

NEC HPC Systems and Future

26/10/2004

NEC

Toshiyuki Furui

NEC's approach to HPC business

Computer Business Strategy

SX-8

Slide 3

World's technology/quality leader


Technology Leader Products

Advanced Technology & High Quality

Innovative technology for
supercomputer drives
evolution of other products



SX
Supercomputer



ACOS mainframe



NX
Unix Server




f i
Server




iStorage



IPF
Server



IPF
Blade Server



IA
Blade Server

Volume Sales Products

Low Cost & High Quality



Express
PC Server



PC



Special Purpose
Servers



DVD Drive



LTO

World's HPC leader

■ Covers wide range of products for HPC to meet every requirement (Right Product for Right Application)

- Vector supercomputers : SX series
- IPF* scalar servers : SMP server TX7 series
Blade Server Express5800/1020Ba
- EM64T* PC clusters : Express5800 parallel clusters
- Global File System : File sharing solution among various platforms (Collaboration System)

■ Higher weight on vector supercomputer development

- Crucial product for actual HPC fields
- Keep overwhelming sustained performance

■ Provides open scalar servers timely to the markets

- Rich line-ups with IPF / EM64T processors
- Flexible solution in cooperation with SX series

IPF :Intel® Itanium® Processor Family
EM64T :Intel® Extended Memory 64 Technology

