PROWESS

Technical Research Report

OpenRadioss on AWS®: HPC Workload Testing

Prowess Consulting ran up to 92% more OpenRadioss Ford® Taurus 10M simulations per day by choosing Amazon Web Services® (AWS) C6i instances with Intel, as compared to AWS C6g instances with AWS® Graviton2.

Executive Summary

OpenRadioss is a finite element solver for dynamic event analysis. It simulates how materials interact and respond to outside influences. This can help engineers model events like car crashes, a bridge straining under a heavy load, or even a phone dropping on the floor. As a high-performance computing (HPC) workload, OpenRadioss requires powerful CPUs, large amounts of memory, high-speed interconnects, and fast storage.

In a study sponsored by Intel, Prowess Consulting ran OpenRadioss on Amazon Web Services® (AWS®) instances powered by both Intel® Xeon® processors and AWS Graviton® processors—both compute-optimized C6 and C7 generations. We chose OpenRadioss for this analysis because it has similar infrastructure requirements to data analytics, AI, and other HPC workloads, and our findings can be extrapolated to these workloads. We tested two OpenRadioss models—the Chrysler® Neon 1M model and the Ford® Taurus 10M model—both of which are popular among engineers.

Our approach was to identify which AWS instances offer the best value. We know development teams must balance the project-management variables of schedule, cost, and quality. We focused on the tradeoffs between schedule and cost by measuring performance and performance per dollar.

Oftentimes, instances are chosen based on core count and hourly instance pricing. We took a different approach, including the cost of engineering time. We also explored instances with different vCPU values, in addition to configuring instances with 64 physical cores.

We measured performance as the number of simulations that an engineer could run in one eight-hour day, and we called this variable "simulation throughput." Prowess Consulting found that engineers can achieve up to 92% higher simulation throughput by choosing the AWS C6i instances [C6i.32xlarge] with Intel shown in Table 1, as compared to the AWS C6g instances [C6g.16xlarge] with Graviton2. For this specific comparison, we set the 64-vCPU C6i Intel instance OpenMP® thread to one to match the Graviton2 processor, which is single-threaded. With this configuration, the 64-vCPU C6i instance operates with 64 physical cores, which matches the physical core count of Graviton2. We used our research findings to create a checklist for maximizing project value.

Checklist to Maximize Project Value

_Identify instances that maximize performance as defined by your project.

_Compute the project cost using fully burdened costs that include engineering hours.

_Pull in the project schedule by using instances that deliver higher performance per dollar.

_Look for instances that make it easy to scale.

Table 1 | Performance (simulations per day) on the OpenRadioss Chrysler® Neon 1M and Ford® Taurus 10M models

	Chrysler® Neon 1M Model	Ford® Taurus 10M Model
C6 generation	The daily OpenRadioss Neon 1M simulation throughput on the 64 physical-core C6 instance with Intel [C6i.32xlarge] is **1.92x higher** than the 64-vCPU C6 instance with Graviton2 [C6g.16xlarge].	The daily OpenRadioss Taurus 10M simulation throughput on the 64 physical-core C6 instance with Intel [C6i.32xlarge] is **1.92x higher** than the 64-vCPU C6 instance with Graviton2 [C6g.16xlarge].
C7 generation	The daily OpenRadioss Neon 1M simulation throughput on the 64 physical-core C7 instance with Intel [C7i.48xlarge] is **1.61x higher** than the 64-vCPU C7 instance with Graviton3E [C7gn.16xlarge].	The daily OpenRadioss Taurus 10M simulation throughput on the 64 physical-core C7 instance with Intel [C7i.48xlarge] is **1.40x higher** than the 64-vCPU C7 instance with Graviton3E [C7gn.16xlarge].

To understand the total costs of these high-performing instances, we measured performance per price. We defined performance per price as the fully burdened cost per simulation times the number of simulations per day. Our fully burdened cost included engineering costs.

