

# Select Google Cloud N2 Virtual Machines to Run Your Genomics Workloads More Efficiently and Cost-Effectively

## Compared to Other Google Cloud Virtual Machines (VMs) the N2 VMs Featuring 2<sup>nd</sup> Gen Intel® Xeon® Scalable Processors Reduced the Time and Cost to Execute Genomics Analytics Toolkit Tasks by as much as 25%

Genomics workloads demand a great deal of compute resources. If you're shopping for public cloud VMs, it's important to investigate the performance of different options. The right choice can dramatically reduce the amount spent running these workloads. To illustrate the ways VMs can vary in terms of performance and cost, Intel conducted benchmark testing on clusters of three Google Cloud VM series: N2 Standard and C2 Standard, both featuring 2<sup>nd</sup> Gen Intel Xeon Scalable processors, and N1 Standard with 1<sup>st</sup> Gen Intel Xeon Scalable processors.

The Cromwell on Google Cloud benchmark measured how quickly these VMs executed 24 tasks in the Genomics Analytics Toolkit (GATK) application. Test results showed that the N2 Standard cluster performed the tasks in less time than both other clusters, which resulted in a lower cost. For efficiency and cost-effectiveness, run your genomics workloads on N2 Standard VMs featuring 2<sup>nd</sup> Gen Intel Xeon Scalable processors.

### N2 Standard VM Clusters Outperformed N1 Standard VM Clusters and Cost Less

As Figure 1 shows, in our testing the N2 VM cluster featuring 2<sup>nd</sup> Gen Intel Xeon Scalable processors completed genomic tasks in the GATK application in 75% of the time the N1 Standard VM cluster needed. When your VM cluster completes a set of tasks in less time, you pay for less VM uptime, which helps your bottom line. This efficiency translates to savings: the cost of executing the tasks on the N2 VM cluster would be 21% lower than on the N1 VM cluster.

#### Relative Time and Cost to Complete Genomics Tasks N2 vs. N1




Figure 1. Relative time and cost to complete genomics tasks. Lower numbers are better.

Genomics

**Google Cloud N2 VM Clusters Completed Genomics Tasks in as Little as 75% of the Time**  
*vs. N1 Standard and C2 Standard VM Clusters*

**Google Cloud N2 VM Clusters Completed Genomics Tasks at as Little as 78% of the Cost**  
*vs. N1 Standard and C2 Standard VM Clusters*

## N2 Standard VM Clusters Outperformed C2 Standard VM Clusters and Cost Less

As Figure 2 shows, in our testing the N2 VM cluster featuring 2<sup>nd</sup> Gen Intel® Xeon® Scalable processors completed genomic tasks in the GATK application in 78% of the time the C2 Standard VM cluster needed. This efficiency translates to savings: the cost of executing the tasks on the N2 VM cluster would be 22% lower than on the N1 VM cluster.




Figure 2. Relative time and cost to complete genomics tasks. Lower numbers are better.

### Conclusion

Because genomics analysis applications are so compute intensive, selecting a cloud VM with robust performance is vital. In our testing, we showed Google Cloud N2 Standard VMs featuring 2<sup>nd</sup> Gen Intel Xeon Scalable processors reduced both the time and cost to complete genomics tasks, compared to N1 Standard and C2 Standard VMs.

### Learn More

To begin running your genomics workloads on Google Cloud N2 virtual machines with 2<sup>nd</sup> Gen Intel Xeon Scalable processors, visit [https://cloud.google.com/compute/docs/general-purpose-machines#n2\\_machines](https://cloud.google.com/compute/docs/general-purpose-machines#n2_machines).

All tests by Intel on GCP/us-central1 on 9/10/2020-10/9/2020. All tests: CentOS 7, Input Data Set 30X Coverage Human Whole Genome Sequence (WGS); NA12878, Workload GATK Best Practices Pipeline for Germline Variant Calling with pre-processing, GATK 4.0.10.1, Genomics Kernel Library (GKL) 0.8.6, Cromwell 52, BWA 0.7.15-r1140, Samtools 1.3.1. Tools in <https://hub.docker.com/r/broadinstitute/genomes-in-the-cloud/>: us.gcr.io/broad-gotc-prod/genomes-in-the-cloud:2.4.3-1564508330, us.gcr.io/broad-gatk/gatk:4.0.10.1. Workflow defined: <https://github.com/gatk-workflows/gatk4-genome-processing-pipeline/blob/master/WholeGenomeGermlineSingleSample.wdl>. Run Iterations:3. VM details: N1 machine series: Intel Xeon 1<sup>st</sup> Generation Scalable Processor (Skylake): N1-standard-2: 2 vCPUs, 7.5GB RAM; N1-standard-16: 16 vCPUs, 60GB RAM. N2 machine series: Intel Xeon 2<sup>nd</sup> Generation Scalable Processor (Cascade Lake): N2-standard-2: 2 vCPUs, 8GB RAM; N2-standard-16: 16 vCPUs, 64GB RAM. C2 machine series: Intel Xeon 2<sup>nd</sup> Generation Scalable Processor (Cascade Lake): C2-standard-4: 4 vCPUs, 16GB RAM. C2-standard-16: 16 vCPUs, 64GB RAM.



Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://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 0822/JO/PT/PDF US001

