SX-Aurora TSUBASA

June 25, 2018

Shigeyuki Aino
NEC Corporation

History of Vector computing

NEC has always provided high sustained performance by vector supercomputer SX series

Vector Engine (PCI card)

Performance

Hardware innovations

Earth Simulator 3

Earth Simulator 2

Earth Simulator

SX-ACE

SX-9

SX-8

SX-7

SX-6

SX-5

SX-4

SX-3

SX-2

Software innovations

1990

2000

2010

Packed vector technologies accumulated over 30 years into PCI card

Project Aurora

Vector Accelerator Card

✓ NEC's 30 years vector technology is packed into Vector Engine card

Project Aurora
(x86 server + Vector Engine)

PCIe

Vector Engine
(PCIe card)

- Compact and flexible
- Hybrid architecture (standard x86 + Vector)
- Economically deliver supercomputer technology

New Architecture

- **SX-Aurora TSUBASA = Standard x86 + Vector Engine**
- **Linux + standard language (Fortran/C/C++)**
- **Enjoy high performance with easy programming**

SX-Aurora TSUBASA Architecture

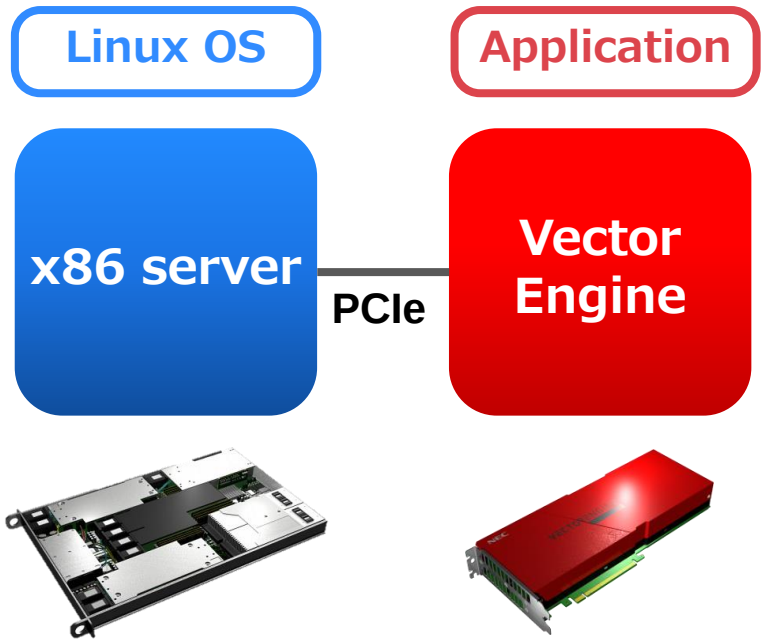

Hardware

- Standard x86 server + Vector Engine

Software

- Linux OS
- Automatic vectorization compiler
- Fortran/C/C++
  → No special programming like CUDA

Interconnect

- InfiniBand for MPI
- VE-VE direct communication support

Easy programming (standard language) → **Automatic vectorization compiler** → **Enjoy high performance**

New Architecture

Hybrid architecture combining Vector Processor with x86 Processor
1. SX-Aurora = x86 server + Vector Engine (VE)
2. VE capability is provided on x86/Linux environment
3. Infiniband Interconnect support