Prowess Consulting found that the AWS instances with Intel that deliver the highest number of simulations per day cost up to 7.4% more than the corresponding AWS instances with Graviton. The actual incremental cost per day is $36.60.[1] Considering this low cost, Prowess Consulting recommends using AWS 64 physical-core C6i and C7i instances with Intel for HPC, analytics, and AI workloads when schedule is a priority.

More About OpenRadioss

OpenRadioss is an open-source finite element analysis (FEA) software package designed to simulate and analyze the behavior of structures under a wide range of loading conditions. It's based on the Altair® Radioss® solver, a commercial FEA software package developed by Altair Engineering.[2] These models help engineers understand what happens when a vehicle crashes—that is, how different parts of a vehicle deform, how the occupants are affected, and how safety features perform.

We used both OpenRadioss Chrysler Neon 1M and Ford Taurus 10M element models for our testing. The Neon 1M model simulates the crash of a vehicle with one million finite elements, while the Taurus 10M model simulates the crash of a vehicle with 10 million finite elements. The Neon 1M model is designed for testing an HPC cluster with a low number of nodes or for testing on a single compute server. The Taurus 10M model is designed for scalability testing on HPC clusters with large numbers of nodes.

Instance Selection

Amazon® Elastic Compute Cloud (Amazon EC2®) provides a wide selection of instance types optimized to fit different use cases. We selected compute-optimized instances (designated by the letter C) that offer high-performance processors. See Table 2 for our selections.

Table 2 | Instance families used in testing

Instance family	C6i	C7i	C6g	C7gn[3]
Processor	3rd Gen Intel® Xeon® Scalable processors	4th Gen Intel Xeon Scalable processors	Arm®-based AWS® Graviton2 processors	Arm®-based AWS Graviton3E processors with up to 200 Gbps of network bandwidth

We selected four C6 AWS instances, as shown in Table 3, and four C7 AWS instances, as shown in Table 4. Our intent was to measure performance and performance per price for the C6 generation instances (comparing instances powered by 3rd Gen Intel Xeon Scalable processors to instances powered by Arm®-based AWS® Graviton2 processors) and the C7 generation instances (comparing instances powered by 4th Gen Intel Xeon Scalable processors to instances powered by Arm-based Graviton3E processors).

For one of these instances, we selected the 64-vCPU C6 instance with Graviton2 (Arm v8), the largest compute-optimized instance available with Graviton2. Our goal was to exercise this instance's vCPUs fully. For a comparable Intel instance, we selected the 64-vCPU C6 instance with the Intel Xeon Platinum 8375C processor. We also selected a 96-vCPU C6i instance to gain additional insight into the effect of adding more vCPUs, and we configured a C6i.32xlarge instance to have 64 physical cores.

Table 3 | AWS® C6 instances used to test OpenRadioss

	64-vCPU C6i	96-vCPU C6i	64 Physical-Core C6i (Configured for One Thread)	64-vCPU C6g
Model name	C6i.16xlarge	C6i.24xlarge	C6i.32xlarge	C6g.16xlarge
CPU	Intel® Xeon® Platinum 8375C CPU			AWS® Graviton2 (Arm® v8) CPU
Number of virtual CPUs (vCPUs)	64	96	Not applicable (N/A); configured for 64 physical cores	64
Memory	128 GB	192 GB	386 GB	128 GB
High-performance block-storage service used with Amazon® Elastic Compute Cloud (Amazon EC2®)	Amazon® Elastic Block Store (Amazon® EBS)			
Amazon EBS volume type	gp2			
Operating system (OS)	Red Hat® Enterprise Linux®			Ubuntu®
OS version	9			22.04
OS kernel	5.14.0-284.11.1.el9_2.x86_64			5.19.0-1025-aws

Our engineers selected the 64-vCPU C7g instance with Graviton3E (Arm v8) so that we could explore the performance of the Graviton3E processor. For a comparable Intel instance, we selected the 64-vCPU C7i instance with the Intel Xeon Platinum 8488C processor. We also selected a 96-vCPU C7i instance to gain additional insight into the effect of adding more vCPUs, and, as with the C6 family, we selected a 96-vCPU C7i instance and configured it to provide 64 physical cores.

