CRAY J98 and CRAY J916 Systems Hardware Overview

HMM-094-B
CRAY J90 Series
Last Modified: April 1998

Figures

Tables

Record of Revision

March 1995

Original printing.

Revision A: June 1995.

Revised to include the master index, glossary, and documentation map for the *CRAY J916 Service Manual Kit*, HMK-101-0.

Revision B: April 1998.

Revised to include the 512-Mword memory options and to include the CRAY J98 system in the title.

CRAY J916 Service Manual Kit

The following figures (refer to Figure 1 and Figure 2) illustrate the contents of the *CRAY J916 Service Manual Kit* and CRAY J90 series videotapes. A list of other related documents is also included. Hardware Publications and Training recommends placing the documents in the kit binder in the order that Figure 1 shows.

The CRAY J98 and CRAY J916 system configurations are included in the CRAY J90 series system configurations. The CRAY J98 system contains a 2 X 2 backplane and is housed in the CRAY J916 cabinet. For documentation purposes, the CRAY J98 system, unless otherwise noted, is functionally equivalent to a CRAY J916 system.

Figure 1. CRAY J916 Service Manual Kit

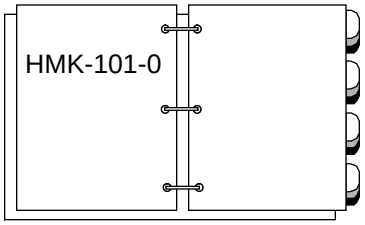

HMK-101-0 CRAY J916 Service Manual Kit

HMM-094-B	Hardware Overview
HMM-118-B	Processors and Memory
HMM-098-D	Input/Output Subsystem
HMM-100-A	Power, Cooling, and Control
HMM-114-B	System Troubleshooting
HMM-073-B	Remote Support
HDM-110-A	Automated Confidence Testing
HDM-099-0	IOS Based Diagnostics
HMM-079-B	Field Replacement Procedures
HMM-116-E	Illustrated Parts Catalog
HGM-016-A	Safety and ESD Guidelines

Figure 2. CRAY J916 Videotapes

CRAY J916 Service Videotapes

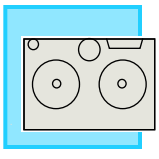

AVT- 00209 CRAY J90 Series Installation

Document Descriptions

The following paragraphs describe all documents in the *CRAY J916 Service Manual Kit*, publication number HMK-101-0. The title and publication number of each document is included, followed by the document description. The publications are listed in the order in which they should be arranged in the kit binder. The *CRAY J916 Service Manual Kit* provides an overview of the CRAY J98 and CRAY J916 systems.

Hardware Overview, HMM-094-B

This document describes the hardware components of the CRAY J98 and CRAY J916 computer systems and provides a brief overview of the system software and diagnostics. It also provides a master index, glossary, and documentation map for the *CRAY J916 Service Manual Kit*, HMK-101-0.

Processors and Memory, HMM-118-B

This document describes the memory modules and processor modules in a CRAY J98, CRAY J916, or CRAY J932 system. The descriptions include ASIC functions, scalar cache, interprocessor communications, processor control, channel communications, and memory addressing.

Input/Output Subsystem, HMM-098-D

This document contains information about the Cray Research, Inc. input/output subsystem that the CRAY J90 series systems use. It includes a system overview, the system theory of operation, and information about the peripheral devices that are connected to the input/output subsystem.

Power, Cooling, and Control, HMM-100-A

This document describes the electrical power, cooling, and control systems of the CRAY J90 series systems. It describes each assembly and its electrical cabling. The document also describes the function of all indicators and switches on the central control unit (CCU) of the mainframe cabinet. It also describes power control during power-up and power-down.

System Troubleshooting, HMM-114-B

This document can be used as a guide to troubleshoot the CRAY J90 series computer systems. Information is provided on system-level troubleshooting, IOS troubleshooting, CPU and memory troubleshooting, and peripherals troubleshooting.

Remote Support, HMM-073-B

This document describes the installation and operation of the hardware and software that enable remote support of CRAY J90 series systems. Additional information in the appendix describes how to modify and continue using previously installed remote support equipment.

Automated Confidence Testing, HDM-110-A

This document contains information about automated confidence testing (ACT). ACT is a set of shell scripts and diagnostic tests that detect and isolate hardware failures in CRAY J90 series computer systems.

IOS Based Diagnostics, HMM-099-0

This document provides detailed descriptions of the CRAY J90 series IOS based diagnostics, which include disk drive tests, tape drive tests, mainframe tests, network tests, input/output processor (IOP) tests, input/output buffer board (IOBB) tests, and the boundary scan test.

Field Replacement Procedures, HMM-079-B

This document contains field removal and replacement procedures for field-replaceable components of the CRAY J916 and CRAY J932 computer systems.

Illustrated Parts Catalog, HMM-116-E

This illustrated parts catalog identifies part numbers for the major components, or field-replaceable units, of the CRAY J916 and CRAY J932 computer systems.

Safety and ESD Guidelines, HGM-016-A

This document provides information about hazard statements, work environment hazards, lockout/tagout (LOTO) procedures, and refrigeration reclamation and certification. Information about electrostatic discharge (ESD) prevention is also provided.

