CRAY J932 Hardware Overview

HMM-193-A
CRAY J90 Series Systems
Last Modified: April 1998

Figures

Tables

Record of Revision

June 1995

Original printing.

Revision A: April 1998

Revised to include the 512-Mword memory option availability.

Introduction

The CRAY J932 computer system, shown in Figure 1, 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.

Figure 1. Basic CRAY J932 Computer System

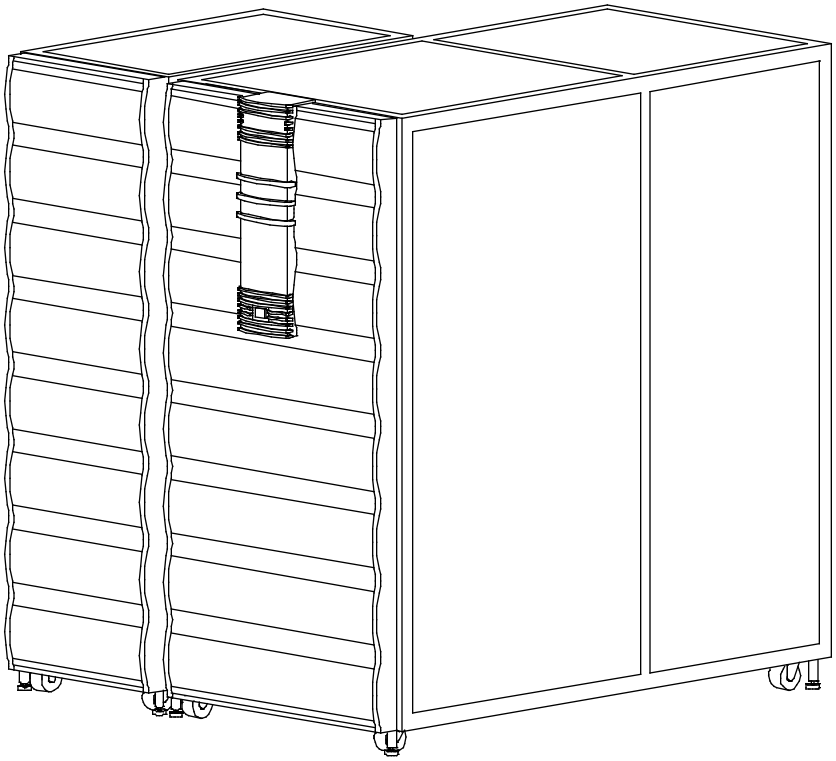

The CRAY J932 is a multicabinet system. The basic CRAY J932 system includes a mainframe cabinet and an I/O cabinet. The mainframe cabinet houses the processor and memory modules along with the system clock. The processor modules contain the CPU components, and the memory modules contain the memory components.

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 2 for the locations of the memory and processor modules, peripherals, and power, cooling, and control components.

