

Save on Time and Cost to Run Genomics Workloads with C3 High-CPU VM Clusters

Compared to a cluster of N2 Virtual Machines (VMs) with Older Processors, a C3 High-CPU VM Cluster featuring 4th Gen Intel® Xeon® Scalable Processors Reduced Time and Cost to Execute Genomics Tasks

Genomics workloads require considerable compute resources. For organizations looking for cloud VMs to suit their needs, it's important to compare the performance of different options. A wise investment can reduce both the time and cost to complete tasks. But with many options on the market, including those with newer processor technologies, how can you know which solution will be a good fit for genomics workloads?

We compared two types of Google Cloud VM clusters: one comprising N2 custom VMs with 2nd Gen Intel® Xeon® Scalable processors and one comprising C3 high-CPU VMs featuring 4th Gen Intel Xeon Scalable processors. We used the Cromwell benchmark to measure how quickly the VM clusters executed 22 tasks in the Genomics Analytics Toolkit (GATK) application. According to results, the C3 high-CPU VM cluster completed tasks in less time than the N2 custom VM cluster, which also reduced cost. This performance indicates that C3 high-CPU VMs enabled by 4th Gen Intel Xeon Scalable processors can deliver better efficiency and cost-effectiveness for running genomics workloads.

Save Time Completing Genomics Tasks

Figure 1 shows the relative time to execute genomics tasks in the GATK application. Compared to the N2 cluster with older processors, the C3 high-CPU cluster featuring 4th Gen Intel Xeon Scalable processors completed tasks in as little as 89 percent of the time of the N2 cluster with older processors. That's a reduction of 11 percent.

Relative Time to Complete Genomics Tasks

Relative Runtime | Less time is better

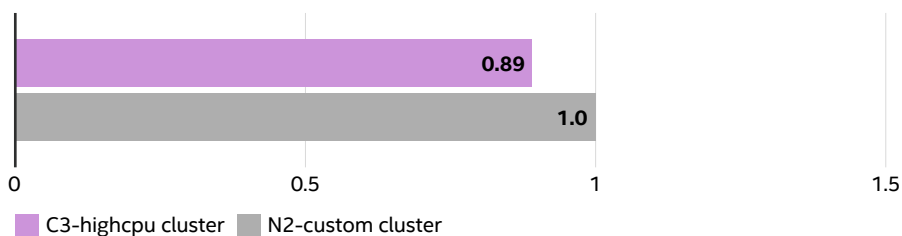




Figure 1. Relative time to complete 22 genomics tasks in the GATK application. Less time is better.

**Genomics**



Reduce Time to Complete Genomics Tasks by 11% with C3 High-CPU Clusters
vs. N2 Custom VM Clusters



Save 28% of the Cost to Complete Genomics Tasks with C3 High-CPU Clusters
vs. N2 Custom VM Clusters

Reduced Cost to Complete Genomics Tasks

The less time a cluster requires to complete a workload, the less VM uptime a customer has to pay for. As Figure 2 shows, the C3 high-CPU VM cluster reduced the cost to complete the same 22 GATK tasks previously mentioned. Completing the tasks with C3 high-CPU VMs featuring 4th Gen Intel[®] Xeon[®] Scalable processors took 72 percent of the cost of N2 VMs with older processors. In other words, they presented a cost savings of 28 percent.

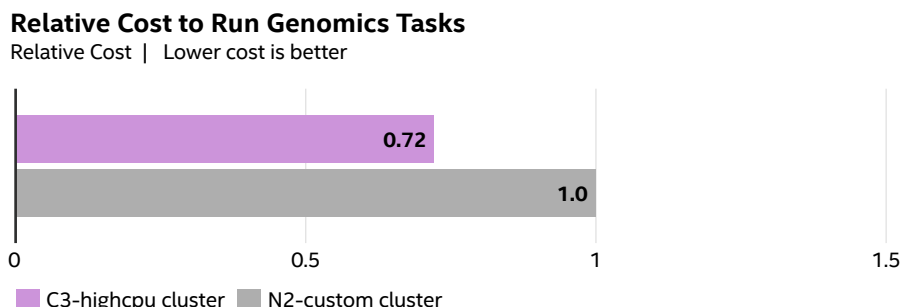


Figure 2. Relative cost to complete 22 genomics tasks in the GATK application. Lower cost is better

Conclusion

Selecting a high-performing cloud VM for your compute-intensive genomics workloads can be key in saving both time and cost. In tests, a cluster of Google Cloud C3 high-CPU VMs featuring 4th Gen Intel Xeon Scalable processors took less time to execute a set of 22 GATK tasks than a cluster of N2 VMs with older processors. With this reduction, the C3 high-CPU VM cluster also reduced the cost to execute the GATK tasks compared to the N2 VM cluster.

Learn More

To begin running your genomics workloads on Google Cloud C3 high-CPU VMs, visit https://cloud.google.com/compute/docs/general-purpose-machines#c3_series

All tests by Intel on Google Cloud us-central1. C3 details: #vCPUs: {4, 8, 22}; Number of Instances: 181; Iterations and result choice: Three; Median; CPU: Intel Xeon Scalable 8481C; Memory/Instance(GB): {8, 16, 44}; Storage/Instance(GB): {3-373}; NW BW/Instance (Gbps) (read/write): 32; Storage BW/Instance(read/write) (Mbps): {240, 800}; OS: Linux; N2 details: #vCPUs: {2, 16}; Number of Instances or VMs: 181; Iterations and result choice: Three; Median; CPU: Family 6 Model 85; Memory Capacity / Instance (GB): 2-16; Storage per instance (GB): {3-373}; Network BW per instance (Gbps) (read/write): 10-32; Storage BW per VM (read/write) (Mbps): {240, 1200}; OS: Linux; Workload and version (all tests): https://github.com/broadinstitute/warp/releases/tag/Who-leGenomeGermlineSingleSample_v3.1.10; Libraries: GATK 4.3.0.0, GKL 0.8.8, Cromwell 85, Samtools 1.3.1, Picard 2.26.10, BWA 0.7.15



Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See above for configuration details. No product or component can be absolutely secure. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Printed in USA 0923/HM/PT/PDF US001

