

SX-6

For Information Contact:

▶ ASIA
NEC HPC MARKETING
PROMOTION DIVISION

7-1 Shiba, 5-chome
Minato-ku, Tokyo 108-8001
Japan
+81-3-3798-9131 phone
+81-3-3798-9132 fax
info@sxsmd.ho.nec.co.jp

▶ NEC High Performance
Computing Europe

Prinzenallee 11
D-40549 Düsseldorf
Germany
+49-211-5369-0 phone
+49-211-5369-199 fax
info@hpce.nec.com

▶ LATIN AMERICA
NEC DO BRASIL S.A. SX-OFFICE

Rua Arabé, 71
CEP 04042-070 V. Clementino
São Paulo SP
Brasil
+55-11-5591-7147 phone
+55-11-5591-7146 fax
ccp@nec.com.br

▶ OCEANIA
NEC AUSTRALIA PTY. LTD. HPCD

635 Ferntree Gully Road
Glen Waverly, VIC 3150
Australia
+61-3-9262-1209 phone
+61-3-9262-1534 fax
info@sx.nec.com.au

NEC SX-6 MULTI-NODE

SCALABLE TO MEET EVEN UTMOST DEMANDS

SOFTWARE

Operating System and Applications

With traditional NEC reliability and the SUPER-UX operating system, now in its 12th year, SX-6 is designed for production sites to perform production computing. All the functions other systems promise, like checkpoint-restart or a robust batch environment, are available today.

The SUPER-UX multi-node kernel is enhanced to recognize a multi-node job class. When a multi-node job enters the system, the kernel will sequence all of the processes across the nodes, initialize the IXS page translation pages for the job, and provide specialized scheduling commensurate with the resources being used. Applications development is made easy through the simplicity of shared memory programming

within a node, and portability is assured by industry-standard MPI and OpenMP support. NEC's PSUITE Integrated Development Environment provides all of the tools and utilities necessary under a single package for project management, editing, compiling, optimizing, and test/debugging. PSUITE is available for cross hosting on all popular workstation class products as well as Linux personal computers to maximize accessibility and development efficiency. The languages, libraries and tools available include Fortran90, OpenMP, C++. Tools and libraries dedicated to the development of parallel and multi-node jobs include MPI, the TotalView debugger and the Vampir/SX performance analysis tool.

CONFIGURATION TABLE

Selected Machine Configurations multi-node Systems

Model Group Name	SX-6/M							
Model Name	1024M128	512M64	256M32	128M16	64M8	32M4	16M2	8M2
CPU								
Number of Nodes	128	64	32	16	8	4	2	2
Number of CPUs	1024	512	256	128	64	32	16	8
Peak Performance ^{*1}	11950GF	5975GF	2987GF	1493GF	746.88GF	373.44GF	186.72GF	93.36GF
Peak Vector Performance ^{**}	9216GF	4608GF	2304GF	1152GF	576GF	288GF	144GF	72GF
Vector Registers	144KB x 1024	144KB x 512	144KB x 256	144KB x 128	144KB x 64	144KB x 32	144KB x 16	144KB x 8
Scalar Registers	64bitsx128x1024	64bitsx128x512	64bitsx128x256	64bitsx128x128	64bitsx128x64	64bitsx128x32	64bitsx128x16	64bitsx128x8
Main Memory Unit								
Memory Architecture	Shared/Distributed Memory							
Max. Capacity	16TB	8TB	4TB	2TB	1TB	512GB	256GB	256GB
Peak Data Transfer Rate	36864GB/s	18432GB/s	9216GB/s	4608GB/s	2304GB/s	1152GB/s	576GB/s	288GB/s
Input/Output Processor								
Max. Number of Channels	16256	8128	4064	2032	1016	508	254	254
Peak Data Transfer Rate	1024GB/s	512GB/s	256GB/s	128GB/s	64GB/s	32GB/s	16GB/s	16GB/s
Internode Crossbar Switch								
Peak Data Transfer Rate	1024GB/s	512GB/s	256GB/s	128GB/s	64GB/s	32GB/s	16GB/s	16GB/s

*1 Peak performance: Summation of total floating point operations per second which can be executed simultaneously in all the floating point execution units (multiply, add/shift and divide pipelines in the vector unit and multiplier and adder/divider in the scalar unit).

** Summation of floating point operations per second which can be executed simultaneously only in vector multiply and vector add pipelines.



EXTREME PERFORMANCE THROUGH VECTOR PROCESSING AND SCALABILITY



By combining the ease of use and efficiency of shared memory with the scalability of distributed memory systems, NEC created one of the most powerful computing systems available today: The NEC SX-6 multi-node, which scales up to configurations with 11 TeraFlops (TF) of peak performance^{*1}, 16 Ter-

aByte (TB) of main memory and 1 TeraByte per second of inter-node communication bandwidth. Vector supercomputers have always provided the absolute highest performance available. High-performance high-bandwidth memory and powerful processors deliver maximum performance on applications and

commercial robustness for production sites.