Figure 2. CRAY J932 Chassis Layout

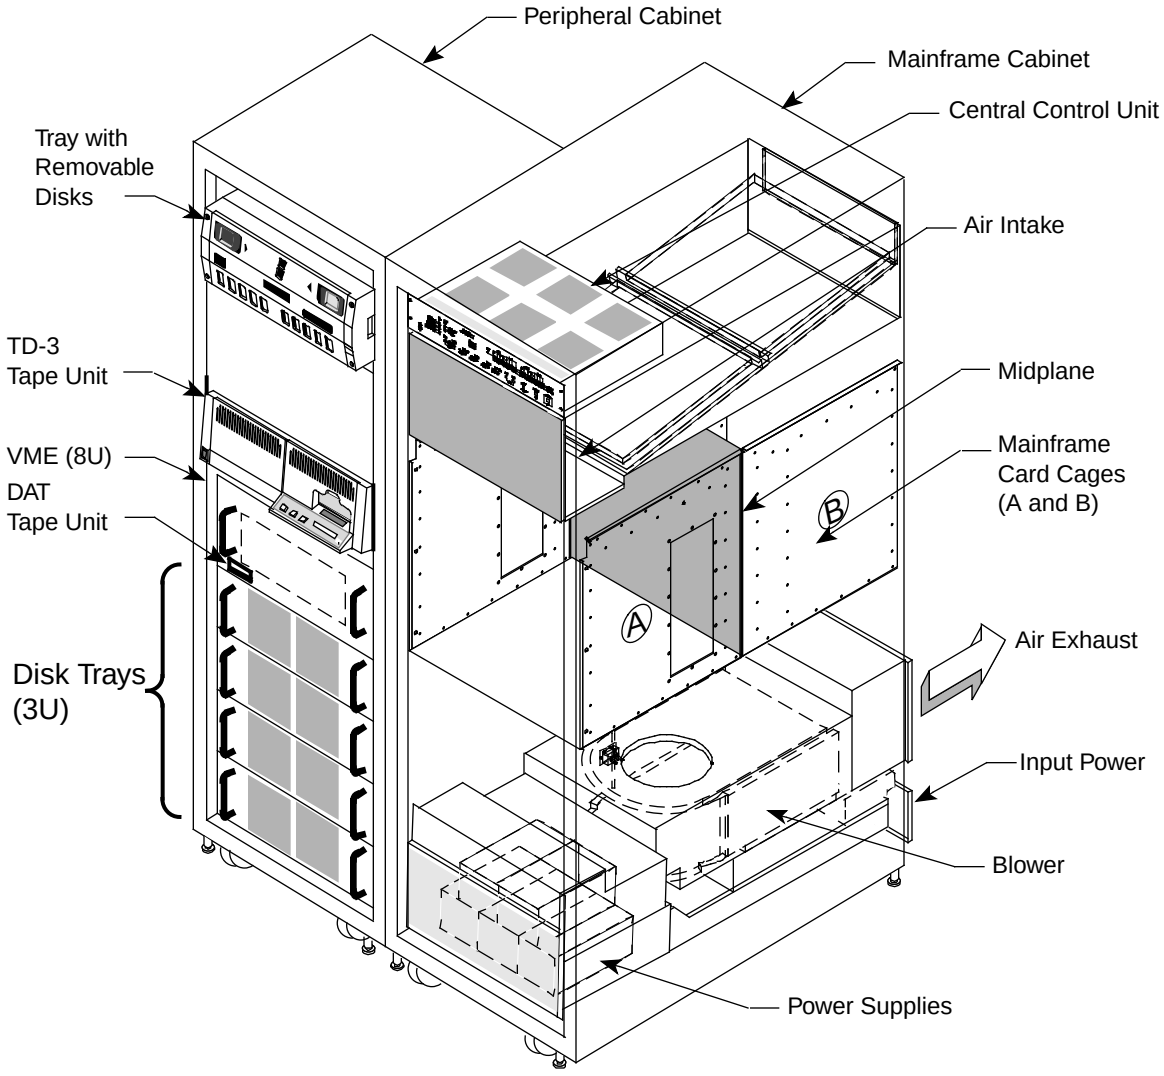

CRAY J932 systems include from 16 to 32 10-ns CPUs with 4 CPUs contained on a single processor module. The memory options include 2,048; 4,096; 8,192; 16,384; or 32,568 Mbytes of memory. Refer to Table 1 for the CRAY J932 system configurations.

The CRAY J932 system is available in an 8 X 8 midplane configuration only. The configurations in Table 1 refer to the maximum number of processor modules and memory modules that the midplane can accommodate. All of the memory module slots are always occupied; however, all of the processor module slots may not always be occupied. Each product name listed refers to the maximum number of processor modules and memory modules that the midplane accommodates for that particular model.

The CRAY J932 system is air cooled and does not require a computer room environment or special power or cooling arrangements.

