Instead, a customer can use the Ethernet port on the OWS to connect the network.

Figure 22. Back View of Support System Chassis

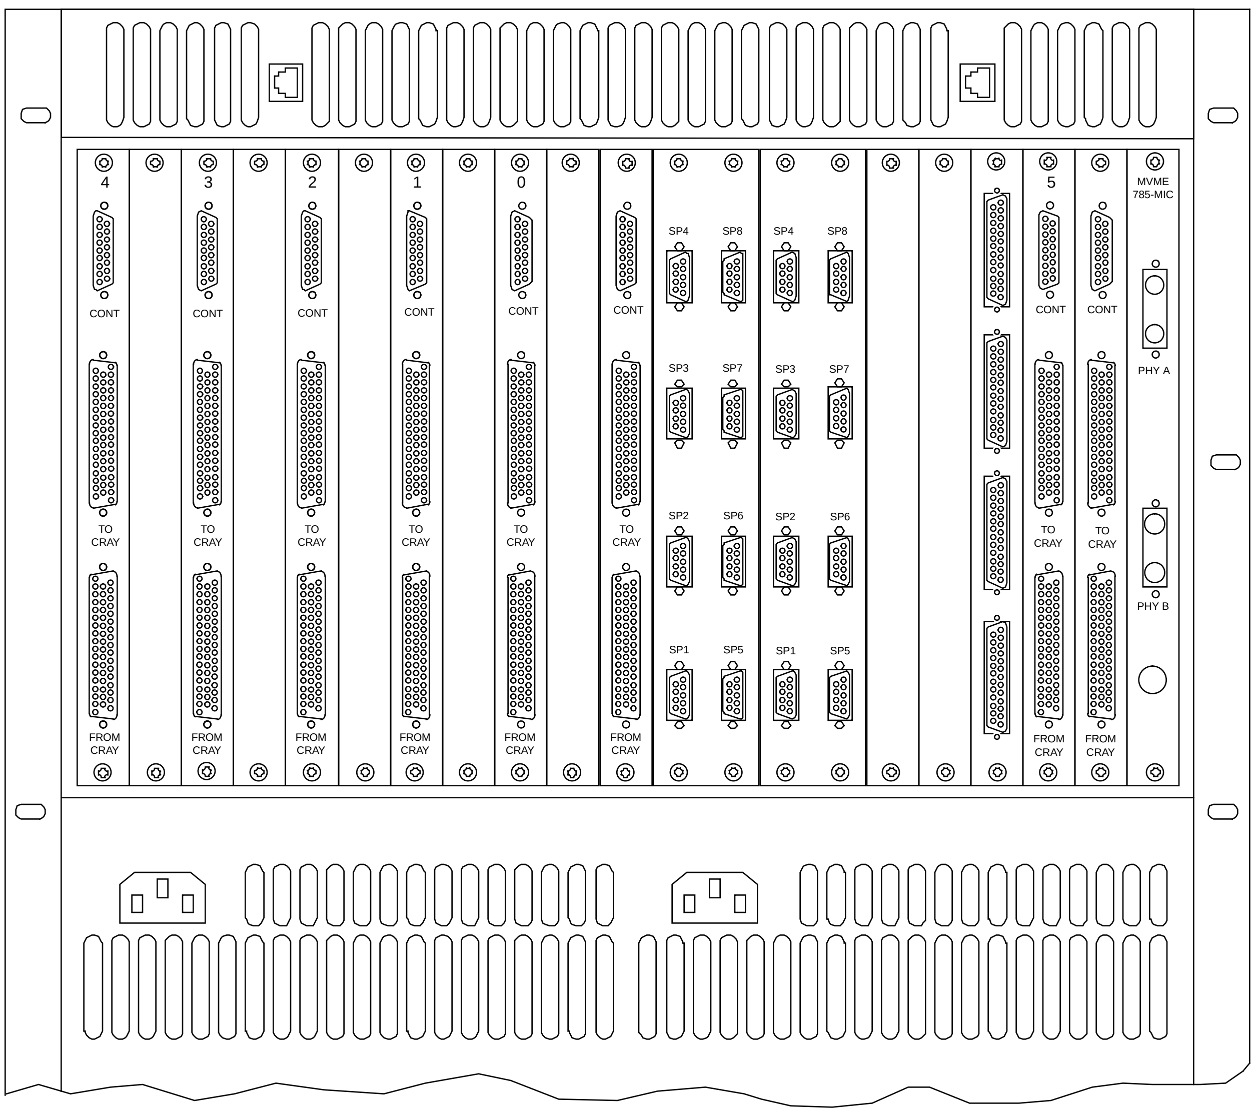

MWS and OWS Workstations

Cray Research uses Sun Microsystems, Inc. SPARCstation 5 workstations as the operator workstation (OWS-T) and the maintenance workstation (MWS-T) with CRAY T90 series systems. The MWS-T and the OWS-T consist of the same subassemblies and peripherals and are physically identical. For this reason, they have the same part number. Refer to Figure 23.

For customers concerned about security, Cray Research also offers a removable disk option for the MWS-T. This option consists of a base unit and a disk module case that plugs into the base unit. With this option installed, the customer can remove the hard drive from the MWS-T and place it in a secure storage area when the hard drive is not in active service.

NOTE: The MWS-T, though connected to the CRAY T90 series system, is the property of Cray Research and is intended exclusively for the use of authorized service personnel.

The FDDI node processor enables the MWS and OWS workstations and the 4-mm tape drive connected to the MWS to communicate with VMEbus-based devices and channels in the support system chassis.

Figure 23. MWS/OWS/SWS Workstation

- 110-MHz SPARC CPU
- 17-in. (408 mm) Color Display
- 535-Mbyte SCSI-2 Hard Drive
- 2.1-Gbyte SCSI-2 Hard Drive
- Internal CD-ROM Player (SCSI-2 Interface, Double-speed)
- Type 5 Keyboard and Optical 3-button Mouse
- 32-Mbyte DRAM memory
- FDDI Node Processor Board
- Deskside 4- to 8-Gbyte, 4-mm Digital Tape Drive

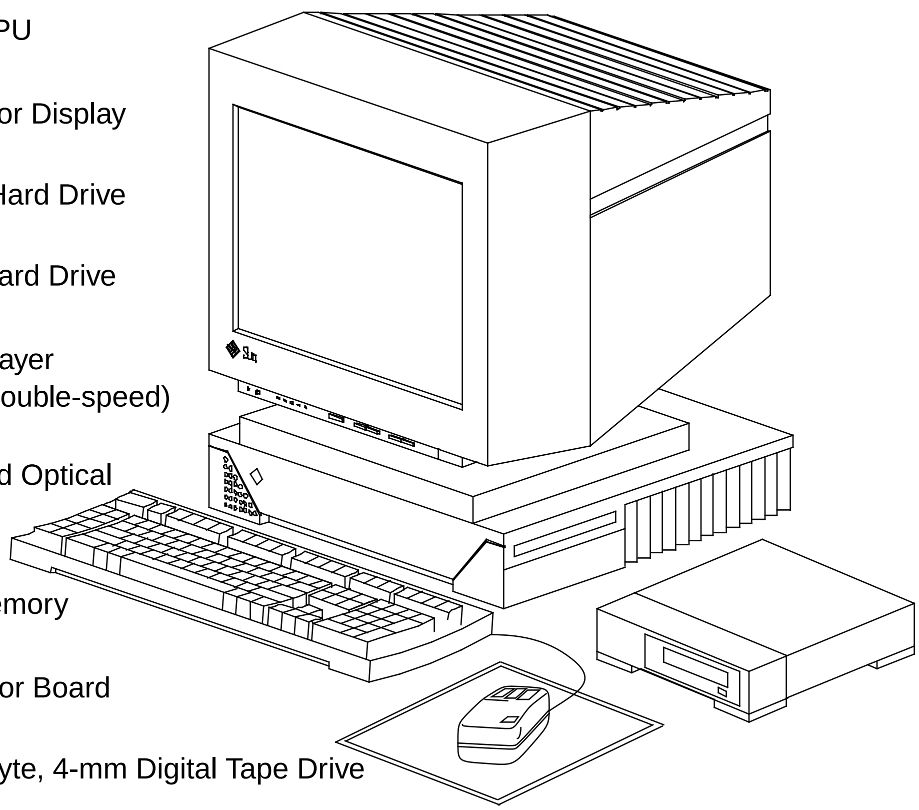

OWS Functions

The OWS is dedicated for the customer operator's use; the MWS is dedicated for use by Cray Research service personnel. However, in the event that one of these workstations fails, the other workstation could be used as a substitute until a replacement workstation is installed.

