NOR-TECH

People Friendly Technology

Highlights of NASA Technical Report on Performance and Cost of Cloud v. On-Premises HPC Solutions[1]

Full NASA Study Available at: https://www.nas.nasa.gov/assets/pdf/papers/NAS_Technical_Report_NAS-2018-01.pdf

Executive Summary

NASA's High-End Computing Capability (HECC) Project is periodically asked if it could be more cost effective through the use of commercial cloud resources. To answer the question, HECC's Application Performance and Productivity (APP) team undertook a performance and cost evaluation comparing three domains: two commercial cloud providers, Amazon and Penguin, and HECC's in-house resources—the Pleiades and Electra systems.

All runs on HECC resources were faster, and sometimes significantly faster, than runs on the most closely matching Amazon Web Services (AWS) resources. In all cases, the full cost of running on HECC resources was less than the lowest possible compute-only cost of running on AWS.

Significant Findings

1. To run the full set of the NPBs, AWS was 5.8-12 times more expensive than HECC, depending on the processor type used. For the full-sized applications, AWS was in the best case 1.9 times more expensive.
2. The NPB runs at POD were 4.7 times more expensive than equivalent runs at HECC. The full sized applications were 5.3 times more expensive.
3. Tightly-coupled, multi-node applications from the NASA workload take somewhat more time when run on cloud-based nodes connected with HPC-level interconnects; they take significantly more time when run on cloud-based nodes that use conventional, Ethernet-based interconnects.
4. Commercial clouds do not offer a viable, cost-effective approach for replacing in-house HPC resources at NASA.

Conclusion

Results show that large applications with tightly coupled communications perform worse on cloud resources than on similar resources at HECC. In addition, per-hour use of cloud resources is more expensive than the full cost of using similar resources at HECC. Taken in combination, the data leads to the conclusion that:

"Commercial clouds do not offer a viable, cost-approach for replacing in-house HPC resources at NASA."

[1] NAS Technical Report NAS-2018-01;Evaluating the Suitability of Commercial Clouds for NASA's High Performance Computing Applications: A Trade Study. S. Chang, R. Hood, H. Jin, S. Heistand, J. Chang1, S. Cheung, J. Djomehri, G. Jost, D. Kokron; NASA Advanced Supercomputing Division, NASA Ames Research Center

Study's Performance and Cost Comparison[2]

Benchmark	NCPUS	# of HECC Skylake nodes	# of AWS c5.18xlarge (Skylake) instances	HECC time (sec)	HECC full cost	AWS time (sec)	AWS Oregon compute cost
bt.D	256	7	8	79.83	$0.16	146.73	$1.00
bt.D	1024	26	29	20.52	$0.15	92.08	$2.27
cg.D	256	7	8	34.3	$0.07	614.2	$4.18
cg.D	512	13	15	17.24	$0.06	650.58	$8.29
cg.D	1024	26	29	13.51	$0.10	719.22	$17.73
ep.D	256	7	8	4.82	$0.01	4.93	$0.03
ep.D	512	13	15	2.44	$0.01	2.46	$0.03
ep.D	1024	26	29	1.24	$0.01	1.19	$0.03
ft.D	256	7	8	27.41	$0.05	526.08	$3.58
ft.D	512	13	15	14.84	$0.05	349.26	$4.45
ft.D	1024	26	29	8.05	$0.06	243.06	$5.99
is.D	256	7	8	2.64	$0.01	71.5	$0.49
is.D	512	13	15	1.45	$0.01	48.28	$0.62
is.D	1024	26	29	0.84	$0.01	44.72	$1.10
lu.D	256	7	8	57.58	$0.11	121.6	$0.83
lu.D	512	13	15	32.37	$0.12	106.79	$1.36
lu.D	1024	26	29	17.17	$0.13	99.3	$2.45
mg.D	256	7	8	8.18	$0.02	39.99	$0.27
mg.D	512	13	15	3.36	$0.01	15.56	$0.20
mg.D	1024	26	29	1.82	$0.01	16.44	$0.41
sp.D	256	7	8	101.52	$0.20	211.48	$1.44
sp.D	1024	26	29	20.32	$0.15	172.02	$4.24
Total Cost					**$1.50**		**$60.98**
Estimated AWS spot cost (30% of on-demand cost)							**$18.29**