Related Hardware Publications

CRAY J90 Series System Programmer Reference, CSM-0301-0B0

This manual describes the configurations, hardware architecture, and functions of CRAY J90 series systems, including the CRAY J916 system. This manual is written to help system analysts and system programmers write and debug programs to run faster and more efficiently.

Videotape Descriptions

The following CRAY J90 series system videotape is available:

CRAY J90 Series Installation, AVT-00209

This video provides the user with a basic understanding of a CRAY J90 series system installation.

Introduction

The CRAY J98 or CRAY J916 computer system, shown in Figure 3, is a multiprocessor system with efficient scalar and vector processing capabilities, large memory, high memory bandwidth, and efficient input/output (I/O) capabilities. This system is capable of configurations that support a customer's price and performance needs, and its modular design enables the customer to perform separate upgrades of the central processing unit (CPU), memory, and input/output subsystem (IOS) components.

The basic system includes a mainframe cabinet and an I/O cabinet. The mainframe cabinet houses the processor and memory modules along with the system clock. The processor modules contain the CPU components, and the memory modules contain the memory components.

Figure 3. CRAY J98 or CRAY J916 Computer System

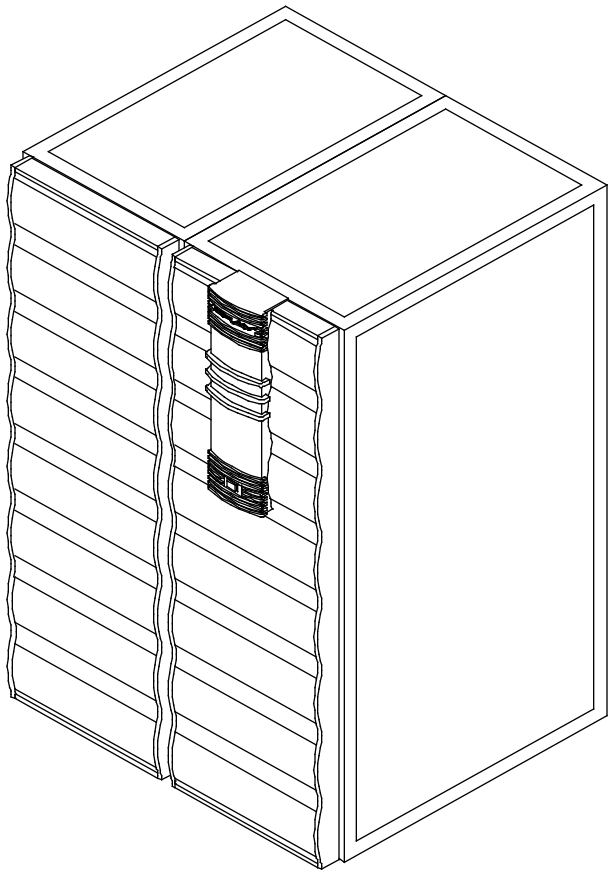

The I/O cabinet includes the IOS and peripherals. An I/O cabinet can house up to four VERSAbus Modular Eurocard (VME) IOSs, and customers may order up to three additional I/O cabinets for more IOS and peripheral capabilities. Refer to Figure 4 for the locations of the memory and processor modules, peripherals, and power, control, and cooling components.

