Orchestrating a brighter world NEC

NEC SX-Aurora TSUBASA
for your better application performance

24-September, 2021

Masaoka, NEC

Agenda

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

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

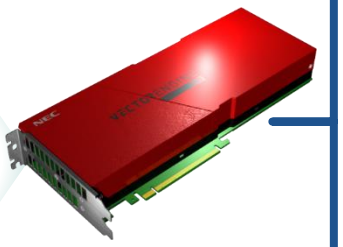

Tower Type

Rack Type

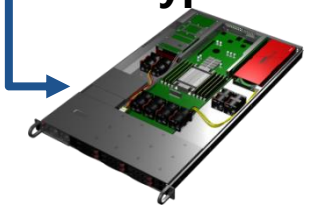

Downsizing of super computer realized by NEC's Technology.

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/Python 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.

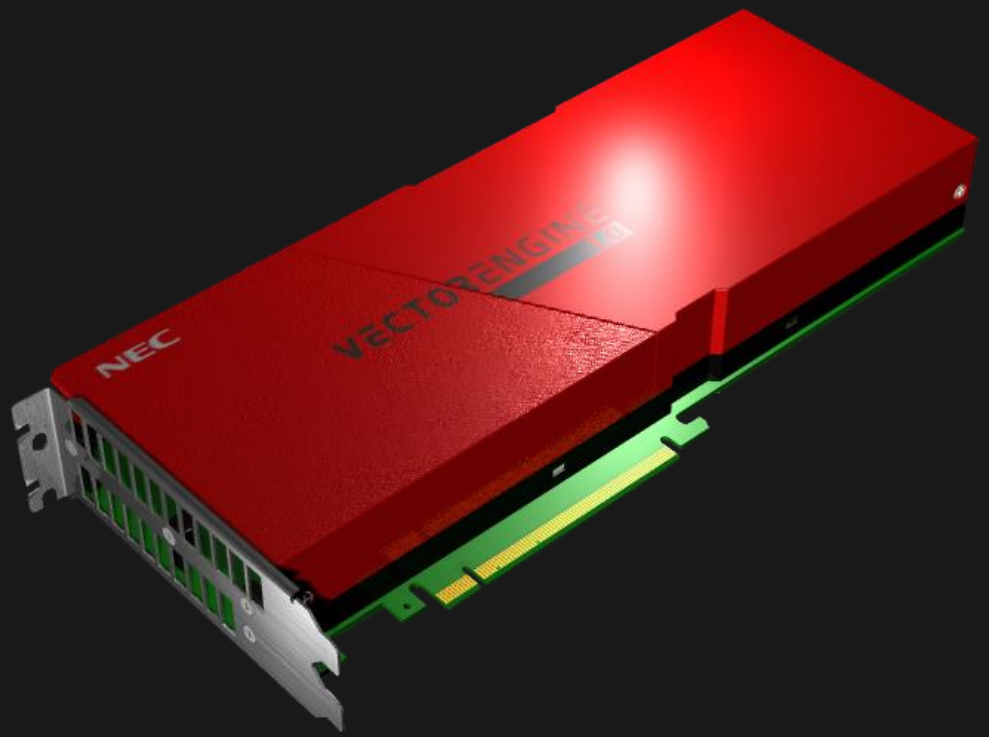

- Vector processor (8/10 cores)
- 1.53TB/s memory bandwidth
- 48GB memory
- 2.45/3.07TF performance (double precision)
  4.91/6.14TF performance (single precision)
- A variety of execution modes
- Standard programming with C/C++/Fortran/Python
- Power consumption < 300 W

