System Overview
(CRAY SV1™ Series Systems)

108-0196-001

Silicon Graphics Private

SiliconGraphics
Computer Systems

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

Requests for copies of service publications should be directed to:

SILICON GRAPHICS, INC.
Customer and Professional Services - Logistics
1100 Lowater Road
P. O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Comments about this publication should be directed to:

SILICON GRAPHICS, INC.
Customer and Professional Services
Service Publications and Training
890 Industrial Blvd.
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

System Overview

108-0196-001
CRAY SV1™ Series Systems
Last Modified: November 1998

Record of Revision 2
Introduction to CRAY SV1 Series 3
Processor Module 8
    Multi-Streaming Processor (MSP) 9
Memory Module 10
Scalable I/O Architecture 11
    MPNs 12
    SPNs 13
    PC-10 Models 14
Power, Cooling, and Control 15
    CCU 16
    Mainframe Power 16
    Mainframe Cooling 16
    PC-10 CCU Functions 17
    PC-10 Power 17
    PC-10 Cooling 17
Diagnostics 18
Operating Systems 19
System Workstation 20

Figures

Figure 1. Basic CRAY SV1 Series Computer Systems 3
Figure 2. CRAY SV1-1A Series Chassis Layout (Rear View) 5
Figure 3. CRAY SV1-1 Series Chassis Layout (Rear View) 6
Figure 4. CRAY SV1 Series with GigaRing Processor Modules 9
Figure 5. CRAY SV1 Series Memory Module 10
Figure 6. SIO Elements 11
Figure 7. MPN 12
Figure 8. SPNs in an NSR 13
Figure 9. System Workstation 20

Tables

Table 1. CRAY SV1 Series Models 4
Table 2. CRAY SV1 Series Single-cabinet Configurations 7

Record of Revision

November 1998

Original printing.

Introduction to CRAY SV1 Series

CRAY SV1 series computer systems, shown in Figure 1, are multiprocessor systems with efficient scalar and vector processing capabilities, large memory, high memory bandwidth, and efficient input/output (I/O) capabilities. These systems are capable of configurations that support a customer's price and performance needs. The modular design enables the customer to perform separate upgrades of the central processing unit (CPU), memory, and scalable input/output (SIO) components.

Figure 1. Basic CRAY SV1 Series Computer Systems

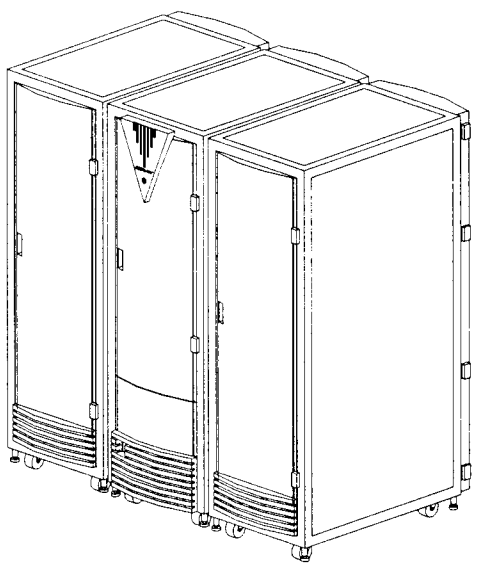

CRAY SV1-1A

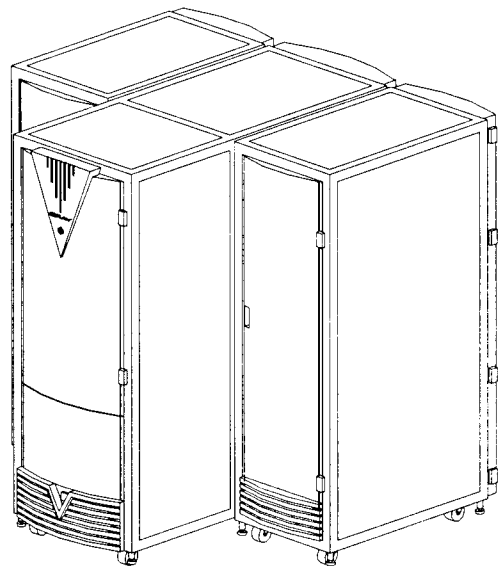

CRAY SV1-1

CRAY SV1 series systems are available in two models: the CRAY SV1-1A model, and the CRAY SV1-1 model. The following example explains the naming conventions for the CRAY SV1 configurations:

CRAY SV1-1A/12-16

- Memory size (Gbytes)
- Number of CPUs
- If "A", system has a 4 X 4 backplane
- Number of mainframe cabinets
- Product name

Each model includes a mainframe cabinet and a minimum of two SIO peripheral cabinets (PC-10s). 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; both reside in a backplane or midplane card cage.