Table 1. CRAY J932 System Configurations

Number of Processor Modules	Product Name	Number of CPUs	Suggested Number of IOSs †	Size of Central Memory (Mbytes)
4	CRAY J932/16-2,048	16	4 to 16	2,048 Mbytes (256 Mwords)
4	CRAY J932/16-4,096	16	5 to 16	4,096 Mbytes (512 Mwords)
4	CRAY J932/16-8,192	16	5 to 16	8,192 Mbytes (1,024 Mwords)
4	CRAY J932/16-16,384	16	5 to 16	16,384 Mbytes (2,048 Mwords)
4	CRAY J932/16-32,568	16	5 to 16	32,568 Mbytes (4,096 Mwords)
5	CRAY J932/20-2,048	20	4 to 16	2,048 Mbytes (256 Mwords)
5	CRAY J932/20-4,096	20	5 to 16	4,096 Mbytes (512 Mwords)
5	CRAY J932/20-8,192	20	5 to 16	8,192 Mbytes (1,024 Mwords)
5	CRAY J932/20-16,384	20	5 to 16	16,384 Mbytes (2,048 Mwords)
5	CRAY J932/20-32,568	20	5 to 16	32,568 Mbytes (4,096 Mwords)
6	CRAY J932/24-2,048	24	5 to 16	2,048 Mbytes (256 Mwords)
6	CRAY J932/24-4,096	24	6 to 16	4,096 Mbytes (512 Mwords)
6	CRAY J932/24-8,192	24	6 to 16	8,192 Mbytes (1,024 Mwords)
6	CRAY J932/24-16,384	24	6 to 16	16,384 Mbytes (2,048 Mwords)
6	CRAY J932/24-32,568	24	6 to 16	32,568 Mbytes (4,096 Mwords)
7	CRAY J932/28-2,048	28	6 to 16	2,048 Mbytes (256 Mwords)
7	CRAY J932/28-4,096	28	7 to 16	4,096 Mbytes (512 Mwords)
7	CRAY J932/28-8,192	28	7 to 16	8,192 Mbytes (1,024 Mwords)
7	CRAY J932/28-16,384	28	7 to 16	16,384 Mbytes (2,048 Mwords)
7	CRAY J932/28-32,568	28	7 to 16	32,568 Mbytes (4,096 Mwords)

† Each peripheral cabinet can accommodate a maximum of four IOSs.

Number of Processor Modules	Product Name	Number of CPUs	Suggested Number of IOSs †	Size of Central Memory (Mbytes)
8	CRAY J932/32-2,048	32	7 to 16	2,048 Mbytes (256 Mwords)
8	CRAY J932/32-4,096	32	8 to 16	4,096 Mbytes (512 Mwords)
8	CRAY J932/32-8,192	32	8 to 16	8,192 Mbytes (1,024 Mwords)
8	CRAY J932/32-16,384	32	8 to 16	16,384 Mbytes (2,048 Mwords)
8	CRAY J932/32-32,568	32	8 to 16	32,568 Mbytes (4,096 Mwords)

† Each peripheral cabinet can accommodate a maximum of four IOSs.

Processor Module Overview

Each CRAY J932 processor module can contain up to 4 CPUs; each CPU includes a scalar and vector processor 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 3 for an illustration of the CRAY J932 processor module.

Each CPU scalar processor contains eight 64-bit scalar (S) registers and a 128-word, 2-way set-associative cache and features a scalar processing rate of 100 million instructions per second (MIPS).

Each CPU vector processor contains eight 64-bit by 64-element vector registers. The peak vector processing rate is approximately 200 million floating-point operations per second.

Each CPU control section contains eight 32-bit address (A) registers, sixty-four 32-bit B registers, sixty-four 64-bit T registers, and eight 32-word instruction buffers that store thirty-two 64-bit words.