Figure 4. CRAY J90 Series I/O Cabinet and Mainframe Cabinet Components

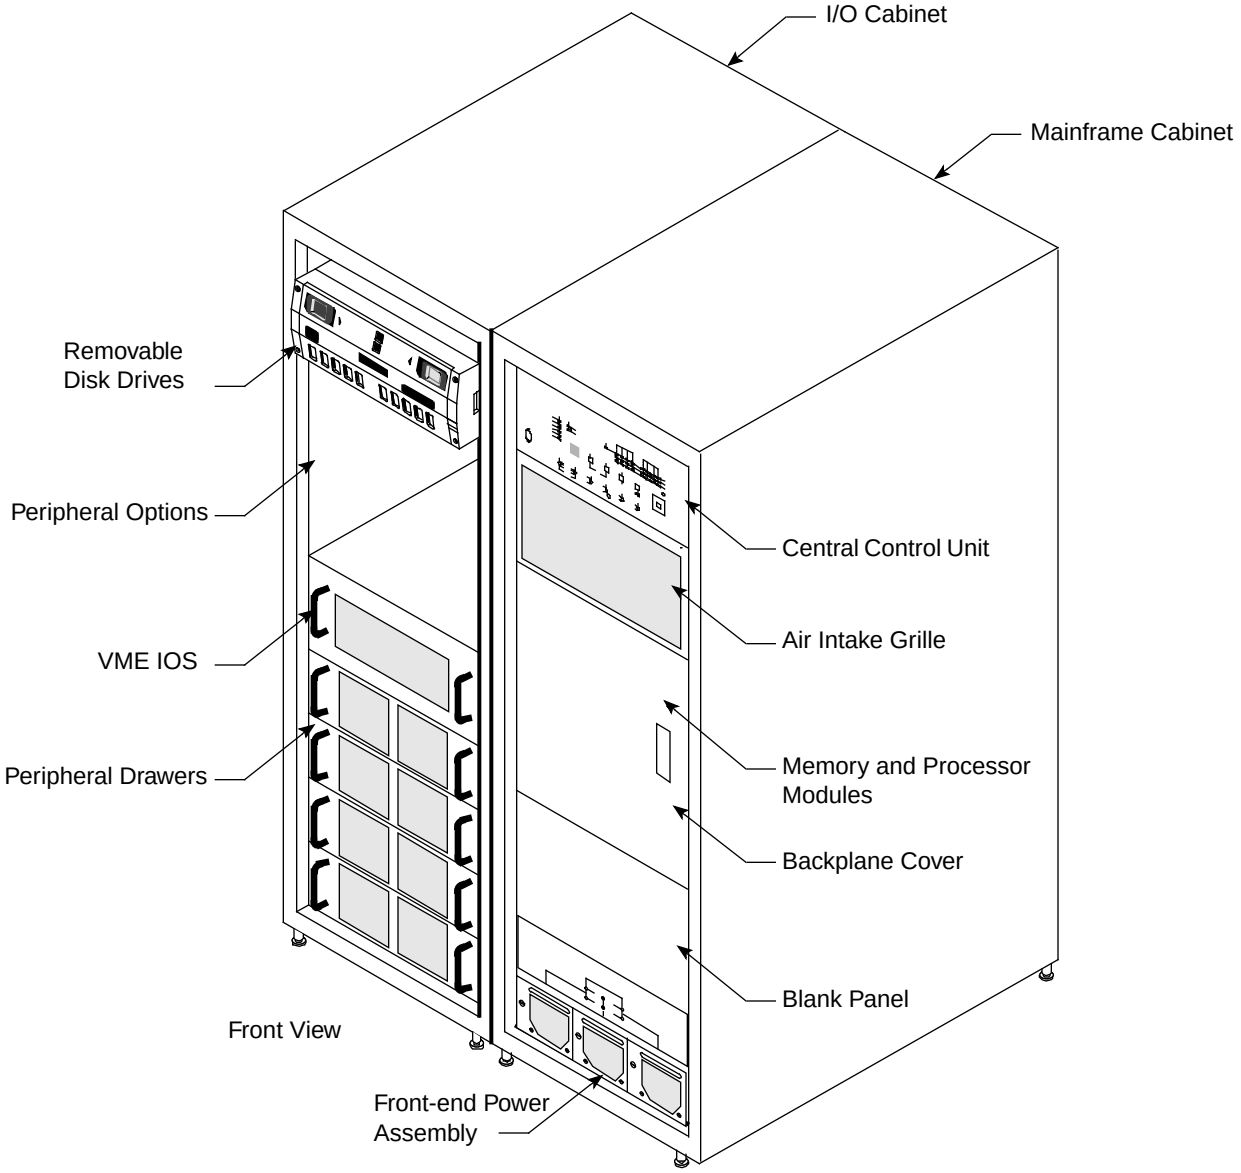

CRAY J98 systems include from 4 to 8 10-ns CPUs with 4 CPUs contained on a single processor. The memory options include 512; 1,024; 2,048; 4,096; or 8,192 Mbytes of memory. Refer to Table 1 for the CRAY J98 (hereinafter referred to as J98) system configurations.

CRAY J916 systems include from 4 to 16 10-ns CPUs with 4 CPUs contained on a single processor module. The memory options include 1,024; 2,098; 4,096; 8,192; or 16,384 Mbytes of memory. Refer to Table 2 for the CRAY J916 (hereinafter referred to as J916) system configurations.

The J98 system is available with a 2 X 2 backplane configuration. The J916 is available in a 4 X 4 backplane configuration. These configurations refer to the maximum number of processor modules and memory modules that the backplane can accommodate. All of the memory slots are always occupied; however, all of the processor module slots may not always be occupied. Each product name, as listed in Table 1, refers to the maximum number of processor modules and memory modules that the backplane can accommodate.

The J98 and J916 systems are air-cooled and do not require a computer room environment or special power or cooling arrangements.

Table 1. CRAY J98 System Configurations

Number of Processor Modules	Product Name	Number of CPUs	Number of VMEs ‡	Size of Central Memory (Mbytes)
1	CRAY J98/4-512	4	2	512 Mbytes (64 Mwords)
1	CRAY J98/4-1,024	4	2	1,024 Mbytes (128 Mwords)
1	CRAY J98/4-2,048	4	2	2,048 Mbytes (256 Mwords)
1	CRAY J98/4-4,096	4	2	4,096 Mbytes (512 Mwords)
1	CRAY J98/4-8,192	4	2 †	8,192 Mbytes (1,024 Mwords)
2	CRAY J98/8-512	5 to 8	2 †	512 Mbytes (64 Mwords)
2	CRAY J98/8-1,024	5 to 8	2 †	1,024 Mbytes (128 Mwords)
2	CRAY J98/8-2,048	5 to 8	2 †	2,048 Mbytes (256 Mwords)
2	CRAY J98/8-4,096	5 to 8	3 †	4,096 Mbytes (512 Mwords)
2	CRAY J98/8-8,192	5 to 8	3 †	8,192 Mbytes (1,024 Mwords)

† This configuration requires additional peripheral cabinets.
‡ Suggested minimum number of VMEs.

