

# Solid-state Storage Device

---



---

**CRAY**

---

---

Cray Research's enhanced SSD solid-state storage device offers large storage capacity, high reliability, and outstanding price/performance. The complete line of SSDs includes five models with memory sizes from 32 million to over 512 million 64-bit words.

Since its introduction in 1982, the SSD has demonstrated its usefulness and reliability in a variety of applications. In conjunction with a CRAY X-MP computer system, significant benefits have been achieved in such applications as seismic processing, reservoir modeling, fluid-flow computation, finite-element analysis, weather forecasting, and other large-scale simulations of physical phenomena.

In the past year, nearly two-thirds of customers ordering a new CRAY X-MP computer system included an SSD in their configuration. The utility and dependability of the

SSD is a continuation of Cray Research's efforts to develop the most powerful computer systems available today.

### **The SSD concept**

The SSD functions as a very high-speed secondary memory and is easily accessible to users via standard Fortran I/O statements. Users reserve and allocate storage on the SSD through the Cray operating systems, COS and UNICOS. Only one COS job control language (JCL) statement or one UNICOS shell statement is required for each dataset or file assigned to the SSD. Because the SSD is a solid-state (that is, non-rotating) memory device, access to data stored on the SSD is not constrained by seek and latency considerations. Access time is essentially zero, and transfer rates are over 100 times faster than conventional disk storage.

### **Performance**

The SSD enhances the performance of the Cray central processing units (CPUs) by significantly reducing I/O wait time. For example, during a

structural analysis problem, a CRAY X-MP computer system configured with an SSD exhibited a 99% decrease in I/O wait time when compared to the same system configured with disk storage. The benchmark showed I/O wait time to be 379 seconds when the CRAY X-MP computer system was configured with disk storage and only 5 seconds when configured with the SSD.

Depending on the Cray computer system being used, transfer rates for the SSD can be 100, 1000, or 2000 Mbytes/sec. Compared to the maximum transfer rate of 10 Mbytes/sec with disk storage, the SSD offers a significant performance improvement.

In addition to the very fast transfer rates, the performance of the SSD can be attributed to an access time of less than 25 microseconds and storage capacities up to 512 million words. All this is possible because of high-density, random-access, metal oxide semiconductor (MOS) memory chips.

---

---

---

## SSD configurations

On CRAY X-MP/4 computer systems, the four largest SSD models connect to the mainframe using two 1000-Mbyte/sec channels. This configuration achieves a maximum transfer rate of 2000 Mbytes/sec. On CRAY X-MP/1 and CRAY X-MP/2 computer systems, the SSD connects to the mainframe using one 1000-Mbyte/sec channel. This configuration achieves a maximum transfer rate of 1000 Mbytes/sec. On a CRAY-1/M or CRAY-1/S computer system, the SSD connects to the mainframe through a high-speed controller. This configuration achieves a maximum transfer rate of 100 Mbytes/sec.

The SSD can also be connected to the I/O Subsystem (IOS). This connection enables data to be transferred between the SSD and the IOS directly, without passing through the central memory of the Cray mainframe.

## Software support

Software support for the SSD is provided by both COS and UNICOS. The following capabilities are common to both systems:

- Ability to copy and save all data from the SSD before a scheduled power-down
- Saved data can be restored at a later time
- Access to the SSD through customary use of standard I/O routines

Under COS, a user job has total control over the allocation of the SSD to each dataset created. COS configures the SSD as a disk device, and access to it is scheduled by the operating system. A user job must request any SSD needs. The job initiates

only after the requested amount of SSD resource becomes available.

COS allows the SSD to be oversubscribed. In an oversubscribed environment, device residence is scheduled on a priority basis (similar to the way central memory is scheduled), with optional manual control by the system operator.

Under UNICOS, the SSD can be configured for three different uses:

- As disk cache, which is used by any disk I/O request issued by any active process*

In this configuration, UNICOS allocates the most frequently accessed disk sectors to the SSD and supports them as cache-resident. This automatically provides SSD access for the most I/O-intensive processes.

---

---

- *As a UNICOS file system accessible through standard read/write routines*

Access is provided to a specific group or user with permission to access the file system. The file system is maintained with standard directory archive utilities.

- *As secondary data segments accessible through special SSD read/write routines*

Secondary data segments are blocks of fast-access secondary memory that are logically attached to process memory. They are device-resident only when the job is in memory.

#### **Physical characteristics**

SSD models 4, 5, 6, and 7 are available as stand-alone units. The SSD-3I is housed in the IOS cabinet.

The stand-alone SSD cabinet consists of four columns arranged in a 90° arc. The power

supplies are hidden by bench-like extensions arranged around the base. The cooling and packaging design of the SSD is based on field-proven enhancements to Cray Research product designs.