Table 1 describes the memory and processor module configurations for both of the CRAY SV1 series models. Figure 2 and Figure 3 show possible chassis layouts of the CRAY SV1-1A and SV1-1 series systems.

Table 1. CRAY SV1 Series Models

Model	Backplane or Midplane	Memory Modules	Processor Modules	CPUs
CRAY SV1-1A	4X4	4	2 to 4	8 to 16
CRAY SV1-1	8X8	8	2 to 8	8 to 32

The backplane/midplane configurations listed in Table 1 refer to the maximum number of processor modules and memory modules that the backplane or midplane can accommodate. All of the memory slots are always occupied; however, all of the processor module slots may not always be occupied.

The SIO architecture is a multicabinet (PC-10), air-cooled system that provides scalable high-performance and high-resilience I/O support for CRAY SV1 series computer systems. SIO data and control information transmits over GigaRing channels. The flexibility of the GigaRing channel architecture enables multiple system configurations with system functionality and performance that are appropriate to the needs of the customer.

The PC-10 is available in two models: the PC-10A and the PC-10B. Typically, CRAY SV1 mainframes use the PC-10A. The PC-10A and the PC-10B have different warning and control system (WACS) features. Refer to the "Scalable I/O Architecture" section in this document for more information about SIO and PC-10 models.

The CRAY SV1 series systems are air-cooled and do not require a computer room environment.

CRAY SV1 series systems include from 8 to 32 CPUs with 4 CPUs contained on a single processor module. The memory options range from 2,048 Mbytes to 32,768 Mbytes. Table 2 lists the available CRAY SV1 series configurations.

Figure 2. CRAY SV1-1A Series Chassis Layout (Rear View)