SX-Aurora Architecture

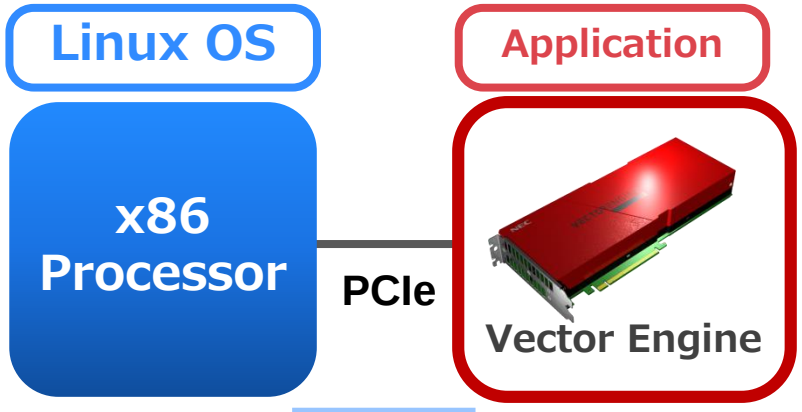

Enables the flexibility

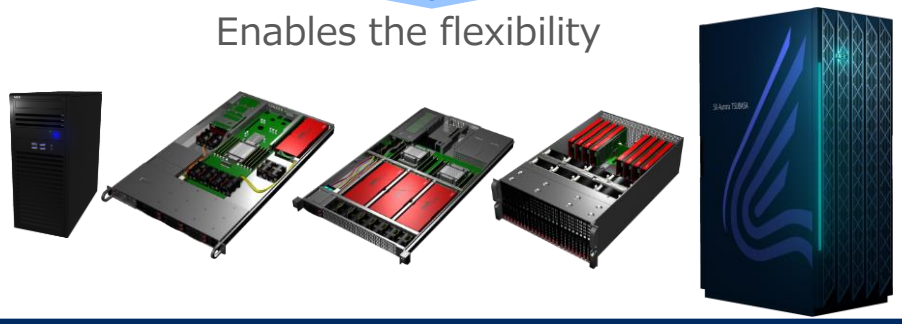

Hardware

- x86 server + VE

Software Environment

- x86 / Linux OS
- Fortran/C/C++ standard programming
- Automatic vectorization by proven vector compiler

NO special programming like CUDA is necessary!

Interconnect

- InfiniBand for MPI

New Values

NEC's Vector technology can invent new Social Values
- as the key to accelerate HPC + AI/Big Data Analytics