The OWS provides a dedicated workstation that customer operators and Cray Research analysts use to operate, administer, and monitor the CRAY T90 series system. The OWS is also used for system boot, dump, clear, and troubleshooting operations and for software support and upgrades.

MWS Functions

The MWS provides a dedicated platform for performing hardware maintenance, monitoring, and support of CRAY T90 series computer systems.

NOTE: The MWS, though connected to the CRAY T90 series system, is the property of Cray Research. The MWS is available exclusively for use by authorized service personnel.

The MWS performs the following functions:

- Provides a system maintenance platform for performing maintenance and running diagnostics on all equipment in the computer system
- Configures the mainframe
- Monitors control system status
- Provides remote maintenance capabilities and a control interface
- Detects and logs system error information
- Provides an OWS-compatible platform
- Provides a platform for retrieving reference documents and for training

IOS/SSD-E Equipment

The following types of IOS/SSD model E chassis are available with CRAY T90 series IO01 systems:

- 600 series with up to two I/O clusters and an SSD-E/128i
- 700 series with up to eight I/O clusters and four SSD-E sections (up to 2,048 Mwords of SSD memory)
- 800 series with up to three I/O clusters and one SSD-E section (512 Mwords of SSD memory)

A CRAY T916 system can support up to two IOS chassis, and a CRAY T932 system can support up to four IOS chassis.

Figure 24 illustrates the 600- or 800-series IOS/SSD-E chassis. A 600-series IOS/SSD is shipped with CRAY T94 systems. As options, a 700- or 800-series IOS/SSD can be used. For CRAY T94 systems, an HEU-T90 is used to cool the 600- or 800-series IOS/SSD chassis.

CRAY T916 and CRAY T932 systems are usually configured with a 700-series IOS chassis. An HEU-WC1 is used to cool a 700-series IOS/SSD that is used with any CRAY T90 series systems. For CRAY T916 and CRAY T932 systems that use 800-series IOSs, an RCU-8 or RCU-10 must be used because the HEU-T90 does not provide sufficient heat-rejection capacity.

The 600- and 800-series cabinets have the same outer appearance. The 600-series IOS/SSD-E shipped with the CRAY T94 system uses 10 slots to house two I/O clusters and an SSD-E/128i. The 800-series IOS/SSD-E contains 32 slots to house three I/O clusters and a 1-section SSD-E.

Figure 24. 600- or 800-series IOS/SSD-E Chassis

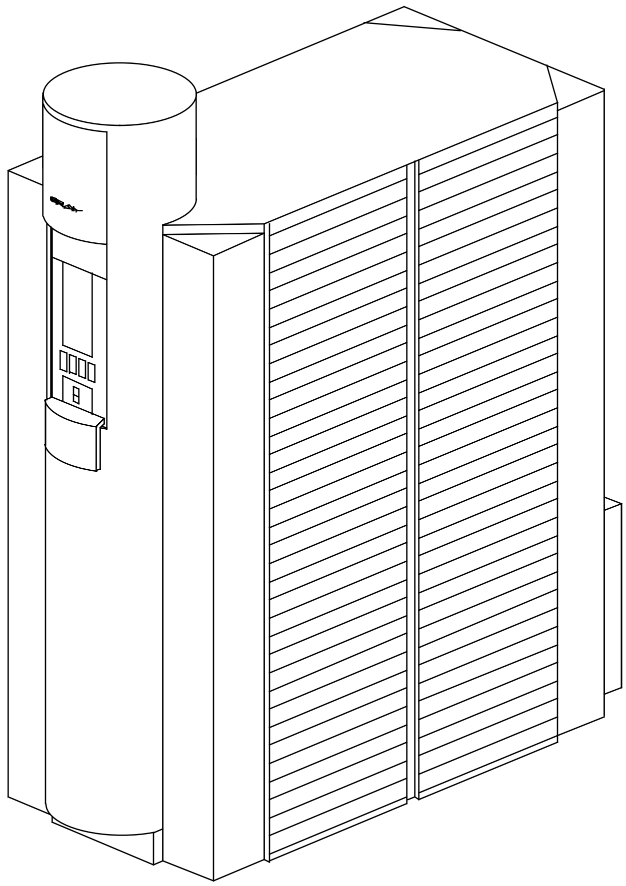

IO02 (Scalable I/O and GigaRing Channel) System Equipment

The scalable I/O (SIO) subsystem is a single-cabinet or multicabinet air-cooled architecture that provides scalable high-performance and high-resilience I/O support for CRAY T3E, CRAY T90 with IO02, and CRAY J90se series systems. Each SIO cabinet has a unique 8000-series serial number.

The basic elements of the SIO architecture are:

- PC-10 peripheral cabinet
- GigaRing channel
- Input/output node (ION)
- System workstation (SWS)

PC-10 Peripheral Cabinet

The SIO peripheral cabinet, called the PC-10 (Figure 25), is a 19-inch rack-mount cabinet that houses all SIO subracks that contain I/O nodes and peripheral devices. There are two models of the PC-10: PC-10A and PC-10B. CRAY T90 series SIO systems use the PC-10B.

All subracks are designed to use room air for cooling. The PC-10 has 36 standard units (SU) of rack space available for various SIO subracks. All PC-10 cabinets must have an integral power subrack that occupies 3 SU, thereby leaving 33 SU of usable customer space for other subracks.

The PC-10B features:

- Integral power subrack
- Power distribution strip (PDS)
- Warning and control system (WACS) monitoring
- Remote power-on/power-off
- Remote reset capabilities

Figure 25. PC-10 Cabinet

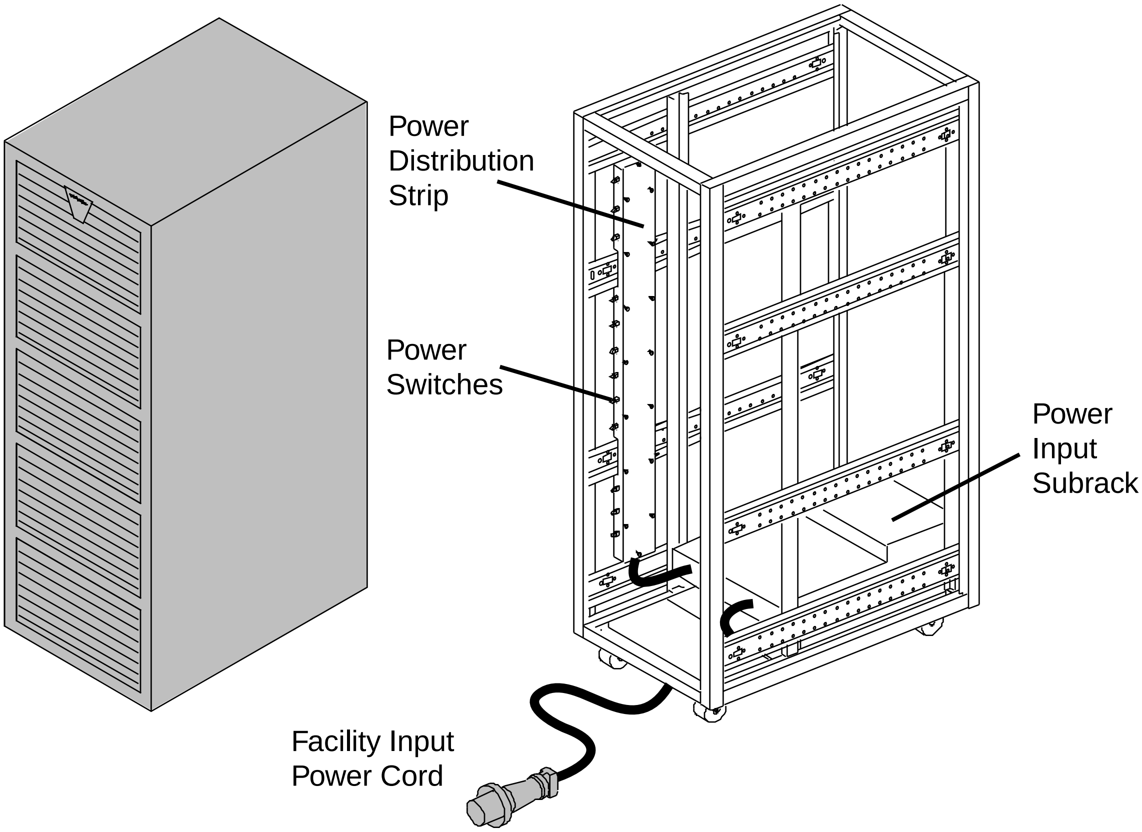

GigaRing Channel

The GigaRing channel design provides a common I/O channel architecture for Cray Research systems; the channel transfers data between IONs. The GigaRing channel consists of a pair of 600-Mbyte/s channels that are configured as counter-rotating rings. The two rings form a single logical channel with a maximum bandwidth of 1.2 Gbytes/s. Protocol overhead lowers the channel rate to 940 Mbytes/s.