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

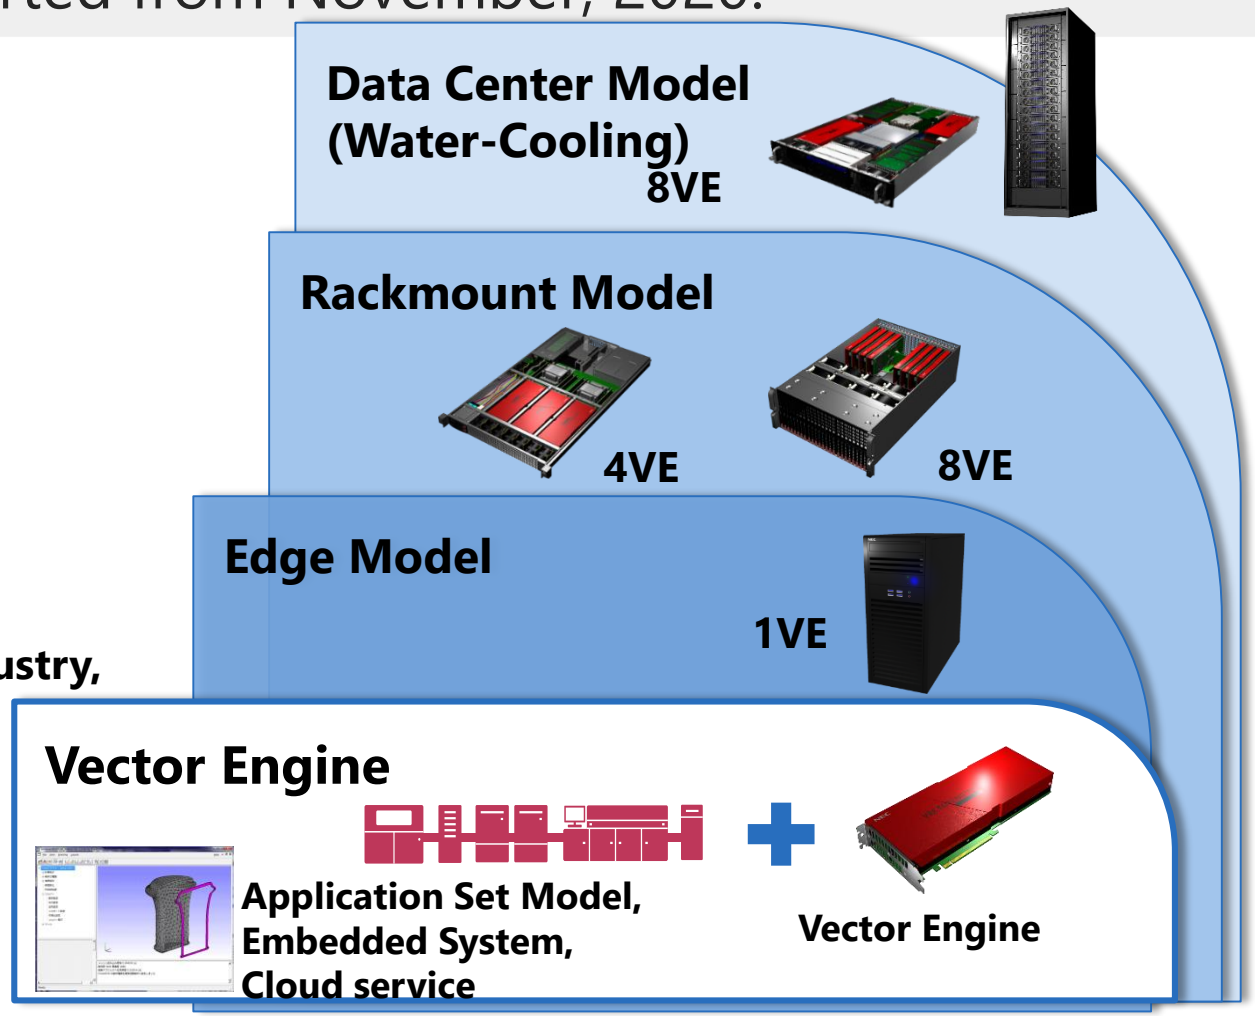

Architecture of SX-Aurora TSUBASA

POINT 2 Ease of Use

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

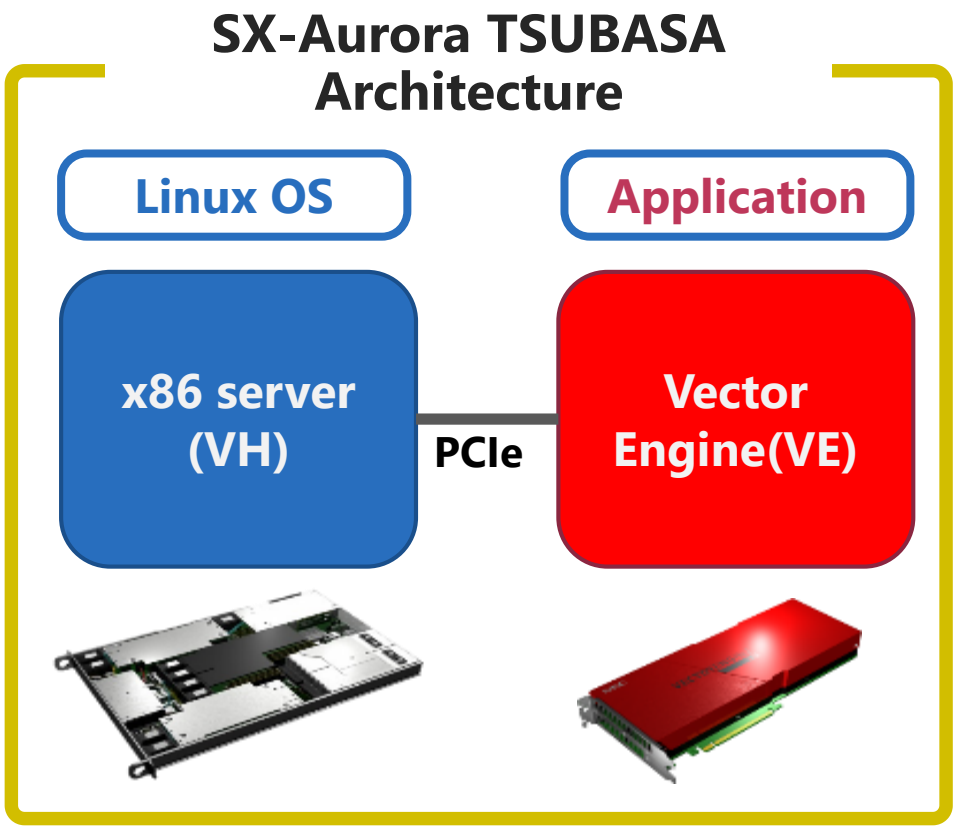

Hardware

- VH(Standard x86 server) + Vector Engine

Software

- Linux OS
- C/C++/Fortran/Python
- Automatic vectorization compiler

Interconnect

- InfiniBand for MPI
  - ✓ VE-VE direct communication support

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

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

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	39	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	54	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%

https://www.top500.org/lists/hpcg/2021/06/

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

NQSV: Switch-over to the urgent job

Job scheduling for the urgent job

- An urgent job can be run immediately suspending normal jobs.
- The normal jobs are resumed once the urgent job is finished.

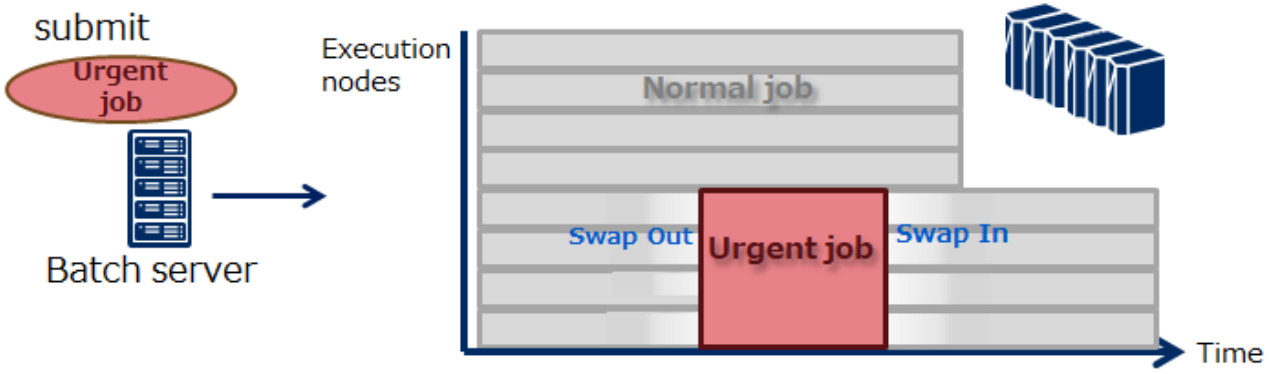

Implementation for the memory swapping

- The VE memory region needed for the urgent job is obtained by swapping out a part of the VE memory used for the normal job to VH.
- The target to be swapped-out can be chosen from VH memory or a file.
  - The performance of swapping-out is better in case VH memory is selected.

Note: Some parts of VE memory used by MPI communication etc. cannot be swapped-out.

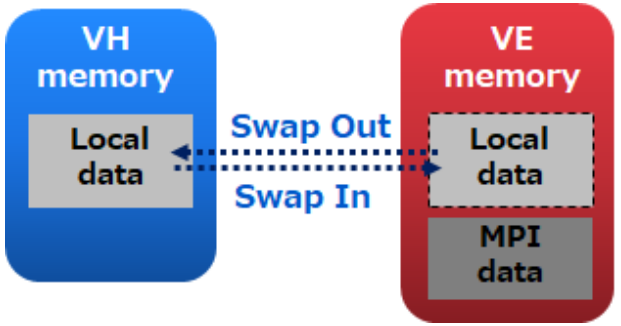

NLCPy : NumPy-like API Accelerated with Aurora

Just by replacing the module name,
Python scripts using NumPy can utilize VE computing power.

Features

- NumPy-like library
  - Provides a subset of NumPy's API
- Highly optimized library
  - Uses optimized library for VE (BLAS and ASL).
  - Provides various vectorized operations.
- Open-source library
  - Licensed BSD License(3-clause).
  - Published on GitHub and PyPI.

```
import nlcpy as np

a = np.random.rand(1000, 1000)
b = np.fft.fft(a)

c = a + b
d = np.linalg.solve(a, c)
```

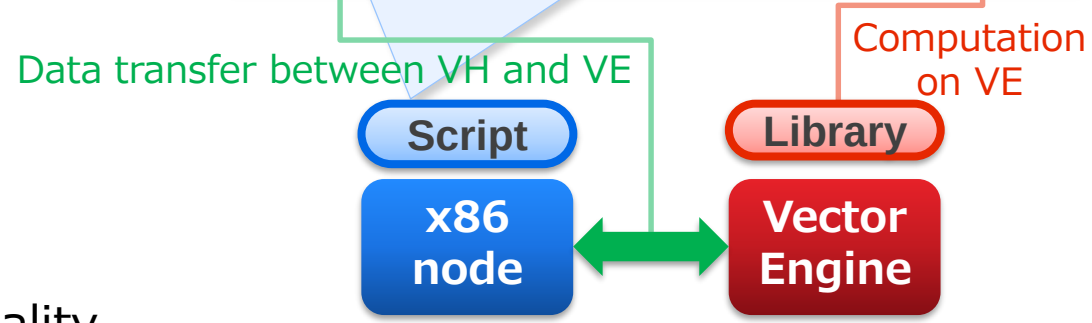

Future Outlooks

- January 2022
  - Supporting "Just-In-Time" compilation functionality.
- April 2022 (Preview Release)
  - Supporting VE-GPU Heterogeneous computing functionality.