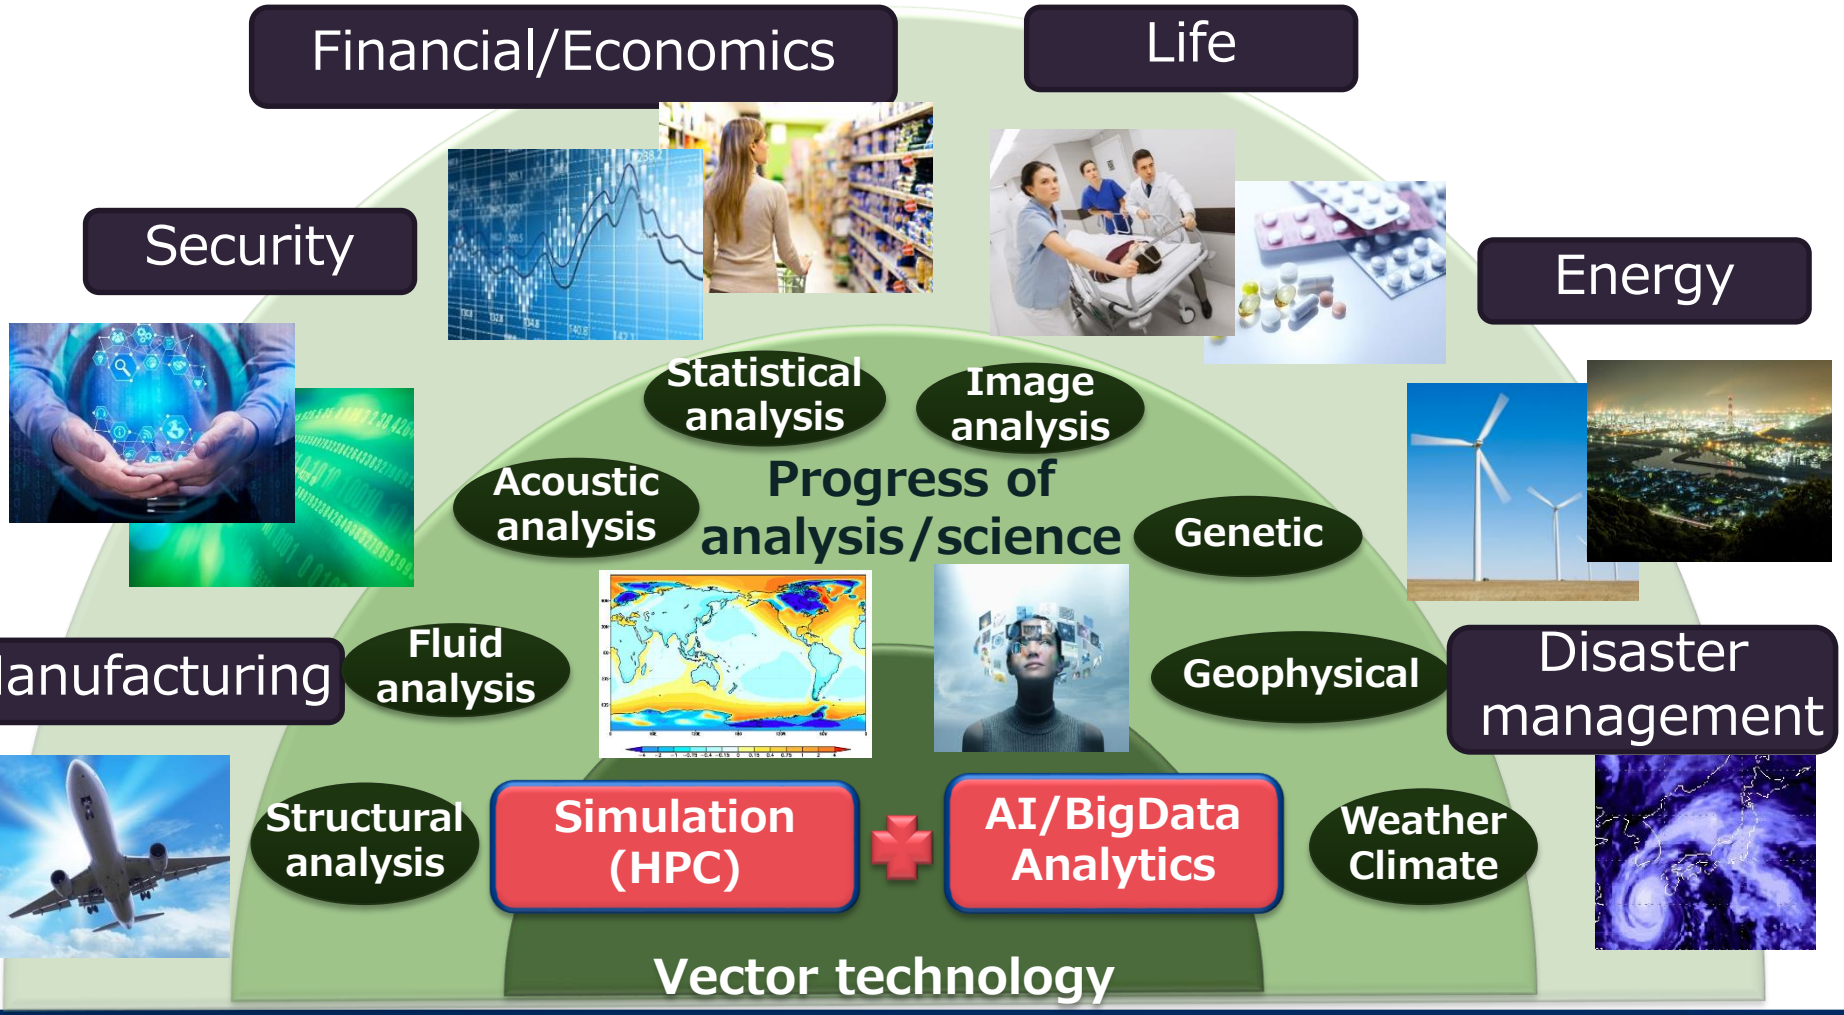

Initial BM results : HPL and STREAM

- **Aurora provides same range HPL performance as SKL**
- **Aurora provides highest memory bandwidth**