GigaRing technology enables the connection of multiple clients on a ring-based channel. The GigaRing technology also enables multiple clients to coexist in a common operational and service environment, and it offers enhanced flexibility in configuring distributed systems. For example, a CRAY T90 series (IO02) computer system and a CRAY T3E computer system can attach to the same ring channel and share a common SWS to support this environment.

The GigaRing channel has a maximum distance of 36 ft (11 m) between IONs. A fiber-optic extender (FOX-1) can be used to increase this distance up to 650 ft (200 m). Two FOX-1 subracks are mounted in each PC-10B to provide the optical extension of the GigaRing channel.

IONs

An ION is any device that connects directly to the GigaRing channel. IONs connect the mainframe and disk, tape, and network peripheral devices to the GigaRing channel. A variety of IONs are available to support a wide range of connectivity requirements.

There are three types of IONs:

- Single-purpose node (SPN)
- Multipurpose node (MPN)
- Mainframe node

SPNs support specific channel interfaces and/or devices, while MPNs provide an SBus-based interface to support industry-standard I/O channels. Mainframe nodes are the GigaRing node chips within the mainframe that interface the mainframe I/O to the GigaRing channel. The GigaRing node chip in the CRAY T90 series mainframe resides on the IO02 module and logically connects the GigaRing channel to a single MPN-1 inside the PC-10 cabinet.

To provide greater reliability and resiliency, SPNs and MPNs are *hot swappable*. This means that power supplies, logic modules, and other components that are designated as field-replaceable units can be replaced without powering down or interrupting the entire SIO subsystem.

SWS Workstation

The system workstation (SWS) is the point of first-level access, control, and status of a CRAY T90 series (IO02) computer system that is attached to a GigaRing channel. The SWS is the focal point for all system operations and management and, at the same time, serves a second role of managing all service activities.

The presence of a single SWS departs from the traditional Cray Research practice that provides an operator workstation (OWS) for the customer and a maintenance workstation (MWS) for the service provider. The SWS combines the functionality of the OWS and the SWS onto a single platform, which is owned by the customer. The customer has access to the functionality of the online concurrent maintenance tools and service software formerly restricted to the MWS. The capability of a single workstation to support both operations and service is a cost-reduction measure that benefits Cray Research and its customers.

The SWS connects to the client system through a series of independent Ethernet and serial connections typically made to the SIO. The Ethernet connection is used for diagnosis, control, and configuration. The serial connection is used for monitoring. An Ethernet port on the back of the SWS is reserved for the customer's Ethernet network.

SWS Optional Components

The following SWS customer options are available; each component is a field replaceable unit (FRU).

- An 8- or 16-serial-port Quad Ethernet SBus card (installed in the SWS).
- An 8- or 16-serial-port *Micro Annex* communications server receives multiple serial data streams from the control system monitoring blocks or peripheral devices and directs them to the SWS, using a single Ethernet connection. A *Micro Annex* server is installed in the PC-10 cabinet if the SWS is more than 50 ft (15m) from the WACS connections, or if additional serial ports are required.

  NOTE: Systems with multiple nodes may require a serial port concentrator. To meet this requirement, an 8- or a 16-serial-port SBus card is typically used. The customer has the option of adding an 8- or a 16-serial-port *Micro Annex* server in place of, or in addition to, the serial port SBus card.

- A 600-dpi HP LaserJet 5MP laser printer. The SWS system does not require a printer. The DAT tape drive is used in place of the printer for sending back diagnostic loops or dumps and returning UNICOS or UNICOS/mk dumps for analysis. Printers that are supported by Cray Research must be connected to the parallel port on the SWS. Currently Cray Research supplies an HP LaserJet 5MP when a customer orders a printer.

- A second 2.1-Gbyte SCSI hard disk drive. A second hard drive can be used for system resiliency. Large system configurations may require a second hard drive.

- A removable disk option for customers concerned about security. This option consists of a base unit and a disk module case that plugs into the base unit. When this option is installed, the customer can remove the hard drive from the SWS and place it in a secure storage area when the hard drive is not in active service.

- Remote support equipment: a 28.8-kBd modem (default) or a 28.8-kBd modem and Telebit Netblazer router (optional).

PC-10 Component Block Diagram

Figure 26 shows PC-10 rack-mounted components and cable connections to the SWS and mainframe.