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. The CRAY J932 channel adapter provides the interface between the Y1 channel and the CPU. The channel adapter is connected to paddle cards that can support up to four Y1 channels (one paddle card for each channel). The paddle cards 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 J932 system includes 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 3. CRAY J932 Processor Module

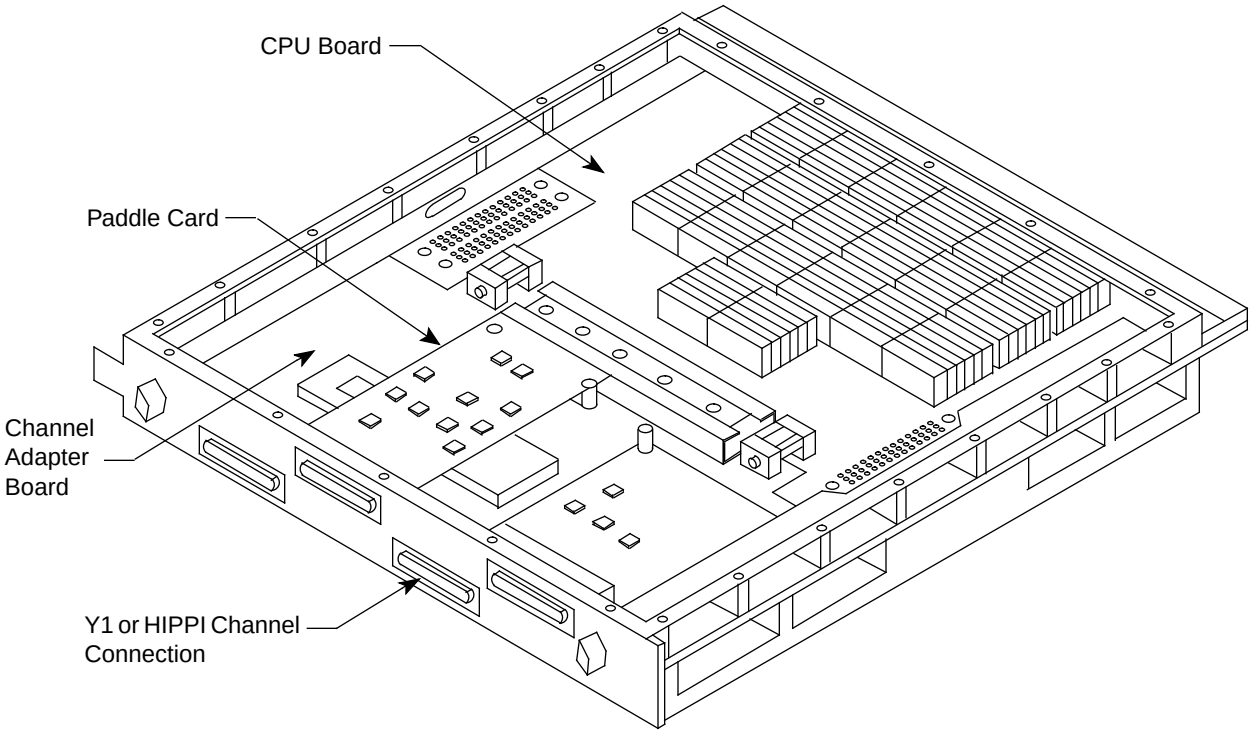

Memory Overview

CRAY J932 memory has a peak bandwidth of 64 words per clock period (CP) (51.2 Gbytes/s). The memory components are distributed across all memory modules; the number of components depends on the configuration. Refer to Table 1 on page 4 for memory configurations.

Dynamic random-access memory (DRAM) chips provide storage for data and correction bits. Three sizes of memory chips are available: 4-Mbit DRAMs, 16-Mbit DRAMs, or 64-Mbit DRAMs. Figure 4 illustrates the memory module hardware. The DRAM chips have a 70-ns access time. Central memory is divided into 8 sections. Each memory section contains 8 subsections; 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 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 4. CRAY J932 Memory Module