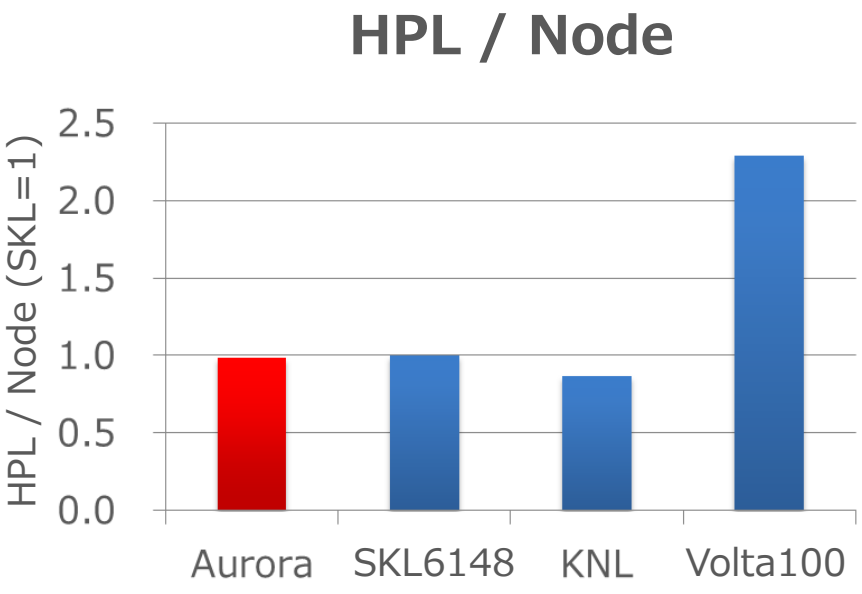

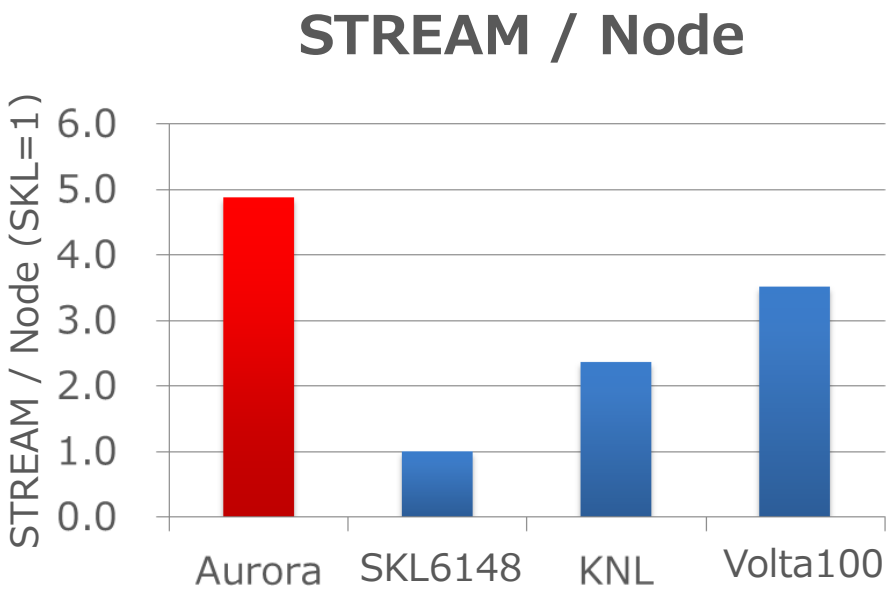

- Aurora is Vector Engine Type 10-B (1.4GHz, 8core)
- SKL is Intel Skylake 6148 Xeon x2/node
- KNL is Intel Knight Landing x1/node
- V100 is NVIDIA Tesla V100 x1/node