Figure 26. CRAY T90 Series SIO Block Diagram

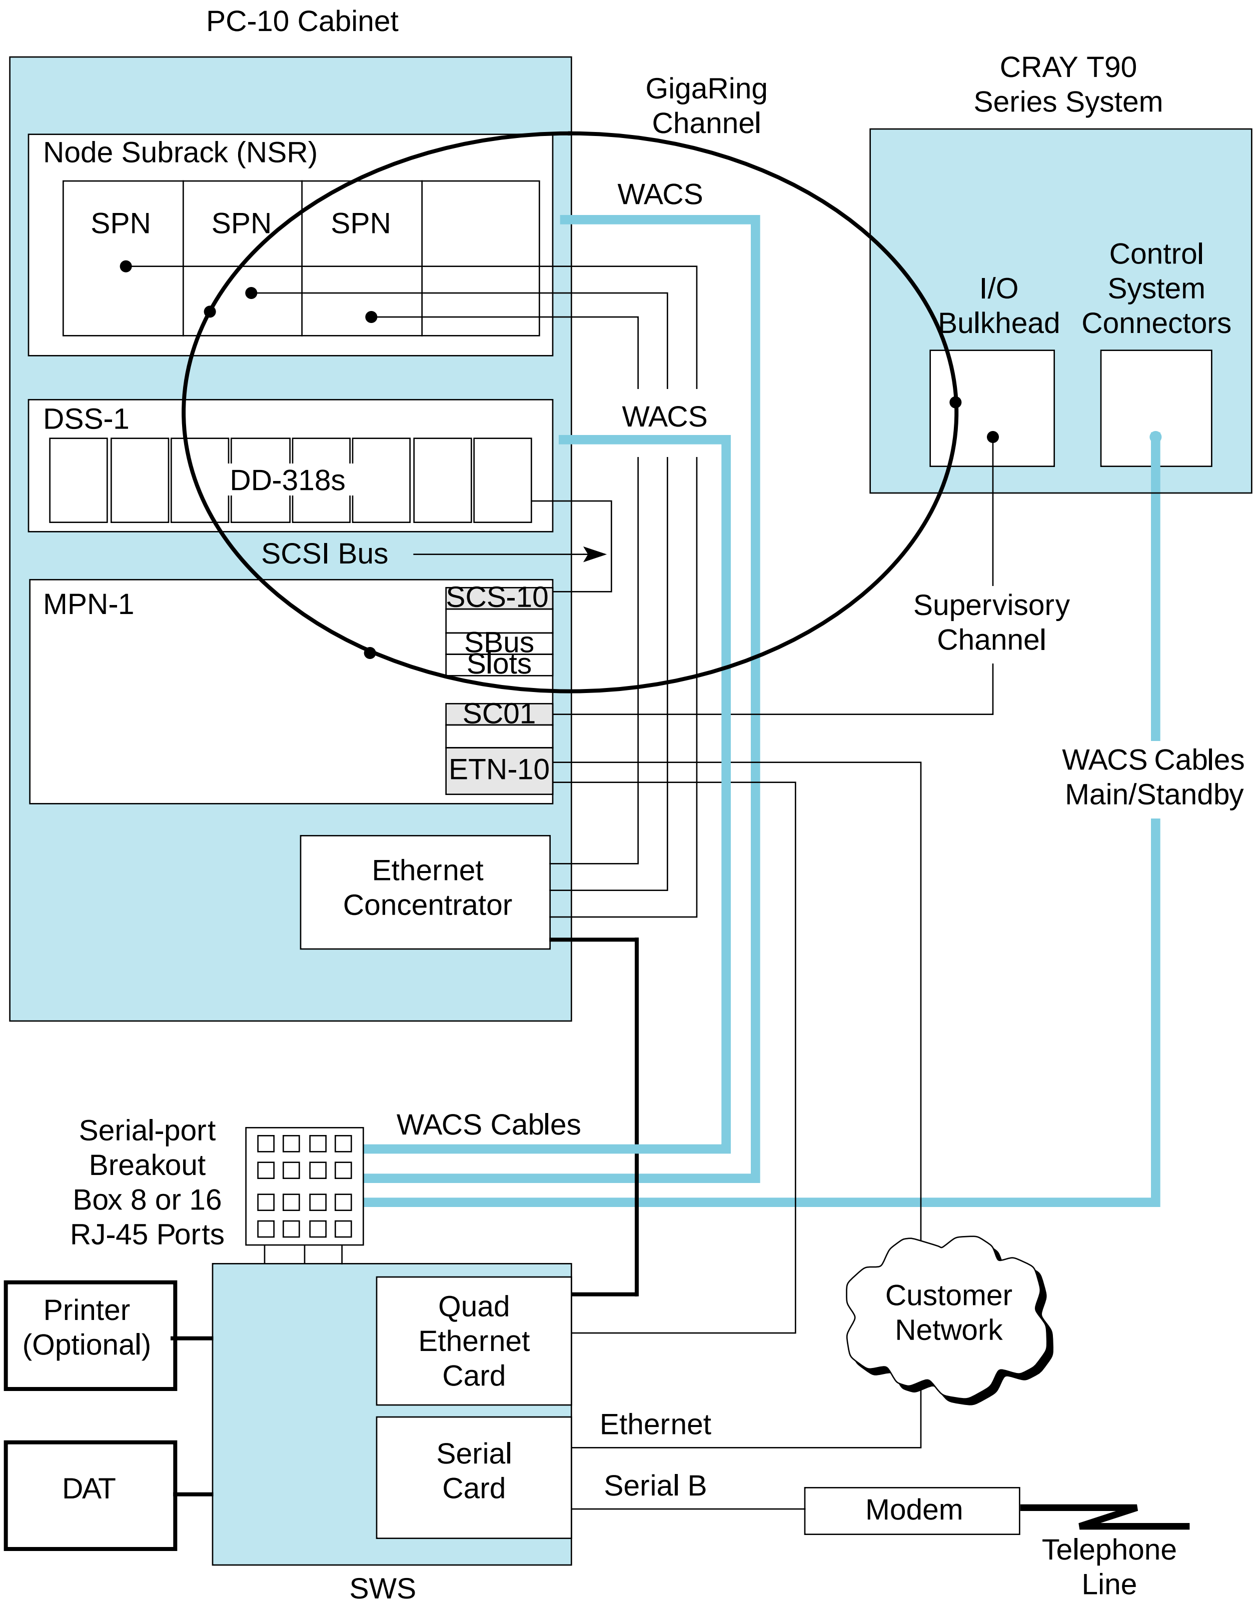

Laser Printer

An HP LaserJet printer can be connected to an MWS port, but it can also be accessed for printing from the OWS. Customers can order printers as optional equipment with SWS systems (printers are not included in the standard SWS configuration). The DAT tape drive can be used in place of the printer on SWS systems for sending back diagnostic loops or dumps and returning UNICOS or UNICOS/mk dumps for analysis. Only printers provided by Cray Research will be supported and connected to the SWS.

Support Channel

Support channel logic provides access to the operating system from the OWS in IO01 systems or from the SWS in IO02 systems. The support channel is used for file system transfers and general administrative functions.

Each IO module contains the logic to provide a single support channel; support channel logic is accessed through the supervisory channel in IO02 systems. CRAY T932 systems can have two support channels; only the primary IO modules (in locations A and M) are cabled to the support channel. CRAY T94 and CRAY T916 systems can each have one support channel.

Support channels are not assigned to a particular CPU for data and control routing as are other I/O channels. The System Configuration Environment (SCE) is used to select the CPU that is used for routing the support channel.

Power and Cooling Equipment

CRAY T90 series computer systems use different power and cooling systems than those used with previous Cray Research computer systems. CRAY T90 series systems use a high-voltage direct-current (HVDC) cabinet that provides power to all components in the system. Motor-generator sets (MGSs) and refrigeration condensing units (RCUs) are not used (an MGS and separate HEU are used if a 700-series IOS/SSD is part of the system).

Unlike previous Cray Research systems, CRAY T90 series systems do not use a refrigerant. This eliminates the requirement for an RCU. Eliminating the RCU has many benefits, which range from reduced power consumption (up to 15% on large computer systems) to increased cooling system reliability.

This subsection describes the following power and cooling components:

- HEU-T90
- HEU-WC1
- Air-cooled cooling unit
- HVDC
- UPS

HEU-T90

Mainframe modules and power supplies are submersed in dielectric coolant. Because the coolant is chemically inert, it does not absorb or dissolve contaminants. This enables the submersed electronic components to operate without being affected by the coolant.

The heat exchanger unit (HEU-T90) routes dielectric coolant through the CRAY T90 series mainframe and IOS/SSD chassis to absorb heat generated by the modules and power supplies. After the dielectric coolant absorbs the heat, it flows back to the HEU-T90 where the heat is transferred to customer-supplied water. Figure 27 shows the HEU-T90 chassis.

Figure 27. HEU-T90

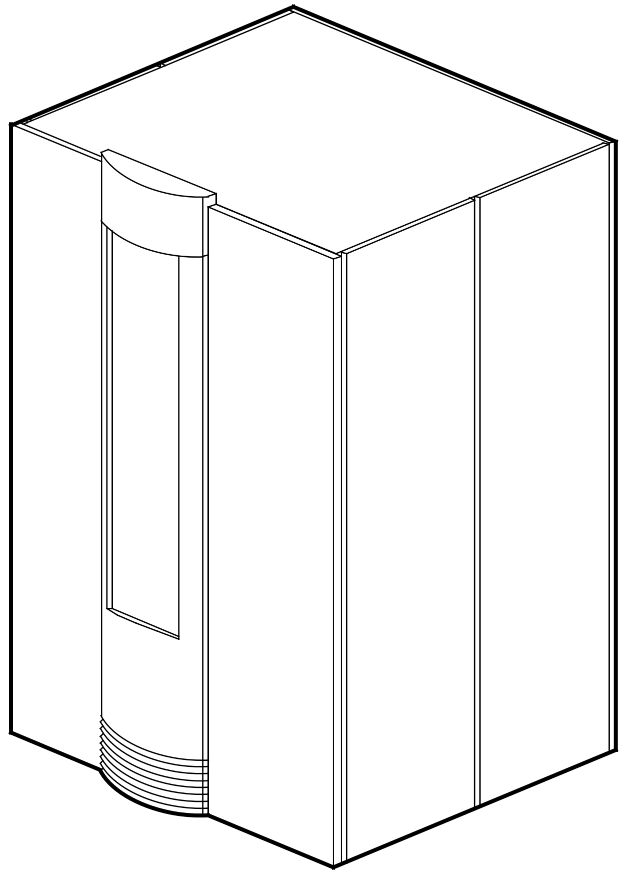

Each CRAY T94 and CRAY T916 system uses a single HEU-T90. Each CRAY T932 system requires two HEU-T90s because of the physical boundary that enables one half of the system to be powered off and drained while the other half continues to function. Each HEU-T90 supports one-half of the chassis.

The HEU-T90 used with CRAY T94 systems has a smaller heat exchanger vessel than the vessels in the HEU-T90s used with CRAY T916 and CRAY T932 systems. Except for this smaller heat exchanger vessel, the HEU-T90s shipped with each CRAY T90 series system are the same.