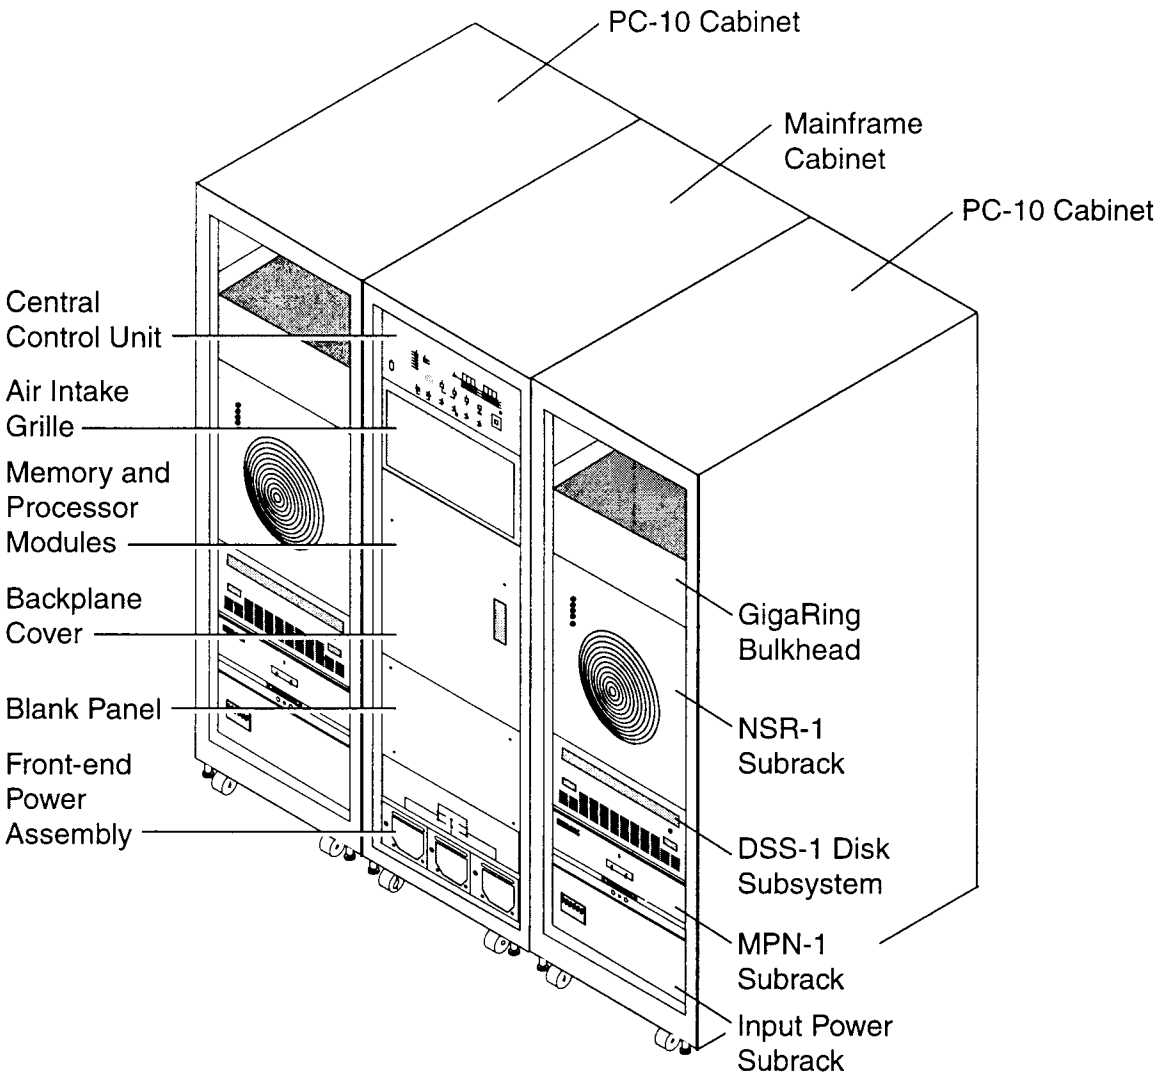

Figure 3. CRAY SV1-1 Series Chassis Layout (Rear View)