Stencil calculation Code Accelerator Interface of NLCPy

- SCA Interface
  - Boosts performance of a wide variety of stencil calculations.
- Example
  - User can easily define arbitrary stencil shapes by specifying relative locations.

```
import nlcpy as vp

a = vp.random.rand(100).reshape(10,10)
b = vp.zeros_like(a)

din, dout = vp.sca.create_descriptor((a, b))
desc_in  = 0.25 * ( din[0,-1] + din[-1,0] + din[1,0] + din[0,1] )
desc_out = dout[0,0]
kern = vp.sca.create_kernel(desc_in, desc_o=desc_out)

for i in range(maxitr):
    a[...] = kern.execute()
```

Laplace Equation

$$\frac{\partial}{\partial x^2}A + \frac{\partial}{\partial y^2}A = 0$$

discretization

finite-difference

$$b_{i,j} = \frac{1}{4}\left(a_{i,j-1} + a_{i-1,j} + a_{i+1,j} + a_{i,j+1}\right)$$

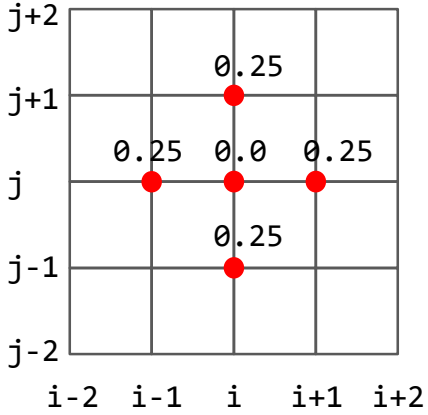

Benchmark Result (2-D XY-axial stencils)

- NLCPy on Vector Engine shows the highest performance.

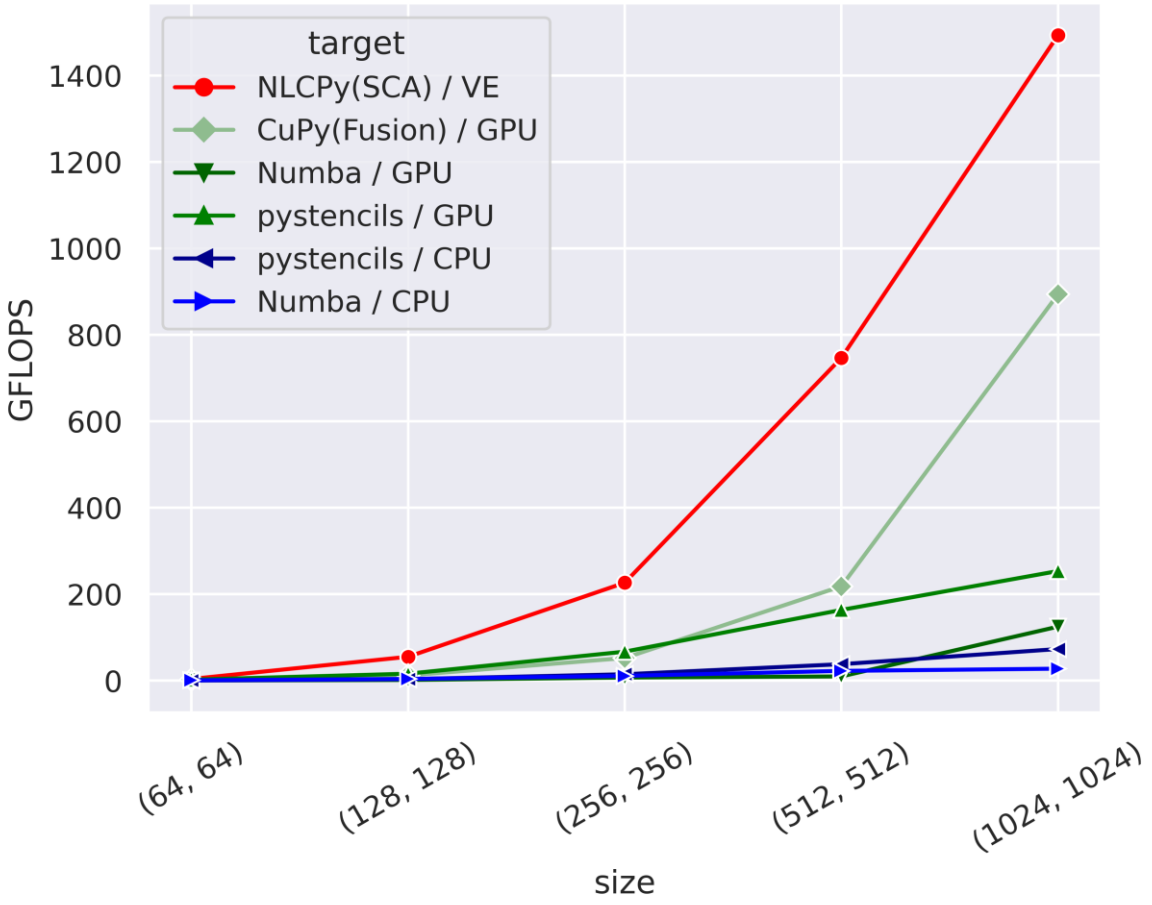

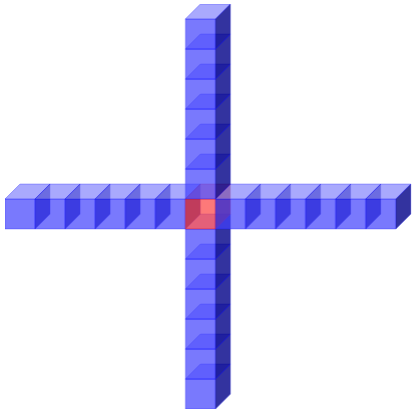

Benchmark Result (3-D XYZ-axial stencils)

- NLCPy on Vector Engine shows the highest performance.

