Orchestrating a brighter world NEC

Introduction of SX-Aurora TSUBASA for ISC21

-PCIe card type vector computer-

June 30th, 2021

NEC Corporation

Agenda

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Announcement of two new functions and card business
4. Value of Vector Engine
5. Roadmap

History of SX Vector Supercomputer

SX-2 1983

Technology:	Bipolar
CPU Frequency:	166 MHz
CPU Performance:	1.3 GFlops
CPU Memory Bandwidth:	10.7 GB/sec

SX-3 1989

Technology:	Bipolar
CPU Frequency:	340 MHz
CPU Performance:	5.5 GFlops
CPU Memory Bandwidth:	12.8 GB/sec

SX-4 1994

Technology:	350 nm
CPU Frequency:	125 MHz
CPU Performance:	2.0 GFlops
CPU Memory Bandwidth:	16.0 GB/sec

SX-5 1998

Technology:	250 nm
CPU Frequency:	250 MHz
CPU Performance:	8.0 GFlops
CPU Memory Bandwidth:	64.0 GB/sec

SX-6 2001

Technology:	150 nm
CPU Frequency:	500 MHz
CPU Performance:	8.0 GFlops
CPU Memory Bandwidth:	32.0 GB/sec

SX-7 2002

Technology:	150 nm
CPU Frequency:	552 MHz
CPU Performance:	8.8 GFlops
CPU Memory Bandwidth:	35.3 GB/sec

SX-8 2004

Technology:	90 nm
CPU Frequency:	1.0 GHz
CPU Performance:	16.0 GFlops
CPU Memory Bandwidth:	64.0 GB/sec

SX-9 2007

Technology:	65 nm
CPU Frequency:	3.2 GHz
CPU Performance:	102.4 GFlops
CPU Memory Bandwidth:	256.0 GB/sec

SX-ACE® 2013

Technology:	28 nm
CPU Frequency:	1.0 GHz
CPU Performance:	256.0 GFlops
CPU Memory Bandwidth:	256.0 GB/sec

Over 35 years experience for High Sustained Performance

Features of SX-Aurora TSUBASA

SX-ACE

（Traditional Super Computer）

Downsize!

SX-Aurora TSUBASA

Vector Engine

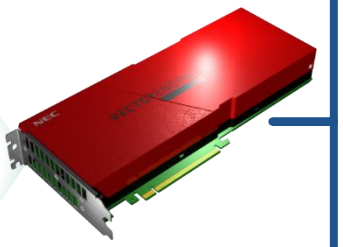

Tower Type

Rack Type

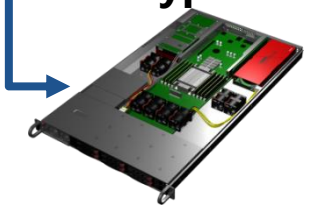

Downsizing of super computer realized by NEC's Technology.

POINT 1 High Memory Bandwidth

Vector technology makes it possible to process multiple and huge data at a time with high memory bandwidth.

POINT 2 Ease of Use

No specialized knowledge is required, AP can be executed only after compiled.
Use C/C++/Fortran to program.

POINT 3 Flexibility

Customer can choose a system which meets their needs. From server type to card specification are all optional,
NEC help customer to maximize the cost performance, to fit all market requirement.

Vector Engine

POINT 1 High Memory Bandwidth

◆ Vector technology is packed into a PCI card.

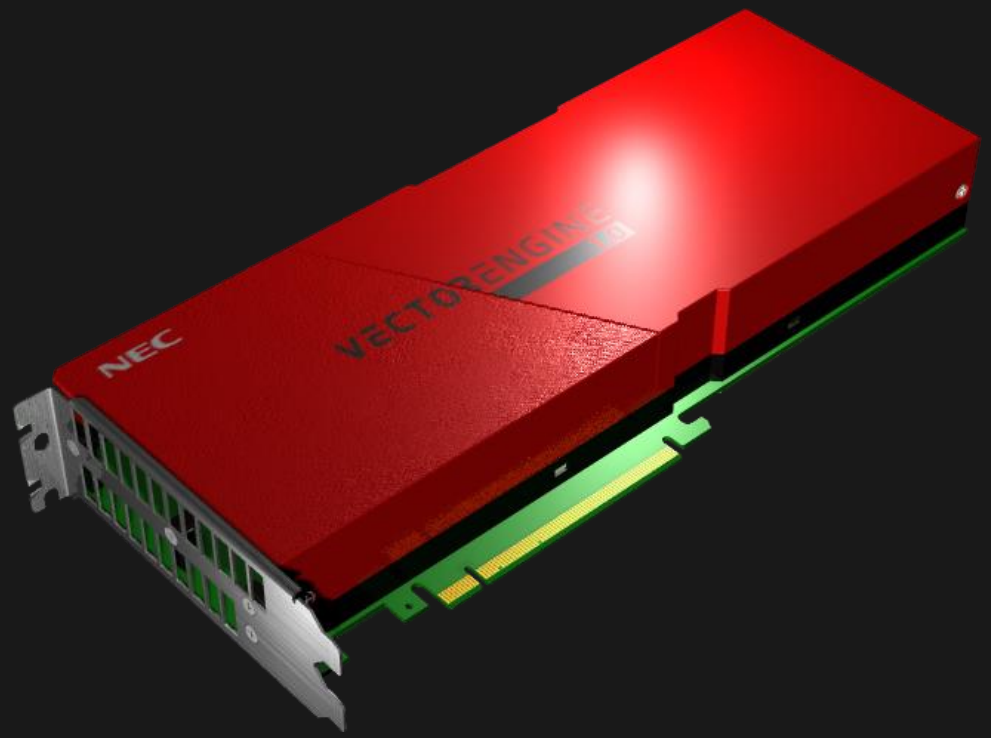

- Vector processor (8/10 cores)
- 1.53TB/s memory bandwidth
- 48GB memory
- 2.45-3.07TF performance (double precision)
  4.91-6.14TF performance (single precision)
- A variety of execution modes
- Standard programming with Fortran/C/C
- Power consumption < 300 W

Architecture of SX-Aurora TSUBASA

POINT 2 Ease of Use

- SX-Aurora TSUBASA = VH + VE
- Linux + standard language (Fortran/C/C++)
- Enjoy high performance with easy programming

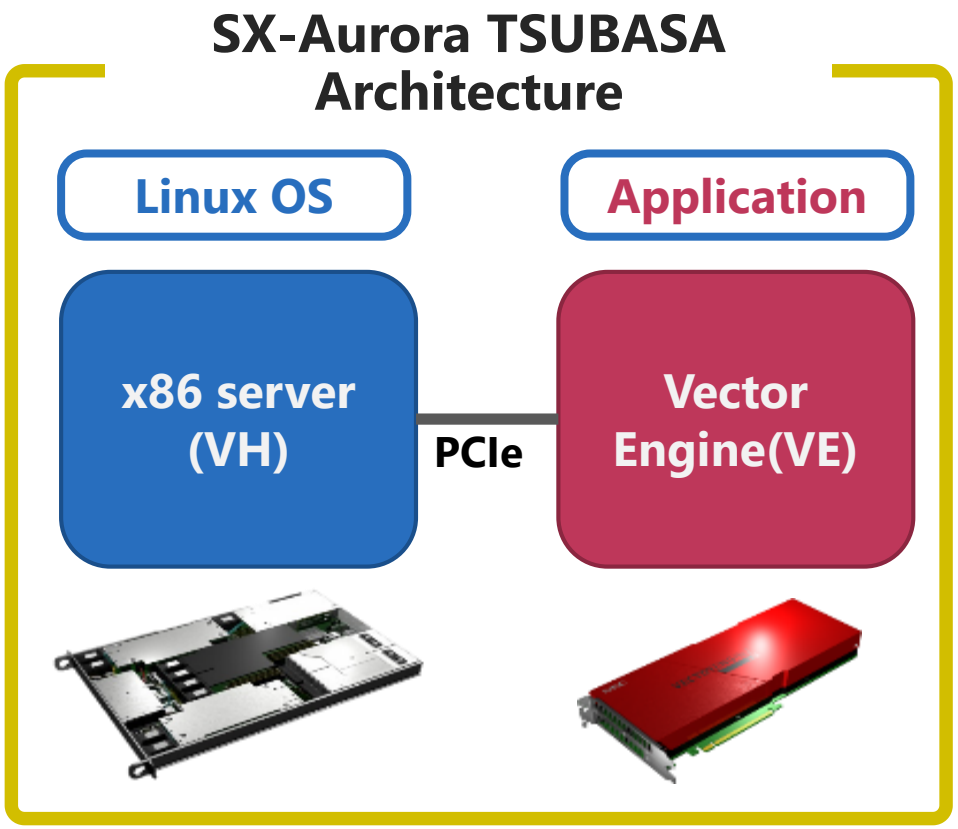

Hardware

- VH(Standard x86 server) + Vector Engine

Software

- Linux OS
- Fortran/C/C++ → Standard language
- Automatic vectorization compiler

Interconnect

- InfiniBand for MPI
- VE-VE direct communication support

Easy programming (standard language) ▶ **Automatic vectorization compiler** ▶ **Enjoy high Performance!**

Lineup of SX-Aurora TSUBASA

POINT 3 Flexibility

Vector Engine supports wide range from desk-side to large-scale Data Centers.
Selling Vector Engine card was started from November, 2020.

Data Center Model
Huge processing in Data Centers

Rackmount Model
Simulation of manufacturing industry, etc.
Use of AI / big data

Edge Model
Simulation of mid-sized manufacturing industry, etc. Laboratory desk-side

Vector Engine
Application acceleration
Embedded use
Cloud service engine

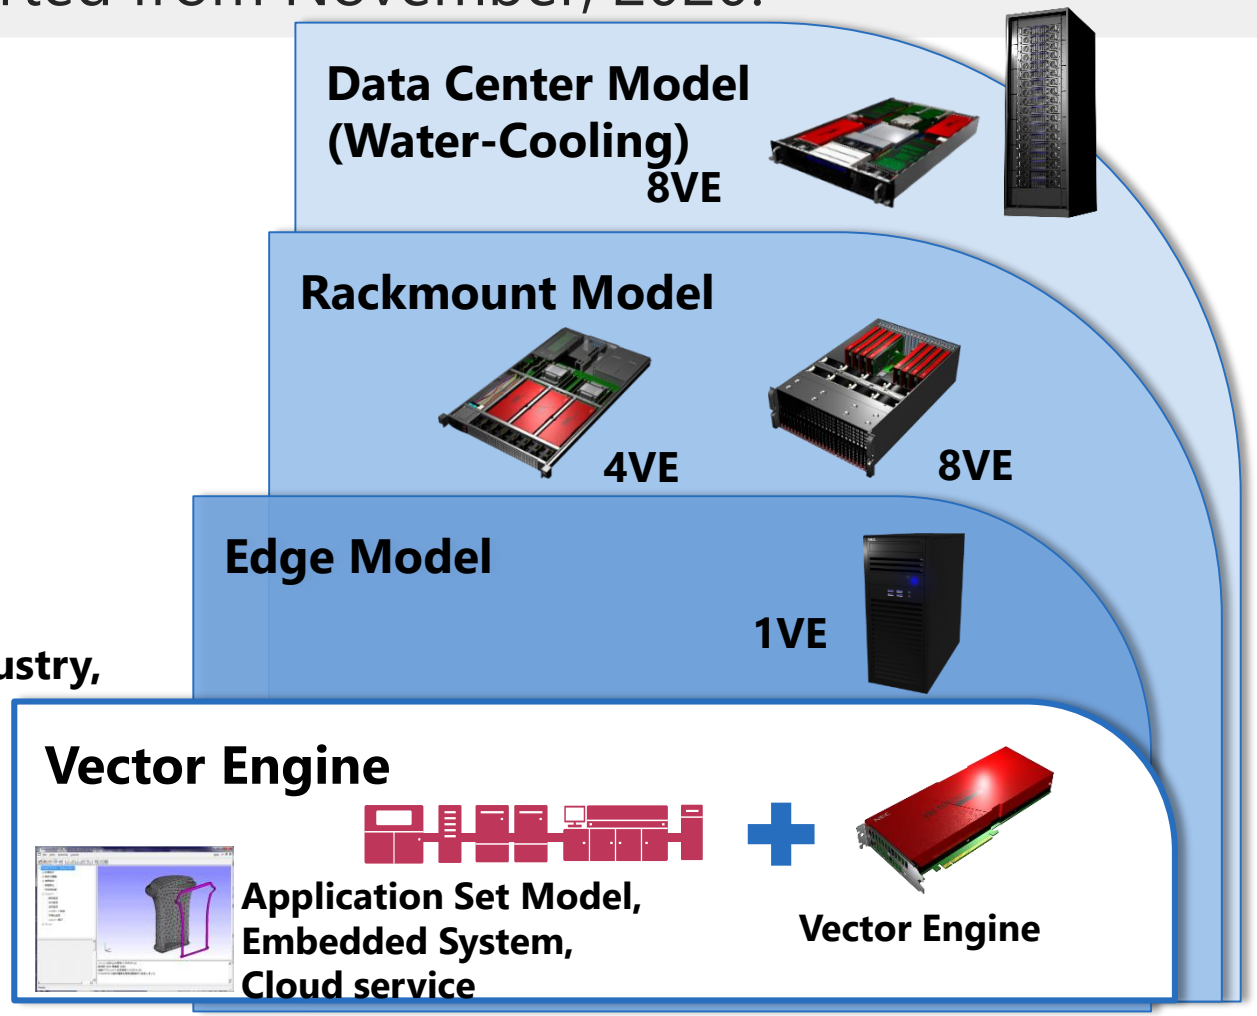

Trusted and Chose by World Famous HPC Centers

Český hydrometeorologický ústav : Weather Climate

Deutscher Wetterdienst : Weather / Climate

NIFS : Fusion Science

National Institute for Fusion Science

Tohoku university : Academic

JAMSTEC : Earth Science

Osaka university : Academic

Performance of large-scale computer system

JAMSTEC Earth Simulator is ranked in TOP10 in the latest HPCG ranking.
High Byte/Flops and high performance single core => High execution efficiency