Control system monitors within the HEU-T90 ensure that temperatures and pressures remain within specific ranges during system operation. If the conditions are out of range, the control system either adjusts valves within the HEU-T90 to compensate for the out-of-range condition or shuts the computer system down to protect the equipment and computer room environment from damage. Refer to the “Control System” for more information.

Level Sensors

Six level sensors monitor the presence of dielectric coolant at specified levels within the CRAY T90 series mainframe. Figure 28 shows the level sensor (LS) locations in a CRAY T94 mainframe. These sensors are located on the same side as the supply and return lines and are used to fill and drain the mainframe. Two level sensors (primary and backup) monitor the presence of dielectric coolant in the top area of the mainframe. Two additional level sensors (primary and backup) monitor the presence of dielectric-coolant in the power supply area of the mainframe. Two more sensors in the drain area and cable drop area monitor the presence of dielectric coolant present in the mainframe during the draining process.

Figure 28. Level Sensor Locations in a CRAY T94 Mainframe

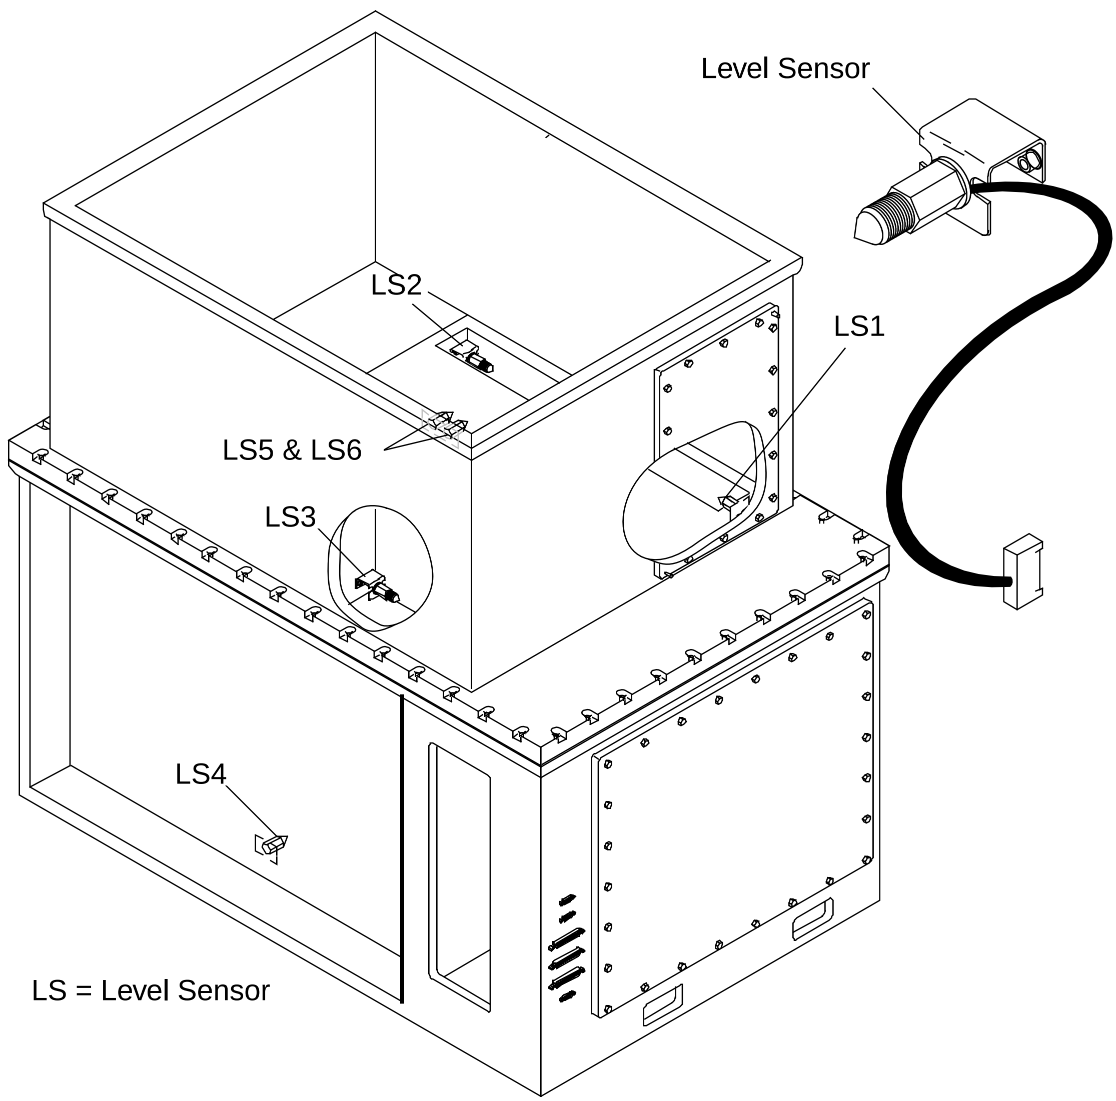

HEU-WC1

CRAY T90 series systems use a water-cooled HEU (HEU-WC1) to cool a 700-series IOS/SSD chassis. This heat exchanger unit is a dielectric coolant-to-water heat exchanger that is similar to the HEU-T90.

The HEU-WC1 has two heat exchanger subassemblies; however, depending upon the heat load, one of the heat exchanger subassemblies can be closed off. The HEU-WC1 control system uses pressure, temperature, and flow sensors to monitor conditions within the IOS/SSD. These sensors send the monitored values to a control board on the HEU-WC1, which connects to the IOS/SSD warning and control system (WACS). The HEU-WC1 was originally designed for CRAY T3D systems.

Air-cooled Cooling Unit

If a CRAY T94 site does not have a chilled-water supply, an air-cooled cooling unit can substitute for a chilled-water supply. The air-cooled cooling unit can either discharge heat into a computer room with an indoor air-heat exchanger or it can discharge heat outdoors using an outdoor-air heat exchanger.

HVDC

Power supplies in CRAY T90 series mainframes operate on 330-Vdc power. Commercial utility power at customer sites may vary from 200 Vac to 480 Vac at frequencies of 50 or 60 Hz. The high-voltage direct-current (HVDC) cabinet transforms any of these ranges of input voltages and frequencies. Figure 29 shows the HVDC systems. The HVDC transforms, rectifies, and regulates input power from the customer’s utility to deliver conditioned 330-Vdc power to system components.

The following HVDC systems are used with CRAY T90 series systems:

- One HVDC-40 (40 kW) for CRAY T94 systems
- One HVDC-160 (160 kW) for CRAY T916 systems
- Two HVDC-160s for CRAY T932 systems