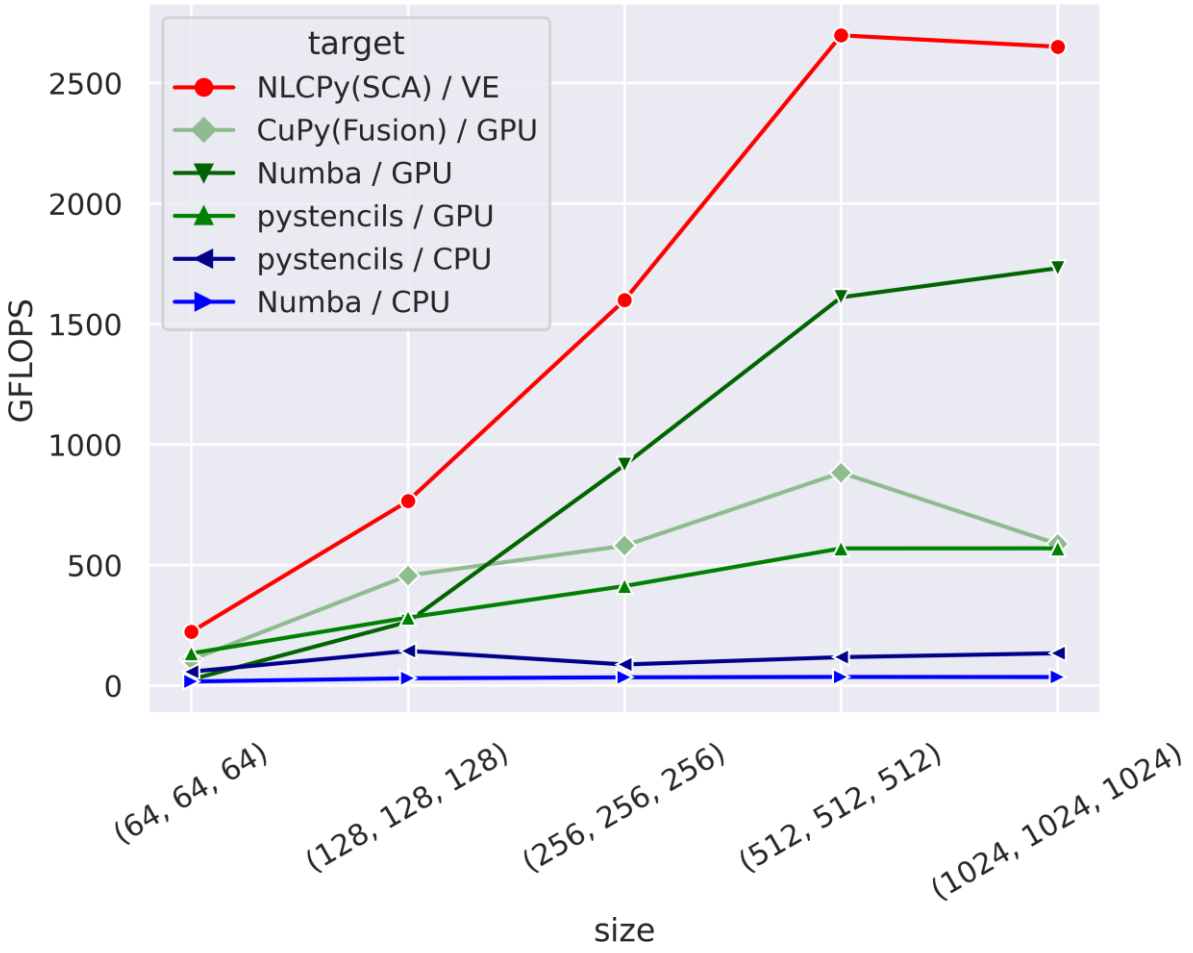

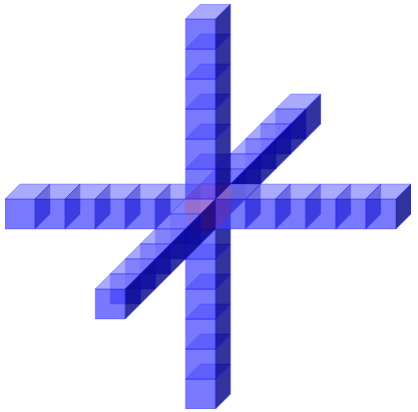

6x6y6za

Benchmark Conditions

- Stencil Shapes and Time Steps
  - 2-D: XY-axial, size 6.
  - 3-D: XYZ-axial, size 6.
  - For 100 time steps.

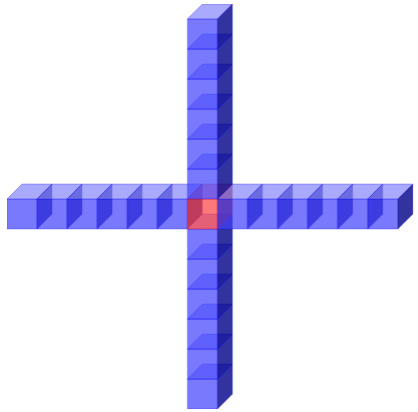

6x6ya

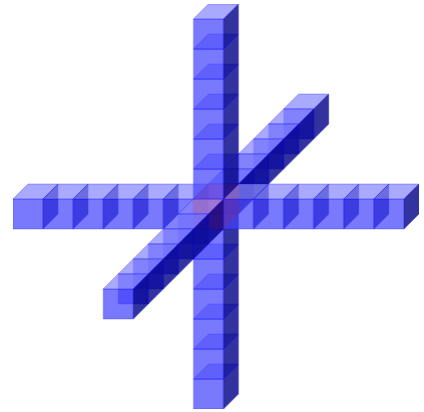

6x6y6za

- Target Libraries
  - Numba:
    - A100 PCI-E 40GB
    - Xeon Gold 6126 x2 (Skylake 2.60GHz, 48 cores)
  - pystencils:
    - A100 PCI-E 40GB
    - Xeon Gold 6126 x2 (Skylake 2.60GHz, 48 cores)
  - CuPy on A100 PCI-E 40GB
  - NLCPy on VE Type 20B (8 cores)

- Data Size
  - 2-D: (64, 64) (128, 128) (256, 256) (512, 512) (1024, 1024)
  - 3-D: (64, 64, 64) (128, 128, 128) (256, 256, 256) (512, 512, 512) (1024, 1024, 1024)

  For each case, single precision was used.