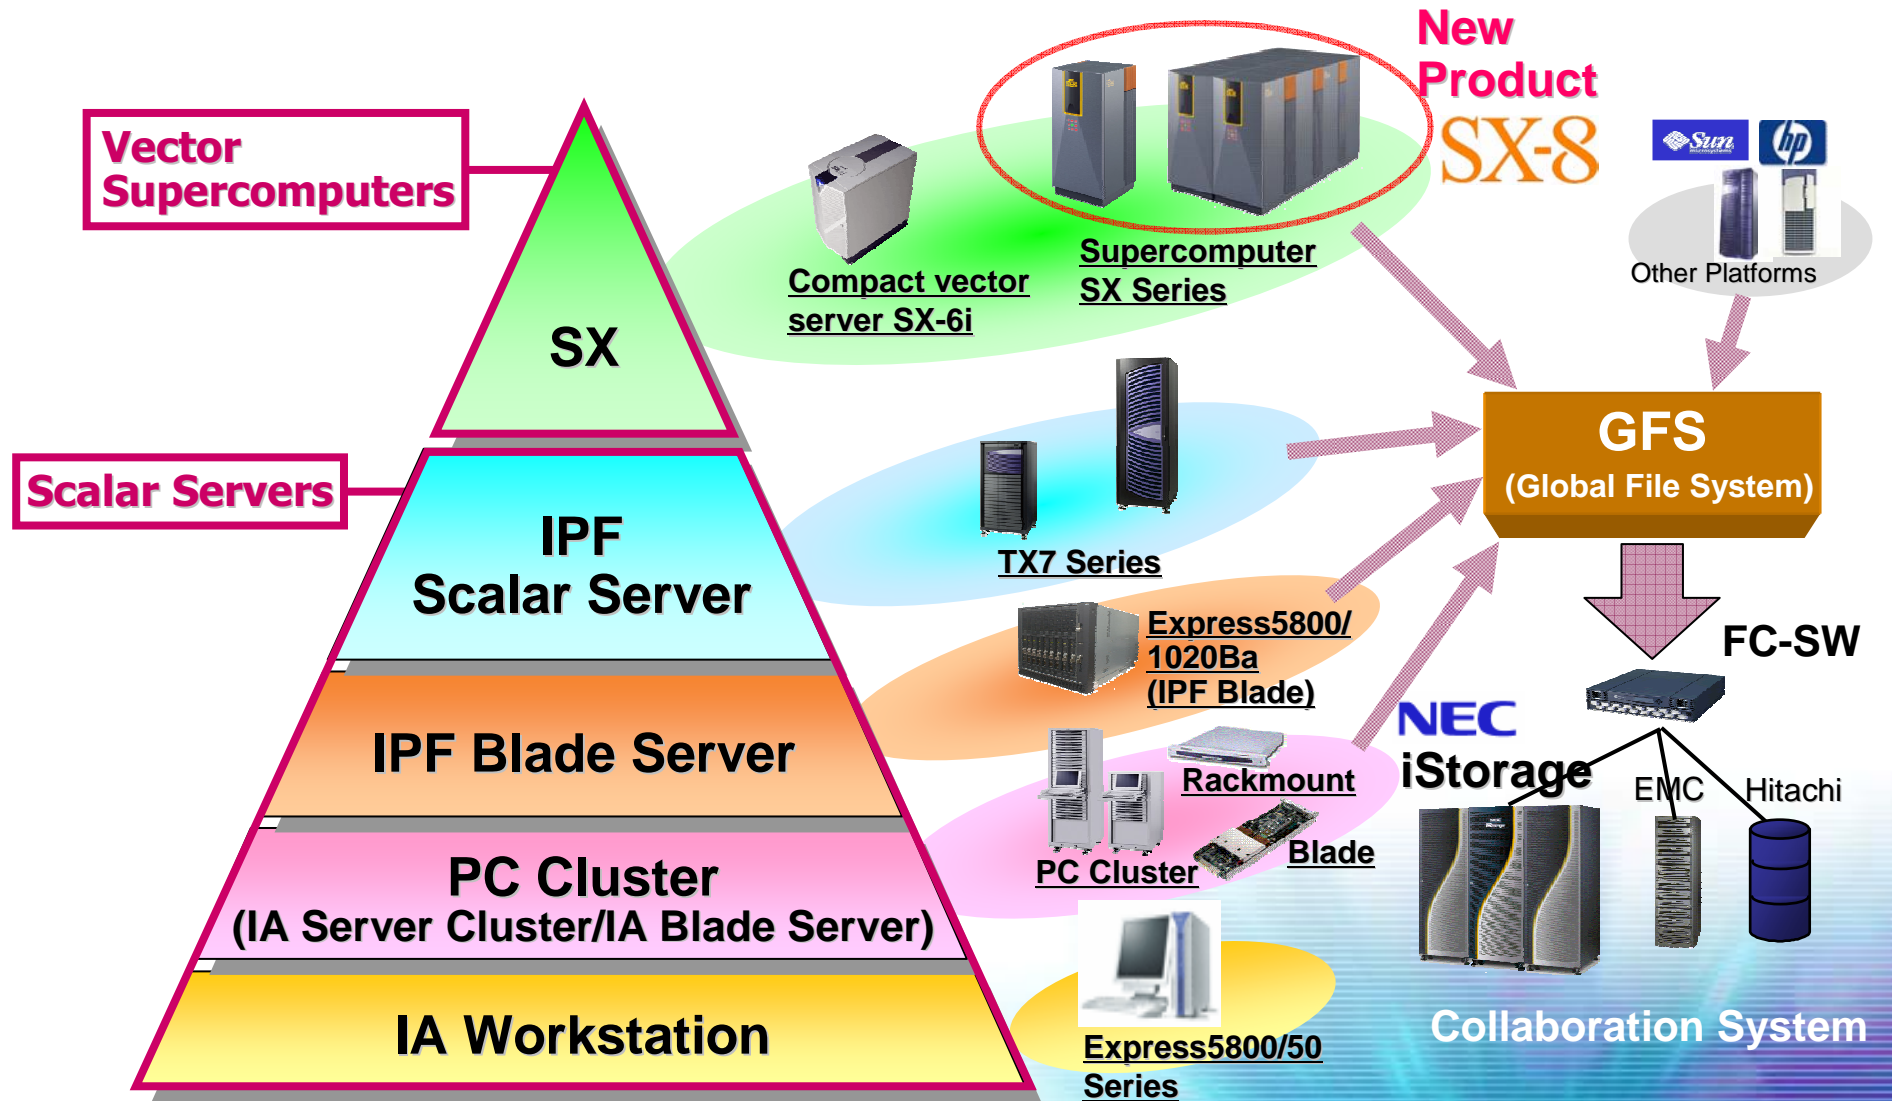
Empowered by Innovation

NEC

HPC Product Lineup

SX-8

Slide 5



Empowered by Innovation

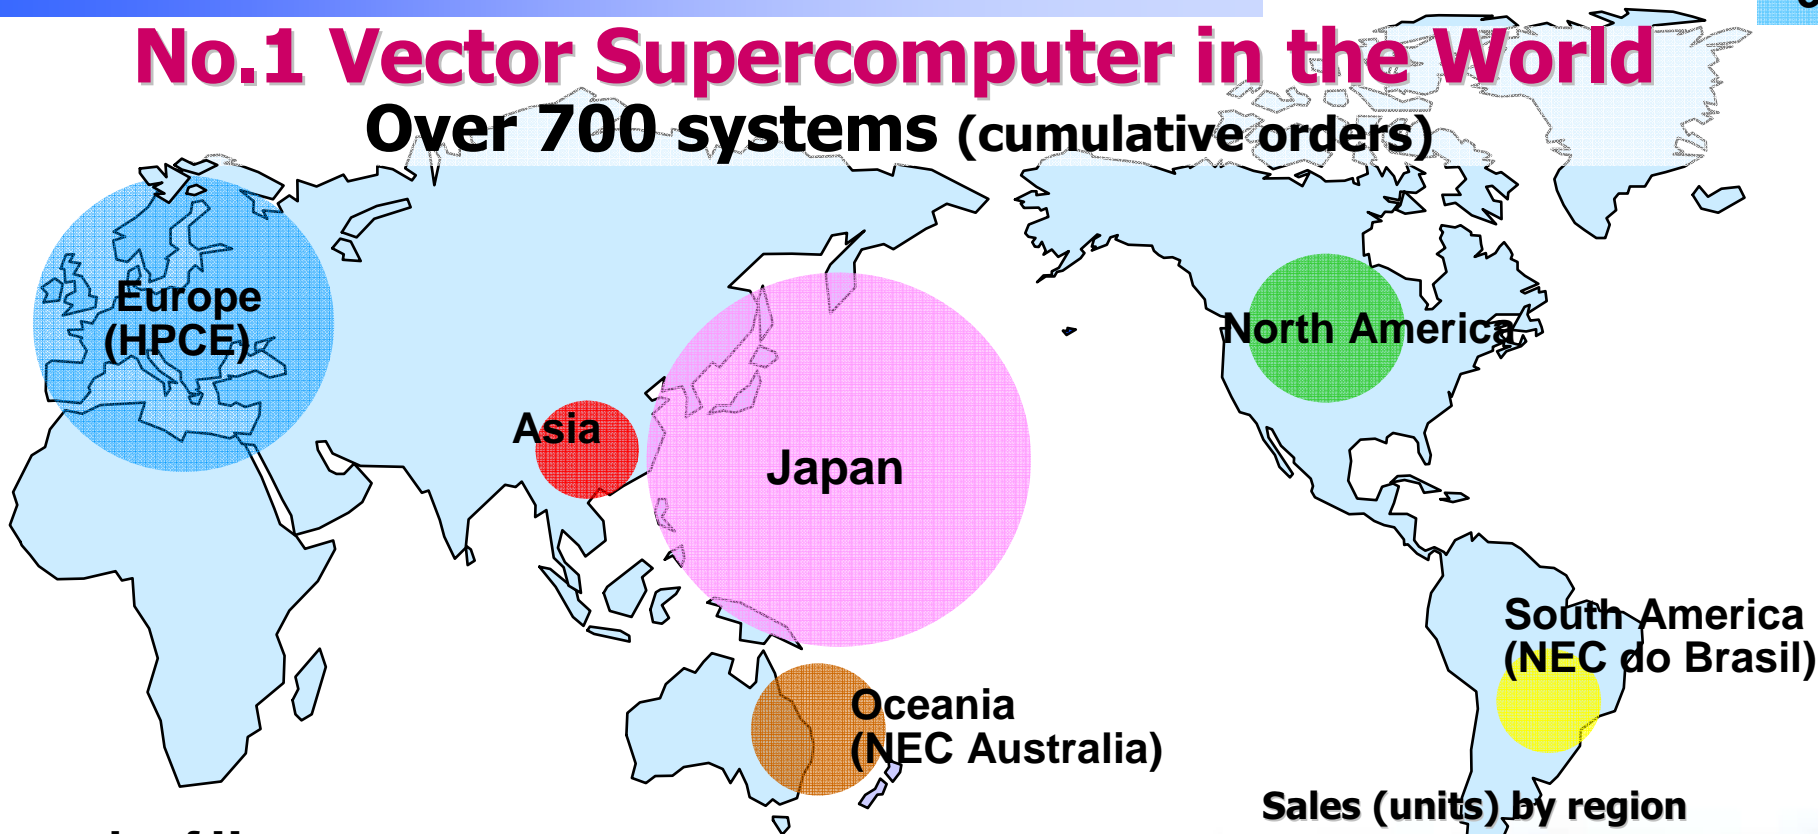
NEC

Sales Status of SX Series

SX-8

Slide 6

No.1 Vector Supercomputer in the World
Over 700 systems (cumulative orders)



Example of Users

Europe

- The Met Office (UK)
- University of Stuttgart (Germany)
- German Climate Computing Center (Germany)
- Volkswagen (Germany)
- Aerospace Laboratory (Netherlands)
- Swiss Center for Scientific Computing (Switzerland)
- CNRS/IDRIS (France)
- Danish Meteorological Institute (Denmark)

Japan

- Osaka Univ./Tohoku Univ. Computer Center
- National Institute for Environmental Studies
- National Institute for Fusion Science
- National Research Institute for Metals
- Nissan Motors
- Obayashi Corporation
- Osaka Gas Co., LTD.
- Mitsui Chemical

Sales (units) by region



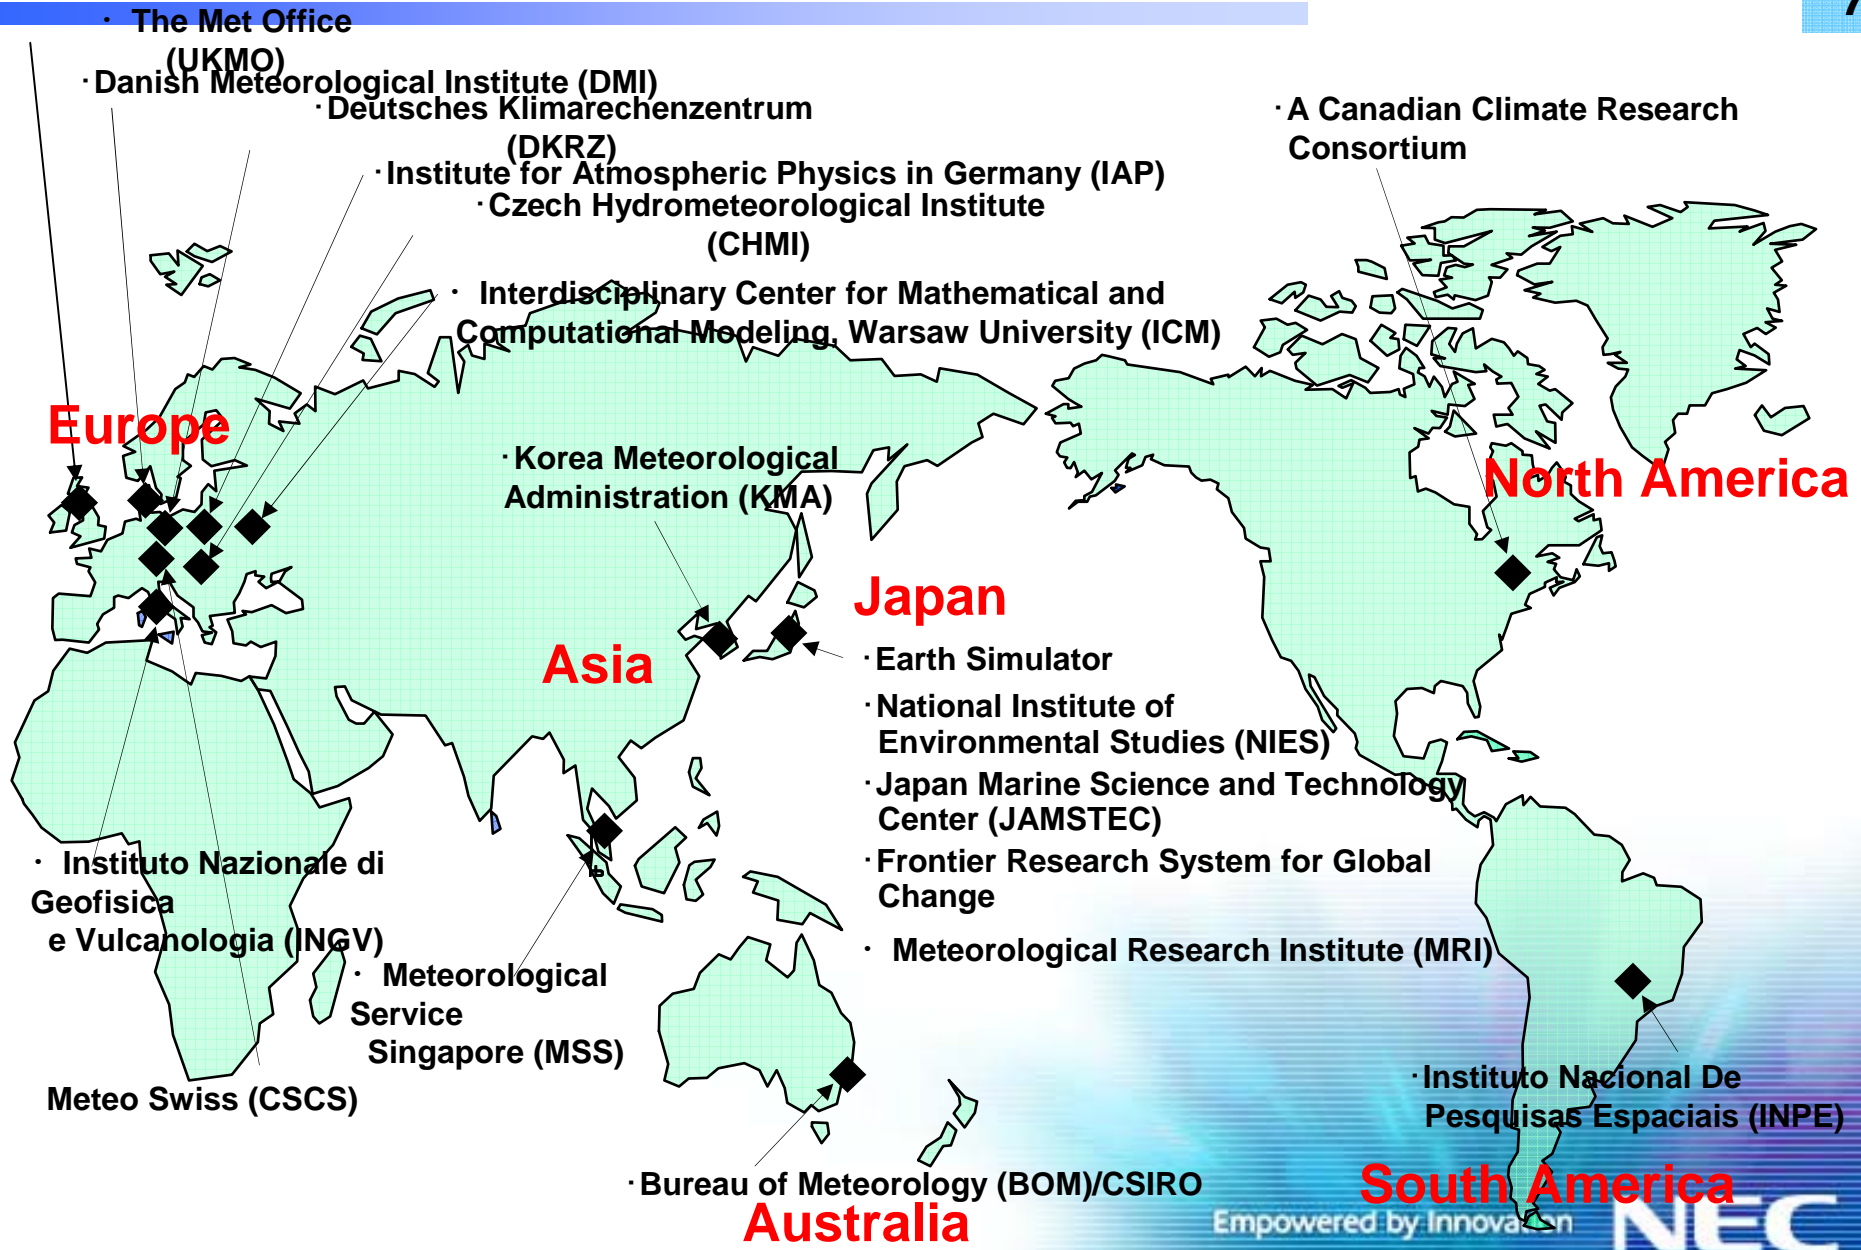
Asia Pacific

- Bureau of Meteorology / CSIRO (Australia)
- Korea Meteorological Administration (Korea)
- Korea Institute of Science, Technology and Information (Korea)
- Meteorological Services of Singapore (Singapore)