JAMSTEC

NIFS
National Institute for Fusion Science

Rank				Cores	HPCG [TFlop/s]	Rpeak [TFlop/s]	Execution efficiency
HPCG	HPL	System	Vendor				
1	1	Fugaku	Fujitsu	7,630,848	16,004.50	537,212.00	2.98%
2	2	Summit	IBM	2,414,592	2,925.75	200,794.88	1.46%
3	5	Perlmutter	HPE	706,304	1,905.44	89,794.48	2.12%
4	3	Sierra	IBM / NVIDIA / Mellanox	1,572,480	1,795.67	125,712.00	1.43%
5	6	Selene	Nvidia	555,520	1,622.51	79,215.00	2.05%
6	8	JUWELS Booster Module	Atos	449,280	1,275.36	70,980.00	1.80%
7	11	Dammam-7	HPE	672,520	881.40	55,423.56	1.59%
8	9	HPC5	Dell EMC	669,760	860.32	51,720.76	1.66%
9	13	Wisteria/BDEC-01	Fujitsu	368,640	817.58	25,952.26	3.15%
10	40	Earth Simulator -SX-Aurora TSUBASA	NEC	43,776	747.80	13,447.99	**5.56%**
11	25	TOKI-SORA	Fujitsu	276,480	614.22	19,464.20	3.16%
12	16	Trinity	Cray/HPE	979,072	546.12	41,461.15	1.32%
13	55	Plasma Simulator	NEC	34,560	529.16	10,510.66	**5.03**%
14	14	Marconi-100	IBM	347,776	498.43	29,354.00	1.70%
15	15	Piz Daint	Cray/HPE	387,872	496.98	27,154.30	1.83%

NEC Network Queuing System V(NQSV) supports Cloud Bursting

Job can be deployed on-premise and burst to the cloud computing system automatically on NQSV infrastructure when the demand for computing power spikes.

- Use the cloud resources on the same UI as on-premise as jobs are submitted to Cloud system on bursting policy
- Reduce the cost for cloud usage by allocating computing resource only as needed
- Select Cloud bursting "yes" or "no" per a job on the user side
- Enable/disable Cloud Bursting function at any time by system administrator

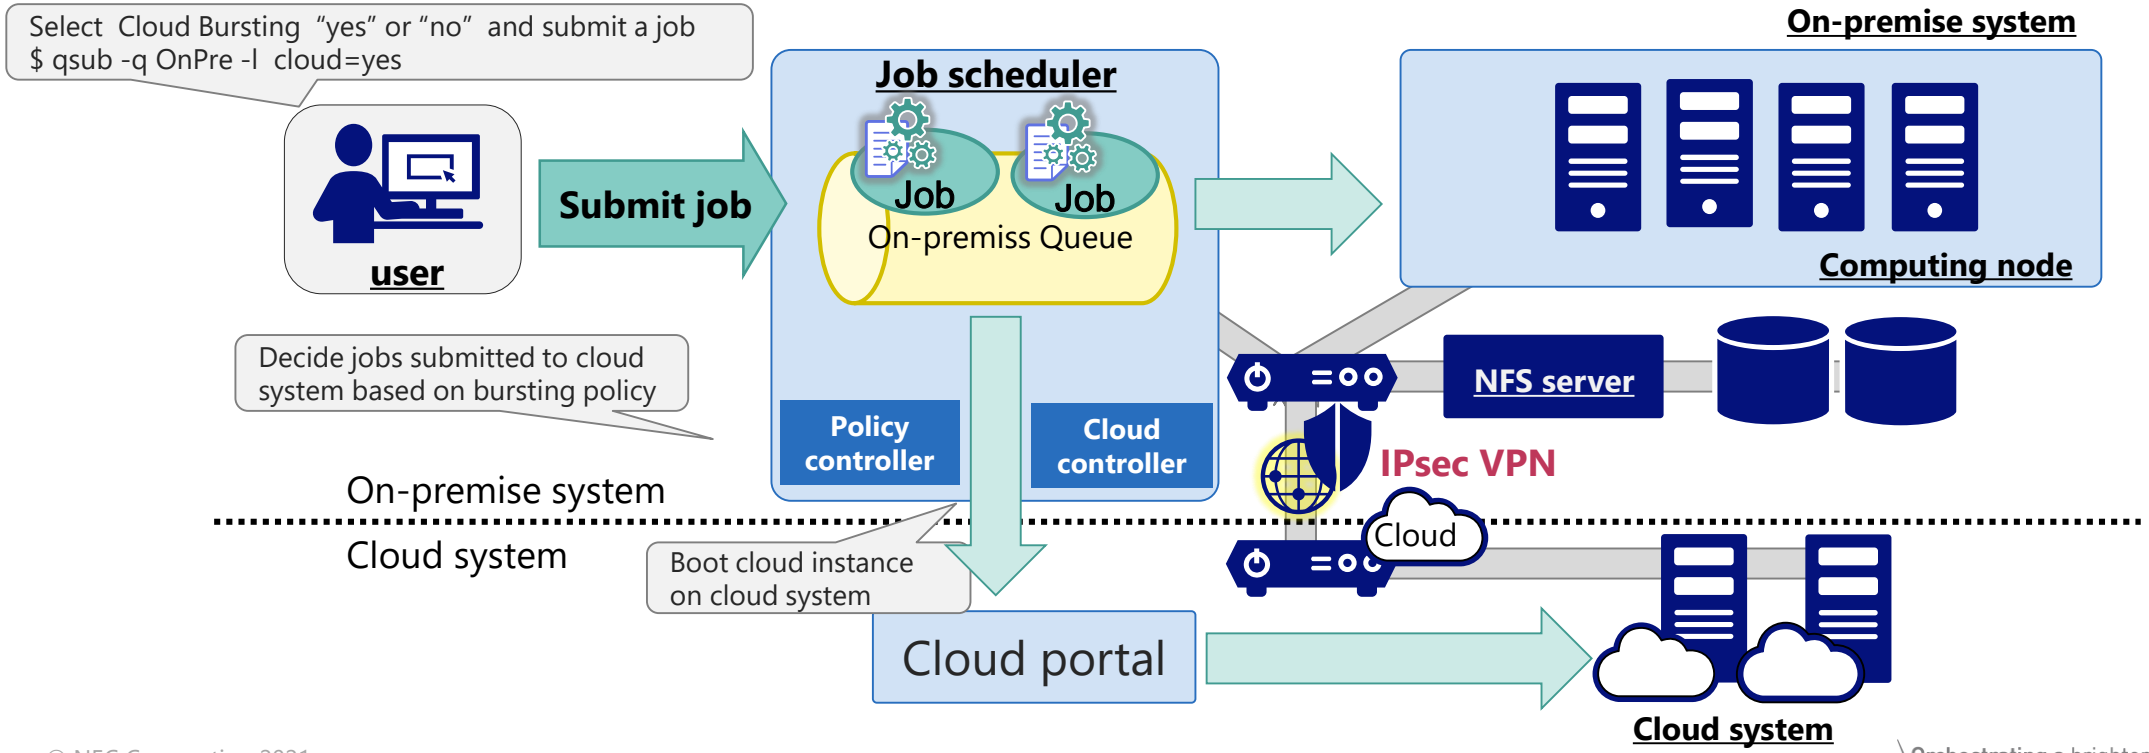

NEC LLVM-IR Vectorizer released in June, 2021

https://www.hpc.nec/forums/topic?id=pA1cPw

Add automatic vectorization feature for VE into clang/flang*[1]/your compilers and create assembler source file including vectorized loops.

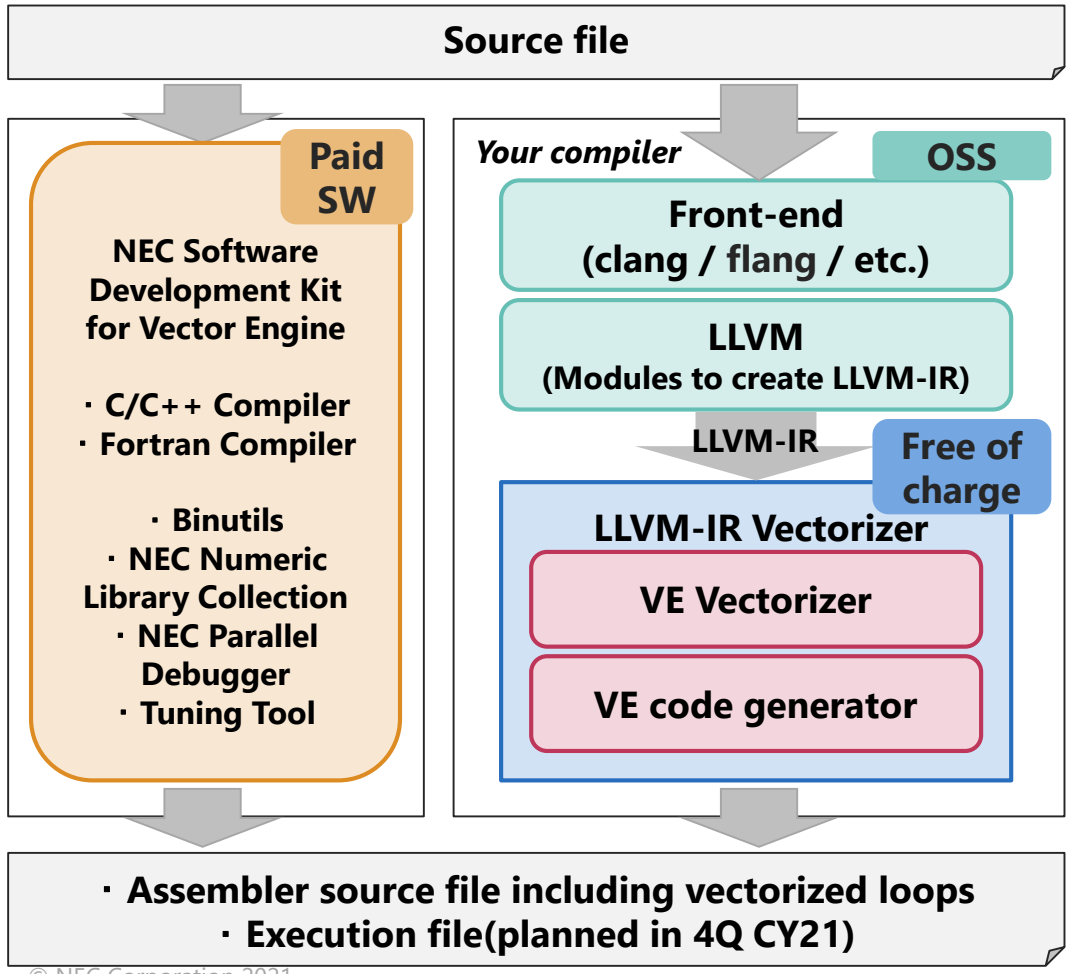

LLVM-IR Vectorizer

*1 will be supported in Q1 CY22.