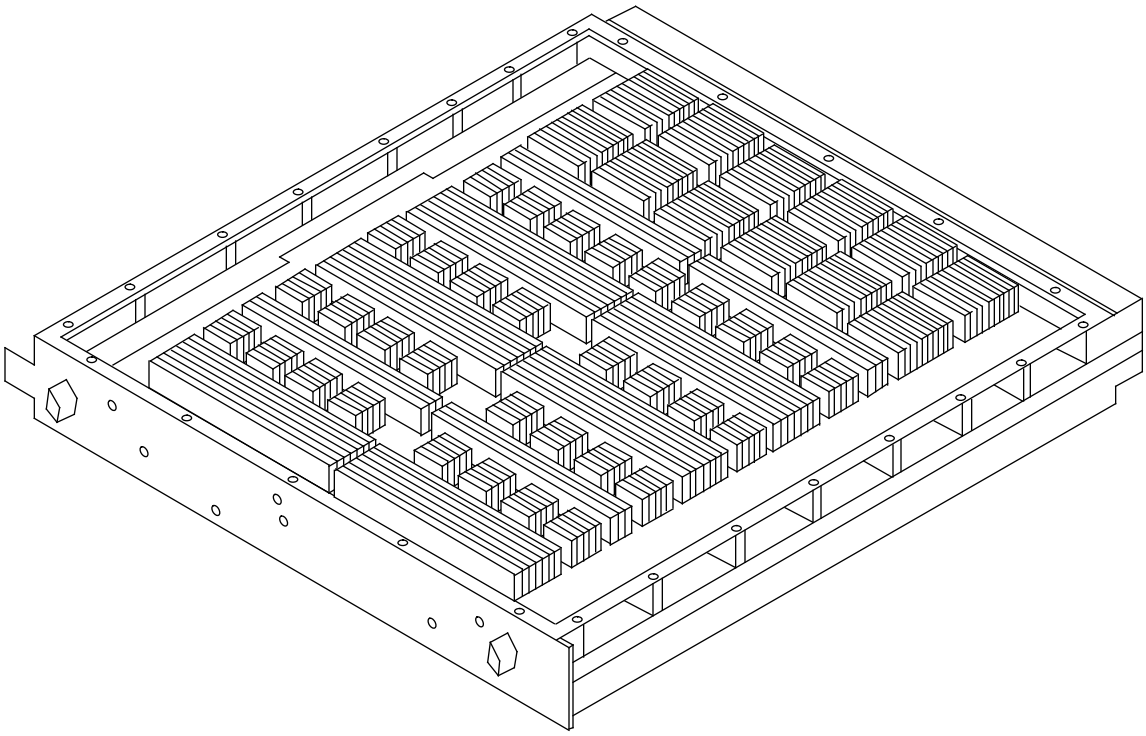

Input/Output Subsystem Overview

The CRAY J932 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 5 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 CRAY J932 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 on the customer's system configuration.

The CRAY J932 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 (FI-2) local area network, both of which use Transmission Control Protocol/Internet Protocol (TCP/IP). The asynchronous

transfer mode (ATM) network controller uses a synchronous optical network (SONET) OC-3C multimode optical interface module. The network interfaces are controller boards that also reside in the IOS. A HIPPI channel provides connections for the external peripheral devices.