Empowered by Innovation

NEC

NEC's Contribution to Meteorological Area



Supercomputer SX series

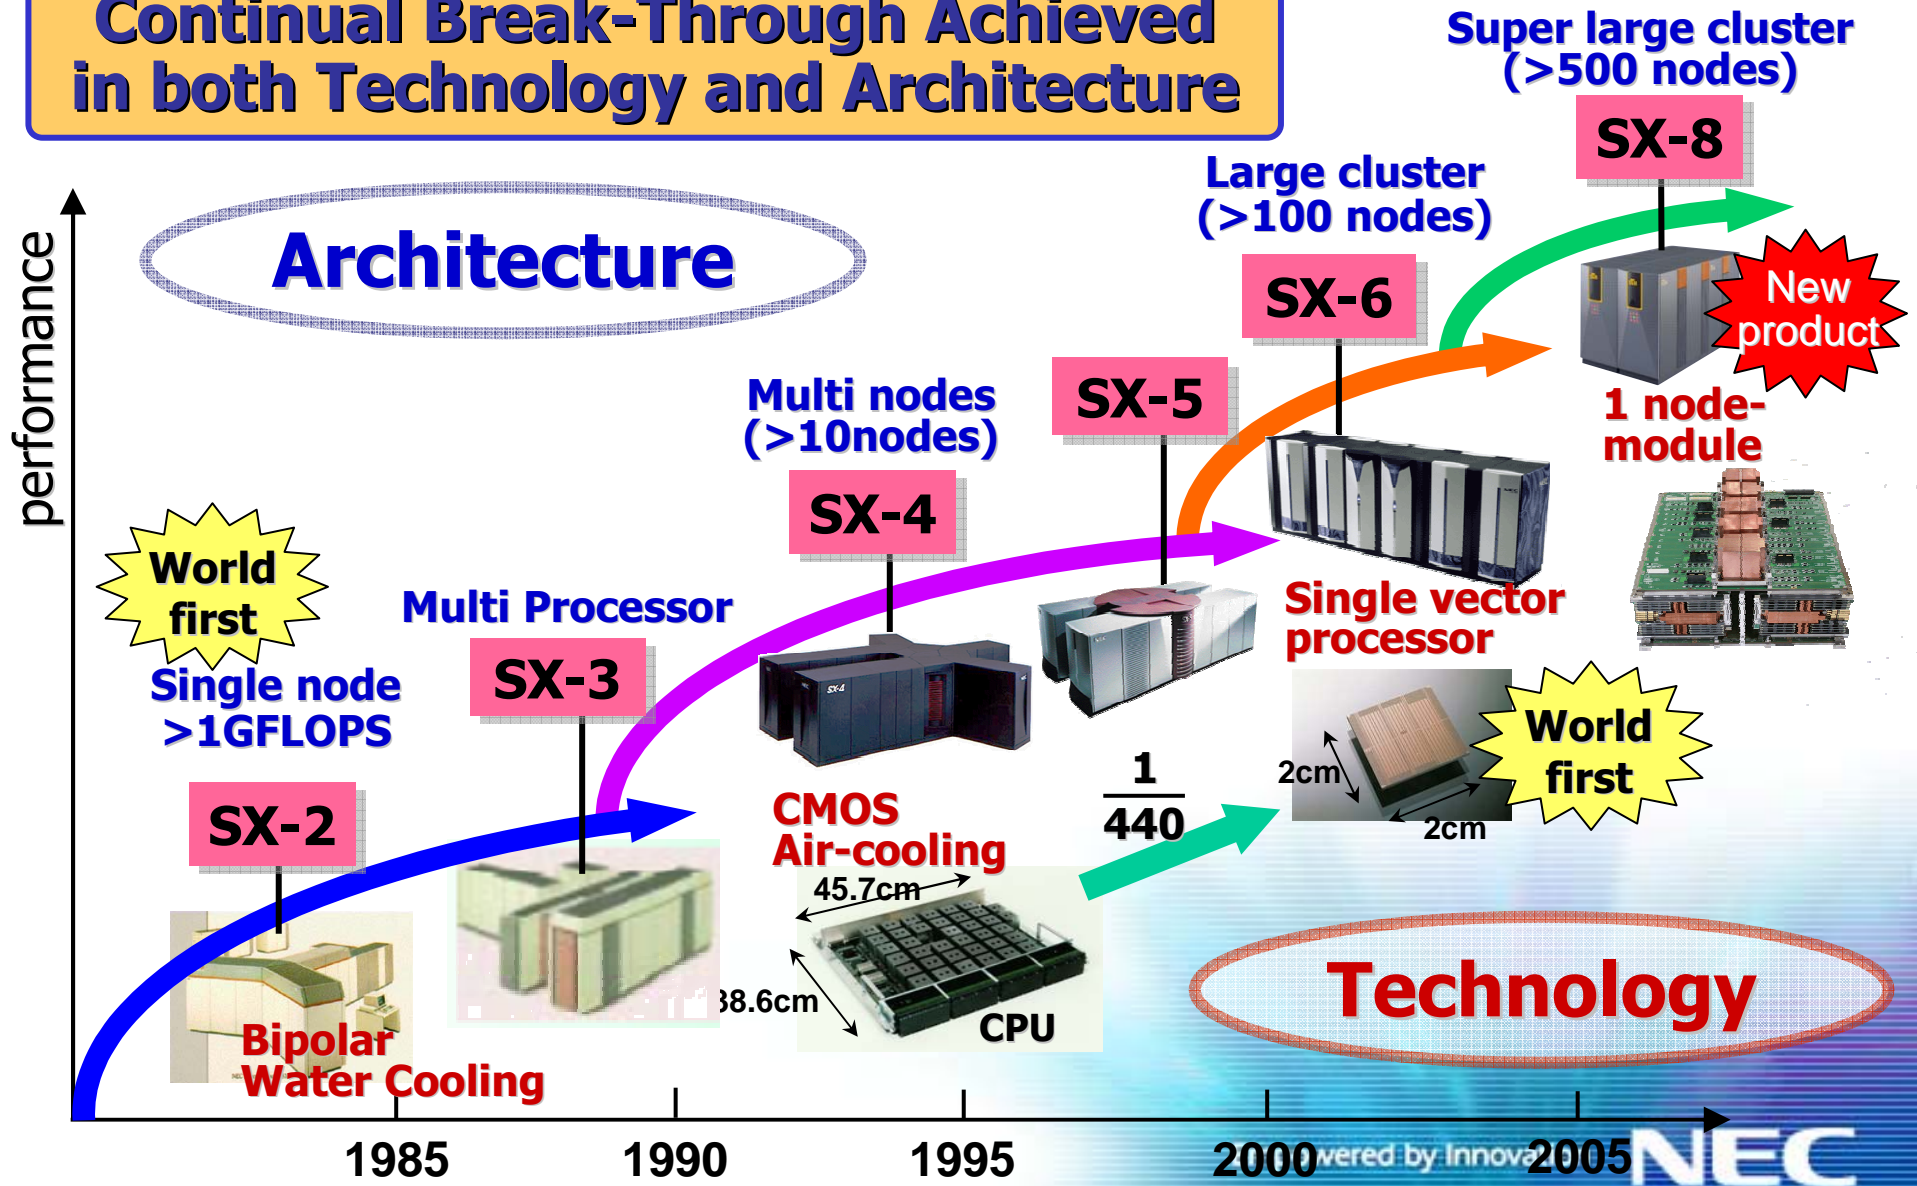
Model SX-8

SX Innovation

SX-8

Slide 9

Continual Break-Through Achieved in both Technology and Architecture



SX-8 development objective

SX-8
Slide 10

- **Provide super high performance platform for high-end computing**
 - Excellent Performance/Price **> 3X (to SX-6)**
 - TCO saving : Space-saving **≒1/4 (to SX-6)**
Power-sa **≒1/2 (to SX-6)**
 - High sustained performance with SX architecture
- **Upward compatible to SX-6 system**
- **Breakthrough by leading-edge technologies**
 - Latest LSI process
 - Low-loss, high-density card packaging structure

SX-8 Product Highlights

1. World's fastest vector supercomputer with maximum performance of 65 TFLOPS

- Very large scale : Up to 512 nodes, 4,096 CPUs (4x to SX-6)
- Very large memory / memory bandwidth: 64TB / 262TB/s (8x to SX-6)
- High speed data transfer between nodes : 8TB/s in total (8x to SX-6)

2. High-density packaging with state-of-the-art technology

- Further-enhanced single-chip vector processor with 16GFLOPS performance
- Leading-edge CMOS technology with 90-nanometer process /copper interconnects
- Single-module node with 128GFLOPS performance (2x to SX-6)

3. Enhanced SUPER-UX / Tuned applications

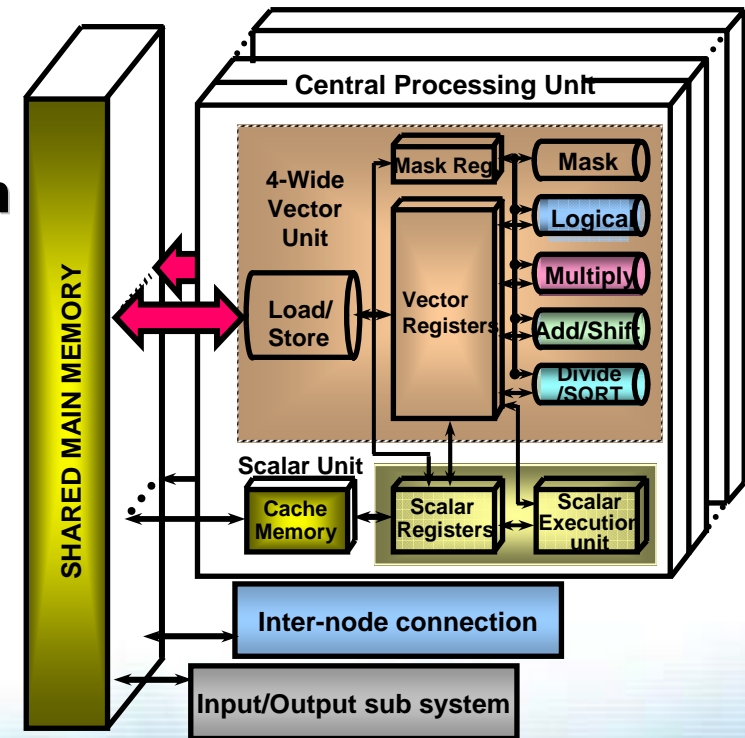
- Proven operating system for SX series enhanced to expand scalability
- A lot of ISV application programs tuned for SX series available

SX-8 Architecture

SX-8

Slide 12

- Upward compatible to SX-5/SX-6
- Vector pipelines
 - 4 logical pipelines : 2GHz
 - 144KB vector register
 - Hardware support of **SQRT** instruction
- Scalar processor
 - 4 way superscalar RISC
- Main memory
 - 2 types of RAMs
 - DDR2-SDRAM : Large capacity 128GB/node
 - FCRAM : High-speed 64GB/node
- Multi node system
 - up to 512 nodes
 - 64 TFLOPS
- Enhanced I/O performance
 - Reduction of I/O overhead by adopting direct CPU control



