



Research Abstract

# Accelerate HPC Simulation Throughput to Reduce Time to Market

Increase productivity by running more OpenRadioss™ simulations per day using Amazon Web Services® (AWS®) instances with Intel® processors.

High-performance computing (HPC) simulations play a crucial role in developing products ranging from automotive components to smartphones to medical devices. When engineers are more productive, companies can reduce project schedules and get their products and services to market faster.

Intel commissioned Prowess Consulting to assess productivity by testing in Amazon Web Services® (AWS®). We ran OpenRadioss™ finite element analysis (FEA) software simulations on AWS instances powered by both Intel® Xeon® and AWS Graviton® processors. We then compared performance and price per performance.

Prowess Consulting noted that AWS Graviton family processors are single-threaded, and Intel Xeon Scalable processors support hyperthreading or two threads per core. To level-set the comparison between AWS instances running on Graviton processors and AWS instances running on Intel® processors, we used AWS settings to configure Intel-based instances to have the same number of cores, 64, and to be single-threaded to match AWS instances with Graviton in AWS instance generations 6 and 7. For details on the instances and the settings, follow our [methodology](#).

We used both OpenRadioss Chrysler® Neon® 1M and Ford® Taurus® 10M element models for 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.

Imagine a typical scenario in which a mechanical engineer is tasked with optimizing the design of a front bumper to best withstand a crash. They set up a finite element model of the car's front section using OpenRadioss Chrysler Neon 1M. The model includes detailed geometry, material properties, and boundary conditions. The engineer then runs the OpenRadioss simulation iteratively, adjusting the variables and re-analyzing the design. When the engineer is finished running approximately 50 simulations, they identify the optimal bumper design.<sup>1</sup> The new design contributes to safer vehicles, reducing injury risk for occupants.

We found that engineers can achieve up to 92% higher simulation throughput by choosing the AWS C6i instances [C6i.32xlarge] with Intel, as compared to the AWS C6g instances [C6g.16xlarge] with Graviton.<sup>2</sup>

## HPC Test Highlights

AWS C7 and C6 instances with Intel® processors delivered higher performance as measured by simulations run per day.

The daily simulation throughput for the OpenRadioss™ Chrysler® Neon® 1M model on AWS C7i is **1.61x higher** than on C7gn.<sup>2</sup>

The daily simulation throughput for the OpenRadioss Ford® Taurus® 10M model on AWS C7i is **1.40x higher** than on C7gn.<sup>2</sup>

The daily simulation throughput for the OpenRadioss Neon 1M model on AWS C6i is **1.92x higher** than on C6gn.<sup>2</sup>

The daily simulation throughput for the OpenRadioss Taurus 10M model on AWS C6i is **1.92x higher** than on C6gn.<sup>2</sup>

We measured performance per price to understand the total costs of these high-performing instances. We defined performance per price as [fully burdened cost per simulation x number of simulations/day]. Our fully burdened cost included engineering costs.

AWS instances with Intel provided the highest number of daily simulations, maximizing engineering productivity. Intel instances cost up to 7.4% more per day (\$36.60) but could run up to 92% more simulations per day (which equates to seven full simulations).<sup>2</sup> This added investment allows engineers to be more productive, providing the opportunity to reduce project schedules and accelerate time to market. Hence, Prowess Consulting recommends using AWS C6i and C7i instances with Intel for HPC, analytics, and AI workloads when schedule constraints are a factor.

## The Right Business Decision

Iterating experiments solves design challenges, and HPC applications allow engineers to run experiments without the time and cost of physical models. Using compute for development isn't a novel idea for engineering teams, but they might not think about how the hardware running their modeling and simulation impacts project schedules.

Assume a project forecast shows 12 months of development time followed by 24 months of revenue-generating product life. When engineering can run more simulations per day, development takes 10 months. Product launch is pulled in by two months. The product gets to market sooner and generates revenue for 14 months, increasing the project's return on investment (ROI).

Look beyond hourly pricing when selecting AWS instances for HPC applications like OpenRadioss or compute-intensive workloads like data analytics and AI. Go with the instance that gives you the highest performance and the best price performance. This approach is the right business decision—especially when getting to market sooner is the priority.

## Learn More

Read the full [technical research report](#) to learn more about Prowess Consulting's testing and results.

<sup>1</sup> The mechanical engineer in this scenario runs approximately 50 simulations: 20 for design exploration, 10 for crash scenarios, 10 for sensitivity analysis, and 10 for optimization iterations.

<sup>2</sup> Prowess Consulting. "OpenRadioss™ on AWS®: HPC Workload Testing." February 2024.



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.

Copyright © 2024 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.

0224/220205