Figure 29. HVDC Cabinets

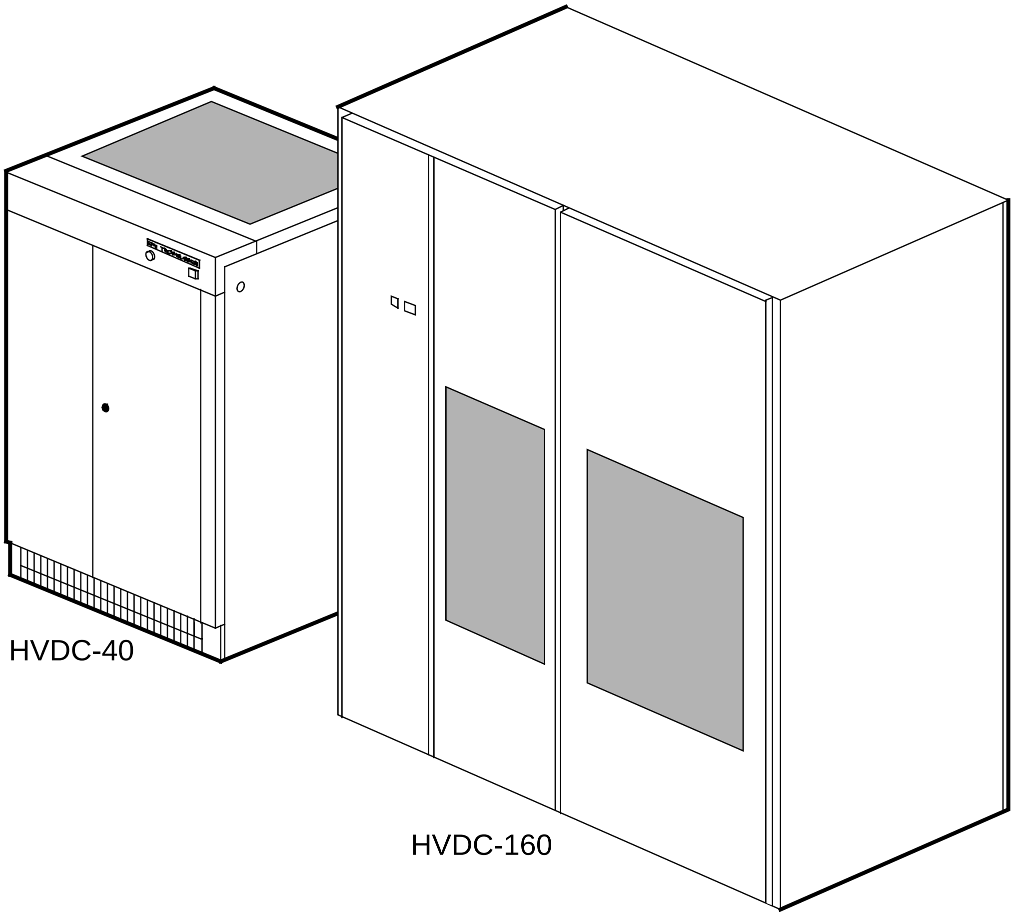

Preshipping HVDC Equipment

Cray Research Site Planning does not plan to preship the HVDC-40 to customers in advance of the rest of the CRAY T94 system. Cray Research will, however, preship the HVDC-160s used with CRAY T916 and CRAY T932 systems.

The HVDC-40 is quality checked with the CRAY T94 system in Systems Test and Checkout (STCO) and the entire system is shipped together. The HVDC-40 wiring consists of one input circuit and five output circuits. With proper preinstallation planning and preparation, the amount of time required to wire the HVDC-40 should be less than 4 hours. The installers should be able to power up CRAY T94 systems on the first day of system installation.

NOTE: A required MG set is preshipped to CRAY T90 series sites that have a 700-series IOS/SSD chassis.

Service and Support for HVDC Equipment

Service for HVDC equipment is provided by Cray Research. This means that Cray Research responds to all maintenance and service calls for this equipment. Cray Research provides all repairs and preventive maintenance. Any HVDC situations that require escalated support should follow existing escalation procedures. The HVDC-40 and HVDC-160 units require minimal preventive maintenance that consists of replacing air filters and verifying operating limits.

Trained Cray Research personnel will complete HVDC start-up as covered by the HVDC vendor's warranty. If Cray Research personnel are not available to provide HVDC start-up procedures, the HVDC vendor will perform the start-up and charge Cray Research based on the vendor's time and material rates.

UPS

If CRAY T94 customers desire power ride-through in excess of 80 milliseconds, Cray Research can provide an uninterruptible power supply (UPS). A UPS is available with enough capacity to power the CRAY T94 mainframe, HEU-T90, IOS/SSD chassis, and peripherals. The UPS system for CRAY T94 systems is an optional piece of equipment that the customer may purchase.

As an option, Cray Research provides a UPS system that attaches to the 160-kW HVDC used with CRAY T916 and CRAY T932 systems. The UPS system is large enough to provide uninterruptible, conditioned power to all Cray Research system components, including the mainframe via the HVDC, workstations, IOS, peripherals, and HEU(s).

The UPS used with an HVDC-160 is not a system-critical device. A bypass circuit on the UPS is available to enable service personnel to perform maintenance on the UPS while the CRAY T90 series system is in operation.

The UPS system used with the CRAY T90 series products is produced by EPE Technologies in the United States of America and Merlin Gerin in Europe. (Both EPE Technologies and Merlin Gerin are subsidiaries of Square D Corporation.)

The optional UPS system supplied by Cray Research provides approximately 2.5 minutes of power ride-through time. Customers may purchase additional battery cabinets to lengthen the ride-through time to the time they desire, with approximately 9 minutes being a reasonable maximum. The UPS uses sealed, maintenance-free batteries that do not require any special environmental considerations. However, the batteries must reside in an ambient room temperature below 77 degrees F (25 degrees C).

Customer Use of Existing UPS Systems

Customers may connect their existing UPS systems to the HVDC-40 used with CRAY T94 systems. Customers cannot directly connect a UPS to an HVDC-160.

Customers may also connect other equipment within their facilities (such as network equipment and chilled water pumps) to the UPS that Cray Research provides, as long as the added load does not exceed the kW limitations of the UPS.

Advantages to Using the HVDC and UPS

There are several advantages to using an HVDC and UPS rather than an MG set. The HVDC increases electrical efficiency and reduces required floor space and weight. The MGs were custom built specifically for Cray Research, which made the MG much more expensive than commercially available products like the HVDC and UPS.

The MG provides conditioned, isolated power only to the mainframe, IOS, and SSD. The other components in the system run on utility power, unless the customer provides conditioned power. The HVDC units provide conditioned, isolated power to the mainframe, HEU, and control system. The UPS significantly increases hold-up times. CRAY T90 series system availability should be increased because of the increased hold-up times and because conditioned power is supplied to all components in the system.

The UPS systems used with CRAY T90 series systems are standard products that are shipped with 5-week lead times. The custom-made MG set required a 26-week lead time.

Control System

The CRAY T90 series control system monitors various conditions within system components and, if necessary, shuts the system down if a monitored condition deviates from its predefined range. The control system is based on GE Fanuc Series 90-70 programmable logic controllers (PLCs) and the GE Fanuc Genius I/O System. Figure 30 is a block diagram of the CRAY T932 control system.

The control system consists of the following components:

- GE Fanuc Series 90-70 programmable logic controllers (PLCs)
- GE Fanuc Genius I/O System and communications system - analog, digital, and 115-Vac I/O remote monitoring blocks
- GE Fanuc Genius I/O System hand-held monitor
- MWS software - `nwacs` program
- Sensors