Empowered by Innovation


NEC

SX-8 Single Node Module

SX-8

Slide 13

- **Up to 8 CPUs/node**
 - Peak Vector Performance(PVP):
16 GFLOPS/CPU
128 GFLOPS/node
- **Symmetric multiprocessing (SMP)**
- **Large Capacity Memory**
 - Up to 128GB
- **Ultra-high memory bandwidth**
 - 64GB/s per CPU
 - Total 512GB/s per node
- **Large I/O throughput**
 - 12.8GB/s per node



Empowered by Innovation

NEC

Large Scale Multi Node System

SX-8

Slide 14

High speed processing of large data with high performance single node, large number of nodes, and high speed interconnects among nodes

Key points for high performance

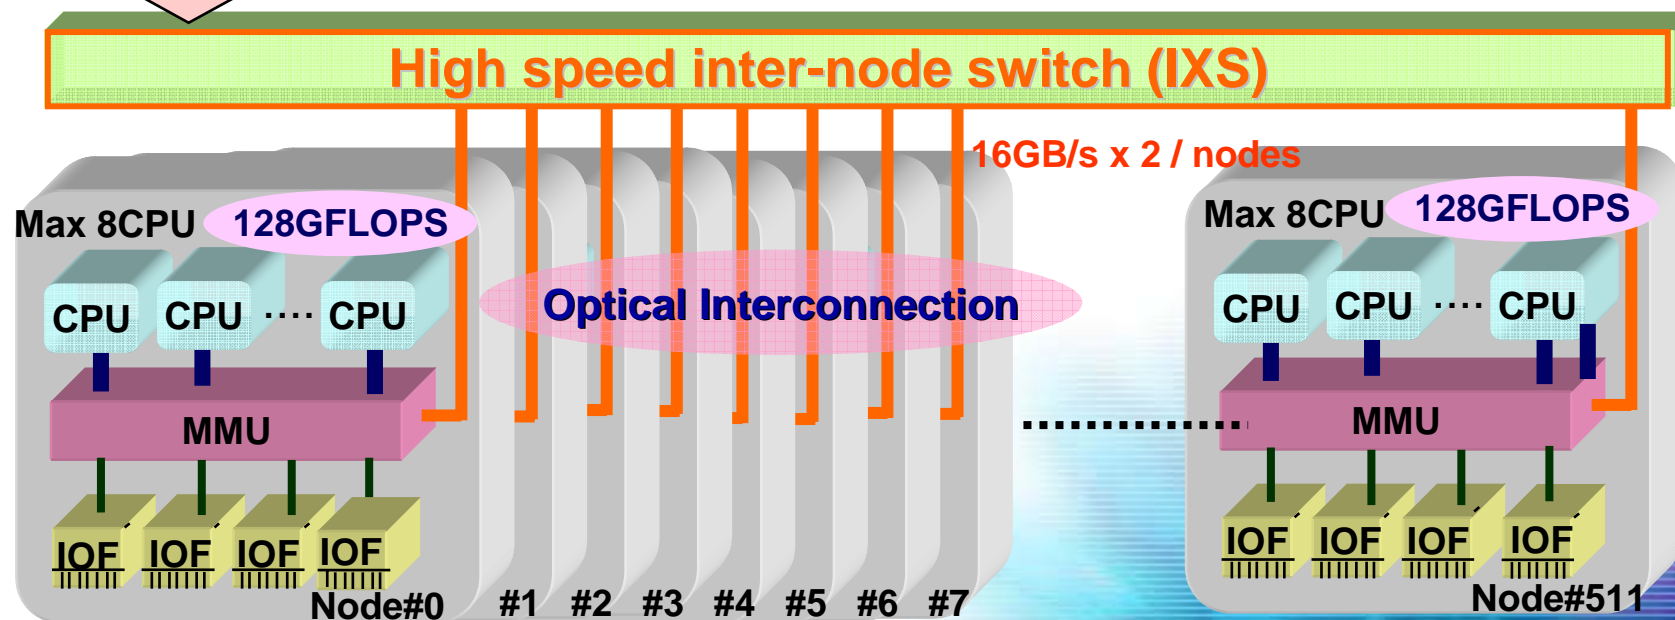
1. Single node performance
2. Maximum number of node
3. Data transfer rate among nodes

Max 128GFLOPS 2x (to SX-6)

Max 512 nodes 4x (to SX-6)

Max 8TB/s
(Peak data transfer rate) 8x (to SX-6)

Very efficient non-blocking switch




Max 512 nodes

Empowered by Innovation

NEC

SX architecture enables optimization for Multi-Grain parallelism

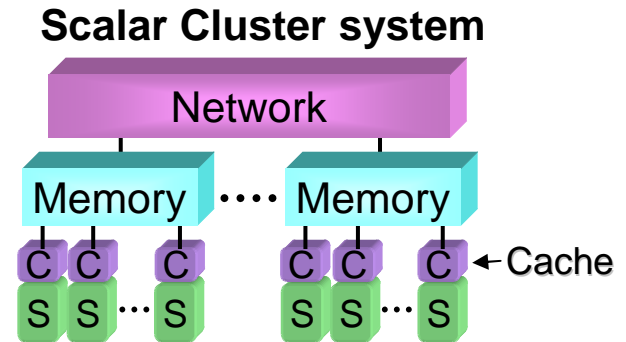
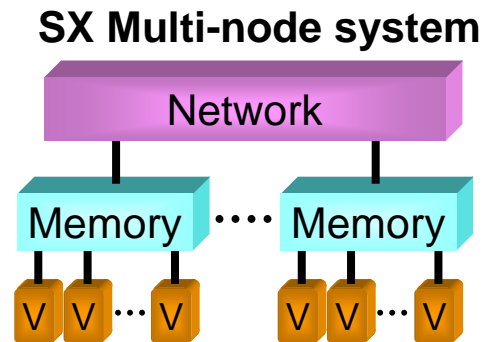
Grain-size	Level	HW architecture	Compiler/Programming
Fine-grain	Instruction	Vector-SIMD	Vectorization
Medium-grain	LOOP	SMP	Microtask/OpenMP
Coarse-grain	Thread	Multi-node Cluster	HPF/MPI



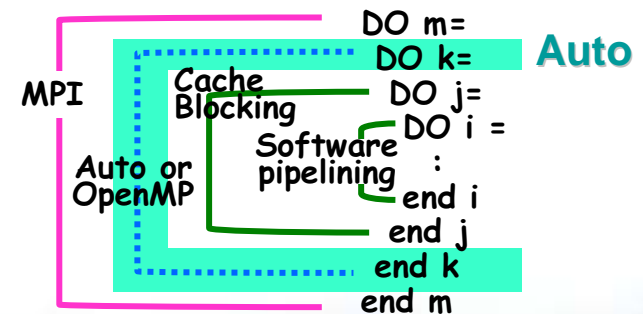
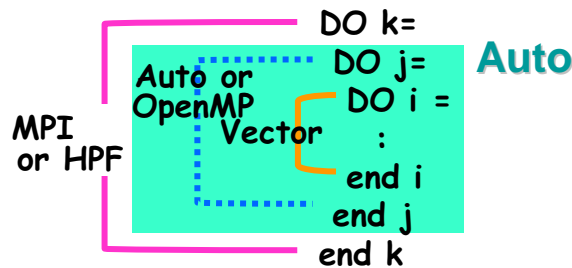
Advantages of SX: Easy Programming !!

- Programming for SX: Easier than Scalar Cluster
- Provides Auto-parallelizing features for all parallel levels

System architecture



Programming technique




	SX Multi-node system	Scalar cluster system
Distributed Memory (Coarse-Grain)	<ul style="list-style-type: none"> • MPI (User) • HPF (User/Compiler) 	<ul style="list-style-type: none"> • MPI (User)
Shared Memory (Medium-Grain)	<ul style="list-style-type: none"> • Auto-parallelization (Compiler) • OpenMP (User/Compiler) 	<ul style="list-style-type: none"> • Auto-parallelization (Compiler) • OpenMP (User/Compiler)
Processor (Fine-Grain)	<ul style="list-style-type: none"> • Auto-vectorization (Compiler) 	<ul style="list-style-type: none"> • Cache Blocking (User) • SW pipelining (User/Compiler)

Advantages of SX: Memory bandwidth

SX-8

Slide 17



	ALU/CPU	FP Reg. Vec Reg.	CBW/CPU	Cache (data)	MBW/CPU	Transfer mode	Memory (Total)
SX-8 (VecProc. @2GHz, 8CPU/Node)	16GF (8op./clk)	144KB (256W x 72Reg.)	←		0.5data/op. (64GB/s)	· block · stride · list	16GB/CPU (128GB/8CPU)
SX-6 (VecProc. @1GHz, 8CPU/Node)	8GF (8op./clk)	144KB (256W x 72Reg.)			0.5data/op. (64GB/s)	· block · stride · list	8GB/CPU (64GB/8CPU)
SX-5 (VecProc. @250MHz, 16CPU/Node)	8GF (32op./clk)	144KB (256W x 72Reg.)			1data/op. (64GB/s)	· block · stride · list	8GB/CPU (128GB/16CPU)
Cray X1e (VecProc. @1.2GHz, 8CPU?/Node)	19.2GF (16op./clk)	64KB (256W x 32Reg.)	2B/op. /CPU	E-cache: 2MB/2CPU?	0.08data/op. (25.6GB/s /2CPU?)	· block · stride · list	4GB/CPU (32GB/8CPU)
IBM p690 (Power4+ @1.9GHz)	7.6GF (4op./clk)	256B (32Reg.)	4B/op.(L1)	L1: 32KB L2:1.5MB/2CPU L3:128MB/8CPU	0.08data/op. (10.6GB/s /2CPU)	· block	16GB/CPU (256GB/16CPU)
SGI ALTix3700 (Itanium2 @1.5GHz)	6GF (4op./clk)	1KB (128Reg.)	4B/op.(L2)	L1: Non for FP L2: 256KB L3: 6MB	0.07data/op. (6.4GB/s /2CPU)	· block	8GB/CPU (512GB/64CPU)