Table 2. CRAY J916 System Configurations

Number of Processor Modules	Product Name	Number of CPUs	Number of VMEs (minimum)	Size of Central Memory (Mbytes)
1	CRAY J916/4-1,024	4	1 to 4	1,024 Mbytes (128 Mwords)
1	CRAY J916/4-2,048	4	1 to 4	2,048 Mbytes (256 Mwords)
1	CRAY J916/4-4,096	4	2 to 4	4,096 Mbytes (512 Mwords)
1	CRAY J916/4-8,192	4	2 to 4	8,192 Mbytes (1,024 Mwords)
1	CRAY J916/4-16,384	4	2 to 4	16,384 Mbytes (2,048 Mwords)
2	CRAY J916/8-1,024	5 to 8	1 to 8 †	1,024 Mbytes (128 Mwords)
2	CRAY J916/8-2,048	5 to 8	1 to 8 †	2,048 Mbytes (256 Mwords)
2	CRAY J916/8-4,096	5 to 8	1 to 8	4,096 Mbytes (512 Mwords)
2	CRAY J916/8-8,192	5 to 8	1 to 8	8,192 Mbytes (1,024 Mwords)
2	CRAY J916/8-16,384	5 to 8	1 to 8	16,384 Mbytes (2,048 Mwords)
3	CRAY J916/12-1,024	9 to 12	1 to 12 †	1,024 Mbytes (128 Mwords)
3	CRAY J916/12-2,048	9 to 12	1 to 12 †	2,048 Mbytes (256 Mwords)
3	CRAY J916/12-4,096	9 to 12	1 to 12 †	4,096 Mbytes (512 Mwords)
3	CRAY J916/12-8,192	9 to 12	1 to 12	8,192 Mbytes (1,024 Mwords)
3	CRAY J916/12-16,384	9 to 12	1 to 12	16,384 Mbytes (2,048 Mwords)
4	CRAY J916/16-1.024	13 to 16	1 to 16 †	1,024 Mbytes (128 Mwords)
4	CRAY J916/16-2.048	13 to 16	1 to 16 †	2,048 Mbytes (256 Mwords)
4	CRAY J916/16-4.096	13 to 16	1 to 16 †	4,096 Mbytes (512 Mwords)
4	CRAY J916/16-8.192	13 to 16	1 to 16	8,192 Mbytes (1,024 Mwords)
4	CRAY J916/16-16.384	13 to 16	1 to 16	16,384 Mbytes (2,084 Mwords)

† This configuration requires additional peripheral cabinets.

Processor Module Overview

Each CRAY J90 series processor module can contain up to 4 CPUs; each CPU includes a scalar and vector functional unit that consists of operating registers, functional units, and a control section. Refer to the document entitled *Processors and Memory*, Cray Research publication number HMM-118-B, for more detailed information on the CPU. Refer to Figure 5 for an illustration of the CRAY J90 (classic) processor module.

Each CPU scalar processor contains a 128-word, 2-way set-associative cache and features a scalar processing rate of 100 million instructions per second (MIPS).

Each CPU has its own scalar and vector functional units. Each CPU contains eight 32-bit address (A) registers, eight 64-bit scalar (S) registers, and eight 64-bit by 64-element vector registers. Each CPU also contains eight 32-word instruction buffers that store thirty-two 64-bit words, sixty-four 64-bit T registers, and sixty-four 32-bit B registers.

The CRAY J90 series CPU features a peak vector processing rate of approximately 200 million floating-point operations per second. Each vector unit contains eight vector registers of 64 elements in Y-MP mode.

A single processor module consists of two printed circuit boards (PCBs): one PCB contains the four CPUs, and a second board contains the channel adapter board. The CRAY J90 series channel adapter board provides the interface between the Y1 channel and the CPU. The channel adapter board is connected to a paddle card that can support four Y1 channels. The paddle card can also support High Performance Parallel Interface (HIPPI) connections. The processor module also includes a power module that provides power for the onboard circuits.

The CRAY J90 series systems include a performance monitor that tracks groups of hardware-related events. The performance monitor is a set of registers used by system analysts to monitor operating software on a system. Eight performance counters track four groups of hardware-related events. These results can be used to indicate the relative performance of a program.

The very large-scale integration (VLSI) chips used in the processor module are application-specific integrated circuits (ASICs) constructed of complementary metal oxide semiconductors (CMOS). These ASICs contain more than 500,000 gates each.