Table 4 | AWS® C7 instances used to test OpenRadioss

	64-vCPU C7i	96-vCPU C7i	64 Physical-Core C7i (Configured for One Thread)	64-vCPU C7gn
Model name	C7i.16xlarge	C7i.24xlarge	C7i.48xlarge	C7gn.16xlarge
CPU	Intel® Xeon® Platinum 8488C CPU			AWS® Graviton3E (Arm® v8) CPU
Number of virtual CPUs (vCPUs)	64	96	N/A; configured for 64 physical cores	64
Memory	128 GB	192 GB	256 GB	128 GB
High-performance block-storage service used with Amazon® Elastic Compute Cloud (Amazon EC2®)	Amazon® Elastic Block Store (Amazon® EBS)			
Amazon EBS volume type	gp2			
Operating system (OS)	Red Hat® Enterprise Linux®			Ubuntu®
OS version	9			22.04
OS kernel	5.14.0-284.11.1.el9_2.x86_64			5.19.0-1025-aws

Simulation Results

After the setup and debug process, we ran three simulations on each instance. Tables 5 and 6 show the total median time in seconds to run the simulations.

Table 5 | Total median time to run a simulation in seconds on C6 instances

	64-vCPU C6i	96-vCPU C6i	64 Physical-Core C6i	64-vCPU C6g
Model name	C6i.16xlarge	C6i.24xlarge	C6i.32xlarge	C6g.16xlarge
Chrysler® Neon 1M total simulation time (seconds)	3,915.94	2,422.07	2,202.54	4,243.35
Ford® Taurus 10M total simulation time (seconds)	20,376.67	12,620.16	10,889.93	20,944.26

Table 6 | Total median time to run a simulation in seconds on C7 instances

	64-vCPU C7i	96-vCPU C7i	64 Physical-Core C7i	64-vCPU C7gn
Model name	C7i.16xlarge	C7i.24xlarge	C7i.48xlarge	C7gn.16xlarge
Chrysler® Neon 1M total simulation time (seconds)	2,922.12	2,714.66	1,732.4	2,791.59
Ford® Taurus 10M total simulation time (seconds)	15,867.52	13,339.09	9,494.96	13,372.9

Reducing Development Schedules

Engineering productivity is critical in competitive business environments. An automotive manufacturer, for example, must quickly complete simulations to meet schedule deadlines. In a simple scenario, engineers might perform two dozen simulations to test different design iterations or assess a particular component's behavior. In a complex scenario, engineers might require hundreds of simulations to explore conditions, failure modes, or design variations.

We measured the number of simulations an engineer could run in one eight-hour workday to assess productivity and how the choice of instance could affect a project schedule. We assumed that the simulations were executed back-to-back without downtime. We also assumed that the last simulation of the day was started by an engineer and then allowed to run to completion without supervision. We calculated simulations per day by taking the number of seconds in one eight-hour day (28,800) and dividing it by the total simulation time, in seconds. Simulations per day are shown in Tables 7 and 8. These results were calculated by dividing 28,800 seconds by each simulation time in Tables 5 and 6.