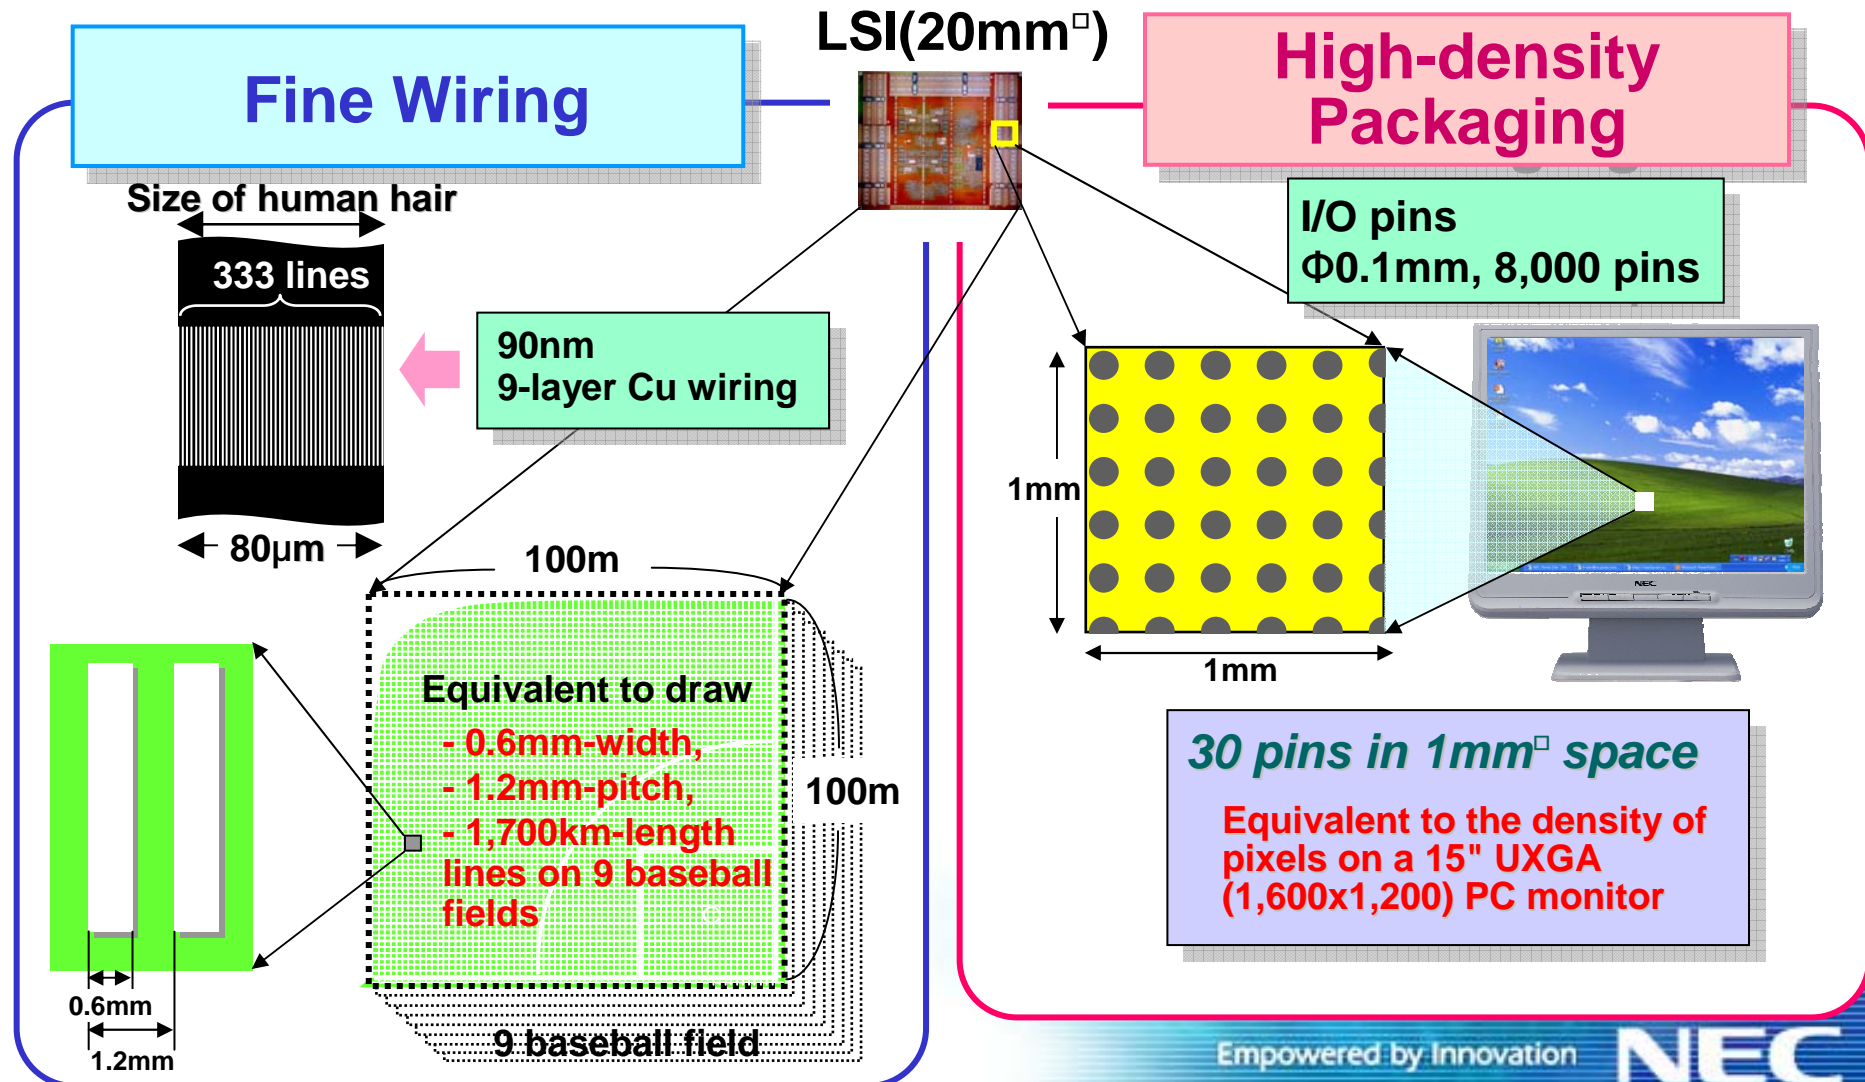
MBW: Memory Bandwidth, CBW: Cache Bandwidth

Empowered by Innovation

NEC

CPU LSI technology

Overwhelming CPU performance achieved with the most advanced LSI technology



Packaging Technology

Node Module


SX-6

64GFLOPS



SX-8


128GFLOPS



**Cubic Volume
/Performance
1 / 5**

- High-density multilayer printed wire board for high-speed signal transmission
 - Eliminated all cable connections to downsize the module
- Packaging density increased by reduction of LSI power consumption and high-efficiency cooling design


SX Single Node Comparison



	SX-5	SX-6	SX-8
Size (W x D x H)(mm)	6320x3200 x1800	1000x1100 x1800	825x850 x1800
Footprint (Including Maintenance area)	30.4m ²	3m ²	2.5m ²
Weight	6950kg	730kg	570kg
Power consumption	90KVA	7KVA	7KVA
Cooling	Air	Air	Air

Excellent Space-saving

Installation space significantly improved with high-density packaging technology




	SX-6 (16 nodes)	SX-8 (8 nodes)
Installation Space*	48.0m ²	13.0m ²

*including maintenance space



SX-8 Sample Configuration System Layout

256-node example

PVP: 32TFLOPS



SX-8 Models



Single-node					Multi-node		
Model	CPU		Memory		Max Channel #*	Max Node #	Inter-node Bandwidth
	#	Max performance	Type	Max Capacity			
A	4 ~ 8	128GF	DDR2	128GB	55	512	16GB/sx2 or 8GB/sx2
			FCRAM	64GB			
B	1 ~ 4	64GF	DDR2	64GB	27	8	8GB/sx2
			FCRAM	32GB			

(*) Excluding one channel for system control

Application Software covering Wide Areas

ISV Application Software

Structural Analysis

- ❖ MSC.Nastran
- ❖ ADINA
- ❖ AKUSMOD
- ❖ PERMAS
- ❖ INDEED
- ❖ OptiStruct
- ❖ OPTISHAPE
- ❖ CDH/AMLS

Computational Chemistry

- ❖ Cerius2/ADF
- ❖ AMBER
- ❖ GAMESS
- ❖ GAUSSIAN
- ❖ GROMOS
- ❖ MM2
- ❖ MNDO94
- ❖ MOPAC

Crash Analysis

- ❖ LS-DYNA
- ❖ PAM-CRASH
- ❖ PAM-STAMP
- ❖ RADIOSS

Fluid Dynamics

- ❖ PAM-FLOW
- ❖ RADIOSS-CFD
- ❖ CFX
- ❖ FIDAP
- ❖ STAR-CD
- ❖ PHOENICS
- ❖ FIRE
- ❖ STREAM
- ❖ POLYFLOW
- ❖ SIMPOREA

Electromagnetic Field Analysis

- ❖ PAM-CEM/CEM-3D

Engine Analysis

- ❖ PAM-VEF

Mathematical Library

- ❖ ASL
- ❖ MSL Fortran90
- ❖ NAG Fortran
- ❖ NAG Fortran90
- ❖ MathKeisan
- ❖ FMSlib
- ❖ BCSLIB-EXT