Initial BM results: HPCG

Performance/power of Aurora shows 7 times better than SKL

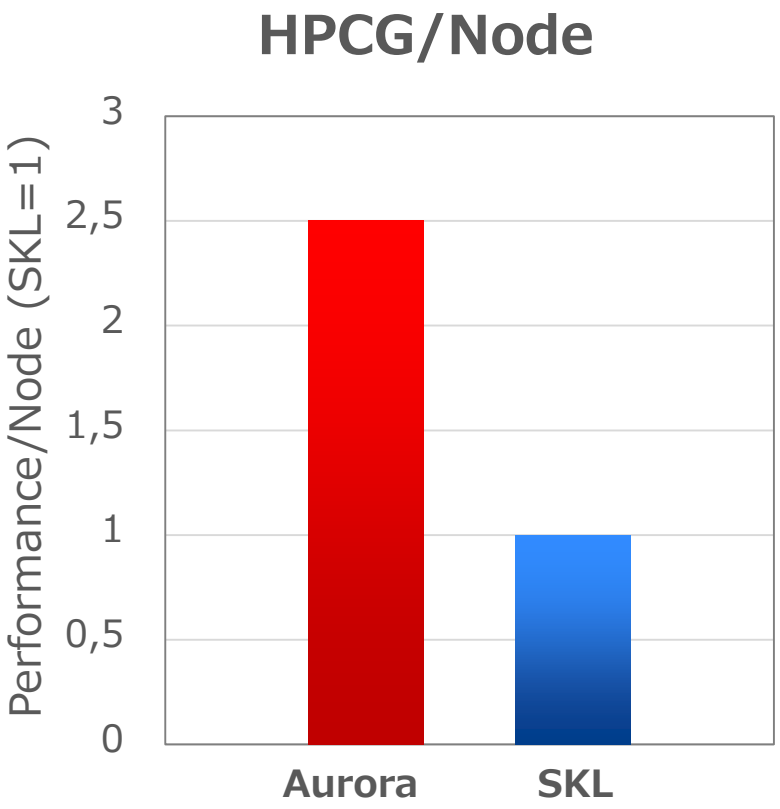

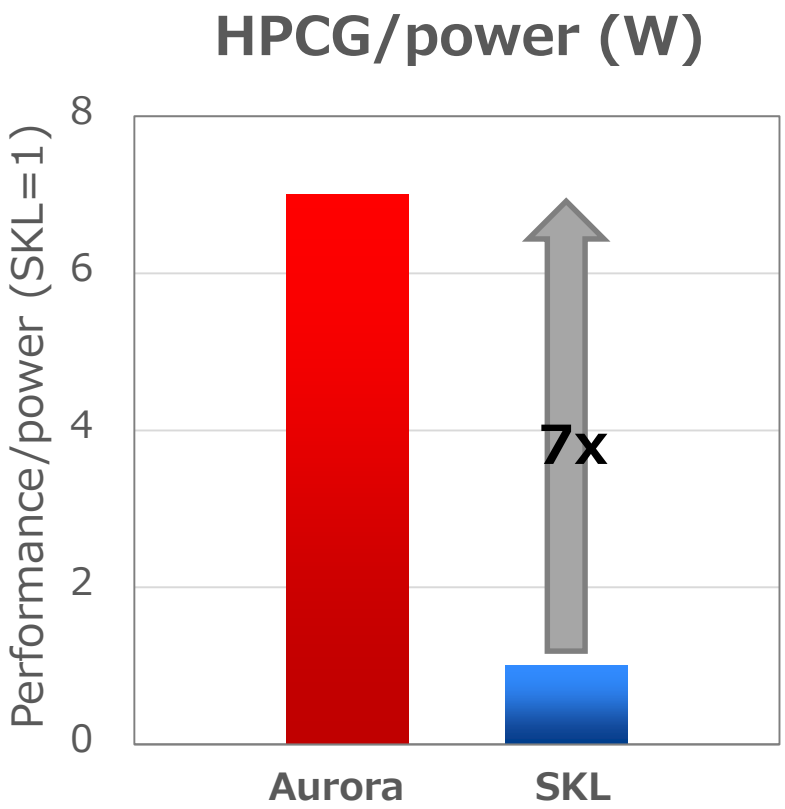

- Aurora is Vector Engine Type 10-B (1.4GHz, 8core)
- SKL is Intel Skylake 6148 Xeon x2/node