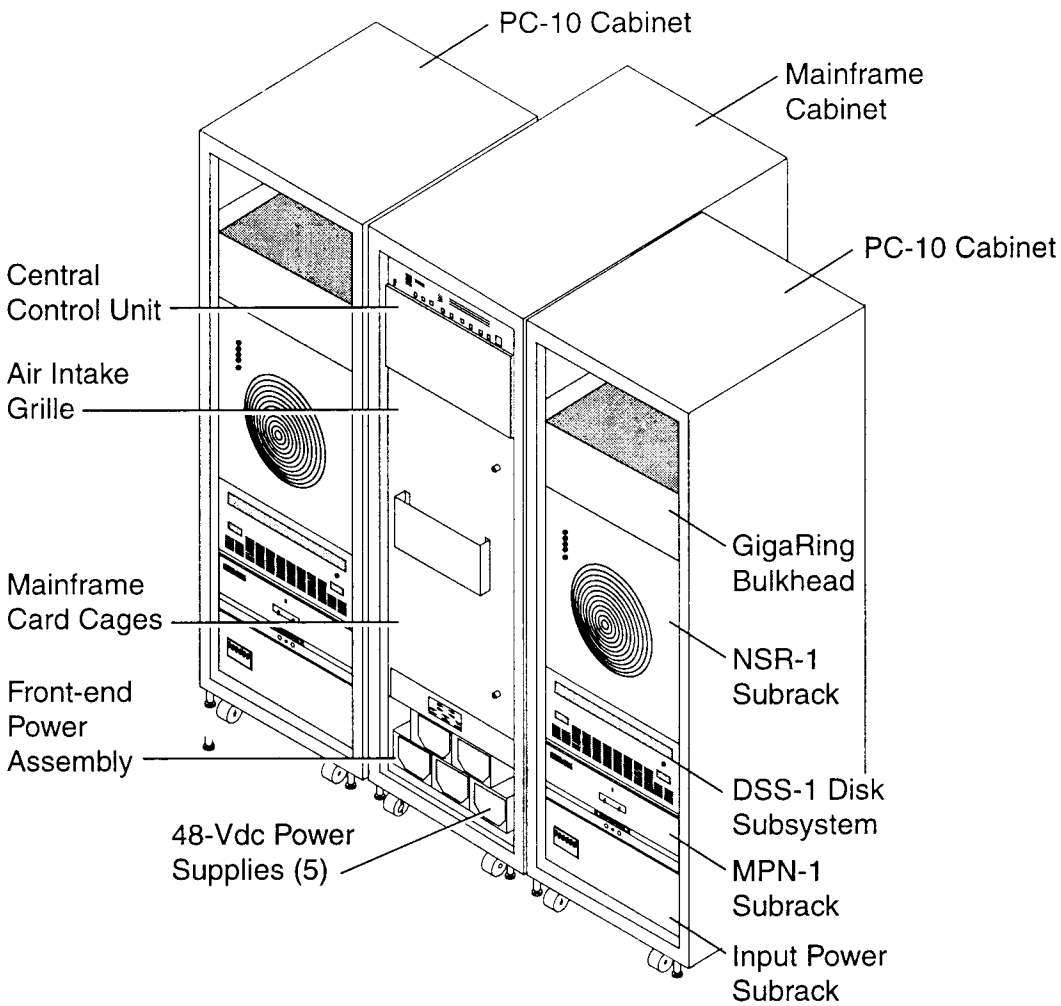

Table 2. CRAY SV1 Series Single-cabinet Configurations

Product Name	Number of Mainframe Cabinets	Total Number of CPUs	Total Memory (Gbytes)
SV1-1A/8-2	1	8	2
SV1-1A/8-4	1	8	4
SV1-1A/8-8	1	8	8
SV1-1A/8-16	1	8	16
SV1-1A/12-2	1	12	2
SV1-1A/12-4	1	12	4
SV1-1A/12-8	1	12	8
SV1-1A/12-16	1	12	16
SV1-1A/16-2	1	16	2
SV1-1A/16-4	1	16	4
SV1-1A/16-8	1	16	8
SV1-1A/16-16	1	16	16
SV1-1/16-4	1	16	4
SV1-1/16-8	1	16	8
SV1-1/16-16	1	16	16
SV1-1/16-32	1	16	32
SV1-1/24-4	1	24	4
SV1-1/24-8	1	24	8
SV1-1/24-16	1	24	16
SV1-1/24-32	1	24	32
SV1-1/32-4	1	32	4
SV1-1/32-8	1	32	8
SV1-1/32-16	1	32	16
SV1-1/32-32	1	32	32

Processor Module

Each CRAY SV1 series processor module contains four 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 *Processor and Memory Components*, Silicon Graphics publication number 108-xxxx-xxx, for more detailed information on the CPU.

Each CPU scalar processor contains eight 64-bit scalar (S) registers and a 256-Kbyte, 128-word, 4-way set-associative cache. Each scalar processor features a scalar processing rate of 400 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 1 GFLOP per CPU. 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 each store thirty-two 64-bit words.

Each processor module contains four CPUs on one printed circuit board (PCB). Processor modules with the SIO GigaRing channel adapter contain two additional PCBs. The CRAY SV1 series GigaRing channel adapter provides the interface between the GigaRing channel and the CPU. The processor modules also include power modules that provide power for the onboard circuits.

Refer to Figure 4 for an illustration of the CRAY SV1 series processor module.

The CRAY SV1 series systems include a performance monitor that tracks groups of hardware-related events. The performance monitor is a set of registers that are used by system analysts to monitor operating software on a system. CPU user instructions allow the user to select and read the performance monitor.

The CRAY SV1 processor modules use application-specific integrated circuit (ASIC) technology. Each module comprises an array of ASICs, which are based on very large-scale integrations (VLSI) complementary metal oxide semiconductor (CMOS) technology. These ASICs contain more than 1.6 million gates each. Each processor module includes 256-Kbyte scalar and vector cache that reduces latency time.