- Includes vectorizer and code generator for VE.
- Inputs LLVM-IR from memory or an IR-file and outputs an assembler source code for VE.
- Applies automatic vectorization to LLVM-IR.
- Has APIs to support compiler directive.
- Provides runtime library including vector mathematical functions (sin, cos etc.)
- Will support flang and MLIR.
- Will be enhanced to create execution file in 4Q CY21.

➢ You can build your own compiler having vectorization feature for VE!
➢ You will enjoy vector computing power without additional SW license fee in the end of this year!

Future of NEC ‘s Vector Supercomputer Business

Develop new markets by downsizing vector supercomputers and selling Vector Engine card through partner sales.

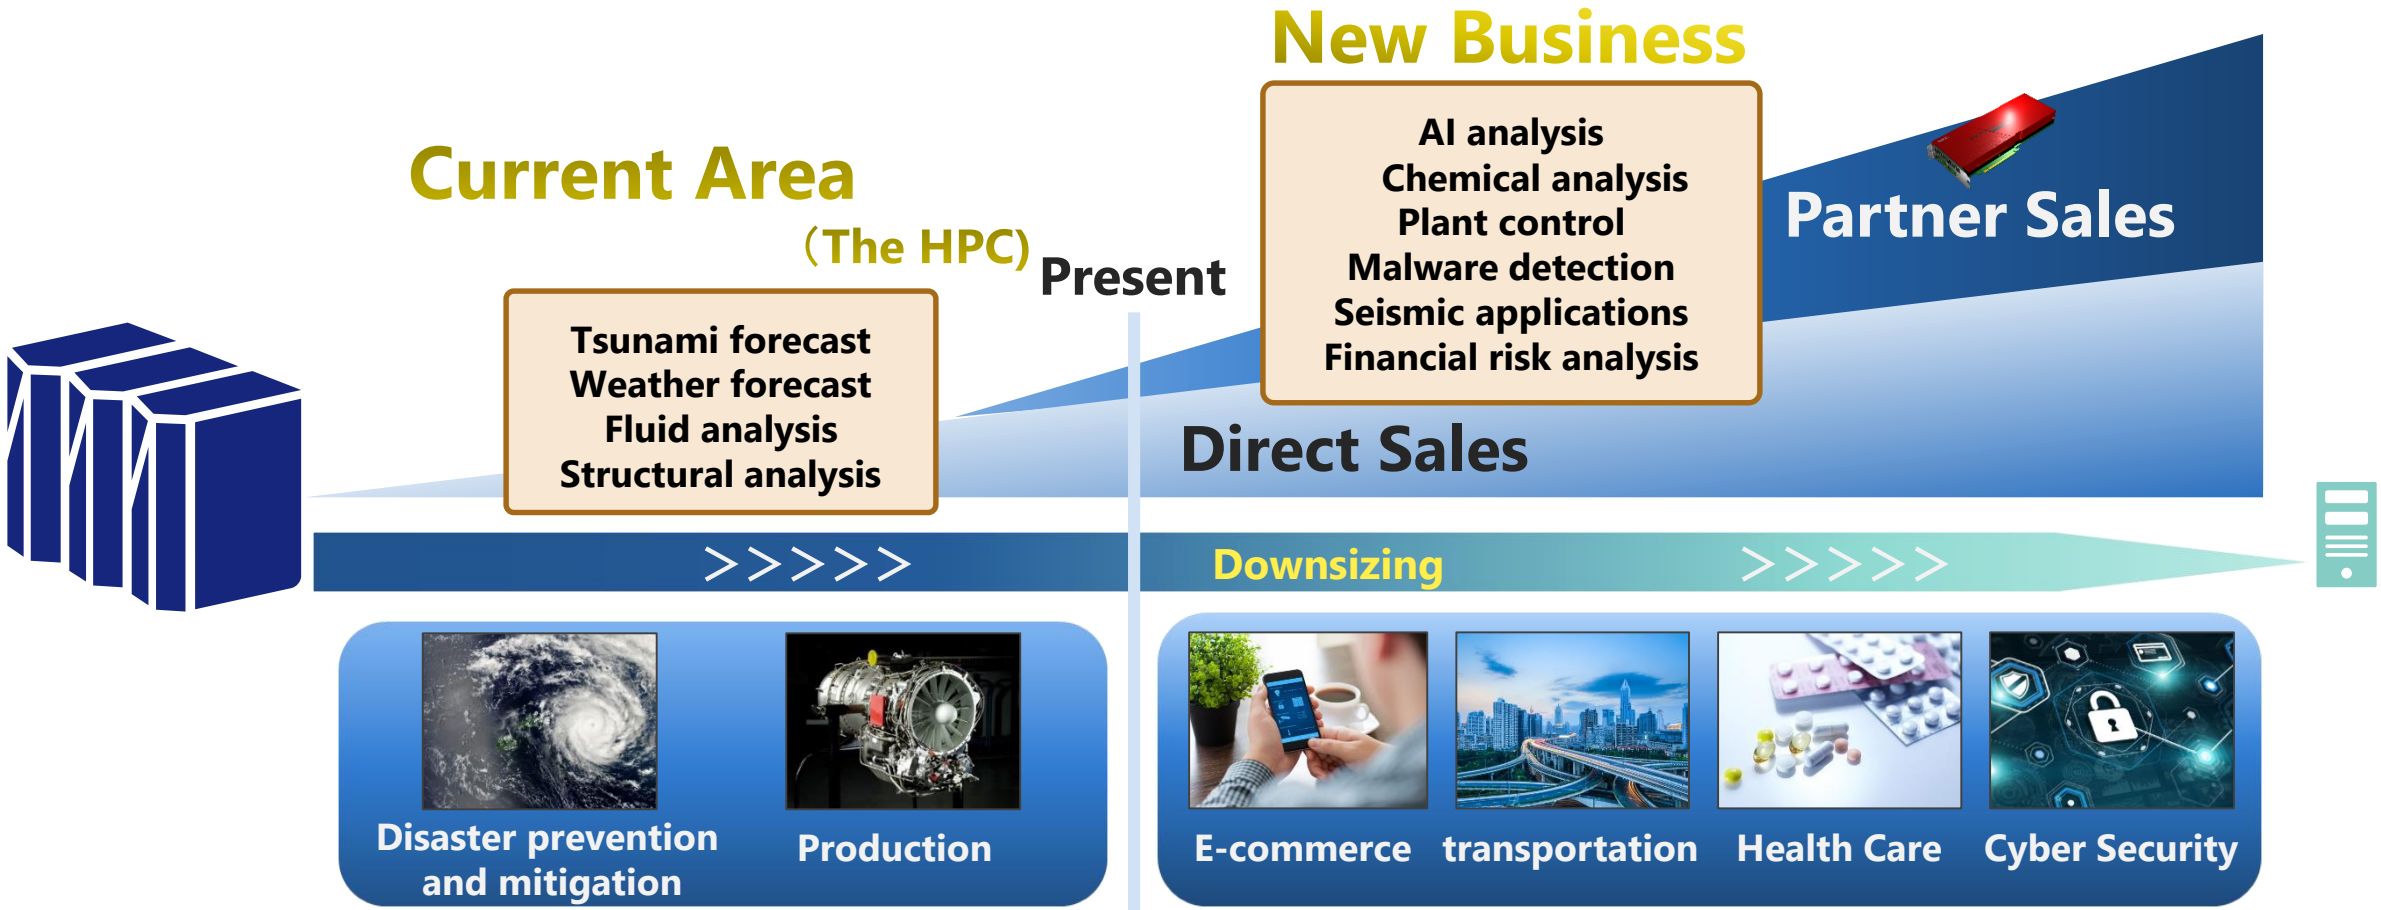

Start card business

Last year NEC announced starting PCIe card selling through system integrators as NEC partners