Table 7 | Productivity calculations (simulations/day) based on simulation time for C6 instances

	64-vCPU C6i	96-vCPU C6i	64 Physical-Core C6i	64-vCPU C6g
Model name	C6i.16xlarge	C6i.24xlarge	C6i.32xlarge	C6g.16xlarge
Simulations/day (Chrysler® Neon 1M)	7.35	11.89	13.07	6.78
Simulations/day (Ford® Taurus 10M)	1.41	2.28	2.64	1.37

Table 8 | Productivity calculations (simulations/day) based on simulation time for C7 instances

	64-vCPU C7i	96-vCPU C7i	64 Physical-Core C7i	64-vCPU C7gn
Model name	C7i.16xlarge	C7i.24xlarge	C7i.48xlarge	C7gn.16xlarge
Simulations/day (Chrysler® Neon 1M)	9.85	10.60	16.62	10.31
Simulations/day (Ford® Taurus 10M)	1.81	2.15	3.03	2.15

Using a 64 physical-core C6 instance with Intel [C6i.32xlarge], our engineers were able to run 92% more simulations per day, compared to the 64-vCPU C6 instance with Graviton2 [C6g.16xlarge].

Other top-performing Intel instances yielded the following results:

- **60% more simulations.** The daily OpenRadioss Neon 1M simulation throughput on the 64 physical-core C7 instance with Intel [C7i.48xlarge] was **61% higher** than the 64-vCPU C7 instance with Graviton3E [C7gn.16xlarge].
- **40% more simulations.** We found that the daily OpenRadioss Taurus 10M simulation throughput on the 64 physical-core C7 instance with Intel [C7i.48xlarge] was **40% higher** than the 64-vCPU C7g instance with Graviton3E [C7gn.16xlarge].

For the full set of results, see the Appendix.

Assessing Performance per Dollar

We included instance price and engineering labor costs in our measurement of performance per dollar. To compute a fully burdened cost per hour to run OpenRadioss simulations, we used AWS on-demand instance pricing and industry-standard engineering salaries (see Tables 9 and 10).

Our detailed assumptions on engineering labor costs are outlined in Appendix Table A1. Appendix Table A3 and Table A4 detail the cost of running simulations over an eight-hour day for each instance.

Reviewing these costs, Prowess Consulting determined that the AWS C6i and C7i instances with Intel that provide the highest number of simulations per day cost up to 7.4% more than the AWS C6g and C7gn instances with AWS Graviton2 and Graviton3E. Given the pricing used for this report, a 7.4% increase equates to $36.60/day.[1]

Analysis

When running HPC workloads, organizations want to complete the simulation project quickly to meet engineering milestones. They also want to minimize the project cost by increasing engineering productivity.

When comparing hourly instance pricing, the lowest price options are the 64-vCPU C6g instance with Graviton2 and the 64-vCPU C7g instance with Graviton3E. However, considering other factors like daily simulation throughput and productivity per dollar led us to approach our instance selection process differently.

In the case of OpenRadioss, an engineer might need to run simulations continuously over days or weeks to solve a problem. In this scenario, looking at simulation throughput and engineering productivity per dollar reveals that a 64 physical-core C6 instance with Intel [C6i.32xlarge] or a 64 physical-core C7 instance with Intel [C7i.48xlarge] can deliver better business results than a 64-vCPU C6 instance with Graviton2 or a 64-vCPU C7 instance with Graviton3E, respectively.

Based on these insights, we found several takeaways:

- Engineers can achieve up to 92% higher performance, as measured by the number of simulations per day, by choosing AWS C6i and C7i instances with Intel, as compared to AWS C6g and C7gn instances with Graviton2 and Graviton3E.
- The fully burdened daily cost of using high-performance AWS instances with Intel to run simulations is up to 7.4% higher than the cost of using AWS instances with Graviton2/Graviton3E. Given the pricing used for this report, a 7.4% increase equates to up to $36.60 per day, which might be a bargain given the performance increase.[1]

Recommendation

Prowess Consulting recommends looking beyond hourly instance pricing when selecting AWS instances for HPC applications like OpenRadioss and other compute-intensive workloads, such as data analytics and AI. Account for engineering time, in addition to hourly instance pricing, to improve business decisions, especially when minimizing schedule is a priority.

Learn More

Discover more about Intel on AWS by visiting "When to select Intel over Graviton."
Give your developers the tools they need to develop in the cloud.
See more research reports by Prowess Consulting.