Figure 4. CRAY SV1 Series with GigaRing Processor Module

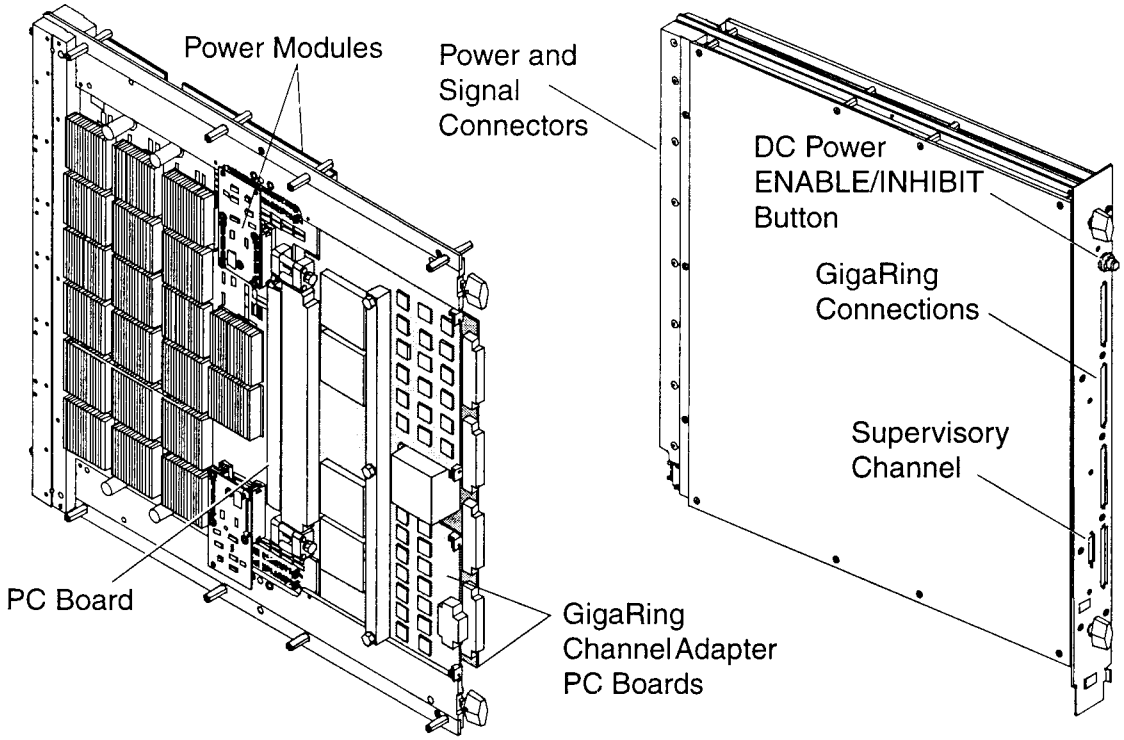

Multi-Streaming Processor (MSP)

Future CRAY SV1 series systems will utilize the MSP. This design includes four tightly coupled two-pipe CRAY SV1 CPUs that are managed through the software. By coupling the four CPUs, a multipipe architecture is created without hardware changes. This MSP has eight vector pipes with a peak performance of 4.0 GFLOPS. The UNICOS kernel automatically selects CPUs based on the optimal machine layout, and the compiler automatically generates the multistreaming code.

Memory Module

Memory components are distributed across all memory modules; the number of components depends on the configuration. CRAY SV1-1A memory has a peak bandwidth of 25.6 Gbytes/s for a 4 X 4 backplane configuration. CRAY SV1-1 memory has a peak bandwidth of 51.2 Gbytes/s for an 8 X 8 midplane configuration. Refer to Table 2 on page 7 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 5 illustrates the memory module hardware. The DRAM chips have a 70-ns access time. Central memory is divided into 8 sections. Each CRAY SV1-1 memory section contains 8 subsections, and each CRAY SV1-1A memory section contains 4 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 it to make simultaneous memory references. A third port is used for I/O and instruction fetch operations.

Figure 5. CRAY SV1 Series Memory Module

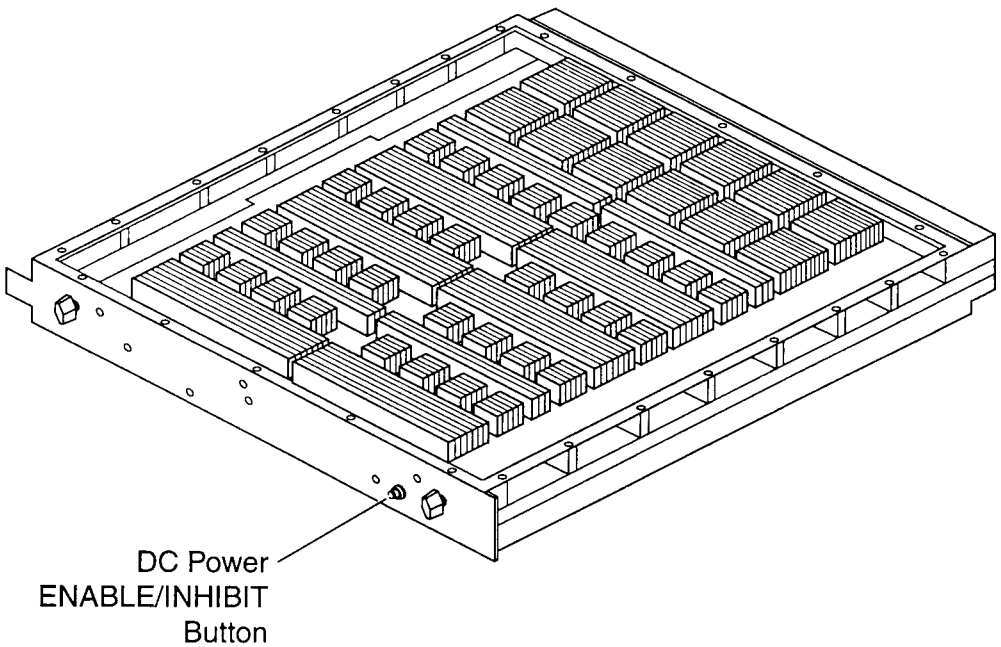

Scalable I/O Architecture

The CRAY SV1 series systems support connections to a variety of disk, tape, and network peripheral subsystems through the SIO architecture. The basic elements of the SIO architecture are the GigaRing channel, input/output nodes (IONs), and the system workstation (SWS).