Figure 5. CRAY J90 Series Processor Module

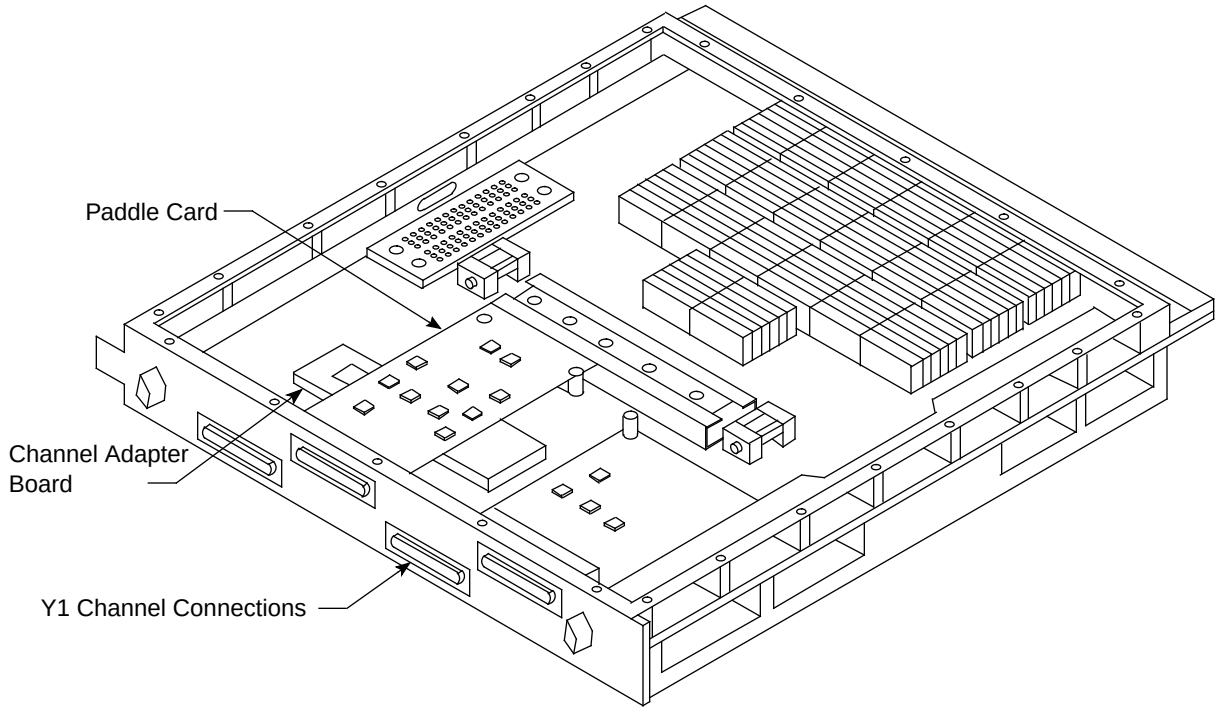

Memory Overview

Memory has a peak bandwidth of 32 words per clock period (CP) (25.6 Gbytes/s) for a 4 X 4 backplane (J916) configuration and 16 words per CP (12.8 Gbytes/s) for a 2 X 2 backplane (J98) configuration. The memory components are distributed across all memory modules; the number of components depends on the configuration. Refer to Table 1 and Table 2 for memory configurations.

Dynamic random-access memory (DRAM) chips provide storage for data and correction bits. Figure 6 illustrates the memory module hardware. The DRAM chips have a 70-ns access time. Central memory is divided into 8 sections. Each memory section in a 4 X 4 system contains 4 subsections, and each subsection contains 16 pseudobanks if fully populated and 8 banks if half populated. Each memory bank can be accessed once every 14 CPs. A system with a 2 X 2 backplane (J98) contains 2 memory subsections. Each memory word consists of 72 bits: 64 data bits and 8 error-correction bits.

Memory is shared among all of the CPUs in the mainframe cabinet. Each CPU has two memory ports that allow two simultaneous memory references from each CPU. A third port is used for I/O and instruction fetch operations.

Figure 6. CRAY J90 Series Memory Module

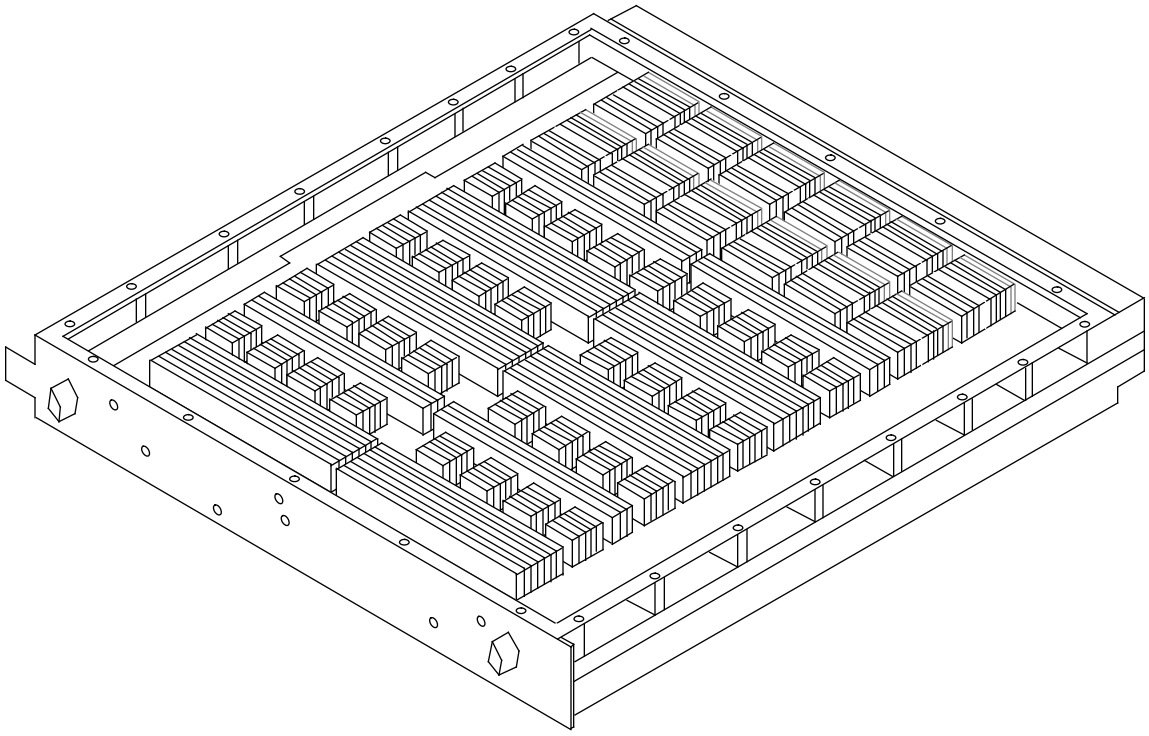

Input/Output Subsystem Overview

The CRAY J98 and J916 I/O cabinet contains a 20-slot VME IOS card cage that is laid out in a 6-4-6-4 slot configuration. The IOS-6 contains four customer-selectable slots, and the IOS-4 contains two customer-selectable slots. Figure 7 is an illustration of the configuration options that are available. The system can support a minimum of one and a maximum of four IOSs per I/O cabinet.