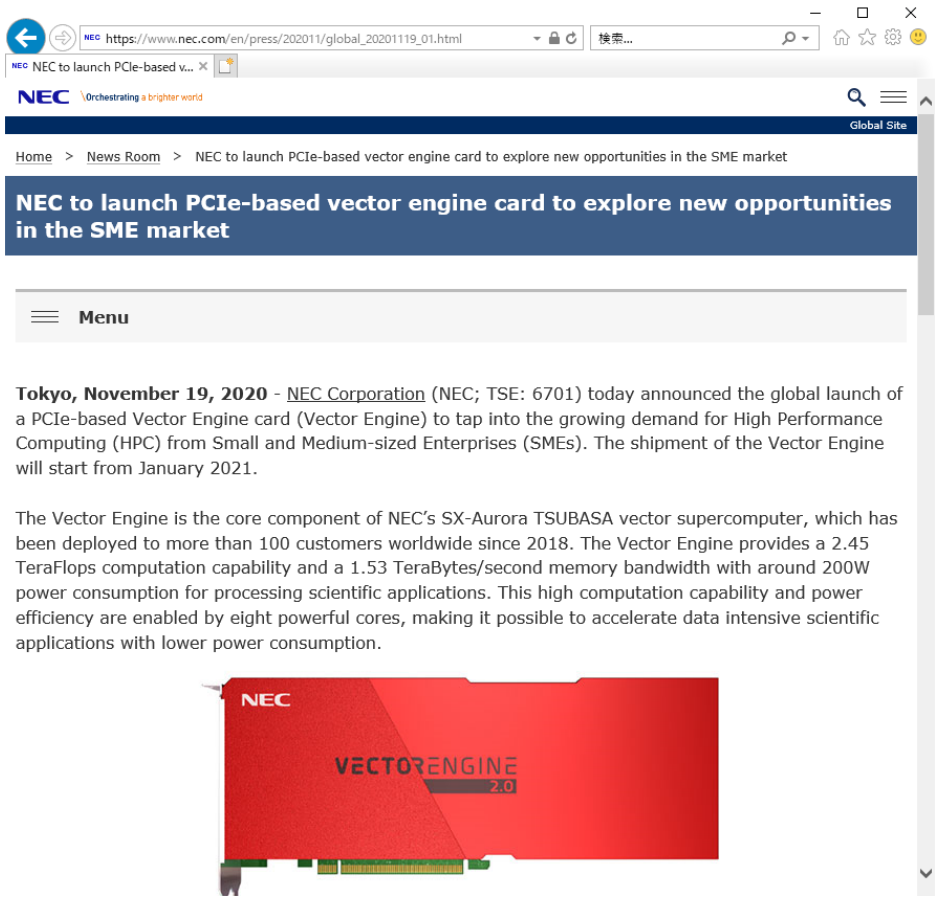

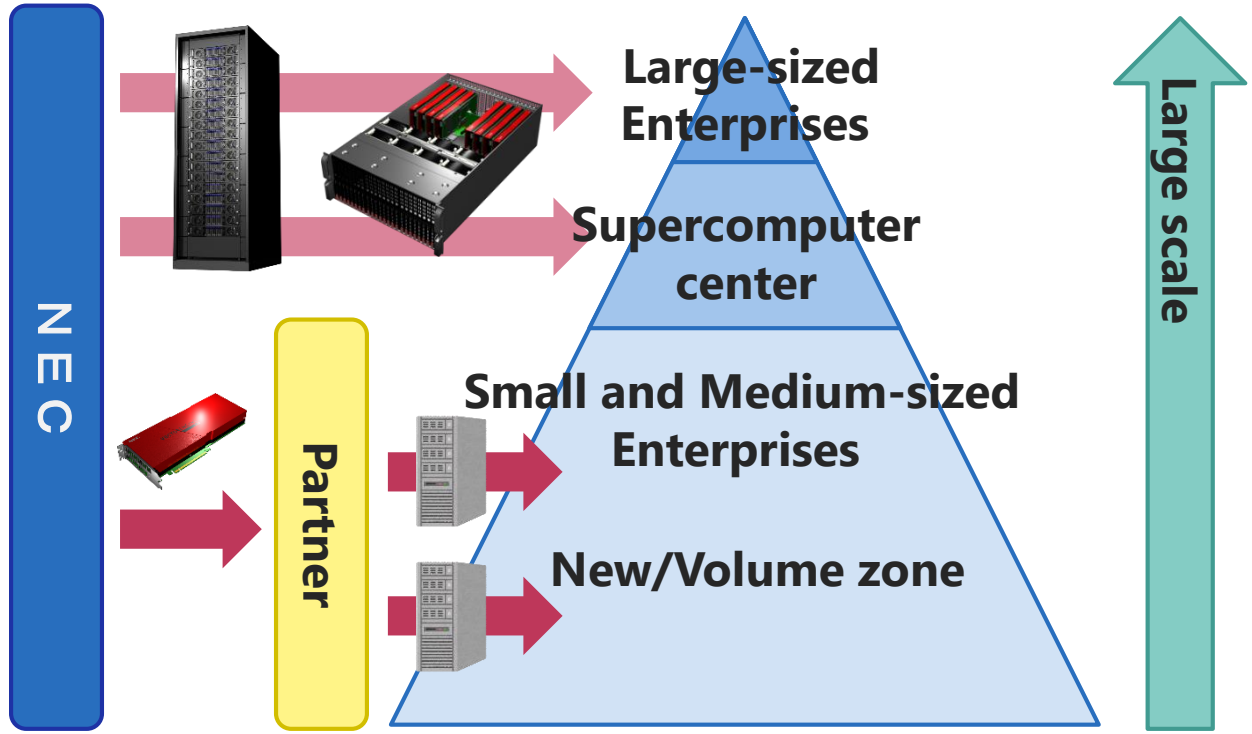

https://www.nec.com/en/press/202011/global_20201119_01.html

Value of Vector Engine

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

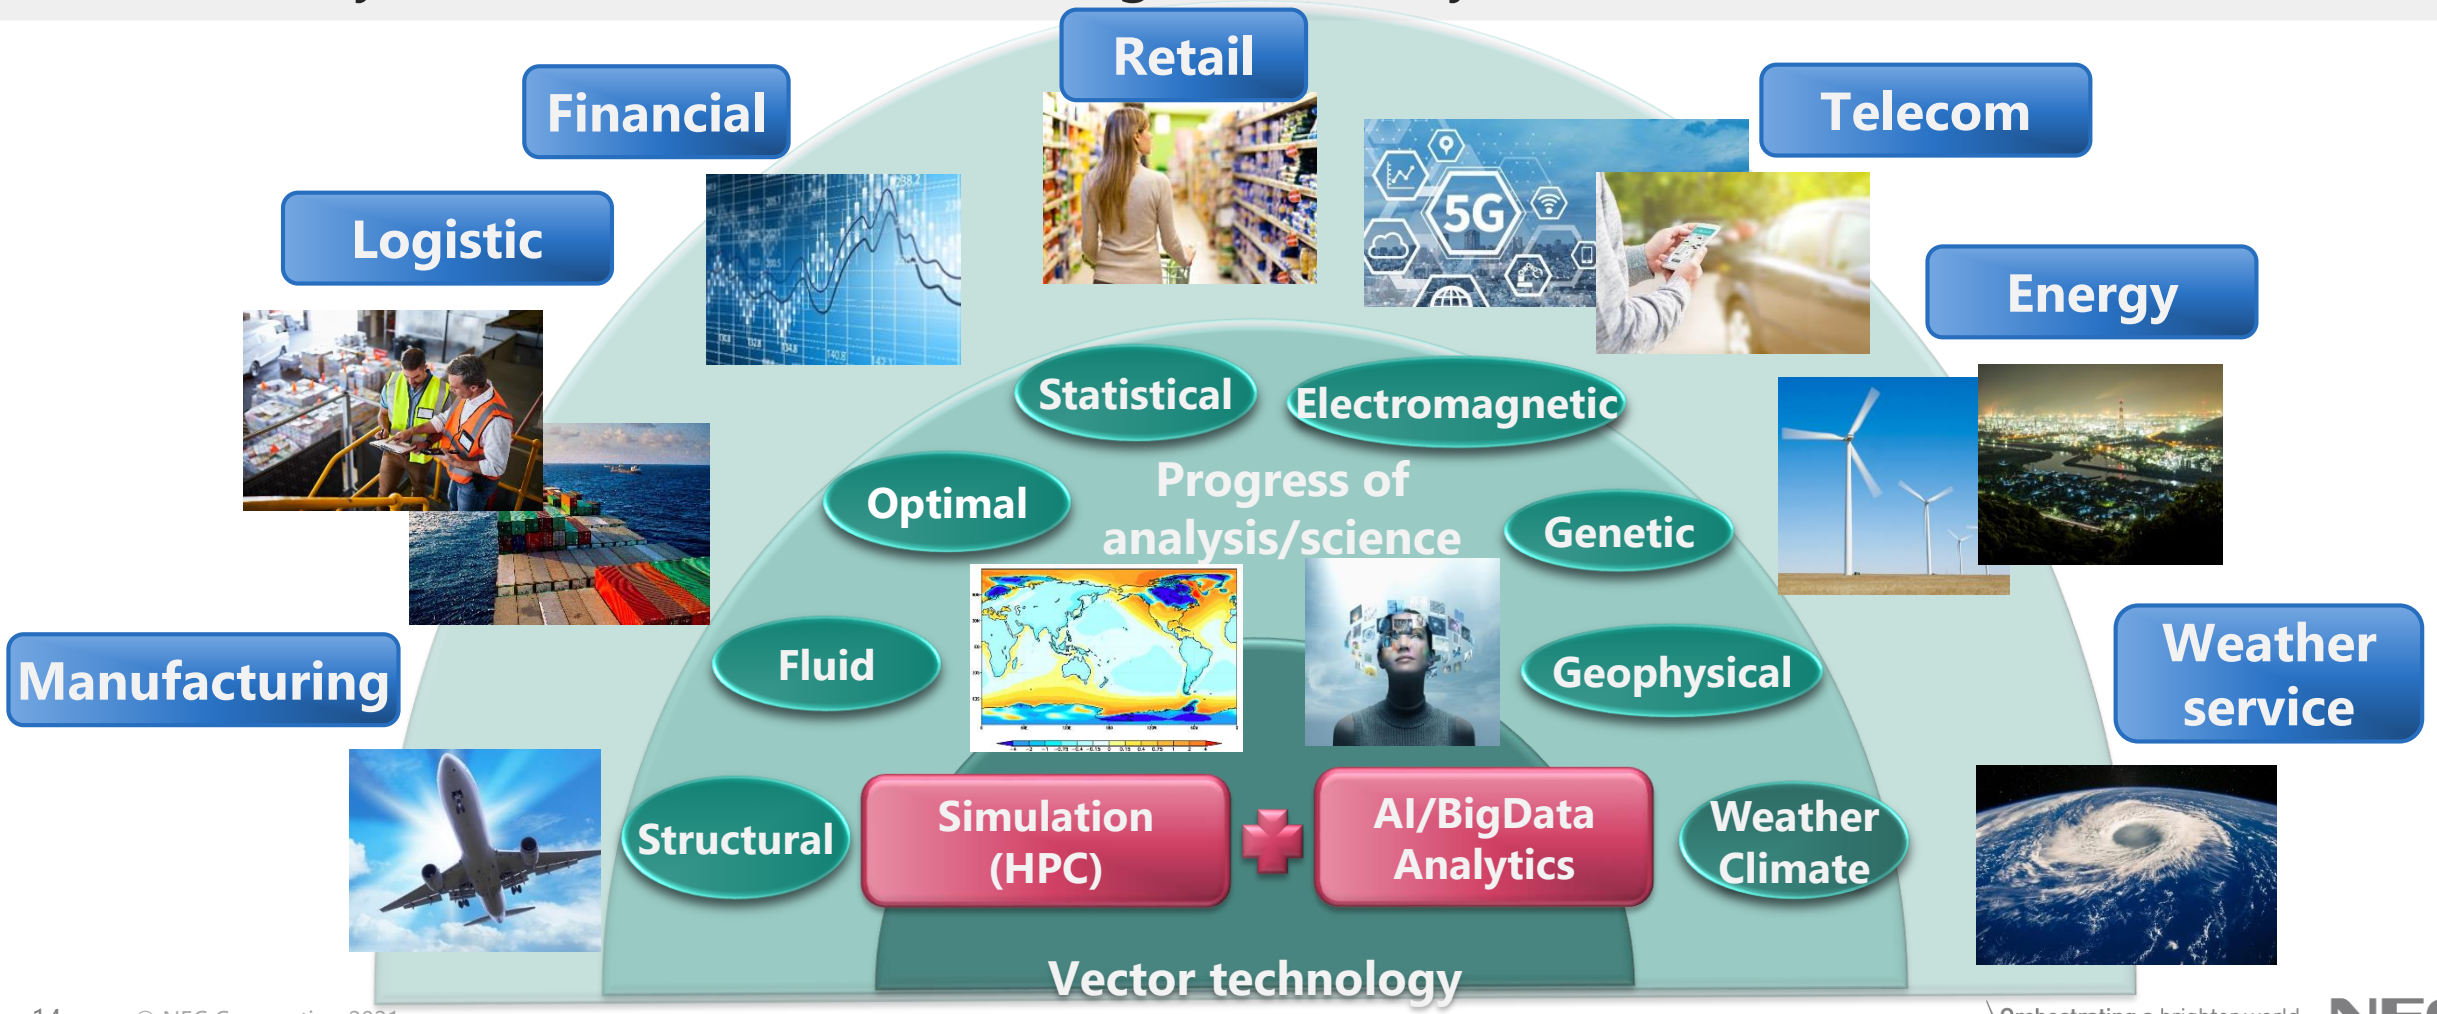

Meteorology

Weather service

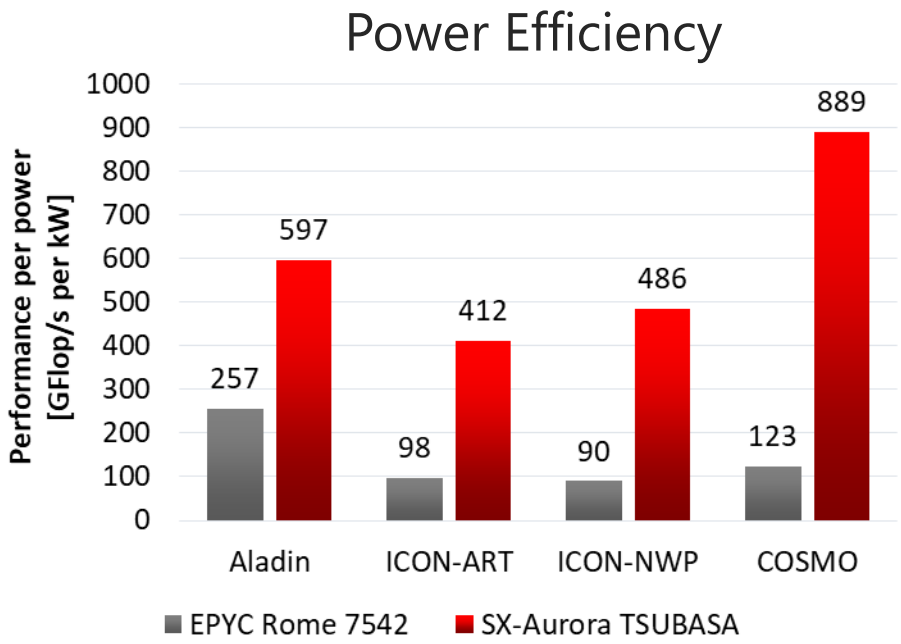

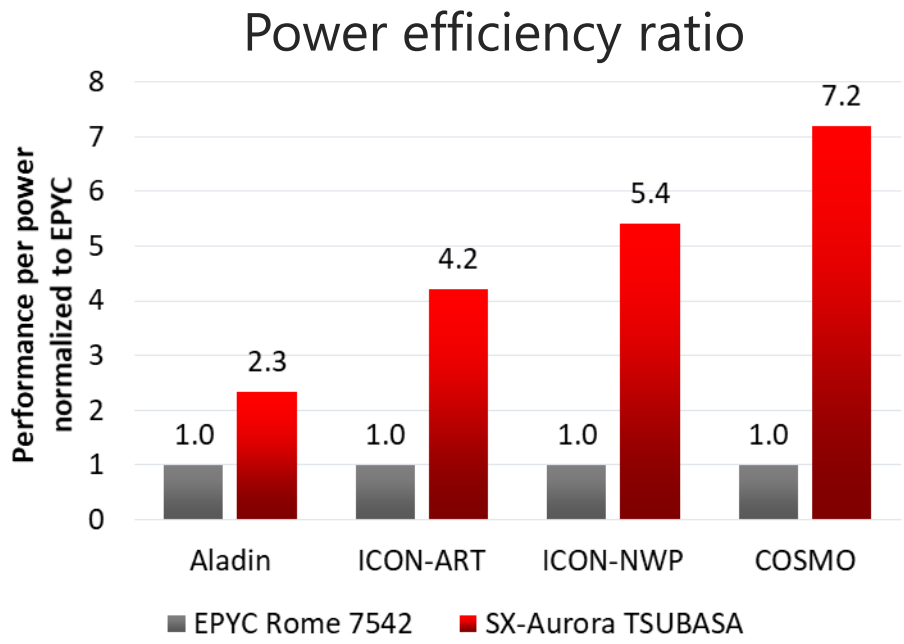

- **EPYC Rome 7542:** EPYC Rome 7542 32 cores/socket, 2.9GHz. 2 sockets per node
- **SX-Aurora TSUBASA:** VE10AE x8 / VH (single socket Rome)
- **ICON-ART:** Status as of 2019 for ICON-ART

- Power supply limitation is one of the big limiting factor of each system size
- Aurora contributes to accelerate meteorology codes within the power limitation
- For the major meteorology codes, Aurora provides 2-7x higher sustained performance with same power consumption

CFD (FDL3DI, developed by US-AFRL)

Manufacturing

The below performance evaluations were conducted at US Naval Research Lab.