Figure 6. SIO Elements

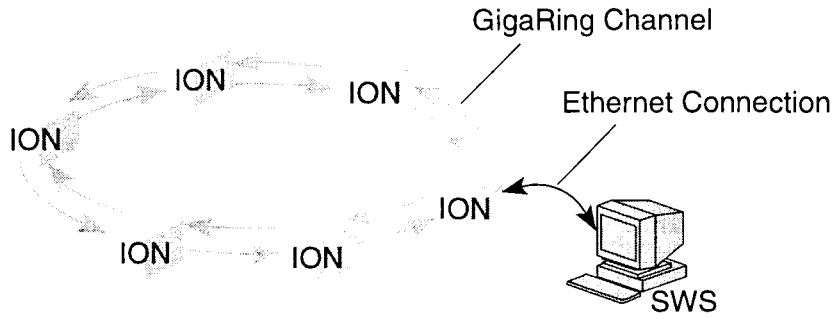

The GigaRing channel is a flexible and scalable channel that standardizes system interconnection and communication across all system peripherals and the mainframe computer. The GigaRing channel consists of a pair of counter-rotating rings, which form a single logical channel.

An ION is any device that connects directly to the GigaRing channel; for example, mainframes, multipurpose nodes (MPNs), and single-purpose nodes (SPNs). Most IONs are independent units that reside in the PC-10 and connect disk, tape, and network peripheral devices to the GigaRing channel.

The SWS serves as the maintenance and operation platform. The SWS provides support for diagnostics, monitoring, serviceability, configuration, and failure isolation through a private Ethernet connection. (Refer to Figure 9 for an illustration of the SWS.)

Disk, tape, and network IONs reside in subracks within the PC-10. The PC-10 also contains the power input subrack and several support subracks. The following subracks are available:

- Multipurpose node subrack (MPN)
- Node subrack (NSR-1)
- Ethernet concentrator
- SCSI disk subracks (DSS)
- Fibre channel disk subrack (DSF)
- Micro Annex communications server
- Fiber-optic channel extender (FOX-1)
- Power input subrack

MPNs

MPNs connect industry standard SBus-based I/O peripherals to the GigaRing channel to provide I/O services for the mainframe node. The SBus controllers reside inside the MPN (shown in Figure 7). CRAY SV1 series systems include the following industry standard SBus controllers:

- ETN-11 Ethernet network interface
- ATM-10 asynchronous transfer mode (ATM) network interface
- FDI-10 fiber distributed data interface (FDDI) network interface
- SCS-10 small computer system interface (SCSI) disk and tape drive interface

Figure 7. MPN

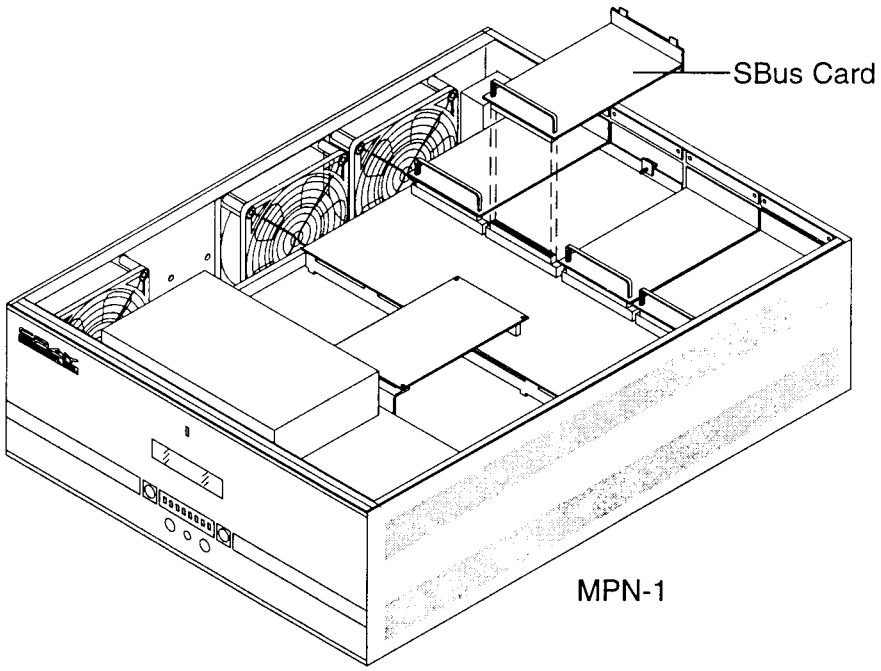

SPNs

SPNs connect specific I/O peripherals to the GigaRing channel to provide I/O services for the mainframe node. All SPNs reside in the NSR (shown in Figure 8). The NSR provides power and cooling for a maximum of four SPNs. The following disk, tape, and network SPNs are available:

- IPN-1 intelligent peripheral interface (IPI-2) disk
- FCN-1 and FCN-2 fibre channel disk
- BMN-1 block MUX tape
- HPN-1 High Performance Parallel Interface (HIPPI)
- HPN-2 High Performance Parallel Interface (HIPPI)
- ESN-1 Enterprise Systems Connection Architecture/390 (ESCON) tape

Figure 8. SPNs in an NSR

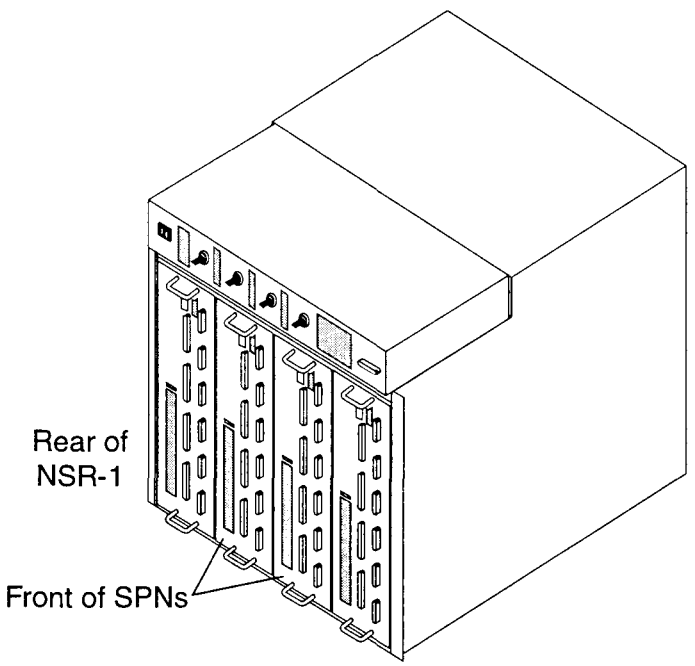

PC-10 Models

When a CRAY SV1 mainframe is connected to PC-10A cabinets, the following PC-10 WACS features are available:

- Single-point hard reset of all MPNs
- Single-point system power-off of the connected PC-10A cabinets
- Fault light notification that a disk subrack fan failed

When a CRAY SV1 mainframe is connected to PC-10B cabinets, the following PC-10 WACS features are available in addition to the features listed for a PC-10A:

- Remote hard reset of all MPNs
- Remote system power-off of individual PC-10B cabinets and NSRs
- Remote monitoring of individual disk subrack fans and power supplies

Power, Cooling, and Control

Each CRAY SV1-1A series mainframe cabinets contain one power connection each; CRAY SV1-1 series cabinets contain two power connections. The PC-10 cabinets also contain one power connection. Each mainframe power connection requires 200 to 240 Vac single-phase power at 47 to 63 Hz. Most of the CRAY SV1 peripherals in the PC-10 cabinet are packaged in 19-in. rackmount subracks that have their own power supplies and cooling fans.

For more information about CRAY SV1-1A series power, cooling, and control, refer to the document entitled *CRAY SV1-1A Mainframe Cabinet Power, Cooling, and Control*, publication number 108-xxxx-xxx. For more information about CRAY SV1-1 series power and cooling, refer to the document entitled *CRAY SV1-1 Mainframe Cabinet Power, Cooling, and Control*, publication number 108-xxxx-xxx. Periodically check the Service Publications and Training Web site at the following URL for availability of these documents: `http://servinfo.csd.sgi.com/`.

CCU

The central control unit (CCU) provides the power system monitoring and control, temperature monitoring and protection, and PC-10 fault monitoring. LEDs on the front panel of the CCU indicate system power and cooling statuses.

A PC-10B peripheral cabinet includes a remote CCU that enables remote monitoring of the system via NWACS software on the SWS. The RS-232 COMM port on the remote CCU is connected to one of the serial ports on the SWS. The NWACS program is used to display the status of the fault LEDs on the CCU. To get an NWACS reading, power down the processor and memory modules, power down the 48-V bus, reset the CPUs, and set voltage margins for the processor and memory modules. You may also power down by resetting the SIO MPNs.