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

Actual Application Performance

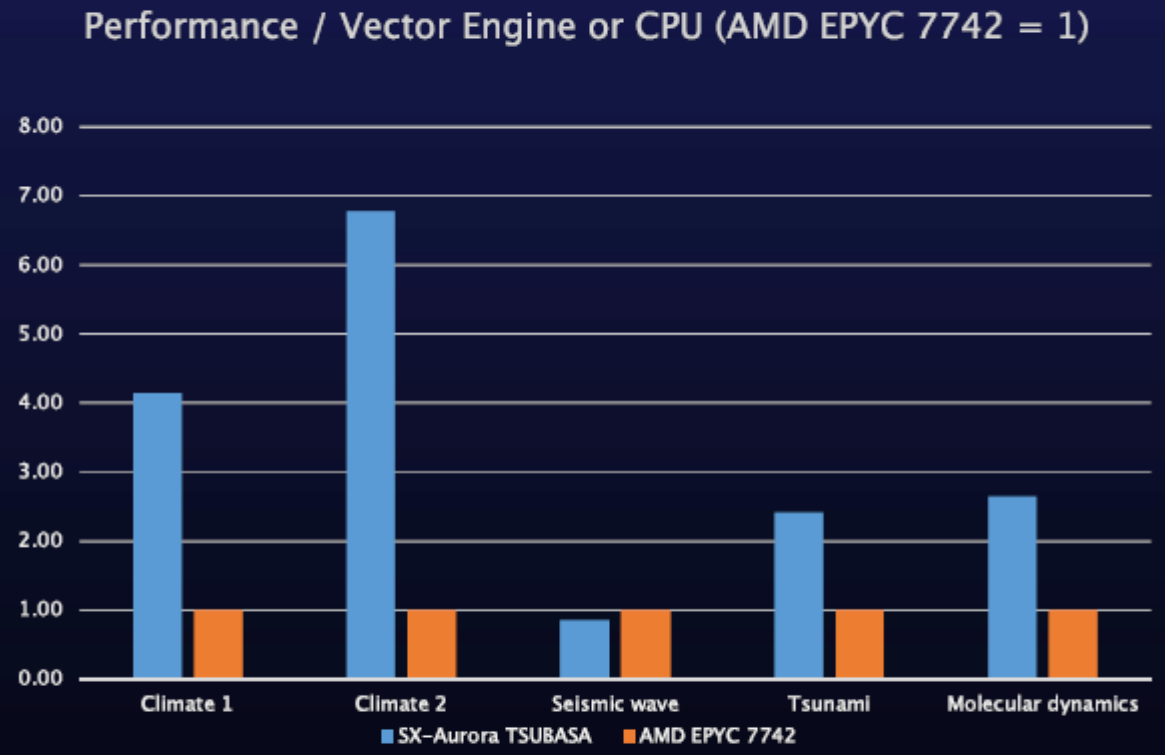

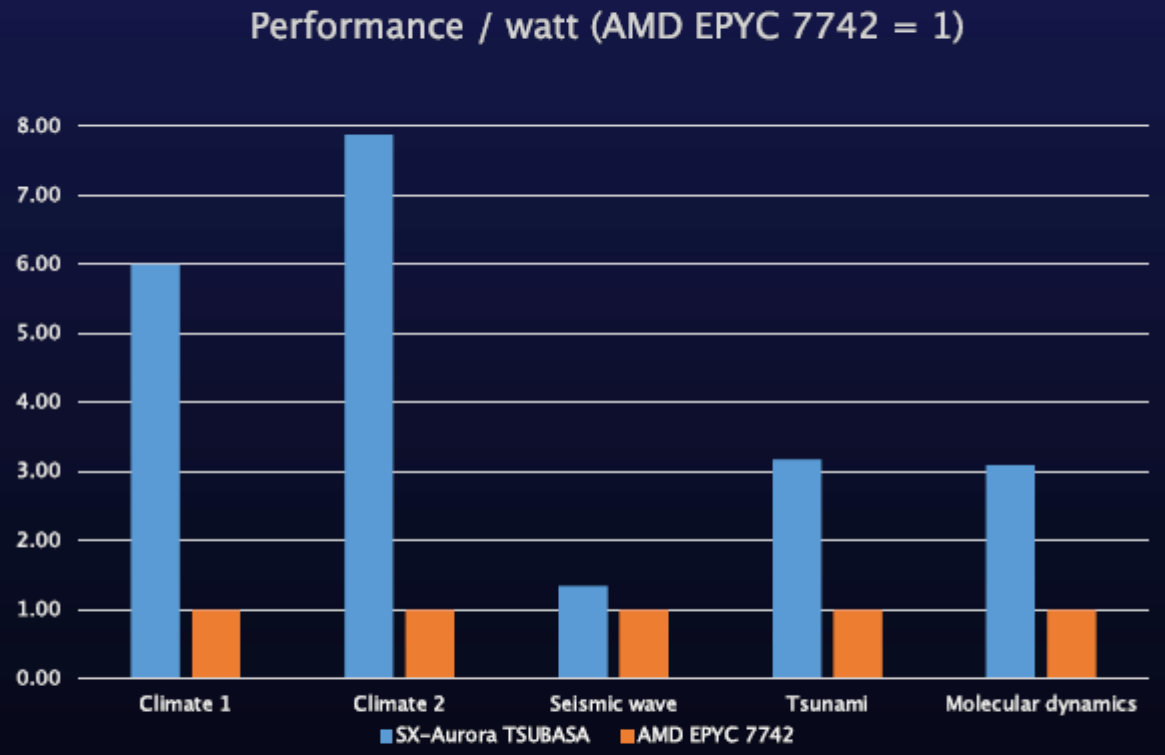

We measured the execution time of some actual applications and maximum power consumption during the execution on NEC SX-Aurora TSUBASA Type 20B (1.6 GHz, 8 core) and AMD EPYC 7742 (2.25 GHz, 64 core). NEC SX-Aurora TSUBASA showed better performance and performance/watt for climate applications in particular.

The task force to port and optimize programs to the new system has started. Its focus includes not only existing programs but also selected programs for GPU.

Meteorology

Weather service

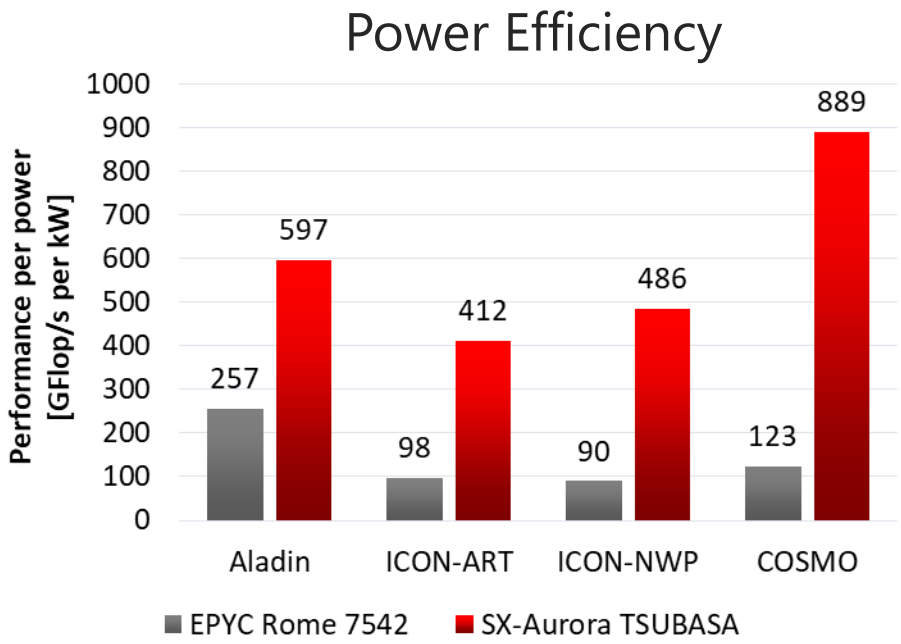

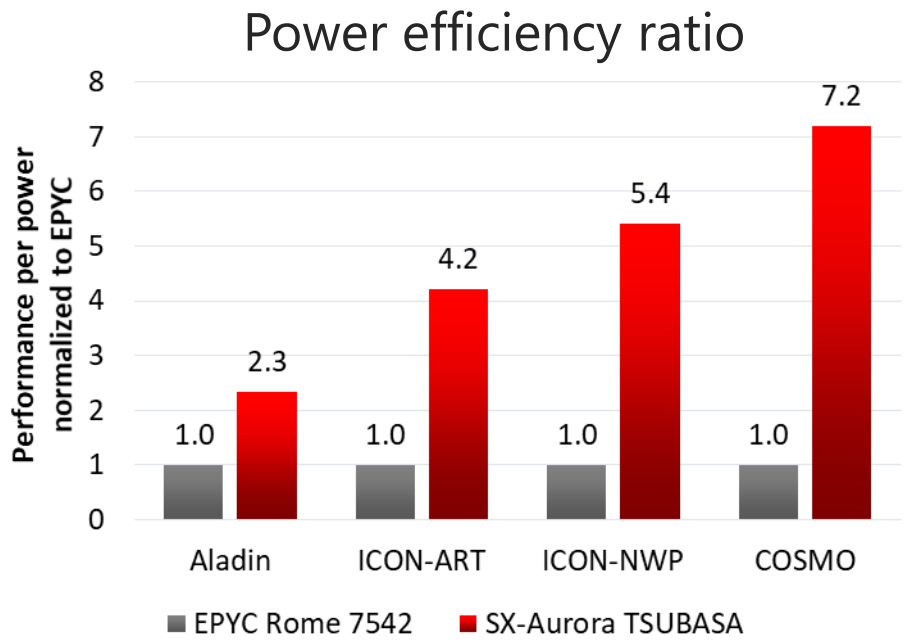

- **EPYC Rome 7542:** EPYC Rome 7542 32 cores/socket, 2.9GHz. 2 sockets per node
- **SX-Aurora TSUBASA:** VE10AE x8 / VH (single socket Rome)
- **ICON-ART:** Status as of 2019 for ICON-ART