The IOS provides fast data transmission between memory and peripheral devices and networks. Data travels from a peripheral device, across a data channel to the device controller and then from the device controller, across the VMEbus to the I/O buffer board (IOBB). From the IOBB, data travels to the mainframe memory through the 50-Mbyte/s data channel.

The 64-bit IOS consists of a SPARC based I/O processor (IOP), an IOBB, and VME-64 based controllers, depending upon the customer's system configuration.

The CRAY J98 or J916 system can serve as a stand-alone system or can be networked into an existing computing environment. The system can be connected to a multiple system network by either an Ethernet (EI-1) controller or a fiber distributed device interface (FDDI-2) local area network, both of which use Transmission Control Protocol/Internet Protocol (TCP/IP). The asynchronous transfer mode (ATM) uses a synchronous optical network (SONET). The network interfaces are controller boards that also reside in the IOS. A HIPPI channel provides connections for the external peripheral devices.

Figure 7. CRAY J98/J916 IOS Configuration Options

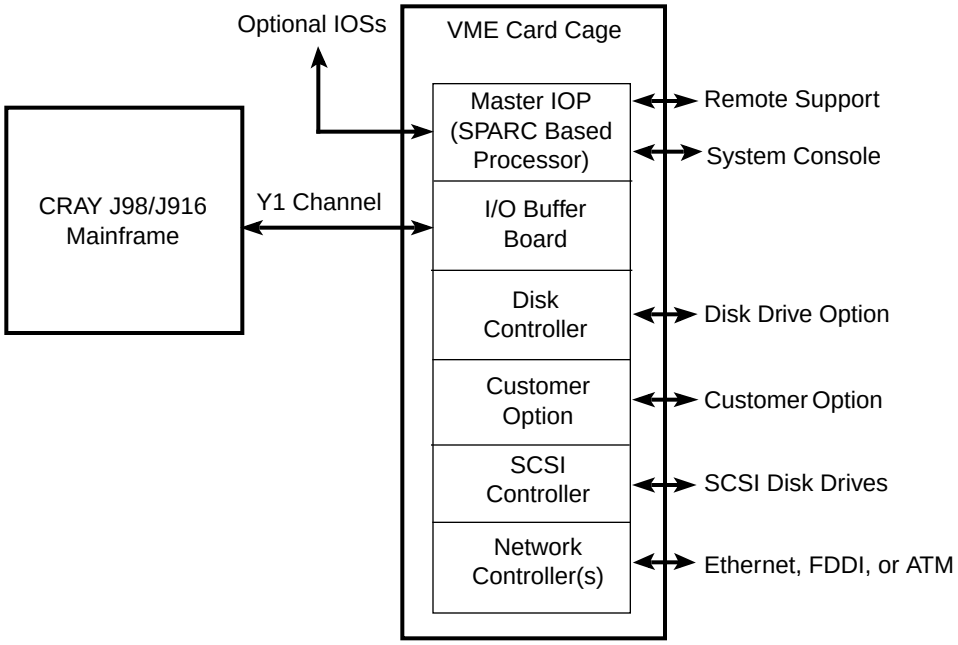

Peripherals

Four disk subsystems are available with the CRAY J98 or J916 system.

- A small computer system interface (SCSI) disk drive subsystem-10 (DDS-10) that includes:
  - 4 SCSI (DD-5S) disk drives
  - A SCSI-3 (SI-3) controller
  - Cables
- A small SCSI disk drive subsystem-1 (DSS-1) that includes:
  - 8 SCSI - DD314/318 disk drives
  - A SCSI-3 (SI-3) controller
  - Cables
- An intelligent peripheral interface (IPI) disk drive subsystem (DDI-10) that includes:
  - 4 DD-5I disk drives
  - A DC-5I disk controller
  - Cables
- A DDR-5 removable IPI disk drive that includes:
  - 2 DD-5I drives
  - An enhanced DC-5I disk controller

External SCSI Peripherals

The CRAY J98 or J916 system offers an optional external SCSI connection for external tape libraries.

Power, Cooling, and Control Overview

Each CRAY J98 and J916 cabinet contains its own power connection and requires 180 to 240 Vac single-phase power at 47 to 63 Hz. Most of the J98/J916 peripherals in the I/O cabinet are packaged in 19-in. rack-mount sliding trays that have their own power supplies and cooling fans.

Refer to the document entitled *Power, Cooling, and Control*, Cray Research publication number HMM-100-A, for more information.

CCU

The central control unit (CCU) system provides the power system monitoring and control, temperature monitoring and protection, and I/O cabinet fault monitoring. LEDs indicate system power and cooling statuses.