Mainframe Power

Mainframe power consumption depends on the customer's system configuration. Refer to *Preparing for a CRAY SV1-1A System Installation*, publication number 007-3864-002, for specific information about CRAY SV1-1A power consumption. Refer to *Preparing for a CRAY SV1 System Installation*, publication number 007-3865-002, for specific information about CRAY SV1-1 power consumption.

The single-phase mainframe AC power enters the system through one (CRAY SV1-1A system) or two (CRAY SV1-1 systems) power cords, which connect to an entry box that contains a circuit breaker and one or two line filters. The power then passes into a front-end power system that converts it to 48 Vdc. The 1700-W power supplies, which are in an n+1 hot-swap configuration, convert the power and provide power factor correction and isolation. The 48-Vdc power is then bused to the backplane or 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 rear area of the mainframe cabinet. This air flows through an inlet filter, through the modules, and then discharges at the lower front of the cabinet through the exhaust plenum. An interlocked damper system that uses forced air cools the processor and memory modules.

PC-10 CCU Functions

The CCU includes a control signal bus that enables the PC-10 to send fan statuses to the mainframe cabinet. Refer to the SIO system document entitled *Warning and Control System (WACS),* publication number HTM-235-A, for information about the LEDs on PC-10 subracks.

PC-10 Power

PC-10 power consumption depends on the customer's system configuration. Refer to *Preparing for a CRAY SV1-1A System Installation*, publication number 007-3864-002, for specific information about PC-10 power consumption in a CRAY SV1-1A system. Refer to *Preparing for a CRAY SV1 System Installation*, publication number 007-3865-002, for specific information about PC-10 power consumption in a CRAY SV1-1 system.

The PC-10 also includes an AC power input subrack that contains a circuit breaker and a line filter. Each node and peripheral subrack contains its own power connection. AC power is distributed to the subracks through a power distribution strip; each subrack plugs into this distribution strip. The Micro Annex communications server and the Ethernet concentrator plug directly into the power input subrack.

PC-10 Cooling

The PC-10 provides only air inlet and exhaust ventilation. Blowers inside the node and peripheral subracks pull the air through the perforated rear door and exhaust it through the front door; no central cooling fan is required. Most node and peripheral subracks include their own fans.

Diagnostics

Note: Periodically check the Service Publications and Training Web site at the following URL for availability of CRAY SV1 series documents: `http://servinfo.csd.sgi.com/`.

Standard UNICOS online diagnostics and utilities as well as offline diagnostics are available for troubleshooting CRAY SV1 series systems. The primary tool for running offline mainframe diagnostics is the automated confidence testing (ACT) menu interface. The CRAY J90 maintenance console (JMC) software can also be used to run offline mainframe diagnostics.

ACT is a suite of shell scripts and C applications that detect and isolate mainframe hardware failures in CRAY SV1 series systems. ACT provides two levels of system testing. The first level of testing in ACT is invoked automatically after the system is powered up and the ION VxWorks operating system 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 *CRAY SV1 Mainframe Troubleshooting*, publication number 108-xxxx-xxx.

SIO diagnostic tools are available in addition to mainframe diagnostics; refer to the *CRAY SV1 Documentation Map*, publication number 108-xxxx-xxx, for a listing of SIO documentation.

Operating Systems

The CRAY SV1 series systems function with two distinct operating systems: the UNICOS operating system for the mainframe and the VxWorks operating system for the SIO architecture.

The UNICOS kernel and configuration files reside on the system workstation. However, once these files have been loaded into central memory and the kernel has been booted, the UNICOS operating system operates 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. UNICOS 10.0.0.3 or later is required for CRAY SV1 series systems.

The VxWorks operating system is a real-time operating system that has been tuned to provide highly efficient I/O control.

System Workstation

The system workstation (SWS) for CRAY SV1 series systems is a Sun Microsystems, Inc. SPARCstation 5 workstation. Refer to Figure 9 for an illustration of the SWS. The CPU unit is configured with a 2.0-Gbyte SCSI hard disk drive, a Quad SBus Ethernet card, and a video monitor SBus card.

Note: This workstation is designed to support CRAY SV1 series system software and hardware maintenance activities. Use of this system workstation for other purposes may cause the system to operate unpredictably.

The SWS provides a CD-ROM drive for loading software and a SCSI disk drive (system disk) that stores the SIO kernel and configuration files, SIO and mainframe utilities and diagnostics, and other maintenance files. An external digital audio tape (DAT) drive connects to the SWS that is used to back up the SWS, to send diagnostic information to Central Repair, and to return UNICOS dumps for analysis.

Figure 9. System Workstation