The SSD-3I resides within the IOS cabinet, using the IOS power supplies and cooling system. Incorporating the SSD-3I into the IOS cabinet helps minimize cost to the user.

#### **Physical characteristics (stand-alone models)**

- Four columns arranged in a 90° arc
- Requires 24 sq ft (2.3 sq m) of floor space
- Uses liquid refrigerant cooling
- Requires 400-Hz power from motor generators

#### **Reliability and maintenance**

Field experience with the SSD has demonstrated its high reliability and availability. The SSD memory is fully equipped with single-bit-error correction, double-bit-error detection

(SECDED) memory protection. Off-line diagnostics aid in maintenance of the unit. The SSD uses the same micro-computer used for maintenance and for logging memory errors on a Cray mainframe. This micro-computer also permits the SSD to be taken off-line from the computer system for diagnosis and repair.

A number of diagnostics are available to field engineers to aid in quickly identifying hardware problems. Further on-site diagnosis of a problem down to the component level can be performed off-line from the mainframe using a Cray Research module tester. This ensures maximum SSD availability and reliability.

#### **Further information**

For additional information about the SSD and other Cray Research products, please contact the nearest Cray Research sales office.

### SSD models and sizes

SSD model	Size in Mwords	Size in Mbytes	Exact size in 64-bit words
SSD-31	32	256	33,554,432
SSD-4	64	512	67,108,864
SSD-5	128	1024	134,217,728
SSD-6	256	2048	268,435,456
SSD-7	512	4096	536,870,912

Exact size in words = Size in Mwords x 1024 x 1024

### Performance comparisons

Configuration	CPU time (seconds)	I/O wait time (seconds)	Total elapsed time (seconds)	
Structural analysis problem*	CRAY X-MP/2 system with SSD-5	165	5	170
	CRAY X-MP/2 system with DD-49 disk drives	165	379	544
Petroleum reservoir modeling problem**	CRAY X-MP/4 system with SSD-5	92	20	112
	CRAY X-MP/4 system with DD-49 disk drives	92	231	323

\* During this benchmark, only one processor of a two-processor CRAY X-MP system was used; only 14 million words of a 128-million-word SSD were used.

\*\* During this benchmark of a 35,000-cell black-oil model, only one processor of a four-processor CRAY X-MP system was used; only 6 million words of a 128-million-word SSD were used.

### SSD transfer rates

CRAY X-MP models	Maximum number of SSD channels	Transfer rate (Mbytes/sec per SSD channel)	Maximum transfer rate (Mbytes/sec)
CRAY X-MP/1	1	1000	1000
CRAY X-MP/2	1	1000	1000
CRAY X-MP/4	2	1000	2000

With a CRAY X-MP/4 computer system, the SSD-31 has a maximum of one channel and, therefore, a maximum transfer rate of 1000 Mbytes/sec. The CRAY X-MP/14se does not support an SSD connection.

---

**CRAY**  
**RESEARCH, INC.**

608 Second Avenue South  
Minneapolis, MN 55402  
612/333-5889  
Telex: 4991729

**Domestic sales offices**

Albuquerque, New Mexico  
Atlanta, Georgia  
Beltsville, Maryland  
Boston, Massachusetts  
Boulder, Colorado  
Chicago, Illinois  
Cincinnati, Ohio  
Colorado Springs, Colorado  
Dallas, Texas  
Darien, Connecticut  
Detroit, Michigan  
Houston, Texas  
Huntsville, Alabama  
Laurel, Maryland  
Los Angeles, California  
Minneapolis, Minnesota  
Pittsburgh, Pennsylvania  
Pleasanton, California  
Seattle, Washington  
St. Louis, Missouri  
Sunnyvale, California  
Tampa, Florida  
Tulsa, Oklahoma

**International subsidiaries**

Cray Asia/Pacific Inc.  
Hong Kong

Cray Canada Inc.  
Toronto, Canada

Cray Research France S.A.  
Paris, France

Cray Research GmbH  
Munich, West Germany

Cray Research Japan, Limited  
Tokyo, Japan

Cray Research S.R.L.  
Milan, Italy

Cray Research (UK) Ltd.  
Bracknell, Berkshire, UK

CRAY, CRAY-1, and SSD are registered trademarks and COS, CRAY X-MP, IOS, and UNICOS are trademarks of Cray Research, Inc. The UNICOS operating system is derived from the AT&T UNIX System V operating system. UNICOS is also based, in part, on the Fourth Berkeley Software Distribution under license from The Regents of the University of California. UNIX is a trademark of AT&T Bell Laboratories.