Investigation of Azure CycleCloud Usage Environment and Considerations on Supercomputer Center-Cloud Collaboration

Toru Nagai [1)], Shuichi Gojuki [2)], Masatoshi Kawai [1)], Takahiro Katagiri [1)], Tetsuya Hoshino [1)]
1) Information Technology Center, Nagoya Univ.
2) Microsoft Japan Co., Ltd.

Abstract We conducted a performance measurement of virtual machines targeting Microsoft Azure with the main purpose of investigating the usage environment of public clouds. Specifically, we used Azure CycleCloud, which is specialized for HPC usage environment and executed various benchmark programs on virtual machines. We report on the usage environment of Azure CycleCloud and the benchmark test results on virtual machines. We also discuss the collaboration between supercomputer systems and public clouds.

Azure CycleCloud

- build HPC environment
  - ‘VM size‘ refers to properties or resources assigned to a virtual machine such as CPU, main memory, GPU, and network bandwidth.
  - Based on the VM size, it orchestrates VMs and storage and deploys a job scheduler-configured cluster to an existing network.
  - It also installs specified software at the environment construction.

HPC environment built with CycleCloud

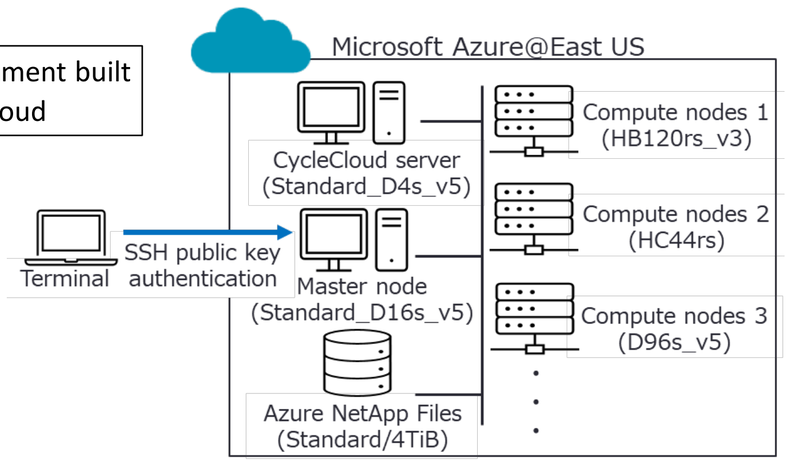

Benchmark test results (1/2)

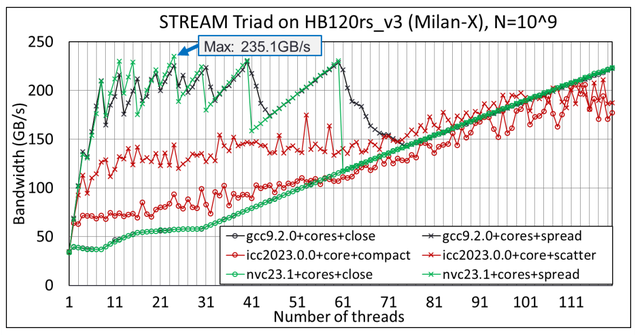

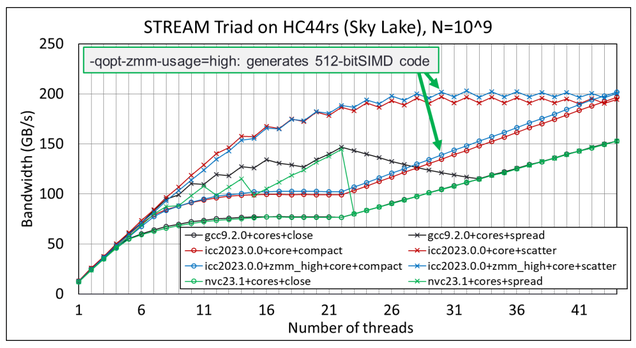

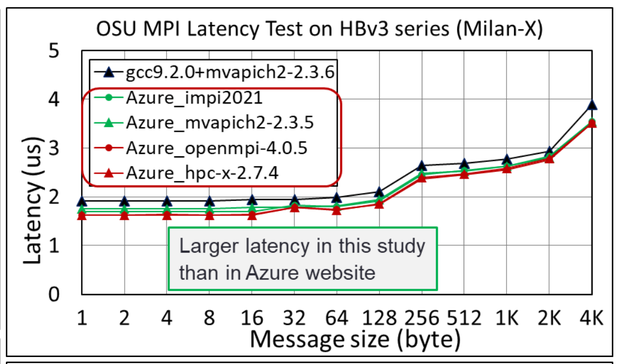

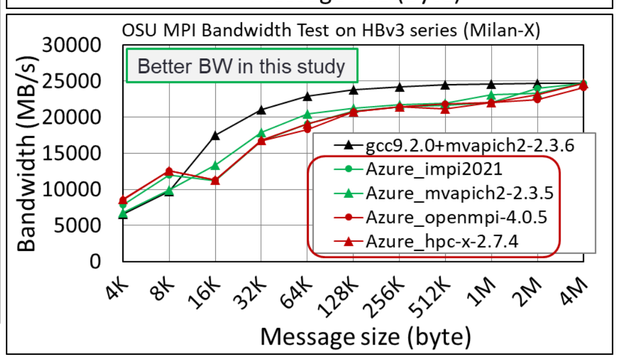

Benchmark test results (2/2)

- BICCG (Krylov subspace method with a block incomplete Cholesky preconditioner) solver as a practical application

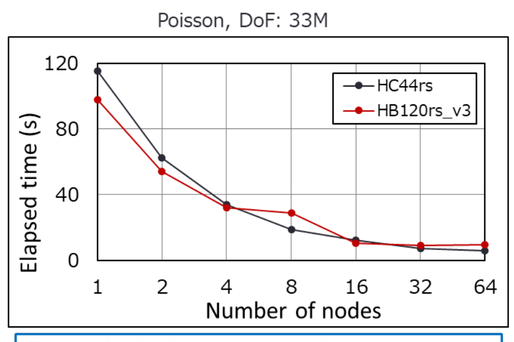

The application is latency-sensitive due to a small message size. Based on the results bellow, it is speculated that HC44rs has a latency advantage over HB120rs_v3.