- 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

1. Features of SX-Aurora TSUBASA
2. World famous HPC Centers utilizing SX-Aurora TSUBASA
3. Aurora System Software
4. Value of Vector Engine
5. Roadmap and the future

Vector Engine 3.0

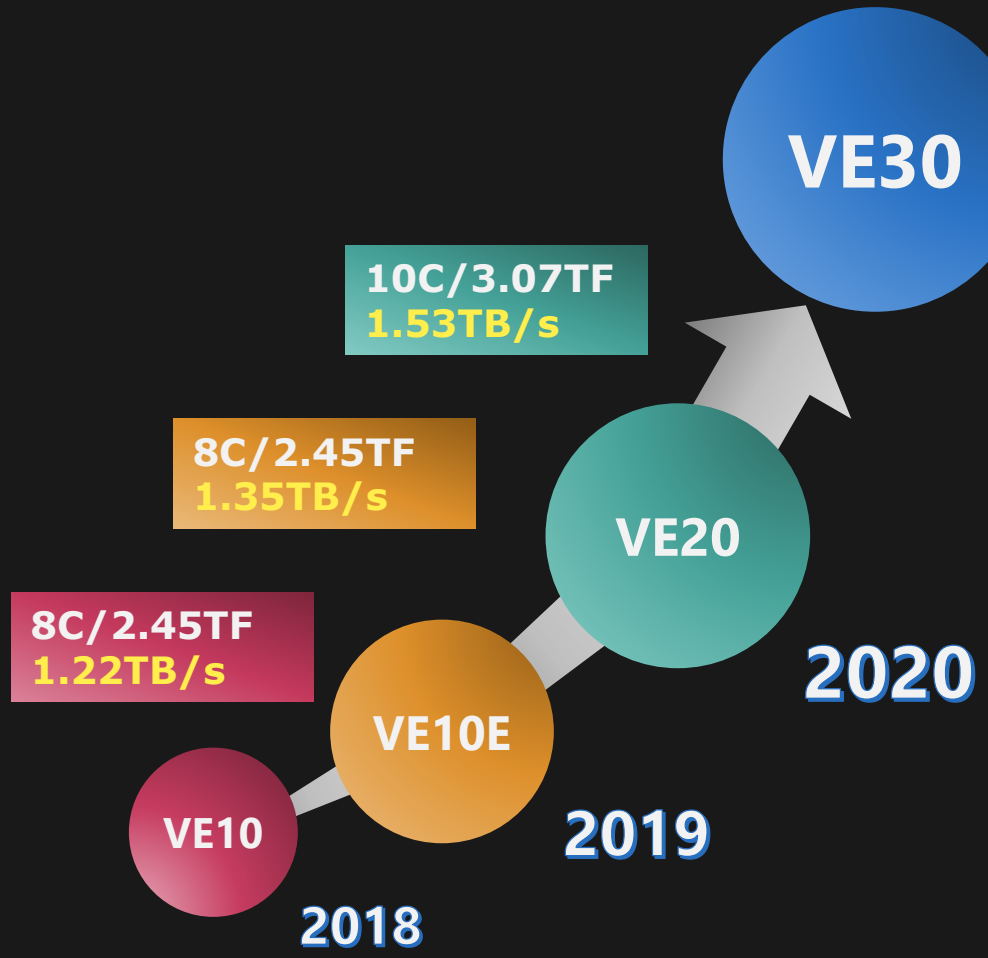

2+TB/s memory bandwidth

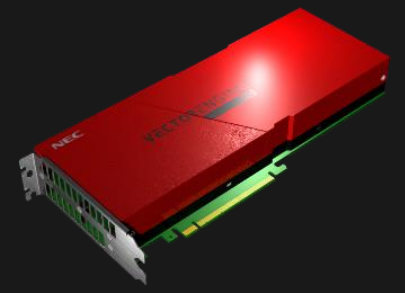

- **Targeting the largest memory bandwidth**
- Inheriting and improving VE/VH architecture
- Higher Flops per processor
- Improved memory subsystem including cache
- Accelerating short vector, and scalar operations
- Maintaining high power efficiency
- Virtual machine support