Now SX-6 breaks new ground, introducing high-end scalable parallel vector supercomputing to the technical server competitive space. With powerful shared-memory nodes and unequaled inter-node communications bandwidth the SX-6 is second to none.

SX-6




SCALABILITY

HARDWARE FEATURES

Model Configuration

A SX-6 node is a complete parallel vector system consisting of up to eight vector processors each with 11.67 GFLOPS of peak performance*. The processors are coupled to an uniform shared main memory of 128 GB capacity. The powerful SMP single-node with uniform shared memory provides very high levels of performance for both capacity and capability requirements. Multi-node configurations with distributed memory scale beyond the limits of a single node. The SX-6 series is compatible with its predecessor SX-5 series. It excels in the total balance of processing performance, memory throughput, input/output performance in much the same way as the former SX series systems did. Existing applications and resources can be easily migrated to the SX-6.



SX-6 Memory module

Single-node System

The SX-6 series single-node models scale up to 8 CPUs, delivering up to 93.36 GF of peak performance*¹ and offering a maximum of 128 GB main memory in shared memory architecture. The memory can stream data to each processor at 36 GB/s for a total memory bandwidth of 288 GB/s.

Multi-node System



From two up to 128 nodes can be connected through NEC's proprietary ultra-high speed Internode Crossbar Switch (IXS). SX-6 series multi-node models cover a range from 8 to 1024 CPUs with a maximum peak performance*¹ of 11 TF. The maximum total system memory scales to 16 TB. Each node is connected via an 8 GB/s channel to the crossbar switch, the maximum throughput of the crossbar being 1TB/s. Although they are built on a distributed memory architecture, multi-node systems still provide a single-system-image. Extreme execution performance can be obtained on a wide range of applications with the powerful single-node systems connected through the IXS.

Ultra-high-speed Vector and Scalar Unit

The vector unit of the SX-6 series processor consists of vector registers and 8 sets of pipelines for logical operations, multiplication, add/shift operations, division, masked operations and load/store. The scalar unit achieves ultra-high-speed performance through a 4-way super scalar design. The combination of the single-chip vector microprocessor with a reduced clock cycle decreases the processing time for each instruction. This leads to the superior short vectors and scalar performances.

Internode Crossbar Switch

The SX-6 series multi-node configurations use an exclusive and proprietary internode crossbar switch (IXS) to connect the individual nodes with an ultrahigh throughput and low latencies network. Eight gigabytes per second of bisection bandwidth is available for each node, for a total of maximum 1024 gigabytes per second internode ultra-high-speed data transfer.




TECHNOLOGY

The Worlds first single-chip Vector Processor

The high gate density possible for the state-of-the-art CMOS technology and LSI design enabled NEC to implement the vector processor on just one chip. This LSI and packaging technology leads to a peak performance*¹ of 11.67 GFLOPS on a single LSI. This ultrahigh integration leads to improved internal latencies and performance in comparison with former generation designs, which used dozens of chips to implement a processor, as well as highly reduced memory latencies by drastically narrowing the distance between memory and processors.

Main Memory Unit

The SX-6 series utilizes ultra-high-speed double data rate synchronous DRAM. Single-node systems have a memory capacity of up to 128 GB and a memory bandwidth 288 GB per second. Multi-node systems have a memory capacity of up to 16 Terabytes and a maximum bandwidth of 36.8 Terabytes per second.



SX-6 single-chip vector processor

IXS Enabled Global Main Memory

For access to memory across nodes the IXS provides page translation tables and global data movement instructions. Because of the characteristics of the IXS block-memory-move instructions offer the highest performance.

All memory is protected by lock-and-key mechanisms under control of the SUPER-UX multi-node operating system. Hardware support tables for memory security are located in both node hardware and IXS hardware. Latencies for internode NUMA memory access are less than most workstation technology NUMA implementations, and the 8 gigabyte per second bandwidth of just a single IXS channel exceeds the entire memory bandwidth of most SMP class systems. This architecture, introduced with the SX-4 series, lends itself to a combination of traditional parallel vector processing (OpenMP, microtasking) combined with message passing (MPI). Message passing alone is also highly efficient on the architecture.

Input/Output Subsystem

The Input/Output subsystem of the SX-6 series can be configured to deliver a bandwidth of up to 8 GB per second on a single-node system. Multi-node systems scale up to an I/O capacity of 1 TB per second.

Global File System

The NEC SX-6 operating system supports SX-GFS, a proprietary global file system. SX-GFS presents a common file system view to the entire NEC SX-6 multi node system. SX-GFS can achieve up to 80% of the performance of a locally attached similar storage device for large I/O requests. NEC offers its Linux-based TX7/i Series machine as a file server for GFS installations.

Ease of Installation

The SX-6 models' power consumption and space requirements have been reduced by 80% when compared with the previous generation of the SX series. The low power consumption allows all models to be fully air-cooled. These two elements contribute to a great reduction of installation costs and complexities.

High Reliability

The usage of highly integrated CMOS technology has led to greatly reduced number of components in a single system. This, in turn, leads to a tremendously improved hardware reliability.