Sustained performance
on 1node and 2x VE10B

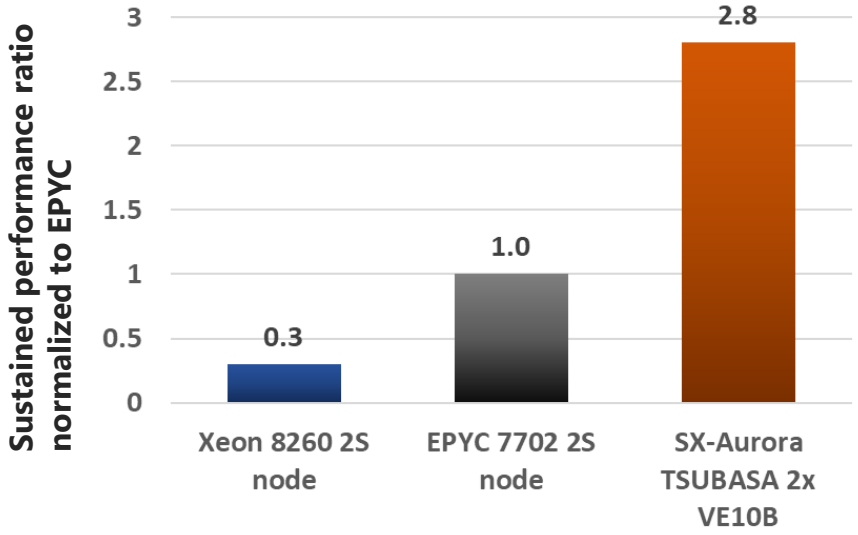

Power efficiency ratio
on 1node and 2x VE10B

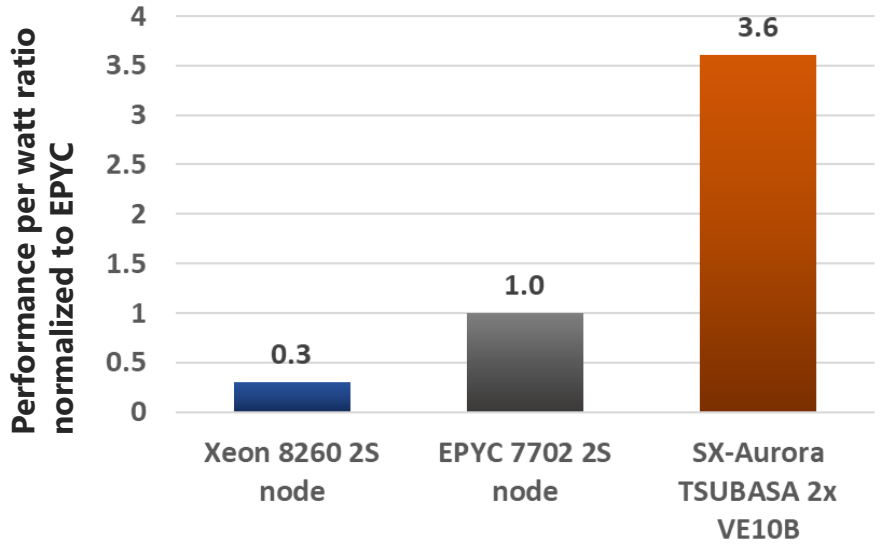

- **FDL3DI:** High-Order Schemes for Navier-Stokes Equations
- **Xeon 8260:** Xeon Cascadelake 8260 24 cores/socket, 2.4GHz, 2 sockets per node
- **EPYC 7702:** EPYC Rome 7702 64 cores/socket, 2.0GHz. 2 sockets per node
- **SX-Aurora TSUBASA:** 2x VE10B / VH (dual socket Xeon)

- Power consumption: 530W/EPYC node(measured), 1,020W/(VH+2x VE10B) (measured), 530W/Xeon(assumption)
- SX-Aurora TSUBASA provides higher performance, and much higher power efficiency than the x86 systems
- Customer's satisfaction with minimum effort for vector tuning without special program language

STAC-A2™ benchmark https://www.stacresearch.com/news/NEC210422

Posted May 12, 2021

Financial

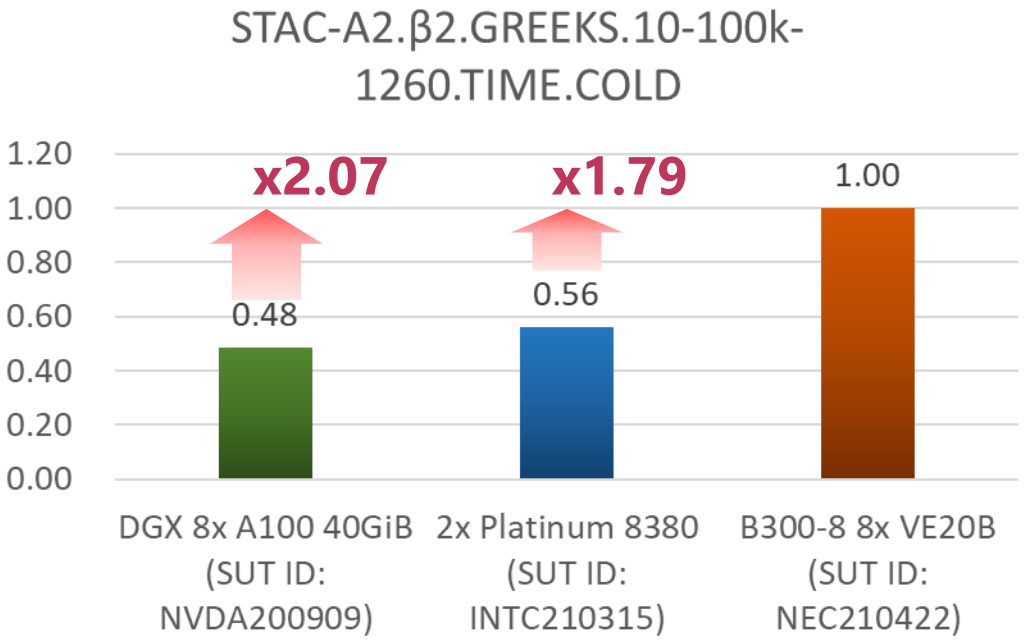

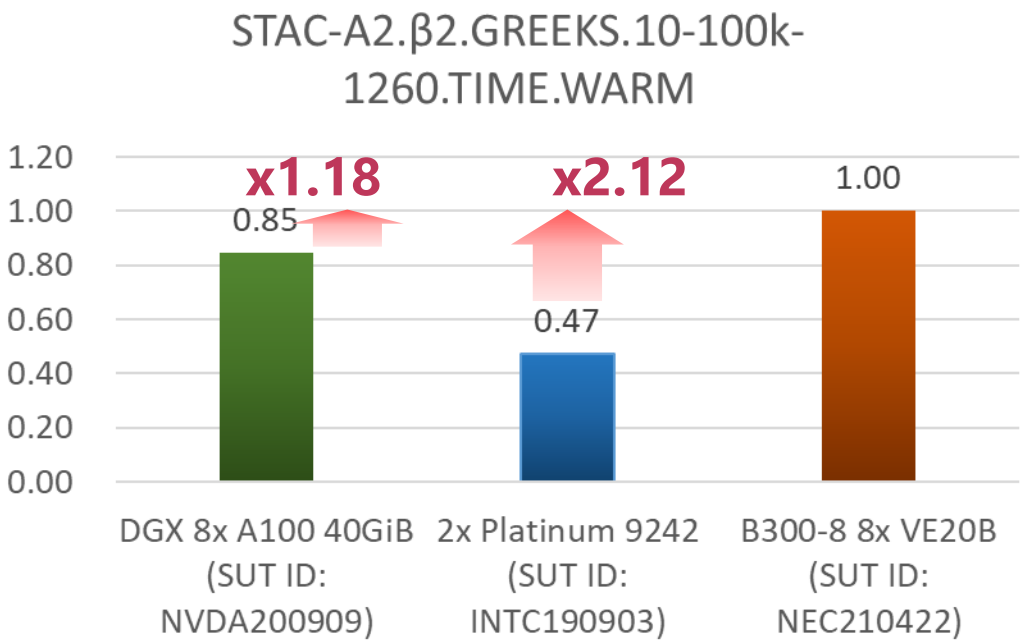

*The above charts illustrate the SKUs with the best reported results from each hardware classification as of May 12. 2021.

- **Large Greeks BM:** Computation of Greeks on American options using Longstaff-Schwartz to price the options in the presence of early exercise with large data size
- **DGX A100:** 8x NVIDIA A100 SXM4 40GiB GPUs per node (https://www.stacresearch.com/news/NEC210422)
- **Xeon 8380:** Xeon Platinum 8380 40 cores/socket, 2.3GHz, 2 sockets per node (https://www.stacresearch.com/INTC210315)
- **Xeon 9242:** Xeon Platinum 9242 48 cores/socket, 2.3GHz, 2 sockets per node (https://www.stacresearch.com/INTC190903)
- **SX-Aurora TSUBASA:** 8x VE20B / Vector Host (dual socket Xeon) (https://www.stacresearch.com/NEC210422)

- STAC-A2 is the technology benchmark standard based on financial market risk analysis.
- Compared to the previous best results for single-server solutions, this (SUT ID: NEC210422) solution was:
  - 79% faster in the cold time for the large Greeks benchmark (STAC-A2.β2.GREEKS.10-100k-1260.TIME.COLD vs. SUT ID INTC210315)
  - 18% faster in the warm time for the large Greeks benchmark (STAC-A2.β2.GREEKS.10-100k-1260.TIME.WARM vs. SUT ID NVDA200909)

Electromagnetic field analysis (OpenFDTD) http://www.e-em.co.jp/OpenFDTD/

Telecom

Performance on 1node and 1VE

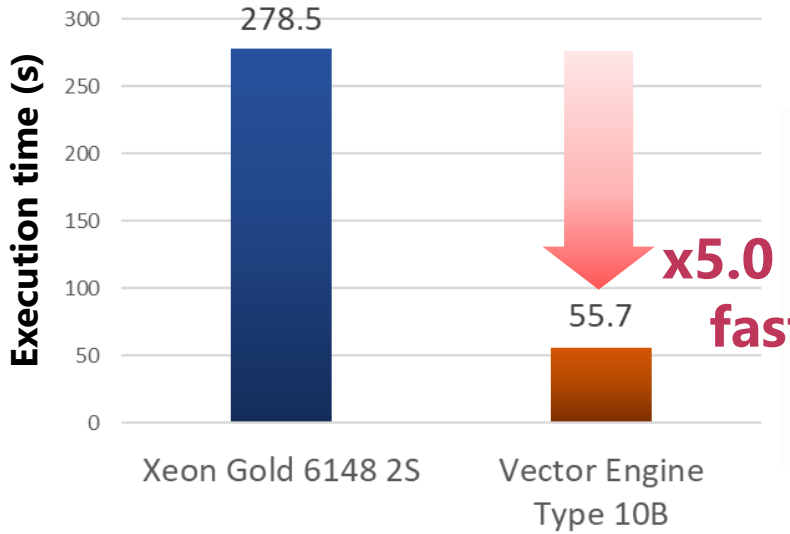

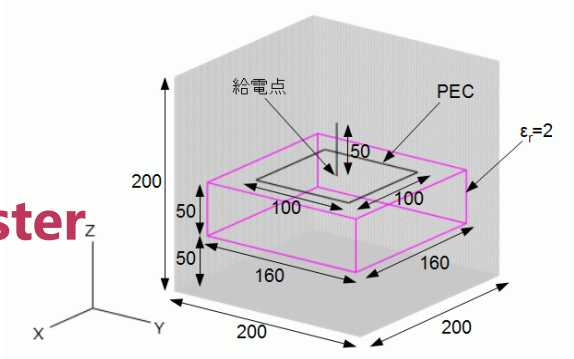

PEC* board is placed on dielectric block.
Monopole antenna stands on the center of PEC board.
*PEC: perfect conductor

*Simulation data provided by EEM

Simulation problem

- **OpenFDTD:** Simulator to analyze electromagnetic field with FDTD, which is provided as free software by EEM
- **BM data:** benchmark500 (# of cell: 500x500x500)
- **Xeon 6148:** Xeon Skylake 6148 24 cores/socket, 2.4GHz, 2 sockets per node
- **SX-Aurora TSUBASA:** VE10B x1 / VH (dual socket Xeon)

- Telecom carriers simulate electromagnetic field and electromagnetic wave propagation for research and development of array antenna for 5G mobile communication. FDTD is one of method for their simulations.
- SX-Aurora TSUBASA provides 5x higher performance than Xeon on OpenFDTD simulator.
- FDTD algorithm is suitable for vector operation, SX-Aurora TSUBASA will contribute to other developments such as automotive millimeter-wave radar, wireless LAN, transmission tower, etc.