Performance of NEC middleware for ML

Frovedis + VE shows over 100x performance compared to Spark + x86

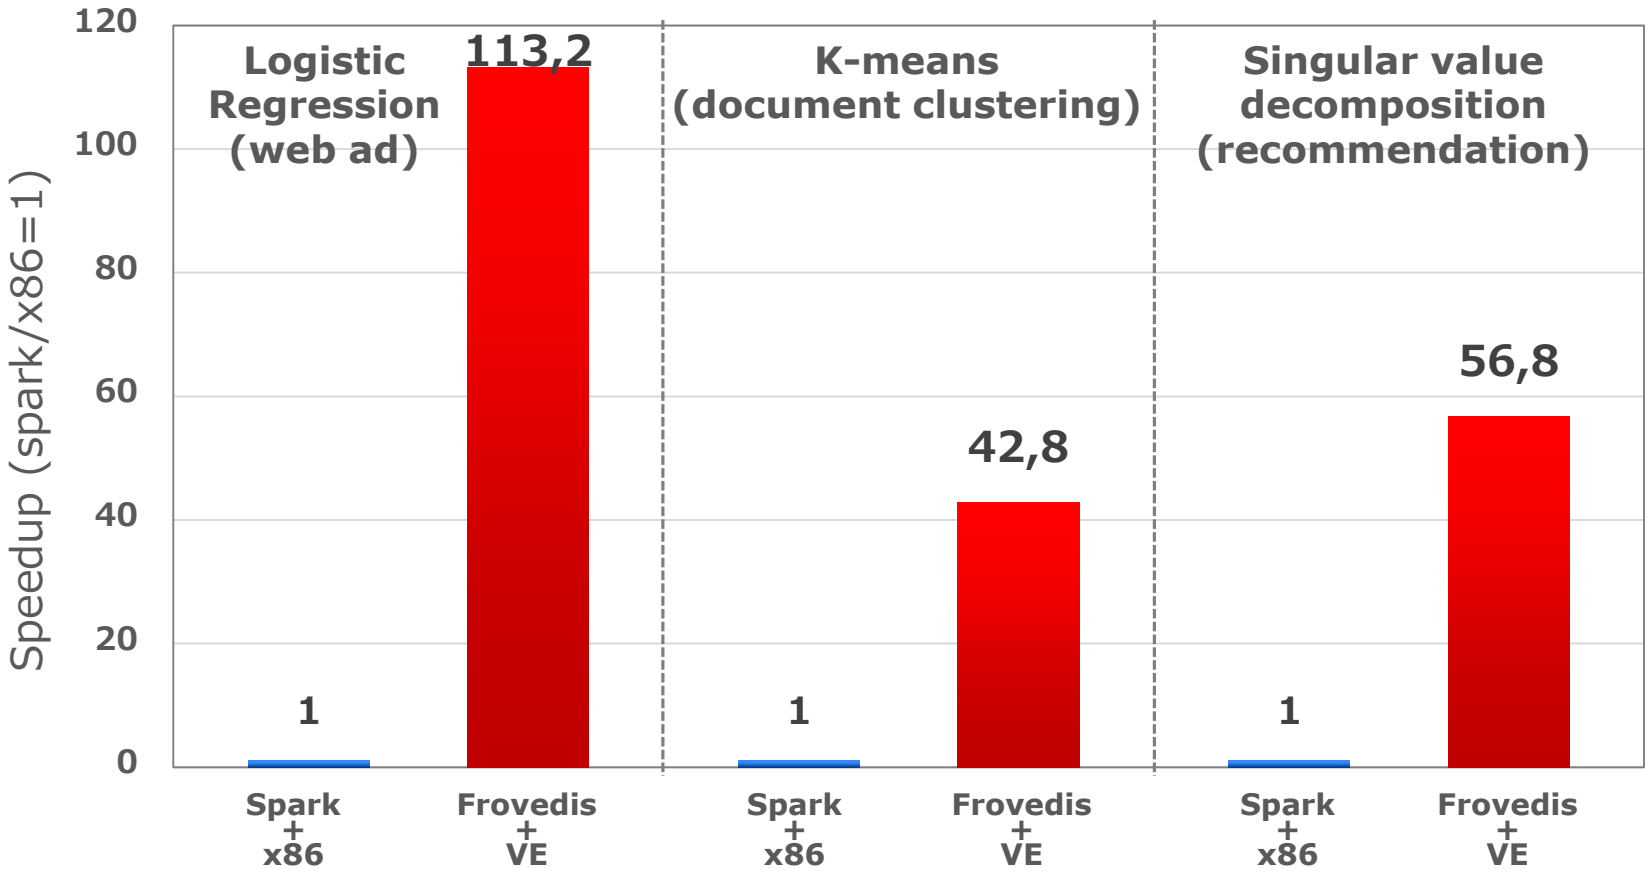

- x86: Intel Xeon Gold 6126 x1 socket
- Aurora: Vector Engine Type 10-B (1.4GHz, 8core) x1
- Performance comparison does not include I/O time

Orchestrating a brighter world

NEC