Appendix

Table A1 | Engineering labor-cost assumptions

Engineering salary/year[4]	Hours worked/year[5]	Engineering rate/hour
$120,191	2,087	$57.5903

Table A2 | AWS® pricing[6]

Instance	Model name	Price/hour ($)
64-vCPU C6i	C6i.16xlarge	2.720
96-vCPU C6i	C6i.24xlarge	4.080
64-vCPU C6g	C6g.16xlarge	2.176
64-vCPU C7i	C7i.16xlarge	2.856
96-vCPU C7i	C7i.24xlarge	4.284
64-vCPU C7gn	C7gn.16xlarge	3.994
64 physical-core C6i	C6i.32xlarge	5.440
64 physical-core C7i	C7i.48xlarge	8.568

Table A3 | Calculating costs for Chrysler® Neon 1M simulations run per day (numbers rounded to the nearest hundredth)

Instance type	AWS® price/hour ($)[6]	Number of hours to run one simulation	Cost to run one simulation	Engineering labor cost per run ($57.59/hr)[7]	Fully burdened cost (instance plus labor)/one simulation	Cost for simulations run per day
64-vCPU C6i.16xlarge	$2.72	1.08	$2.96	$62.64	$65.60	$482.48
96-vCPU C6i.24xlarge	$4.08	0.67	$2.75	$38.75	$41.49	$493.36
64 physical-core C6i.32xlarge	$5.44	0.61	$3.33	$35.23	$38.56	$504.24
64-vCPU C6g.16xlarge	$2.18	1.17	$2.56	$67.88	$70.45	$478.13
64-vCPU C7i.16xlarge	$2.86	0.81	$2.32	$46.75	$49.06	$483.57
96-vCPU C7i.24xlarge	$4.28	0.75	$3.23	$43.43	$46.66	$494.99
64 physical-core C7i.48xlarge	$8.57	0.48	$4.12	$27.71	$31.84	$529.27
64-vCPU C7gn.16xlarge	$3.99	0.77	$3.09	$44.06	$47.75	$492.64

Table A4 | Calculating costs for Ford® Taurus 10M simulations run per day (numbers rounded to the nearest hundredth)

Instance type	AWS® price/hour ($)[6]	Number of hours to run one simulation	Cost to run one simulation	Engineering labor cost per run ($57.59/hr)[7]	Fully burdened cost (instance plus labor)/one simulation	Cost for simulations run per day
64-vCPU C6i.16xlarge	$2.72	5.66	$15.40	$325.97	$341.37	$482.48
96-vCPU C6i.24xlarge	$4.08	3.50	$14.30	$201.89	$216.19	$493.36
64 physical-core C6i.32xlarge	$5.44	3.02	$16.46	$174.21	$190.67	$504.24
64-vCPU C6g.16xlarge	$2.18	5.81	$12.66	$335.05	$347.71	$478.13
64-vCPU C7i.16xlarge	$2.86	4.40	$12.59	$253.84	$266.43	$483.57
96-vCPU C7i.24xlarge	$4.28	3.70	$15.87	$213.93	$229.96	$494.99
64 physical-core C7i.48xlarge	$8.57	2.63	$22.60	$151.89	$174.49	$529.27
64-vCPU C7gn.16xlarge	$3.99	3.71	$14.84	$213.39	$228.77	$492.67

Table A5 | OpenRadioss Chrysler® Neon 1M simulations/day

OpenRadioss Chrysler® Neon 1M Model Simulations/Day				
Performance comparison	AWS® with Intel instance	Intel® Xeon® processor generation	AWS with AWS Graviton® instance	Graviton generation
8% more (C6i)	64-vCPU C6i [C6i.16xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
75% more (C6i)	96-vCPU C6i [C6i.24xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
95% more (C6i)	64 physical-core C6i [C6i.32xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
4.5% less (C7i)	64-vCPU C7i [C7i.16xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
2% more (C7i)	96-vCPU C7i [C7i.24xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
61% more (C7i)	64 physical-core C7i [C7i.48xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E