AI/ML on SX-Aurora TSUBASA

AI/ML that requires memory performance can be well accelerated
Provide frameworks for easy utilization

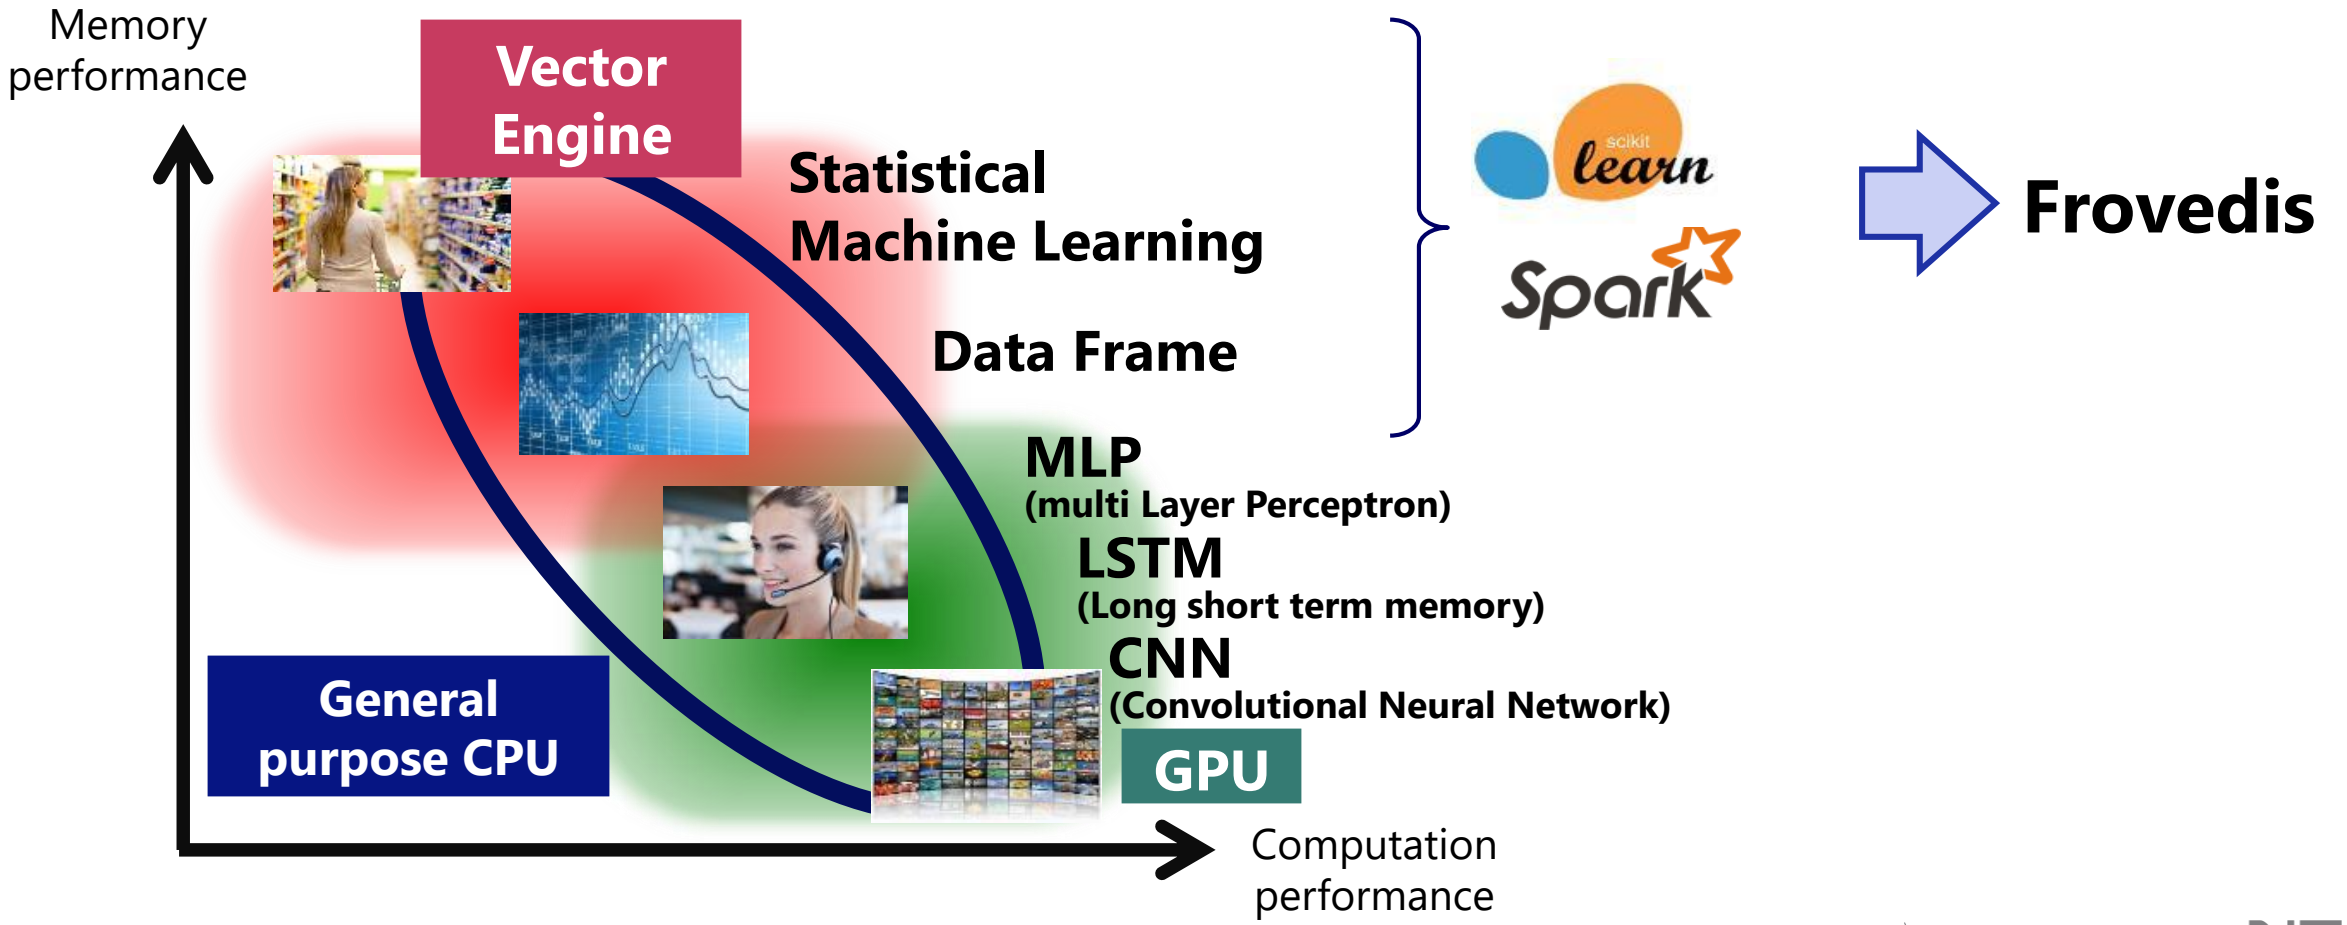

Machine Learning performance

Financial/Retail

Linear SVM

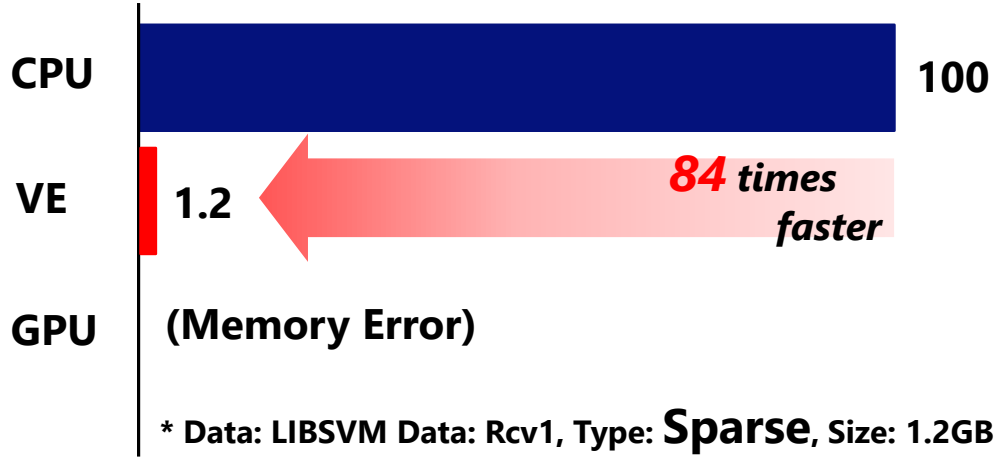

K-means

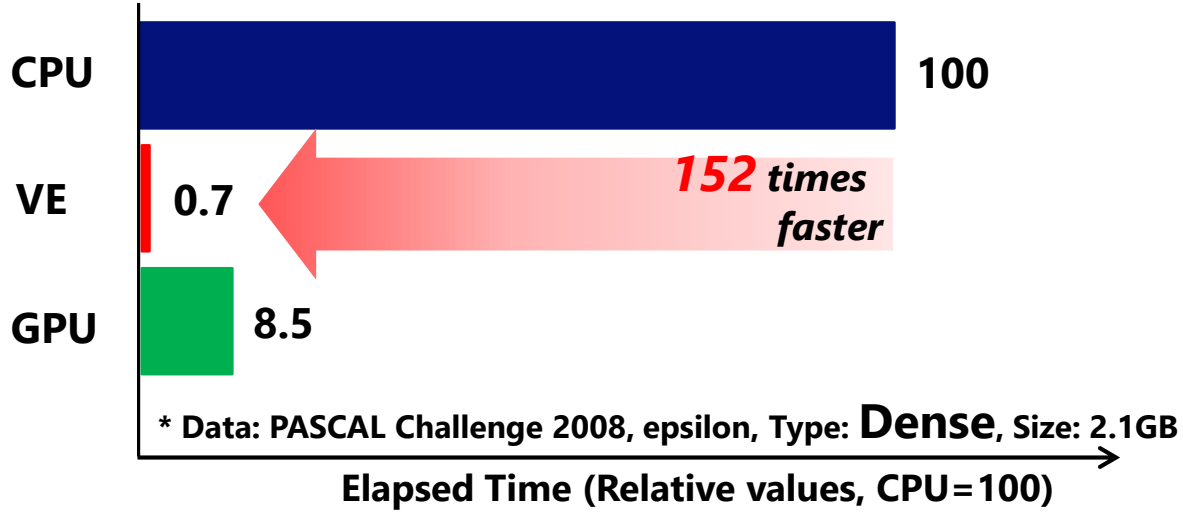

Logistic Regression

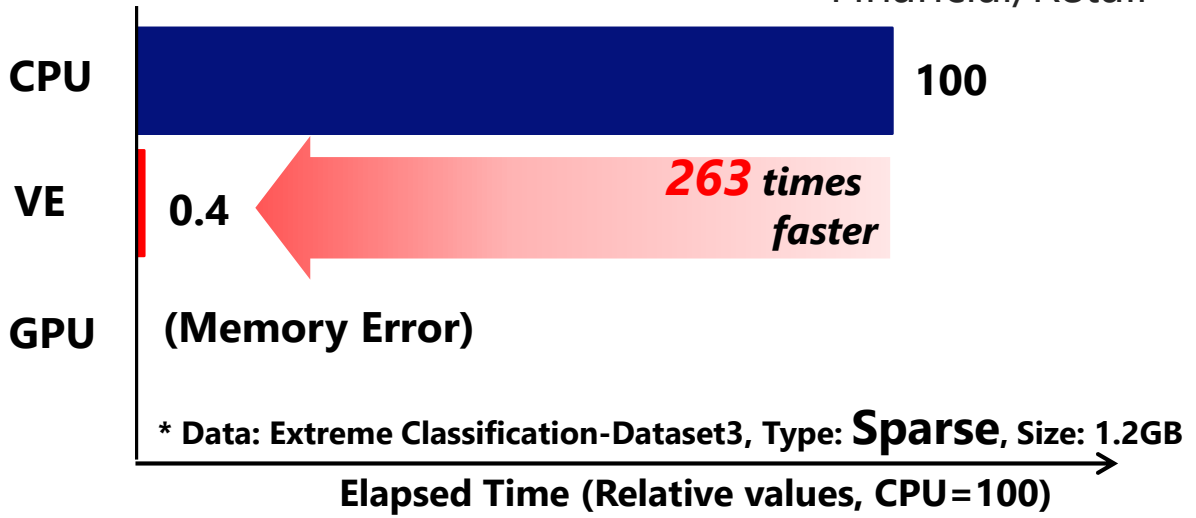

Benchmarking environment

- Hardware

CPU: 2x Intel Xeon Gold 6126 (Skylake)
VE: Type 20B
GPU: A100 40GB

- Software

scikit-learn: 0.23.2
Frovedis: 0.9.8
RAPIDS: 0.17

* GPU BM codes were tuned for V100, and executed on A100 as is.

Frovedis supported algorithms

Implemented with Frovedis Core and Matrix Library

- Supports both dense and sparse data => Sparse data support is important in large scale machine learning

Supported algorithms:

- Linear model
  - Logistic Regression
  - Multinominal Logistic Regression
  - Linear Regression
  - Linear SVM
- ALS
- K-means
- Preprocessing
  - SVD, PCA
- Word2vec
- Factorization Machines
- Decision Tree
- Naive Byes
- DBSCAN
- Graph algorithms
  - Shortest Path, PageRank, Connected Components