Mainframe Power

Mainframe power consumption depends on the customer's system configuration. Refer to *Preparing for a CRAY J916 System Installation*, Cray Research publication number HR-04072-0C, for specific information about power consumption. However, the maximum power consumption for the mainframe cabinet is 3.79 kVA (3.60 kW).

The mainframe AC power enters the system through a power entry box that contains a circuit breaker and a line filter. The AC power then passes into a front-end power system that converts it to 48 Vdc

using three 1,700-W power supplies, which are in an n+1 hot-swap configuration, with power factor correction and isolation. The 48 Vdc power is then bused up to the backplane; the onboard power supplies plug into the backplane.

Mainframe Cooling

A single high-performance, motorized impeller blower provides cooling for the mainframe chassis. It pulls air from the upper front face of the mainframe cabinet. This air flows through an inlet filter, through the modules, and then is discharged at the lower rear of the cabinet through the exhaust plenum. The processor and memory modules are cooled by an interlocked damper system that uses forced air to cool the modules. Refer to Figure 8 for a diagram of the mainframe cabinet power and cooling components.

Figure 8. Mainframe Cabinet Power and Cooling Components (Side View)

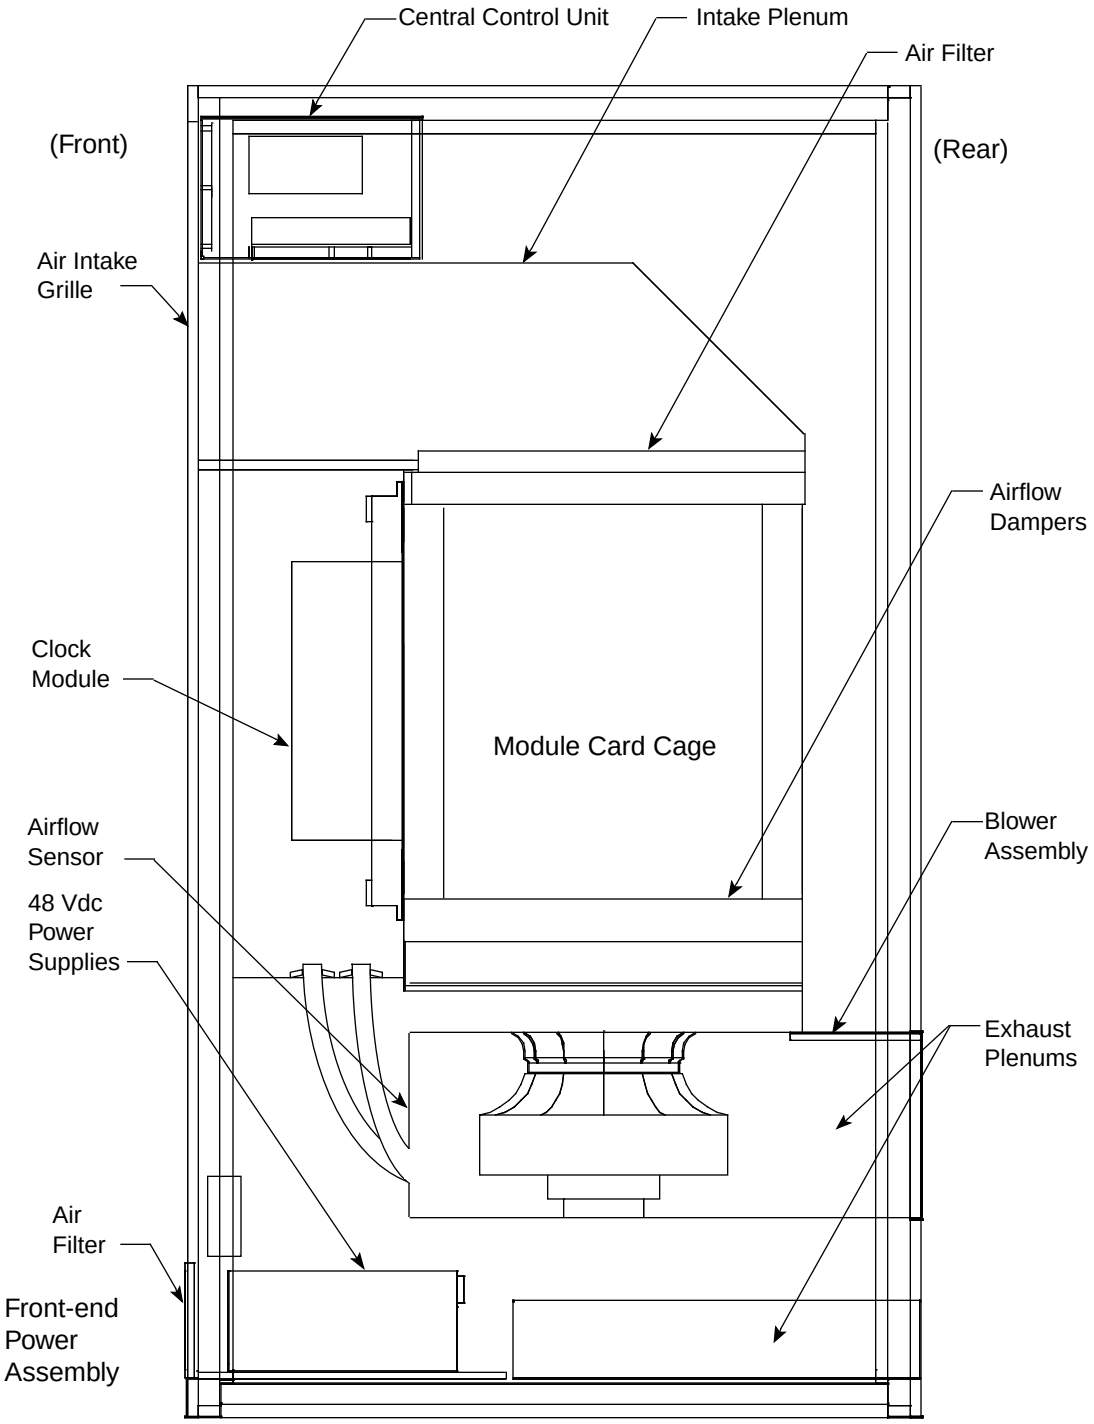

I/O Cabinet CCU Functions

The CCU includes a control signal bus that enables voltage and fan statuses to be sent from the I/O cabinet to the mainframe cabinet. Refer to *Power, Cooling, and Control*, Cray Research publication number HMM-100-A, for information about the LEDs on the I/O cabinet CCU.

I/O Cabinet Power

I/O power consumption depends on the customer's system configuration. Refer to *Preparing for a CRAY J916 System Installation*, Cray Research publication number HR-04072-0C, for specific information about power consumption. However, the maximum power consumption for the I/O cabinet is 4.42 kVA (4.20 kW).

The I/O cabinet also includes an AC power entry box that contains a circuit breaker and a line filter. Each peripheral tray contains its own power connection. AC power is distributed to each peripheral tray through a power distribution rail; each peripheral tray plugs into this distribution rail.

I/O Cabinet Cooling

The IOS VME card cage includes its own n+1 hot-swappable cooling fans and power system. Each peripheral tray is cooled by air that moves from the front of the tray to the rear; no central cooling fan is required. Some peripheral trays may include their own fans.

Diagnostic Overview

Automated confidence testing (ACT) is a suite of shell scripts and diagnostic tests that detect and isolate hardware failures in CRAY J90 series systems. ACT provides two levels of system testing. The first level of testing, ACT power-up, is invoked automatically when the system is powered up and the IOS kernel is loaded. The power-up tests isolate defective system components and record or display the error information. The second level of testing, the ACT menu system, provides a menu-driven interface that is used to select and run specific diagnostics. Each level can be used by on-site, remote-support, or escalated-support service personnel to troubleshoot the system. For more information on ACT, refer to the document entitled *Automated Confidence Testing*, Cray Research publication number HDM-110-A.