Table A6 | OpenRadioss Ford® Taurus 10M simulations/day

OpenRadioss Ford® Taurus 10M Model Simulations/Day				
Performance comparison	AWS® with Intel instance	Intel® Xeon® processor generation	AWS with AWS Graviton® instance	Graviton generation
2% more (C6i)	64-vCPU C6i [C6i.16xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
65% more (C6i)	96-vCPU C6i [C6i.24xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
92% more (C6i)	64 physical-core C6i [C6i.32xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
16% less (C7i)	64-vCPU C7i [C7i.16xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
On par (<1% different)	96-vCPU C7i [C7i.24xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
40% more (C7i)	64 physical-core C7i [C7i.48xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E

Table A7 | Fully burdened cost of running daily OpenRadioss Chrysler® Neon 1M simulations

OpenRadioss Chrysler® Neon 1M Model Simulations/Day				
Performance/ dollar comparison	AWS® with Intel instance	Intel® Xeon® processor generation	AWS with AWS Graviton® instance	Graviton generation
On par (<1% difference)	64-vCPU C6i [C6i.16xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 3.1% higher (C6i)	96-vCPU C6i [C6i.24xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 5.4% higher (C6i)	64 physical-core C6i [C6i.32xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 2% lower (C7i)	64-vCPU C7i [C7i.16xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
On par (<1% different)	96-vCPU C7i [C7i.24xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
Up to 7.4% higher (C7i)	64 physical-core C7i [C7i.48xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E

Table A8 | Fully burdened cost of running daily OpenRadioss Ford® Taurus 10M simulations

OpenRadioss Ford® Taurus 10M Model Performance Per Dollar				
Performance/ dollar comparison	AWS® with Intel instance	Intel® Xeon® processor generation	AWS with AWS Graviton® instance	Graviton generation
On par (<1% difference)	64-vCPU C6i [C6i.16xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 3.1% higher (C6i)	96-vCPU C6i [C6i.24xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 5.4% higher(C6i)	64 physical-core C6i [C6i.32xlarge]	3rd Gen	64-vCPU C6g [C6g.16xlarge]	Graviton2
Up to 2% lower (C7i)	64-vCPU C7i [C7i.16xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
On par (<1% different)	96-vCPU C7i [C7i.24xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E
Up to 7.4% higher (C7i)	64 physical-core C7i [C7i.48xlarge]	4th Gen	64-vCPU C7gn [C7gn.16xlarge]	Graviton3E

[1] For the Ford® Taurus 10M model, the fully burdened cost (hourly instance + fully burdened engineering cost) of running simulations over one-eight hour day on AWS C7gn.16xlarge is $492.67. The fully burdened cost of running simulations over one eight-hour day on AWS C7i.48xlarge is $529.27. The difference is 7.4%, or $36.60. See Tables A2 and A4 for detailed costs.

[2] OpenRadioss. Accessed September 2023. www.openradioss.org/.

[3] Select the C7gn tab to view the C7gn instance.

[4] ZipRecruiter. "High Performance Computing Engineer Salary." Accessed September 2023.

[5] Indeed. "How Many Work Hours in a Year." July 2023.

[6] AWS® pricing from Amazon EC2® on-demand instance pricing page. AWS Region: US East (N. Virginia) US-east-1. Accessed September 2023. In the "instance type" section, select "compute optimized," and in the vCPU section, select "64" and then separately "96," then look for the correct instance lines in the results.

[7] See Analysis in the Appendix. Table A1.

PROWESS

The analysis in this document was done by Prowess Consulting and commissioned by Intel.
Results have been simulated and are provided for informational purposes only.
Any difference in system hardware or software design or configuration may affect actual performance.
Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.

0124/220205