Benchmark	Case	NCPUS	# of HECC Haswell nodes	# of AWS c4.8xlarge (Hasewell) instances	HECC time (sec)	HECC full cost	AWS time (sec)	AWS Oregon compute cost	AWS Gov compute cost
ATHENA++	SBU2	1024	43	57	2268	$14.48	2298	$57.89	$69.68
ATHENA++	SBU2	2048	86	114	1177	$15.03	1374	$69.22	$83.32
ECCO	NTR1	120	5	7	120	$0.09	173	$0.54	$0.64
ECCO	NTR1	240	10	14	65	$0.10	140	$0.87	$1.04
ENZO	SBU2	196	9	11	1827	$2.44	2266	$11.02	$13.26
FVCore	SBU1	1176	49	66	1061	$7.72	1104	$32.20	$38.76
nuWRF	SBU2	1700	71	95	529	$5.58	1302	$54.66	$65.80
OpenFOAM	Channel395	48	2	3	4759	$1.41	7646	$10.14	$12.20
OpenFOAM	Channel395	144	6	8	12547	$11.17	20771	$73.44	$88.39
OpenFOAM	Channel395	288	12	16	10194	$18.16	23013	$162.73	$195.87
Total Cost						**$76.18**		**$472.71**	**$568.96**
Estimated AWS spot cost (30% of on-demand cost)								**$141.77**	
Estimated AWS pre-leasing cost (70% of US-gov cost)									**$398.27**

[2] Appendix: NAS Technical Report NAS-2018-01;Evaluating the Suitability of Commercial Clouds for NASA's High Performance Computing Applications: A Trade Study.

Nor-Tech's Take

1. The number one consideration is workload. If, as with most HPC systems, the overall workload is more than 80%, Nor-Tech calculated the ROI for on-premises computing to be about nine months.
2. Application software is significantly more expensive than hardware. Cloud provider's fees are usually based on time used--the longer it takes to run in the cloud, the more it costs. So the faster the job solves, the less the software cost for that job.
3. Taking into consideration that lifecycles for on-premises systems are a minimum of three years, there is substantial savings in deploying on-premises over cloud even factoring in ongoing power, cooling, and management costs of on-premises computing.
4. Frequency is the most important factor for performance in HPC computing. Generally, lower core count processors have faster frequencies. Cloud providers use high core count processors to improve overall computing density.
5. Some cloud providers only use high speed Ethernet verses the lower latency InfiniBand typically found in HPC solutions. InfiniBand offers a significant performance advance over Ethernet in many HPC applications.
6. Ultimately NASA concluded that commercial clouds do not offer a viable, cost-approach for replacing in-house, on-premises HPC resources.

Why Intel® Xeon® for On-Premises HPC Solutions?

Nor-Tech, an Intel HPC Data Center Specialist and Intel Platinum Partner, integrates Intel's Xeon Scalable processors into high performance servers and clusters. The many benefits of Intel's new 3rd Gen Xeon Scalable processors include:

- Faster time to value with Intel Select Solutions
- Strong, capable platforms for the data-fueled enterprise
- Next-generation platform for cloud-optimized, 5G-ready networks, and next-generation virtual networks
- Breakthrough HPC and high-performance data analytics innovation

Intel's 3rd Generation Xeon Scalable processors with the Intel Optane persistent memory 200 series and Intel Deep Learning Boost (Intel DL Boost) deliver groundbreaking platform innovations for digital transformation, featuring:

- Built-in AI Acceleration: Get faster insights from data-intensive workloads with built-in AI acceleration and massive memory capacity.
- Trusted Protection: Take advantage of multilayer security that helps optimize service delivery and thwart malicious exploits.
- Enhanced Platform: Consistently deliver amazing experiences with hardware-enhanced virtualization across compute, network, and storage.

These new processors are the next step above the previous-generation 4 to 8-socket processor foundation, designed and built for today's AI-infused, data-intensive workloads.

intel Technology Provider Platinum 2020 HPC Data Center Specialist

intel Technology Provider Platinum 2020

Discover all the benefits of Intel® 3rd Gen Xeon® Scalable Processor and download the brief.
https://www.intel.com/content/www/us/en/products/docs/processors/xeon/3rd-gen-xeon-scalable-processors-brief.html

NOR-TECH
People Friendly Technology