Tool

- ❖ TotalView
- ❖ LSF
- ❖ KUBRIX

Graphics

- ❖ AVS
- ❖ AVS/Express
- ❖ EnSight
- ❖ RVSLIB/Server

Petroleum

- ❖ FOCUS



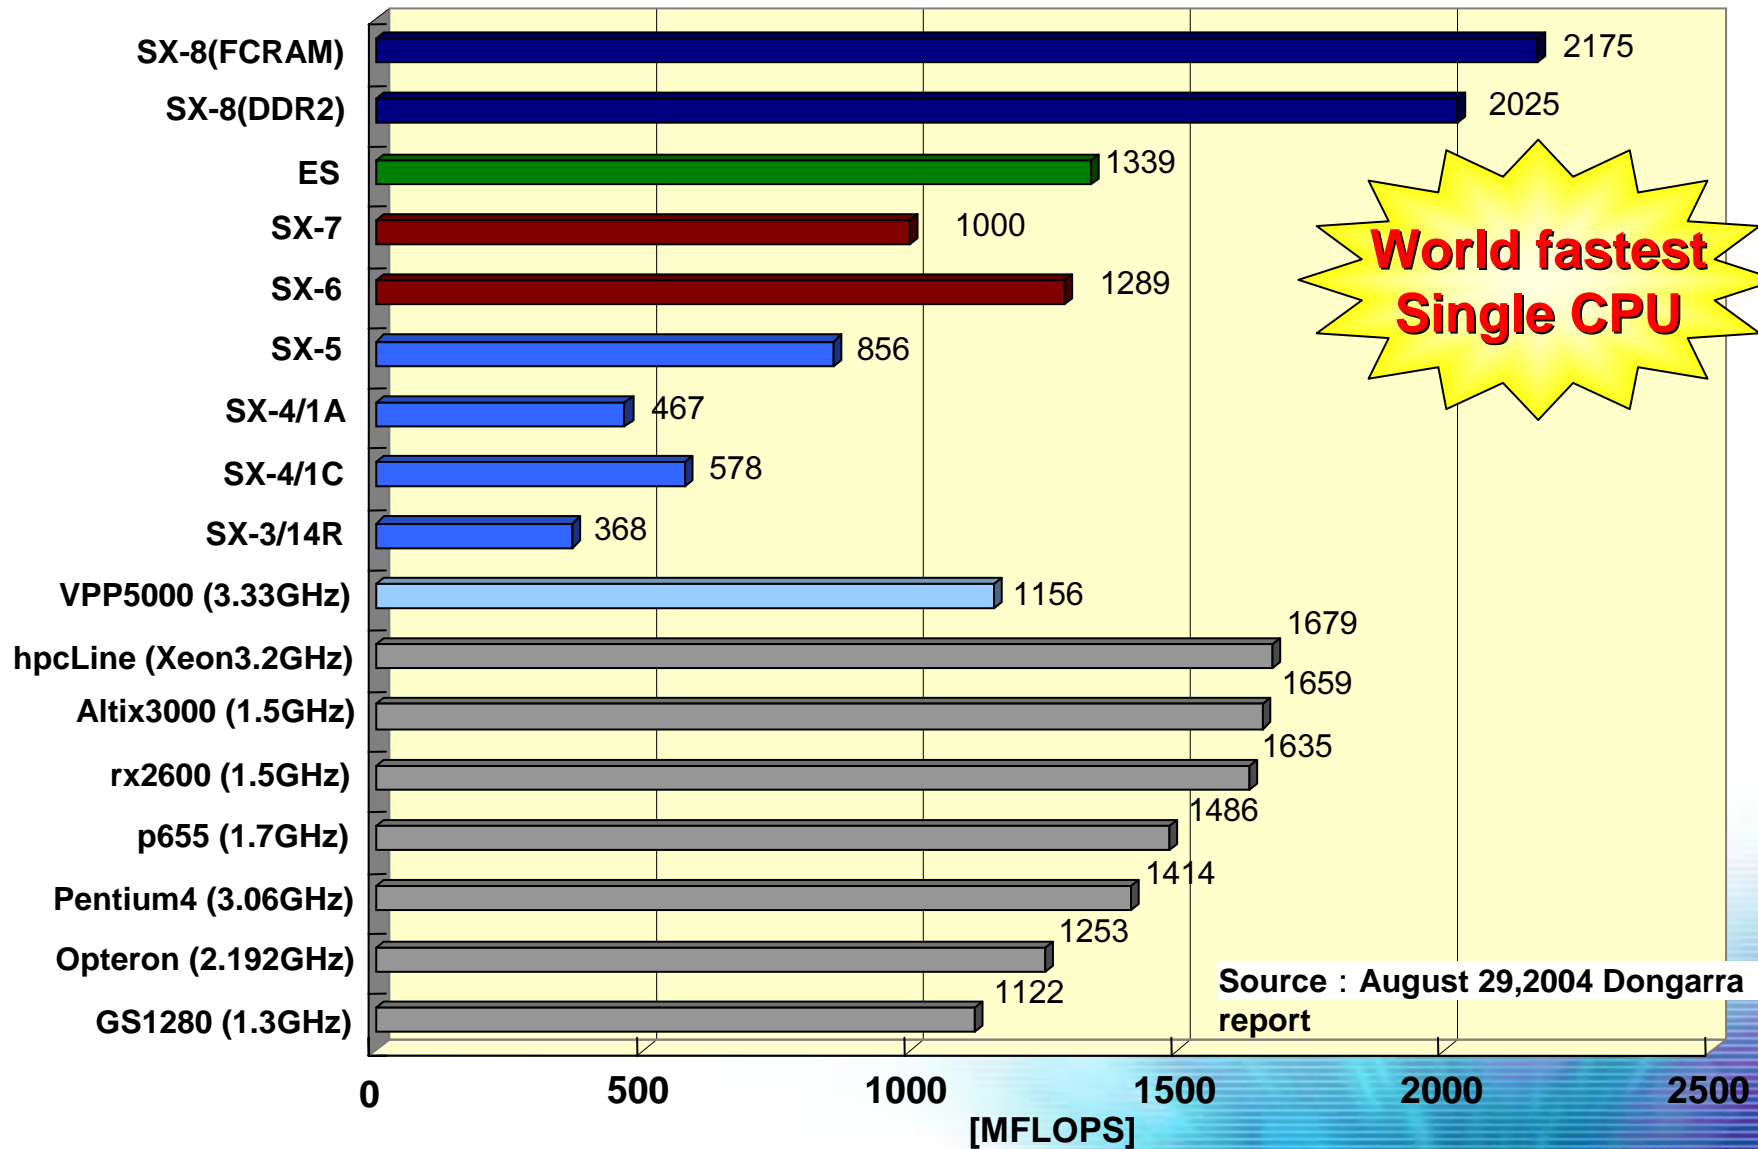
SX-8 Performance

DDR2-SDRAM Model & FCRAM Model

LINPACK 100

SX-8

Slide 26



**World fastest
Single CPU**

Source : August 29,2004 Dongarra report