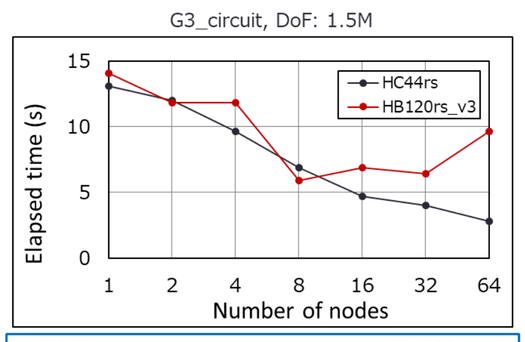

- HC44rs has a latency advantage over HB120rs_v3.
- The variability in HB120rs_v3 in the figure below may have been significantly influenced by the network distance between the nodes deployed.

VM size	CPU	# of sockets	# of cores	# of HW threads	L3 cache size	Main memory (GiB)	InfiniBand
HC44rs	Intel Xeon Platinum 8168 (Sky Lake)	2	44	44	33MB x 2 sockets	352	100Gbps EDR
HB120rs_v3	AMD EPYC 7V73X (Milan-X)	2	120	120	768MB x 2 sockets	448	200Gbps HDR

Supercomputer Center-Cloud Collaboration (1/2)

Case 1: Cloud burst, or users can utilize the system without being aware of whether their submitted jobs are being executed in an on-premises environment or in a public cloud.
Case 2: Users actively utilize the public cloud.

#	Technical requirements	Case 1	Case 2
1	Executables and shell scripts created on premises can be used as is in the public cloud.	OK	OK
2	Input files are automatically transferred to the public cloud.	OK	OK
3	Outputs can be obtained on premise.	OK	OK
4	Results are consistent between the two.	OK	OK
5	Security equivalent to or higher than on premise	OK	OK
6	Accounting information is available.	OK	OK
7	The execution time is approximately the same or less.	NG	OK

- #1-3: It is OK, only if the aligned software execution environments on both the on-premise and public cloud, and each other’s file systems mutually mounted over the network are available.
- #7: Optimization for the cloud environment is difficult in Case 1.
- The user manual is essential in Case 2.

Supercomputer Center-Cloud Collaboration (2/2)

Pricing

Converting one dollar as 145 yen

VM size		CPU	# of HW threads	NVIDIA GPU	Interconnect between nodes	Fee (yen/hour)			
						Pay-as-you-go		Spot	
						East US	East Japan	East US	East Japan
HB120rs_v3		AMD EPYC 7V73X (Milan-X)	120	–	200Gbps HDR	522.0	756.9	52.2	75.7
HC44rs		Intel Xeon Platinum 8168 (Sky Lake)	44	–	100Gbps EDR	459.4	666.1	45.9	66.6
NC96ads_A100_v4		AMD EPYC 7V12 (Milan)	96	A100 x 4	–	2130.3	3088.9	672.3	1202.8
NC24s_v3		Intel Xeon E5-2690 v4 (Broadwell)	24	V100 x 4	–	1774.8	2432.5	826.7	852.6
D96s_v5		Intel Xeon Platinum 8370C (Ice Lake)	96	–	–	668.2	863.0	109.6	87.2
Flow	Type I	Fujitsu A64fx	48	–	Tofu D	31.0			
	Type II	Intel Xeon Gold 6230 (Cascade Lake)	40	V100 x 4	200Gbps EDR	155.1			
	Cloud	Intel Xeon Gold 6230 (Cascade Lake)	80	–	100Gbps EDR	44.3			

- While the usability is not good, it seems possible to use the spot (excluding GPU usage).
- There is also room to consider a plan that allows us to use VMs cheaper than pay-as-you-go by committing to use for a certain period and contracting.

- In Azure, utilization fees reflect the costs of hardware, facility, electricity, and operation expenses. In contrast, with "Furō," only the electricity and part of operation expenses are considered in the usage fees, making it evident that Azure's pricing is higher.
- Let us roughly estimate the hardware, facility, and personnel costs related to "Furō."

Newly considered expenses in “Flow”

	Million yen/yr (%)
Computer rental fee	900.0 (91.6)
Depreciation expense	47.2 (4.8)
Personnel expenses	35.3 (3.6)
Total	982.5 (100.0)

- Although the estimated expenses for depreciation and personnel are highly arbitrary, the computer rental fee accounts for more than 90% of the total.

“Flow” usage fees

	Yen/node･hr		
	Current	Newly calculated	Total
Type I	31.0	72.0	103.0
Type II	155.1	343.9	499.0
On-premise Cloud	44.3	101.1	145.4

- Comparing with the Azure fees, it is found that spot pricing can sometimes be cheaper.
- One thing not considered in the above estimation is the so-called academic discount on the computer rental fee.