Roadmap

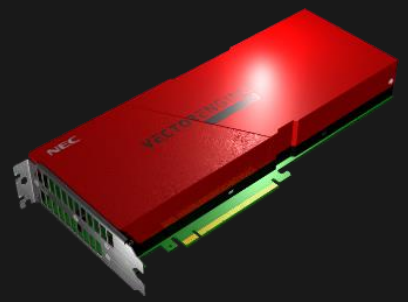

Memory bandwidth / VE

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

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

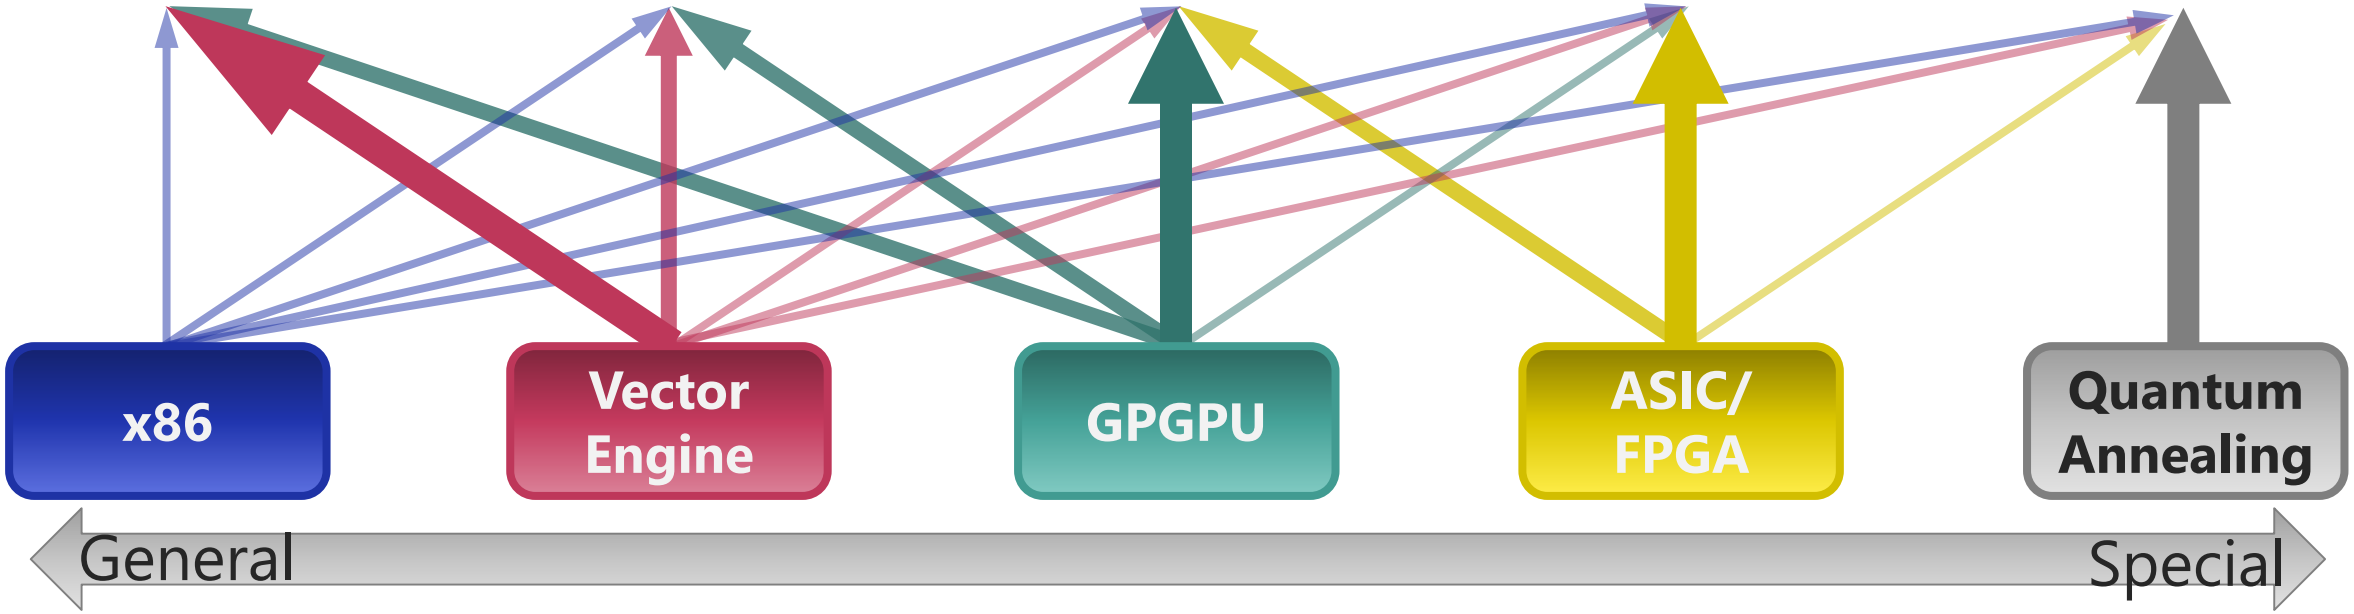

Aurora is ready for your future

- Aurora HPCG performance efficiency is better than the other machines.
  - The Fraction of Peak is over 5%. This will be one of the advantages of Aurora system in your future, especially when you want to build your LARGE computing system.

- Aurora can easily collaborate with the other machines, such as GPGPU, etc.
  - Through interconnect, right now Infiniband
    - Hybrid VE-x86/GPGPU MPI is available right now.

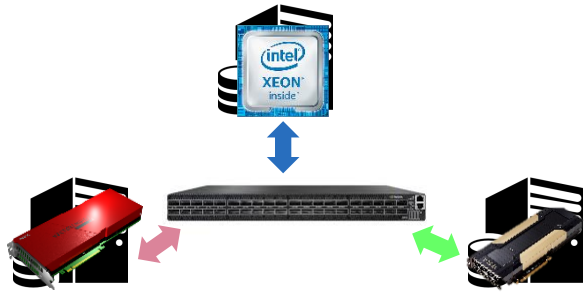

  - In VH (Linux/x86)
    - Not only VEs, but GPGPU, etc. can be installed in the same VH.
      - Hybrid VE-x86 MPI through PCIe is available right now.
      - Hybrid VE-GPGPU MPI through PCIe will be available in 2022.
      - NLCPy/Aurora with GPGPU will be available in 2022.