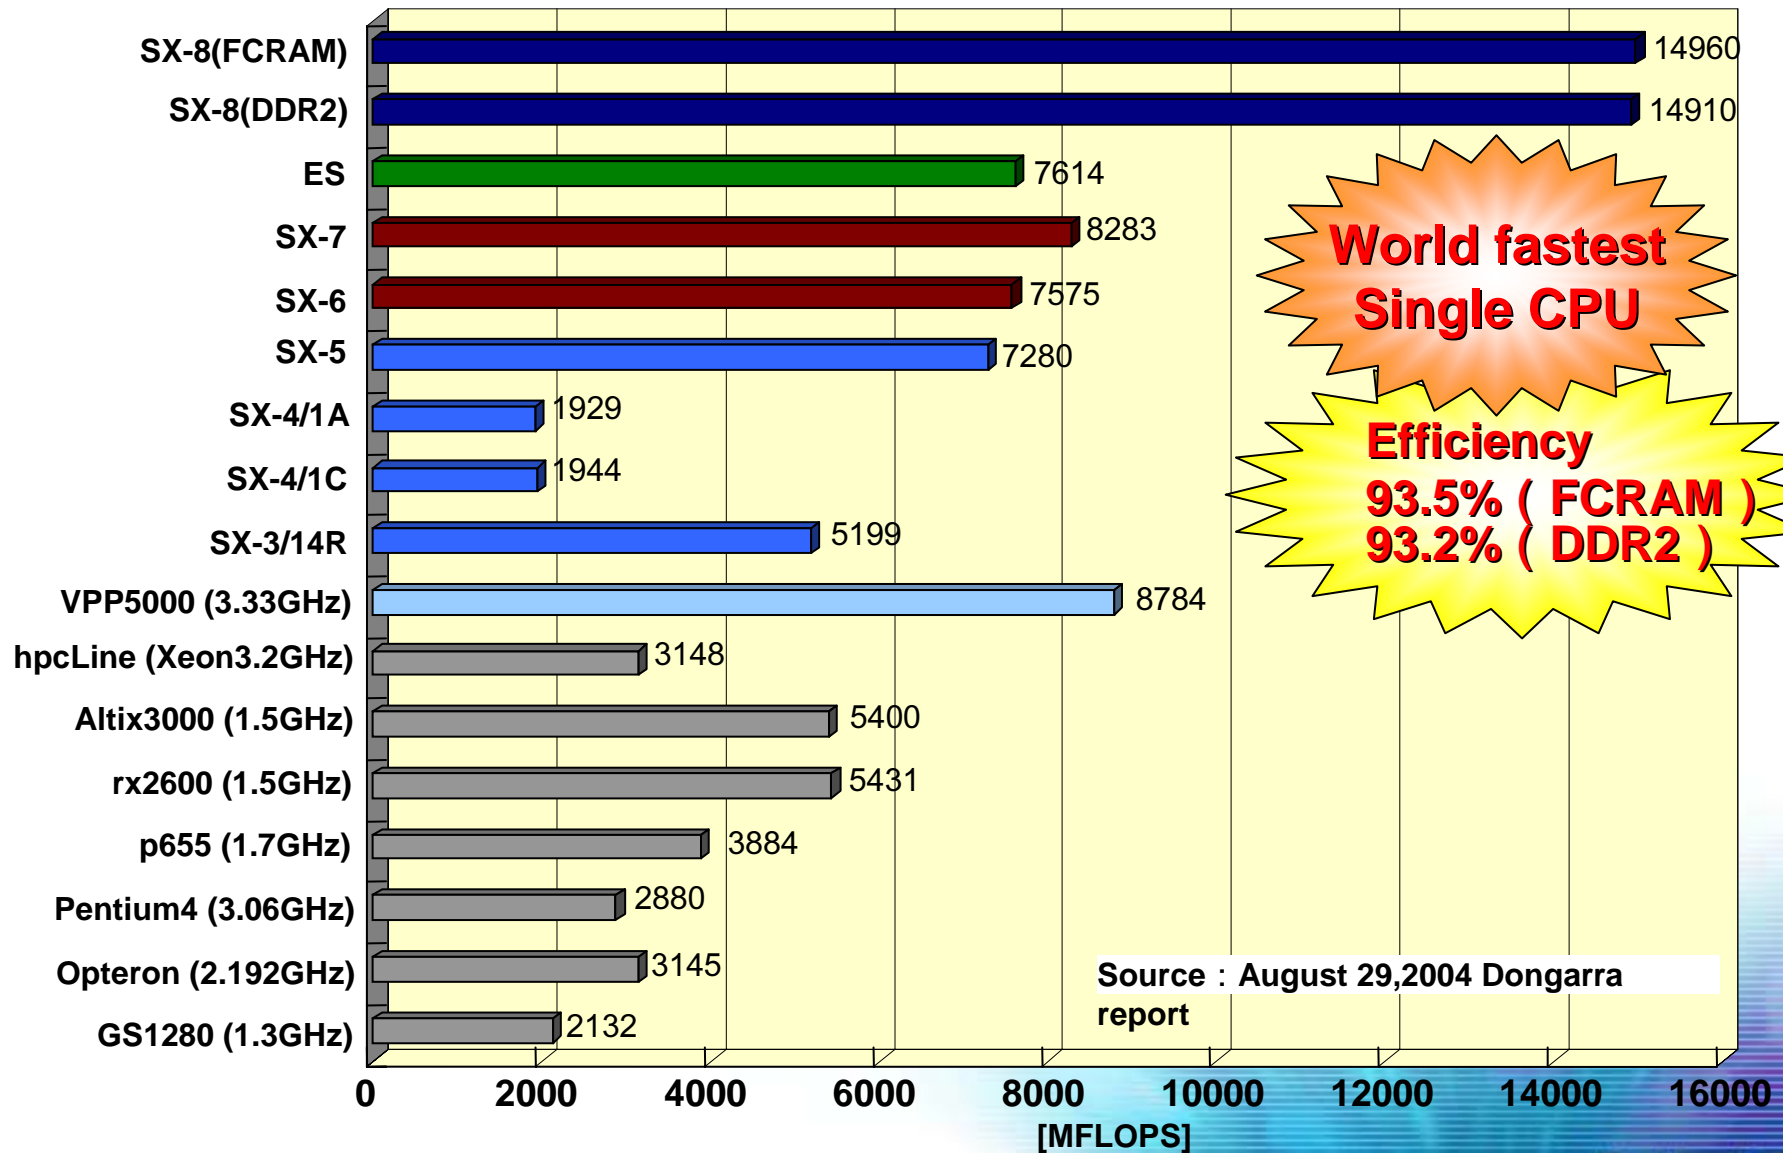
Empowered by Innovation

NEC

LINPACK TPP

SX-8

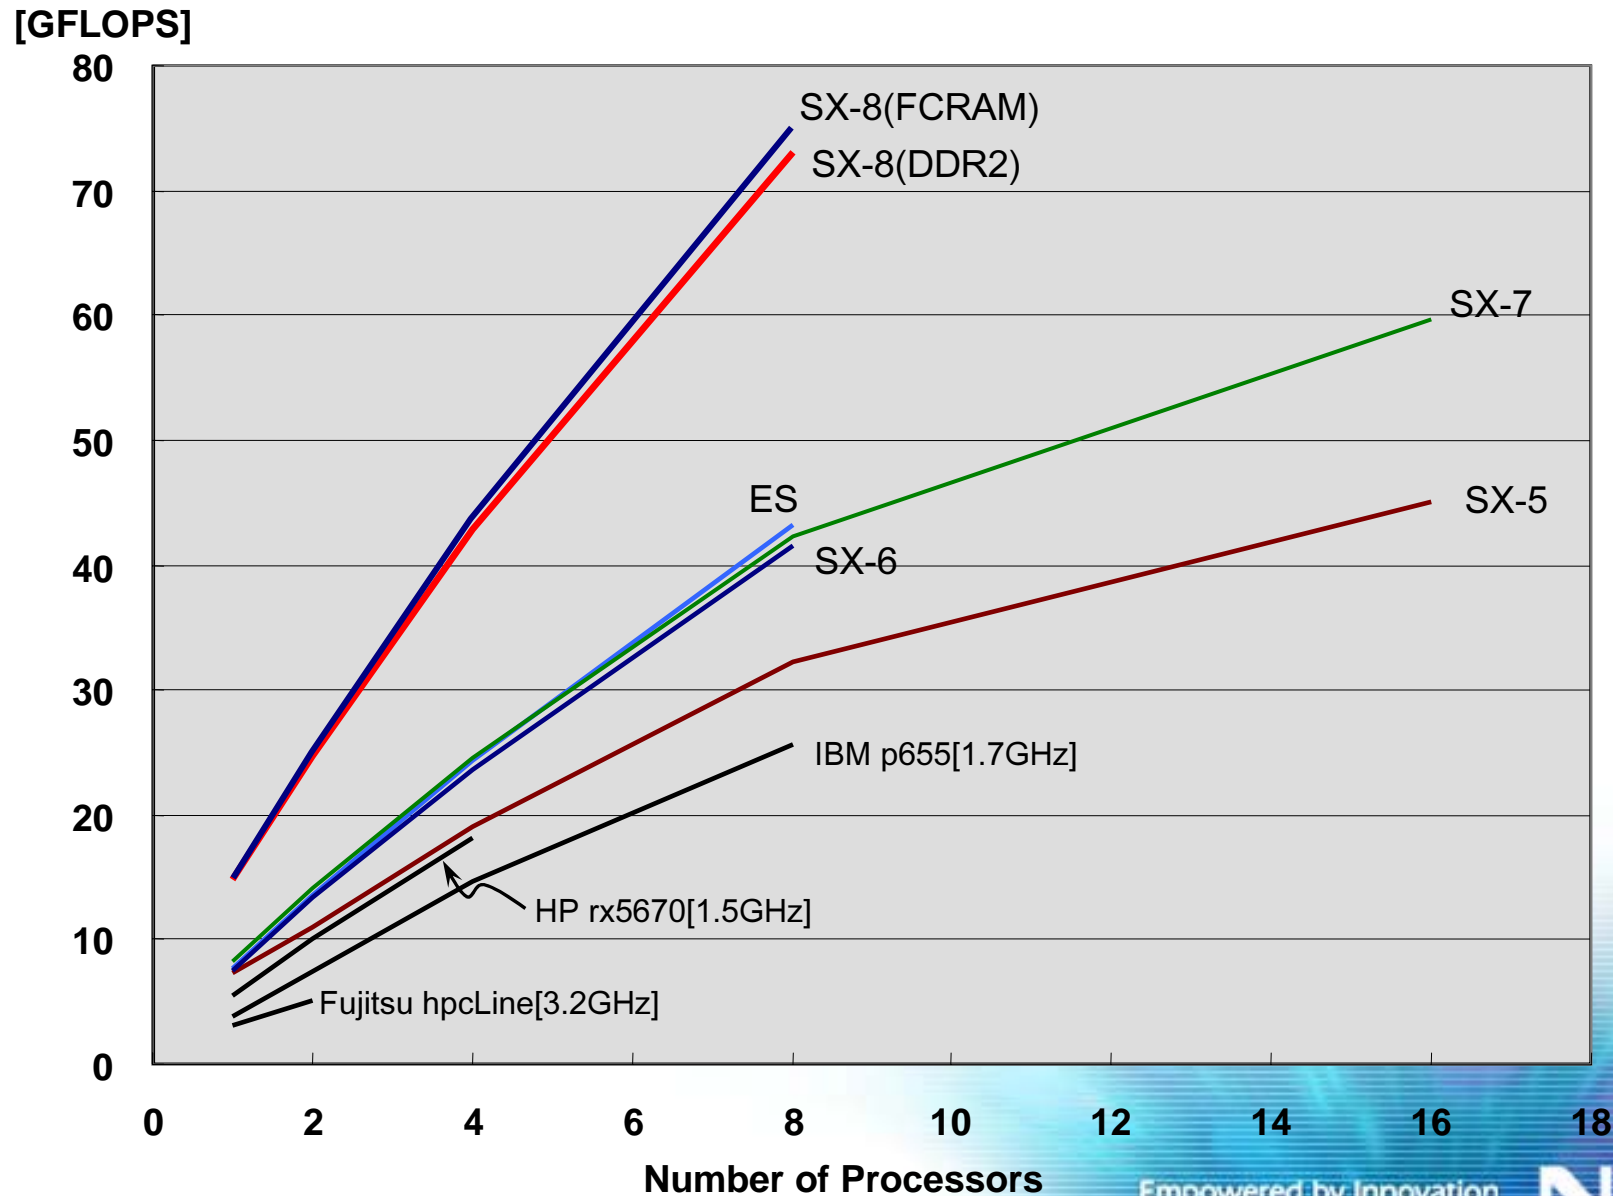
Slide 27



Empowered by Innovation

NEC

LINPACK TPP Scalability

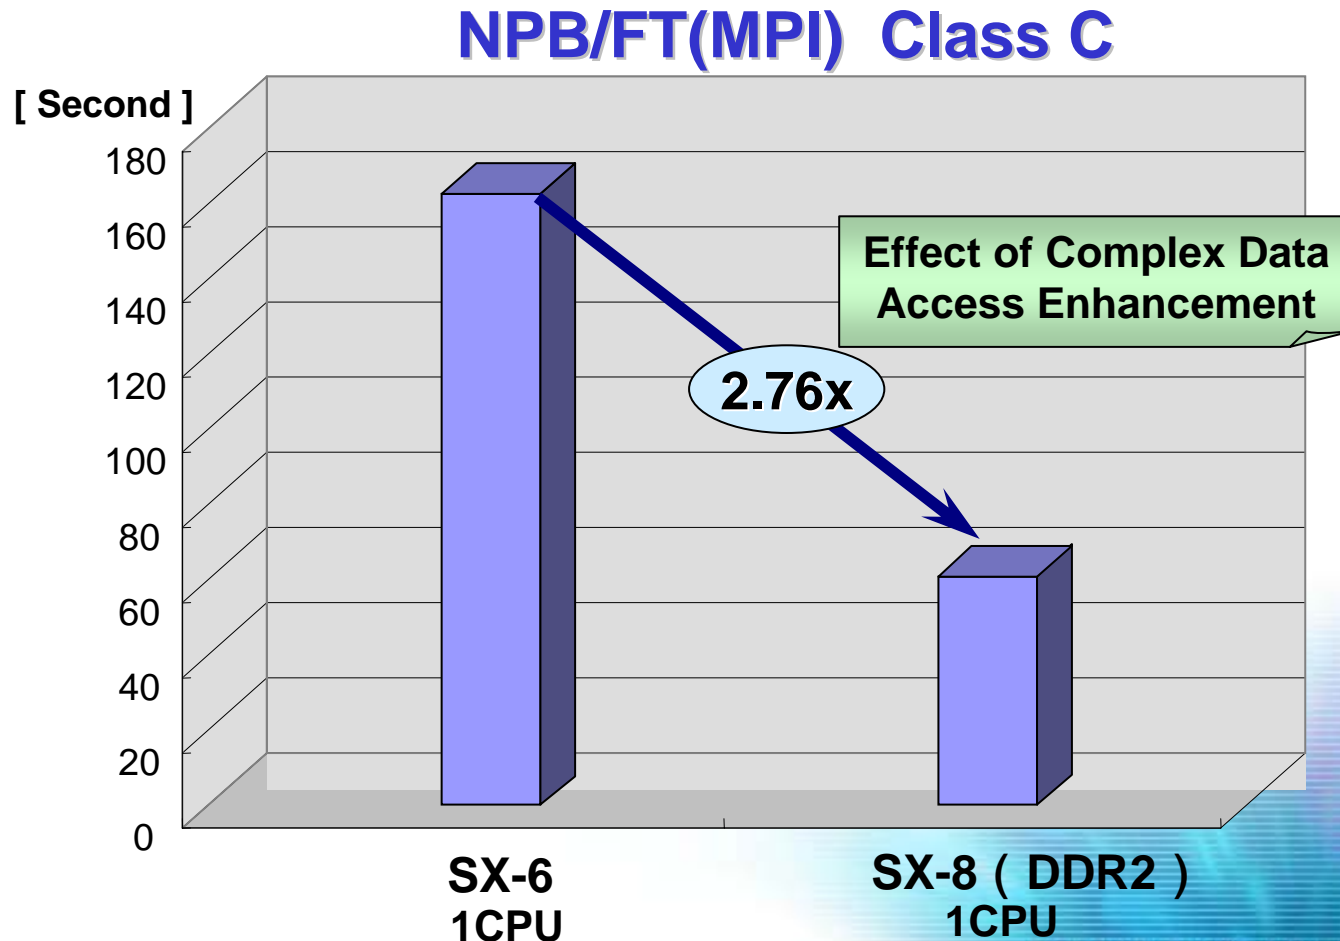


Stride Vector Access

SX-8

Slide 29

- Improved even stride vector access performance (ex. Complex Numbers)