Figure 30. CRAY T932 Control System Block Diagram

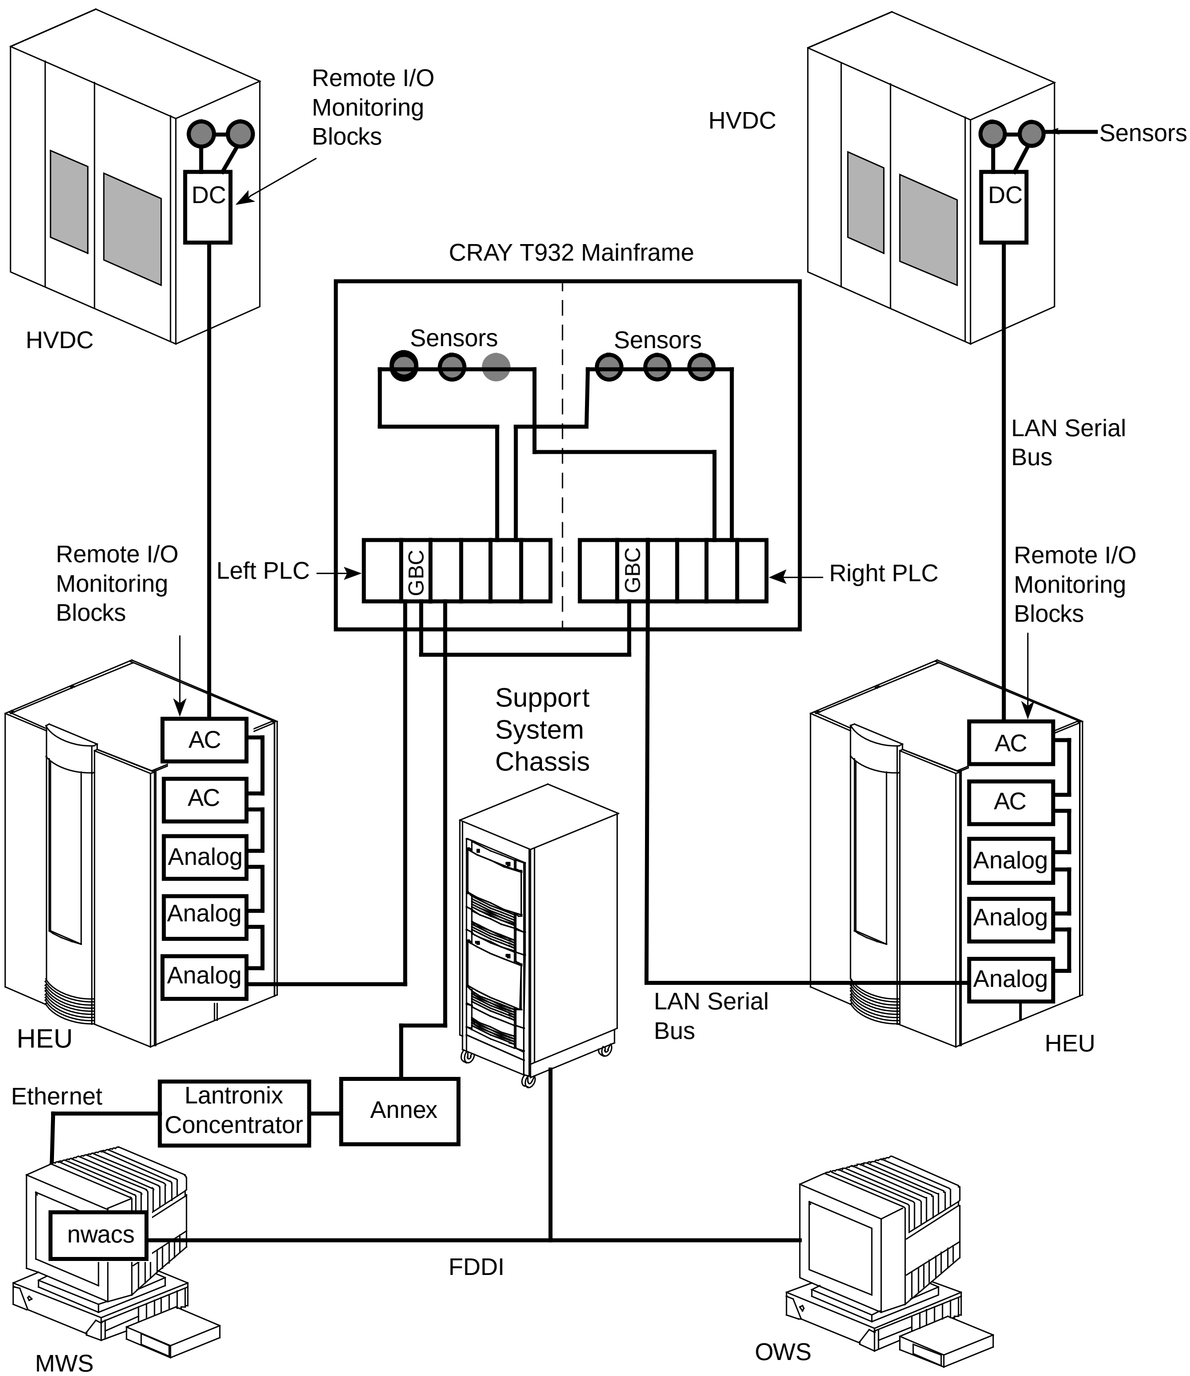

NOTE: A PLC in each side of the CRAY T932 mainframe controls half of the system. If one PLC malfunctions, the other PLC can be used to control the entire computer system. CRAY T94 and CRAY T916 systems each have a primary and secondary PLC, but both systems have only one HEU and one HVDC.

Programmable Logic Controllers

Two programmable logic controllers (PLCs) monitor and control the computer system. Two PLCs (main and standby) are used for system reliability and redundancy purposes. Each PLC is identical except that the redundant PLC does not contain an analog output card, which enables power supply margining. If the main PLC fails, the standby PLC controls the computer system until the main PLC is repaired and returned to operation.

The PLCs and related components (power supply, local CPU, state logic processor, and other modules) mount in a 9-slot VMEbus rear-mount rack as shown in Figure 31. This rack can hold ten PLC components: one power supply and nine other modules.

The control system operates on a local area network (LAN) in which the PLCs are the central components (Figure 30). Each component within the system is daisy chained and communicates through the LAN serial bus.

In CRAY T94 systems, the two PLCs (main and standby) are located next to each other on the right side of the mainframe. In CRAY T916 and CRAY T932 systems, the PLCs are located in the bottom of the chassis next to the I/O bulkhead; the main PLC is in quadrant 1, and the standby PLC is in quadrant 3.

Figure 31. Programmable Logic Controller

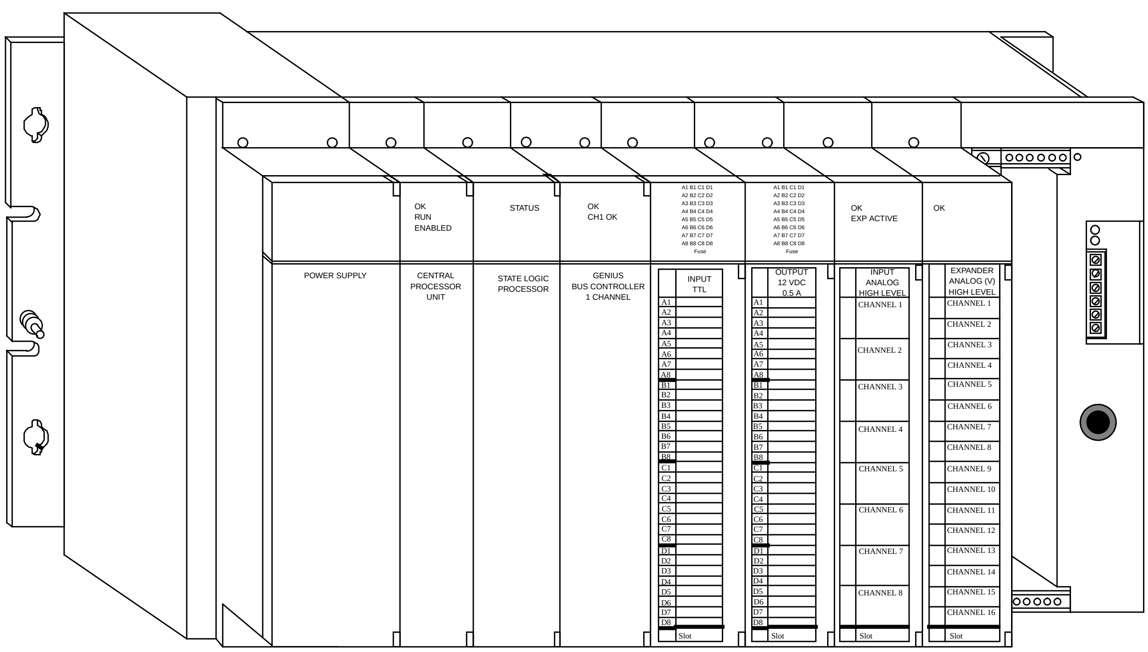

I/O Remote Monitoring Blocks

Both PLCs receive data from the I/O remote monitoring blocks located in the HVDC and HEU units and from the sensors located in the mainframe (Figure 30). However, the remote I/O monitoring blocks receive commands only from the main PLC, unless the main PLC does not report to the remote I/O monitoring blocks for three bus scan cycles. After three scan cycles, the remote I/O monitoring blocks switch over to the standby PLC.

The Genius I/O bus controller (GBC) is a module in the PLC that provides the interface and control to the I/O remote monitoring blocks. The GBC sends I/O information along the VME backplane of the PLC to the CPU module. Any commands from the CPU module in the PLC are routed through the GBC to the remote monitoring blocks.

Sensors

The control system operation relies on the information returned from the sensors, which are located in the mainframe chassis, HVDC, and HEU. Sensors monitor temperature, pressure, power bus voltages, power supply currents, and dielectric-coolant fluid levels and flow rates. The sensor location determines the routing of the sensor information.

The sensors located within the HEU and HVDC route information to the I/O remote monitoring blocks located on the device. The information is routed through the control system to the GBC modules in each PLC.

Status Displaying Software

The MWS or SWS runs a software program called `nwacs` to display the status information of the control system. The `nwacs` program is used to control the system; for example, `nwacs` adjusts voltage levels and dielectric-coolant flow rates, and controls power-up and power-down processes.

Hand-held Monitor

The hand-held monitor (HHM) is used to troubleshoot and configure the control I/O system components. The hand-held monitor plugs into the GBC and I/O remote monitoring block monitor ports. Figure 32 shows the HHM.