- Frequent Pattern Mining
- Spectral Clustering
- Hierarchical Clustering
- Latent Dirichlet Allocation
- Random Forest
- Gradient Boosting Decision Tree (GBDT)

We will support more!

https://github.com/frovedis/frovedis

Background of multi-architecture system -towards Heterogeneous Computing-

Architecture is selected according to characteristics of each of applications.
One of trends in HPC system is hybrid, composed of a variety types of processors.

Scientific calculation

- Weather forecast
- Aerodynamic analysis
- Collision analysis

Statistical processing

- Recommendation
- Demand prediction
- Fraud detection

Image recognition

- Self-driving
- Checking goods
- Cancer diagnosis

Real-time transaction

- Financial transaction
- Face recognition
- Industrial robot

Combinatorial optimization

- Financial portfolio
- Shift schedule
- Delivery planning

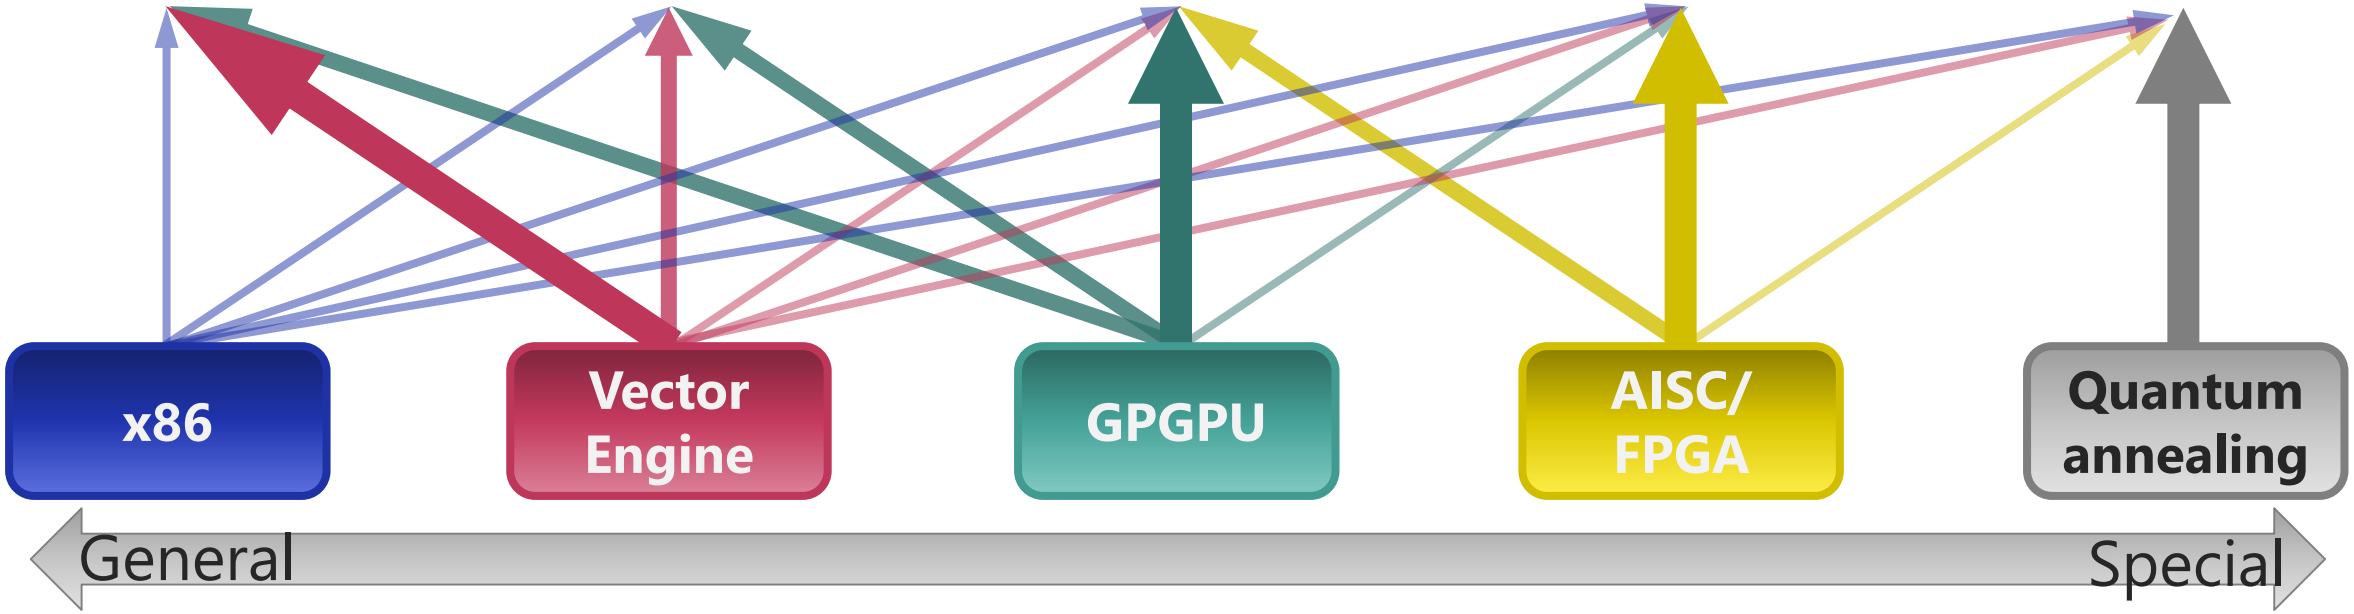

MPI communication on multi-architectural supercomputer

Higher performance by allocating appropriate resources with MPI communication between CPU, GPU and Vector Engine nodes.

A test benchmark execution was successful on JAMSTEC Earth Simulator!