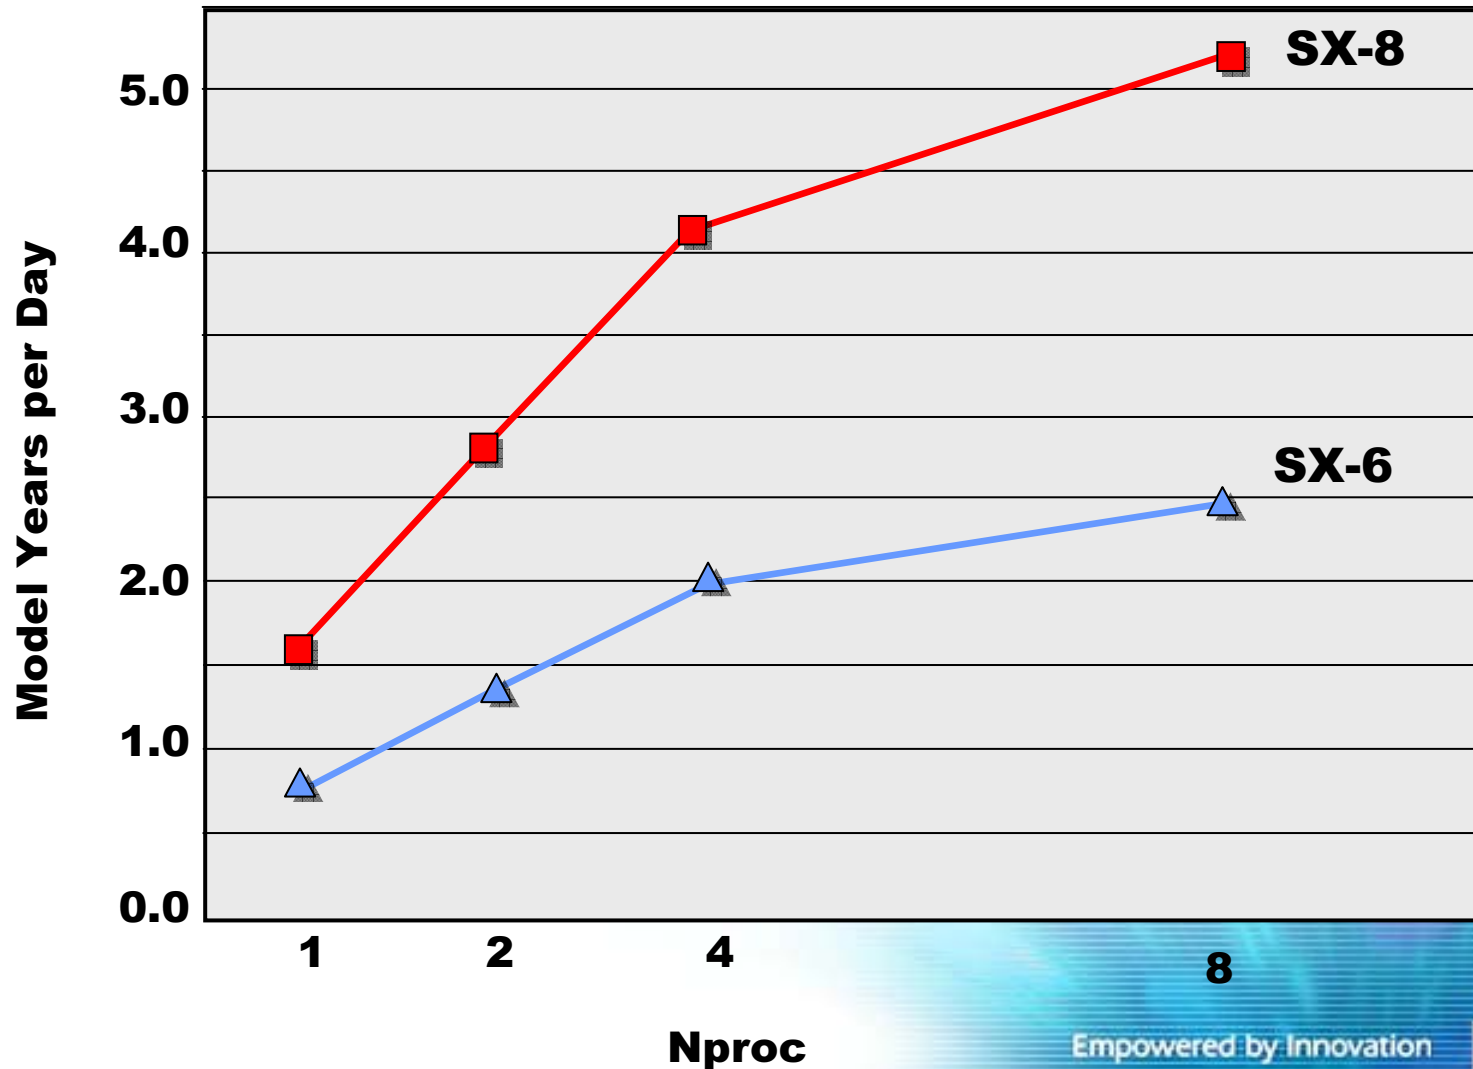


Empowered by Innovation

NEC

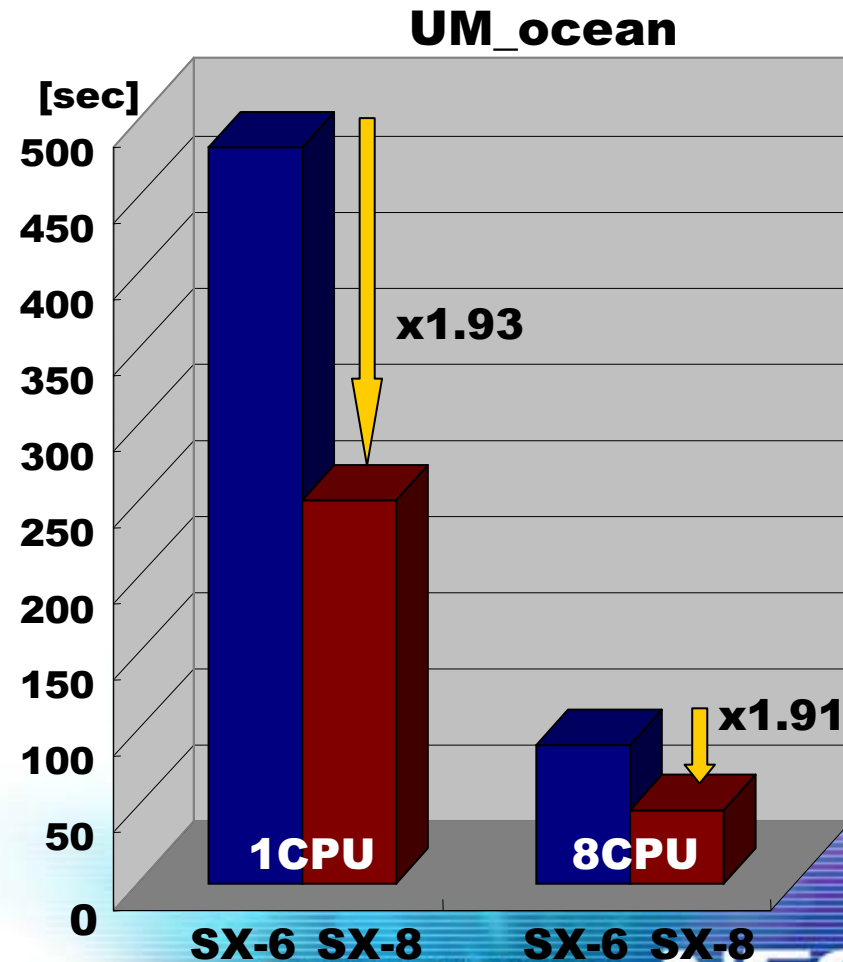
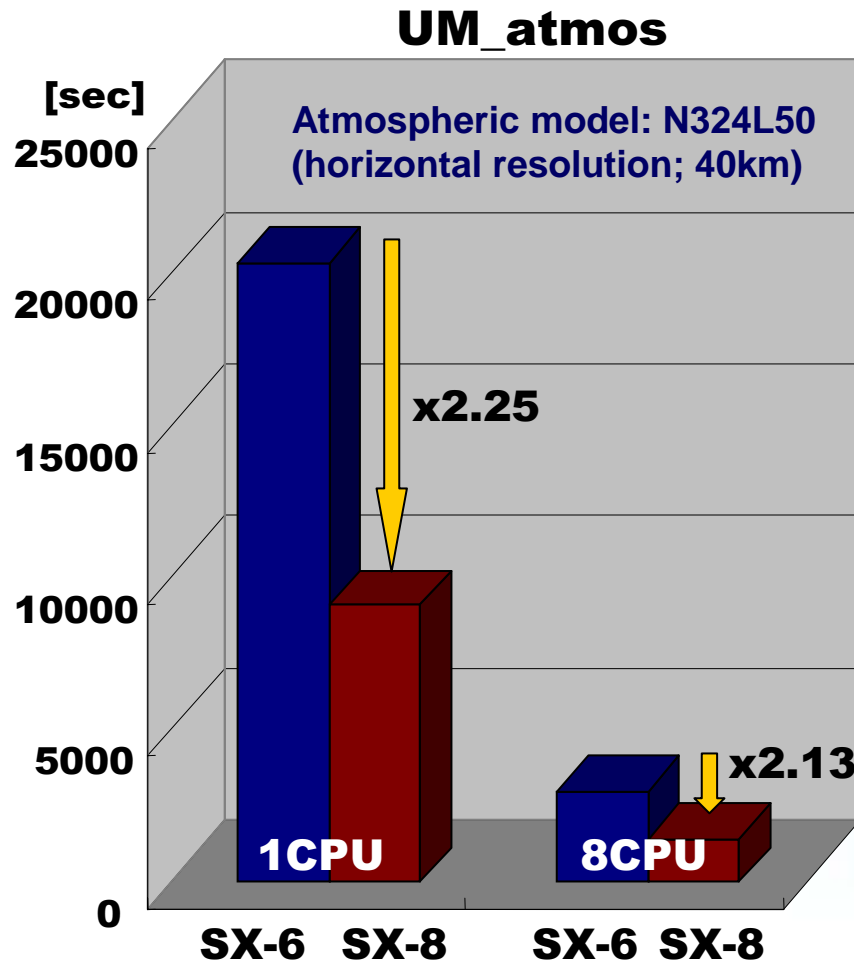
UK Met Unified Model (UM)

Atmospheric model: N48 (horizontal resolution; 270km)



UK Met Unified Model (UM)


Measured on SX-6 and SX-8(FCRAM)
- More than 1.9 times higher on SX-8 over SX-6



NEC will continue to supply powerful vector supercomputers to our worldwide customers in meteorology and climate fields.



Thank You Very Much!





Empowered by Innovation

NEC

Empowered by Innovation **NEC**