The diagnostic tests used by ACT consist of a suite of IOS based quick-look diagnostics and mainframe-based diagnostics. IOS quick-look diagnostics are streamlined versions of menu-driven IOS based diagnostics. IOS quick-look diagnostics are not interactive and are invoked from a shell script or command line. Refer to the document entitled *IOS Based Diagnostics*, Cray Research publication number HDM-099-0, for descriptions of the menu-driven IOS based diagnostics.

There are two types of IOS based tests: menu-driven tests and quick-look tests. The menu-driven confidence and comprehensive tests provide more extensive hardware testing than the quick-look tests. They also provide an interface that enables you to select and run specific sections of each test.

The IOS based diagnostics reside on the system console SCSI disk, and their purpose is to test the IOS internals, CPU, memory, and peripherals that are connected to the IOS.

Use the *System Troubleshooting* document, Cray Research publication number HMM-114-B, as a guide to troubleshooting your system.

Software Overview

The CRAY J90 series systems function with two distinct operating systems: the IOS operating system and the UNICOS operating system. The IOS operating system is a Cray Research proprietary real-time system that resides in the IOP local memory. Although several of the basic IOS commands use the same syntax and arguments as standard UNIX commands, the IOS operating system should not be confused with a UNIX operating system. The IOS kernel and supporting files are stored on the disk drive in the system console.

The UNICOS kernel and configuration files also reside on the system console disk drive. However, once these files have been loaded into central memory and the kernel has been booted, the UNICOS operating system can operate in central memory. The UNICOS supporting files reside on the large-capacity system disk drives.

The UNICOS operating system uses a standard UNIX System V environment. It also complies with the Portable Operating System Interface (POSIX) standard for computer environments. The UNICOS operating system incorporates enhancements that provide more efficient parallel/vector processing, increased security, and more versatile network connectivity.

As part of the UNICOS boot process, the IOS sets the control and register ASICS on the processor and memory modules to a specified initial state. This initialization is done using configuration parameter files. Individual configuration files exist for each processor and memory module in the system. After initialization and after the UNICOS kernel has been loaded and booted, the IOS operating system supports all I/O operations between UNICOS and the VMEbus devices.

System Console

The system console (refer to Figure 9) for CRAY J90 series systems is a Sun Microsystems, Inc. SPARCstation 5 workstation that runs the Solaris 2.x operating system. The system console provides a digital audio tape (DAT) CD-ROM drive for loading software and also contains a SCSI disk drive (system disk) that stores the IOS kernel and configuration files, IOS and mainframe utilities and offline diagnostics, and other maintenance files.

NOTE: This workstation is designed to support CRAY J90 software and hardware maintenance activities. Use of this system console for other purposes may cause the system to operate unpredictably.

Figure 9. System Console