Figure 32. Hand-held Monitor

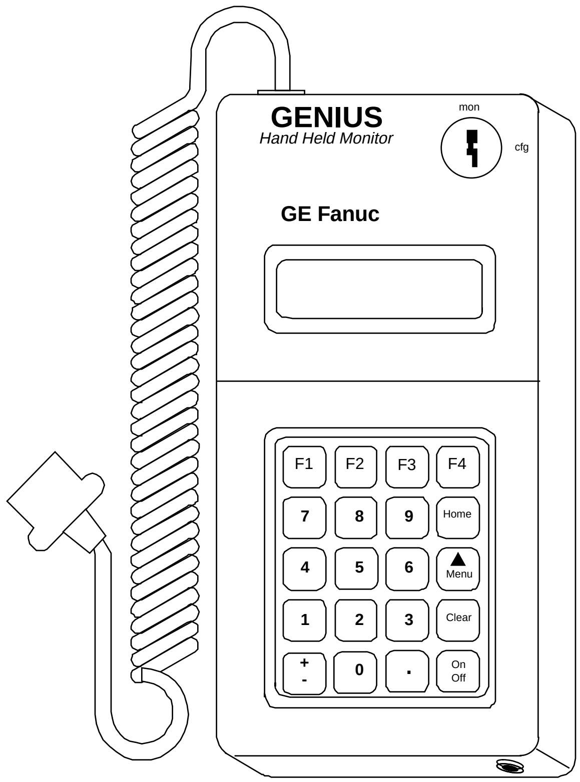

Documentation and Information Sources

The following subsections describe the CRAY T90 series documentation sources.

Online and CD-ROM Documentation

All CRAY T90 series and Scalable I/O documentation produced by Service Publications and Training (SPT) and the service and support and spares plans described in this subsection are available online through the Online Documentation InterfacE (ODIE). This information is also available on CD-ROM for use by field service personnel; the part number for the CD-ROM that includes CRAY T90 series information is CD-VOL1:1.0. The part number for the SIO Architecture Documentation CD-ROM is CD-VOL3:1.0.

The CRAY T90 Series Multimedia Release, part number CD-MMR1:1.0, contains a multimedia job aid that provides quick access to troubleshooting information, field replacement procedures, and a parts catalog for field engineers who maintain CRAY T90 series systems. The troubleshooting information lists nwacs error messages, controller subrack error messages, and miscellaneous errors and links them to a corrective action. The corrective action could result in performing a field replacement procedure. The replacement procedures are supported by color illustrations, video clips, and animations that enhance difficult steps within the procedure.

The procedures include a set-up menu that provides a link to the photographic parts catalog. The parts catalog contains photographs of CRAY T90 series system parts, their part numbers, the product (CRAY T916 or CRAY T932 mainframe, CRAY T94 mainframe, or HEU), and a common name (if applicable). The parts catalog enables the user to search for a part by part name or part number.

The CD-MMR1:1.0 CD-ROM also contains PostScript files (letter size and A size) of the procedures documented in the CRAY T90 Series *Field Replacement Procedures*, publication HMM-111-0, for users who wish to print a hard copy of the procedures.

Service and Support Plan

The Service Planning department writes service and support plans and spares plans for each major Cray Research product. Each plan is an index to the various service and support features for the product. The Service and Support Plans do not replace Cray Research technical instructions. They are intended to provide insight into CRAY T90 series on-site service and support functions.

You can order the *CRAY T90 Service and Support Plan*, part number NPS19945, from Logistics. A PostScript version is available in the `/home/hydra/spe/triton` directory on hydra; you can access this document, and the spares plans described below, in the CRAY T90 series ODIE online documentation set. Each service plan is an index to the various service and support features for the product.

The *CRAY T94 Spares Support Plan* and *CRAY T916 and CRAY T932 Spares Support Plan* are available in the `/home/hydra/spe/triton/spares` directory on hydra. The spares support plans list all field replaceable unit (FRU) parts and describe support levels for available service options.

Hardware Training Courses

T90MC *(CRAY T90 Series Maintenance Course)*. This 15-day course is designed to teach the skills necessary to maintain and troubleshoot CRAY T90 series systems. The course focuses on the use of maintenance tools to maintain CRAY T90 series systems. The course combines theory and practice in the lab.

T90INTR *(CRAY T90 Series System Introduction)*. This 1-day course presents CRAY T90 series system architectural enhancements. The information describes system types, module types, cooling technology, power requirements, and the control system. This course also includes a discussion of the troubleshooting techniques used on a CRAY T90 series system.

T90INTR-VC (*CRAY T90 Series System Introduction-Videoconference).* This 3-hour videoconference course presents CRAY T90 series system architectural enhancements. The information describes system types, module types, cooling technology, power requirements, and the control system. The course also includes a discussion of troubleshooting techniques used on a CRAY T90 series system.

The T90INTR-VC course is available to any site that has videoconference capabilities. The site is responsible for operation of videoconference equipment. Copies of the instructor's overheads are supplied for the class. No other materials are required.

SIO-S (Cray Scalable I/O Architectural Overview). This self-directed learning course introduces students to the scalable input/output (SIO) basic architecture. The course consists of computer-based training (CBT) modules that consist of reading assignments and exercises that evaluate the student's progress. The training modules guide the student from an overview of the SIO architecture, through the logical components and communications, to a general understanding of SIO architecture (hardware and mechanical components, control systems, cooling, and power distribution). This course takes approximately 20 hours to complete.

All technical support and site personnel who enroll in the SIOM course of the CRAY T3E Maintenance (T3EM) course must complete this course first. The information presented in this CBT is not repeated in the SIOM or T3EM classes.

SIO-M (Cray Scalable I/O Maintenance). This 2-week course provides I/O training for Worldwide Customer Service (WCS) personnel who maintain CRAY T90 or CRAY J90se series systems that connect to Scalable Input/Output (SIO) architecture. The course familiarizes students with the maintenance strategy for the SIO architecture. This course emphasizes SIO maintenance and the skills necessary for diagnostic troubleshooting. The diagnostic troubleshooting portion of this course includes a simulator that trains students to identify a failing field-replaceable unit (FRU), and how to configure or reconfigure the SIO architecture and system workstation (SWS). In addition, this course instructs the student how to replace all FRUs for the SIO architecture.

A practical lab provides experience in troubleshooting the SIO architecture, running concurrent diagnostics, configuring the system, and performing FRU procedures.

Software Training Courses

T90WS *(CRAY T90 Series Workshop)*. This 3-day workshop is for Cray Research employees and customers who need an understanding of the CRAY T90 series system architecture and software differences. Students should have knowledge of and experience with UNICOS 8.0. This workshop limits its focus to those features included in UNICOS 8.3 that are specific to a CRAY T90 series system. This workshop is taught at the Cray Research Corporate Training Center in Eagan, Minnesota.

The first day of class, which is taught by an SPT instructor, focuses on system architecture and online diagnostic tools. The rest of the workshop highlights the software environment including start-up/shut-down, configuration, installation, administration, programming, internals, networking, accounting, and other software differences.

Reader Comment Form

Title: **CRAY T90™ Series System Overview** **Number:** **HGM-002-C**

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this document?

______Troubleshooting ______Tutorial or introduction
______Reference information ______Classroom use
______Other - please explain

____

Using a scale from 1 (poor) to 10 (excellent), please rate this document on the following criteria and explain your ratings:

______Accuracy ________________________________
______Organization ________________________________
______Readability ________________________________
______Physical qualities (binding, printing, page layout) ________________
______Amount of diagrams and photos ________________________
______Quality of diagrams and photos ________________________

Completeness (Check one and explain your answer)

______Too much information _____ Too little information ______Correct amount

____
____
____

You may write additional comments in the space below. Mail your comments to the address below, fax them to us at +1 715 726 4353, or E-mail them to us at *spt@cray.com*. When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

____
____
____

NAME ________________________________
JOB TITLE________________________________
E-MAIL ADDRESS ________________________________
SITE/LOCATION ________________________________
TELEPHONE________________________________
DATE ________________________________
[or attach your business card]

CRAY RESEARCH, INC.

Attn: Service Publications and Training
890 Industrial Boulevard
P.O. Box 4000
Chippewa Falls, WI 54729-0078