Figure 5. CRAY J932 IOS Configuration Options

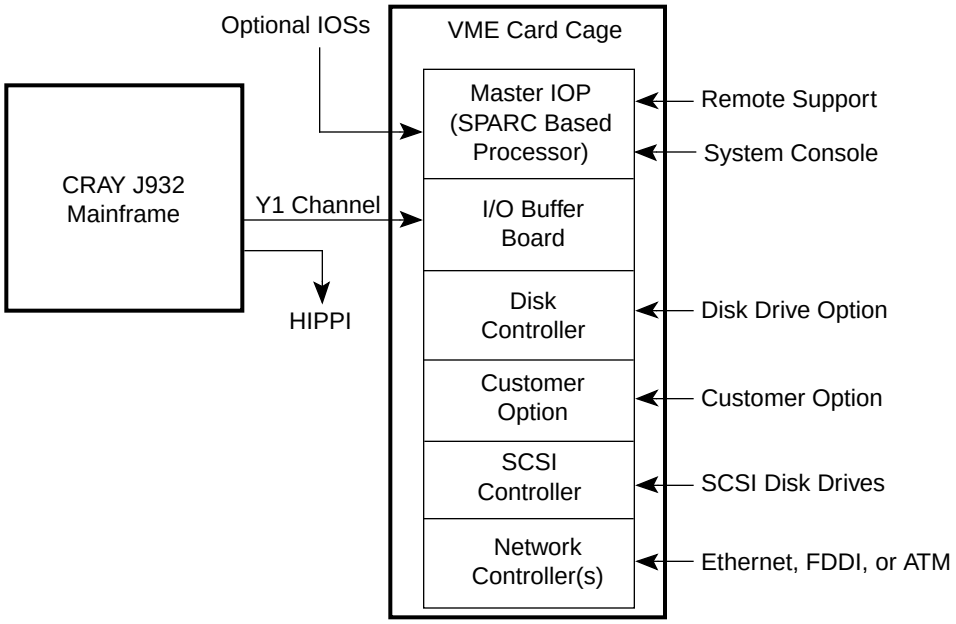

Peripheral Subsystems

The CRAY J932 system supports connections to a number of disk and tape peripheral subsystems.

Disk Subsystems

Four disk subsystems are available with the CRAY J932 system.

- A small computer system interface (SCSI) disk drive subsystem-10 (DDS-10) that includes:
  - 4 DD-5S disk drives or 4 SCSI (DD-6S) disk drives
  - A SCSI-3 (SI-3) controller
  - Cables

NOTE: In a future production phase, the DDS-30 subsystem (that includes 4 DD-6S disk drives instead of 4 DD-5S disk drives) will replace the DDS-10 subsystem.

- 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 removable IPI disk drive (DR-5I) subsystem that includes:
  - 2 DD-5I drives, including canisters
  - An enhanced DC-5I disk controller
  - Cables

Tape Units

Internal cabinet support is available for the optional STK 4220 cartridge drive (TD-3) and standard digital audio tape (DAT) media. (A single cartridge DAT media unit is included with the CRAY J932 system.) The TD-3 unit consists of a controller (SI-3), corresponding tape drive, and connecting hardware. The DAT media unit interfaces directly with the SPARC based IOP.

External Mass Storage

For additional tape storage, an optional SCSI-2 controller (SI-3X) can be used to connect external SCSI-2 based tape drives. There are a number of tape options that may be externally connected to the CRAY J932 computer system. Externally, the limits are 15 tape drives per bus (2 buses per controller) on a 16-bit bus, or 7 drives per bus on an 8-bit bus.

Power, Cooling, and Control Overview

Each CRAY J932 cabinet contains its own power connections (one in an I/O cabinet and two in a mainframe cabinet) and requires 180 to 240 Vac single-phase power at 47 to 63 Hz. Most of the CRAY J932 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 CRAY J932 series system document entitled *Power, Cooling, and Control* for more information.

CCU

The central control unit (CCU) 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 J932 System Installation*, Cray Research publication number HR-04100-0B, for specific information about power consumption. However, the maximum power consumption for the mainframe (processing) cabinet is 7.5 kVA (7.13 kW).

The single-phase mainframe AC power enters the system through two power cords to an entry box that contains a circuit breaker and two line filters. The 200 to 240 Vac power then passes into a front-end power system that converts it to 48 Vdc using five 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 midplane, from which the onboard power supplies (logic power modules) receive power.

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 discharges 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 2 for a diagram of the mainframe cabinet power, cooling, and control components.

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 the CRAY J932 series system document entitled *Power, Cooling, and Control* 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 J932 System Installation*, Cray Research publication number HR-04100-0B, for specific information about power consumption. However, the maximum power consumption for the I/O (peripheral) 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 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 J932 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 J932 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 both the control circuitry and registers on the processor and the memory modules to a specified initial state. This initialization is done with configuration parameters that consist of serial in/serial out (scan latch) data. Individual scan files exist for each processor and memory module in the system. The configuration parameters are stored on the IOS disk drive. 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 for CRAY J932 systems is a Sun Microsystems, Inc. SPARCstation 5 workstation that runs the Solaris 2.x operating system. The system console provides a CD-ROM drive for loading software and 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 J932 software and hardware maintenance activities. Use of this system console for other purposes may cause the system to operate unpredictably.