- 4th generation Earth Simulator system

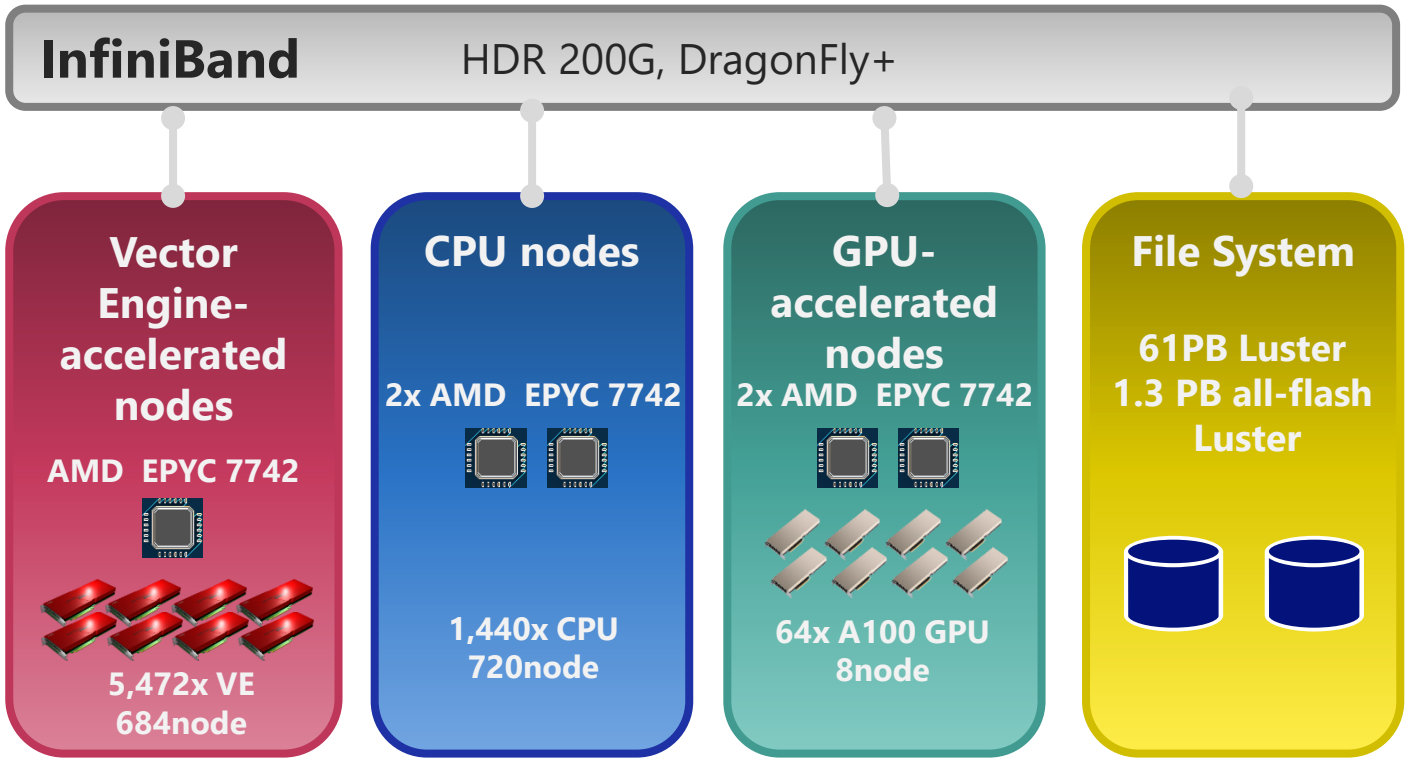

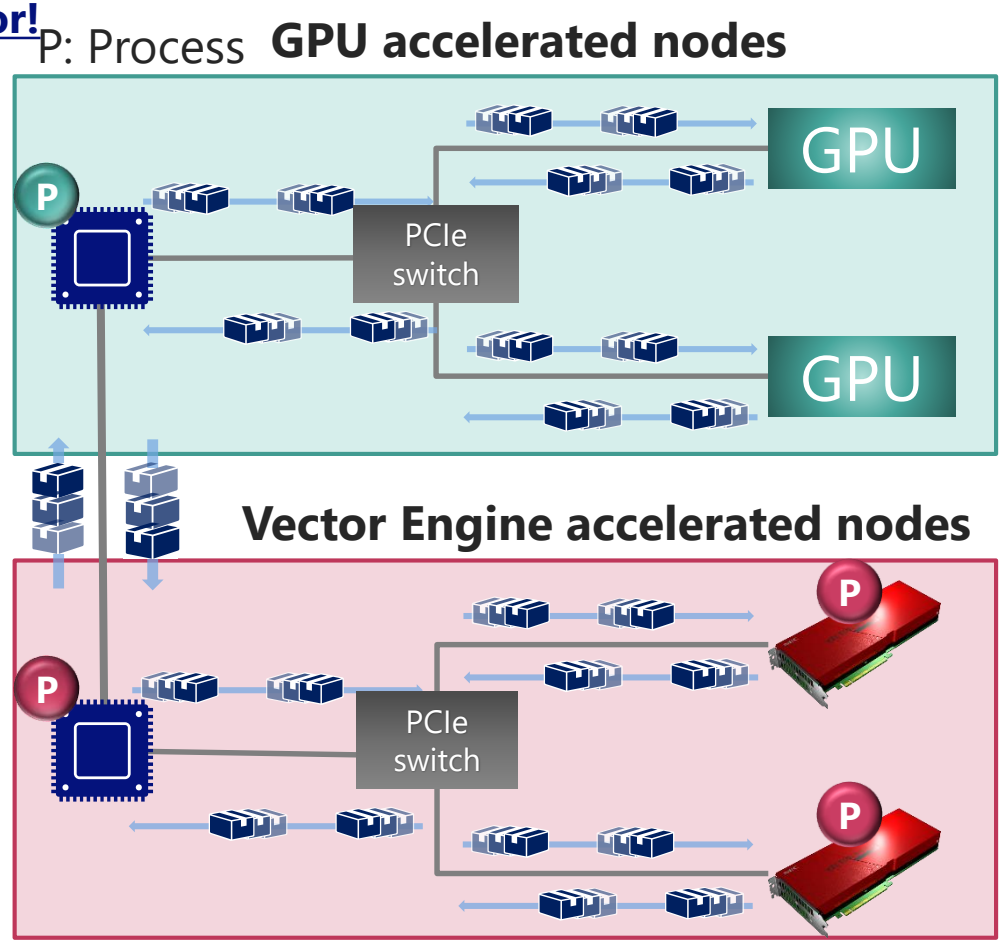

https://www.conferenceharvester.com/uploads/harvester/VirtualBooths/13396/NKBNOCXO-PDF-1-412693%285%29.pdf

Vector Engine 3.0

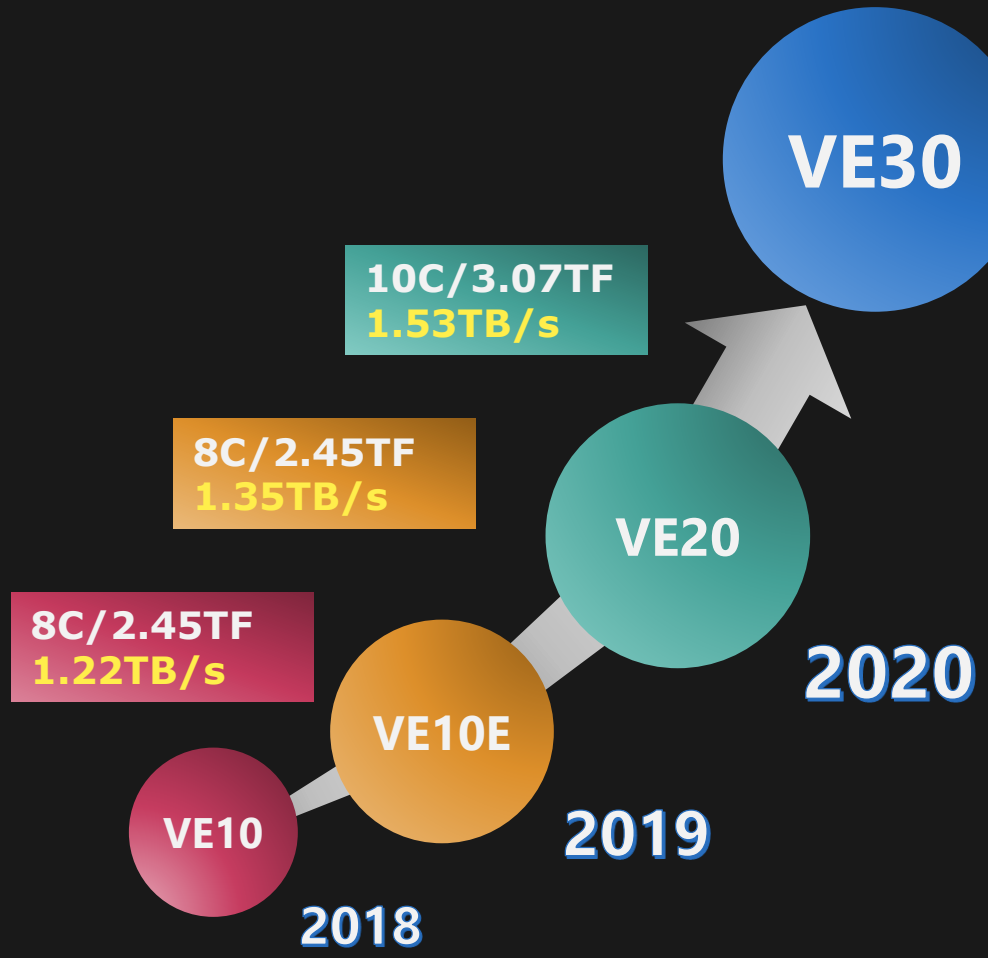

2+TB/s memory bandwidth

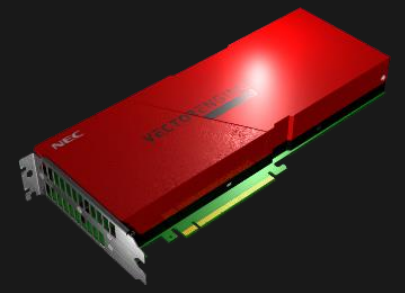

- **Targeting the largest memory bandwidth**
- Inheriting and improving VE/VH architecture
- Higher Flops per processor
- Improved memory subsystem including cache
- Accelerating short vector, and scalar operations
- Adding instructions for AI/ML
- Maintaining high power efficiency
- Heterogeneous computing enhancement
- LLVM-IR Vectorizer(C/C++, Fortran) for VE30
- Virtual machine support

Roadmap

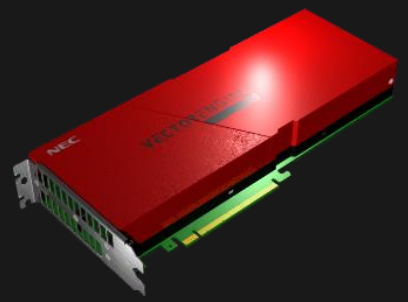

Memory bandwidth / VE

VE10
8C/2.45TF
1.22TB/s memory bandwidth

VE10E
8C/2.45TF
1.35TB/s memory bandwidth

VE20
10C/3.07TF
1.5TB/s memory bandwidth

VE30
2+TB/s memory bandwidth

VE40
2++TB/s memory bandwidth

VE50

2019 2020 2021 2022 2023 2024

Find more information on our website

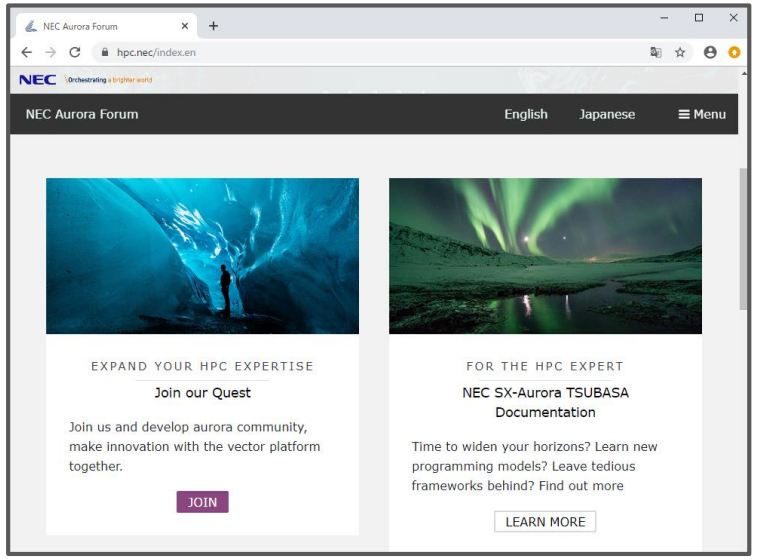

Aurora Web Forum

http://www.hpc.nec

- Latest updates
- Manual, documents
- Bulletin board

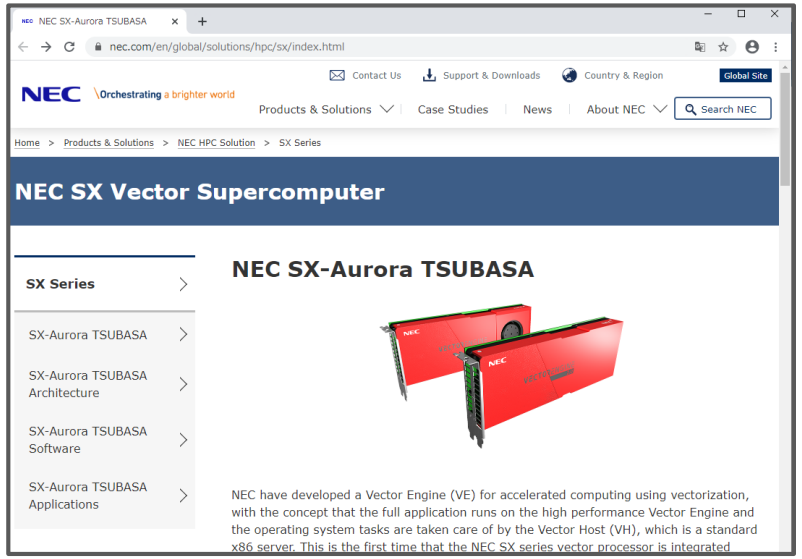

SX-Aurora TSUBASA Website

http://www.nec.com/en/global/solutions/hpc/sx/index.html

- Hardware and software overview
- Supported applications

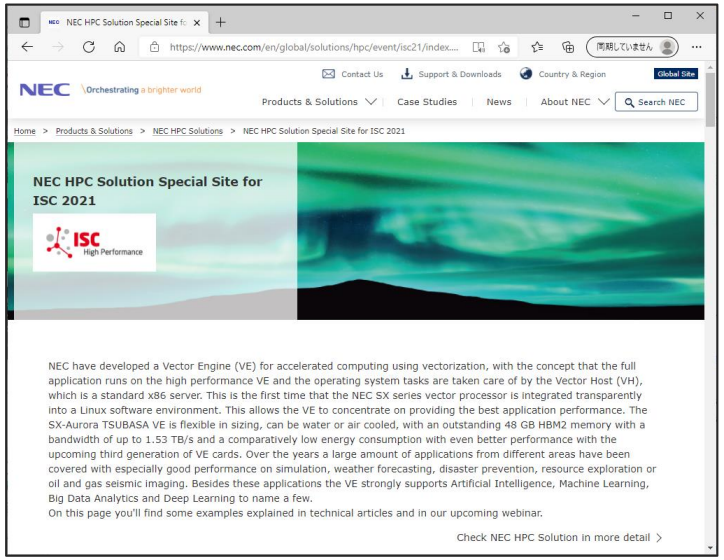

NEC ISC21 website

https://www.nec.com/en/global/solutions/hpc/event/isc21/index.html

- Webinar Jun 09:00 - 12:00- CEST
- Technical articles

✉ info@hpc.jp.nec.com

\Orchestrating a brighter world

NEC creates the social values of safety, security,
fairness and efficiency to promote a more sustainable world
where everyone has the chance to reach their full potential.


Orchestrating a brighter world
NEC