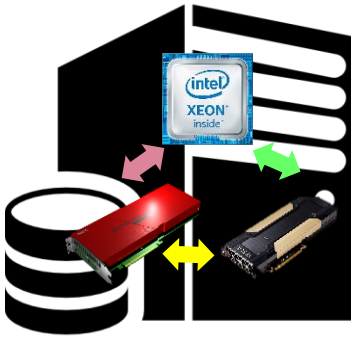

MPI communication on multi-architectural supercomputer

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

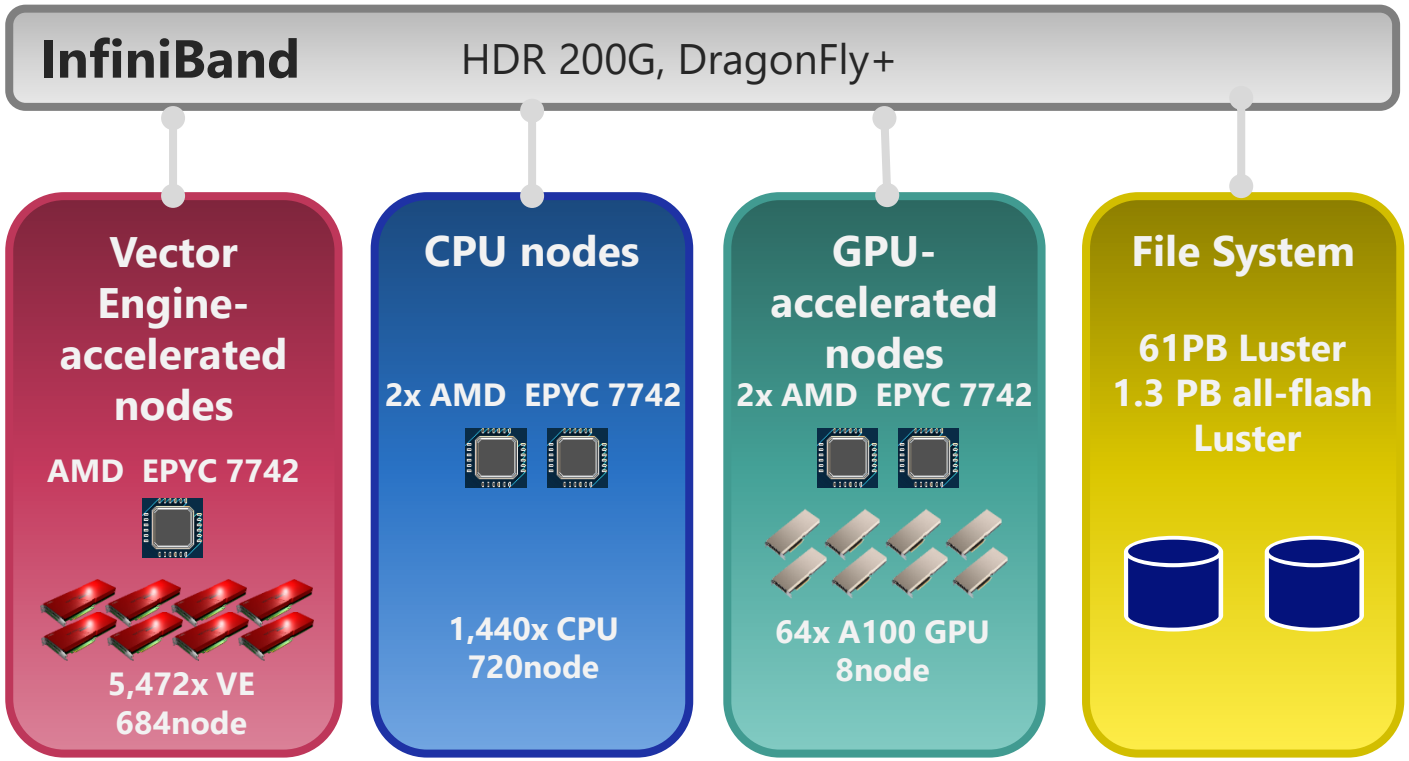

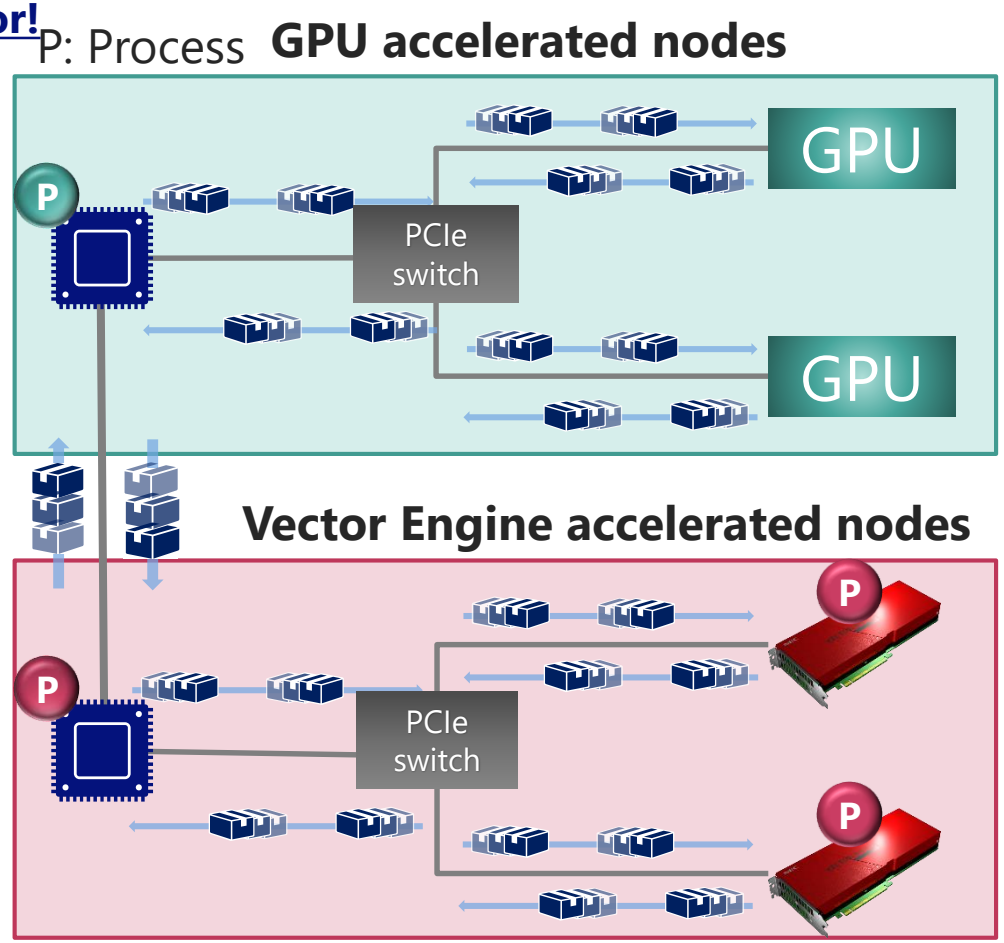

https://www.conferenceharvester.com/uploads/harvester/VirtualBooths/13396/NKBNOCXO-PDF-1-412693%285%29.pdf

NLCPy for Heterogeneous Computing, without MPI

- NLCPy will support VE-GPU heterogeneous computing in April 2022. (Preview)
  - User will be able to choose execution devices just by specifying Python decorators.

```
import nlcpy as vp

@vp.device('VE')
def compute(x, y):
  return x + x * y

x = vp.random.rand(1000)
y = vp.random.rand(1000)
z = compute(x,y)
```

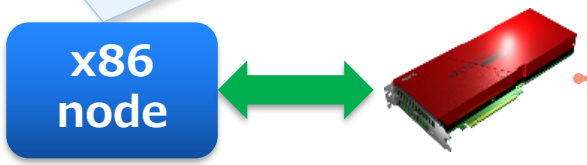


```
import nlcpy as vp

@vp.device('GPU')
def compute(x, y):
  return x + x * y

x = vp.random.rand(1000)
y = vp.random.rand(1000)
z = compute(x,y)
```

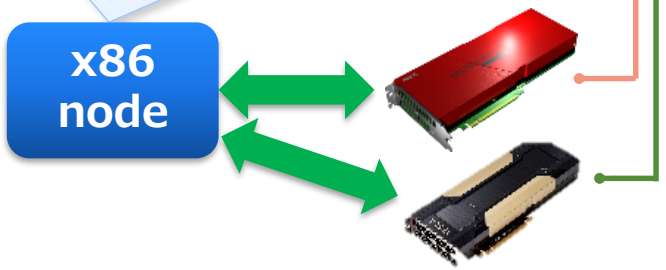

Find more information on our website

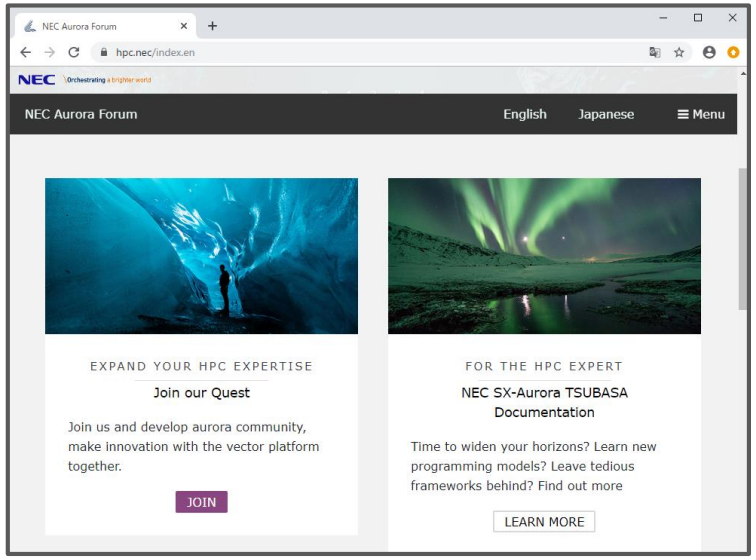

Aurora Web Forum

http://www.hpc.nec

- Latest updates
- Manual, documents
- Bulletin board

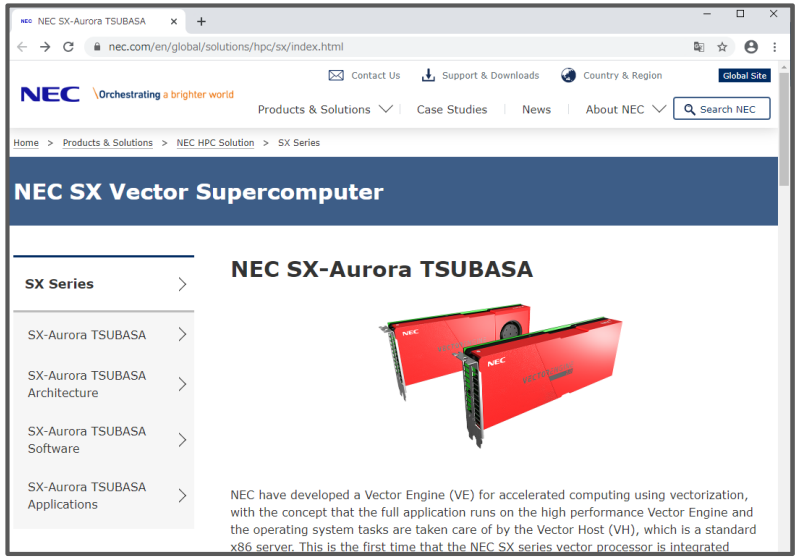

SX-Aurora TSUBASA Website

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

- Hardware and software overview
- Supported applications

✉ 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