DATASHEET

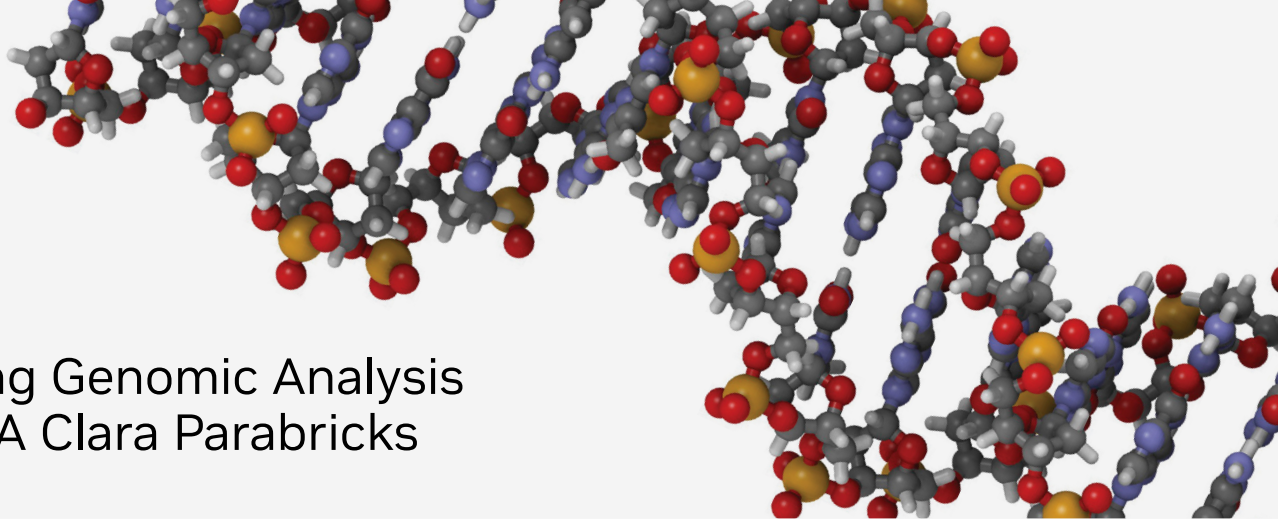

Accelerating Genomic Analysis with NVIDIA Clara Parabricks

NVIDIA Clara™ Parabricks® is an accelerated compute framework for next-generation sequencing data, supporting end-to-end data analysis workflows for DNA and RNA applications. With accelerations of up to 80X for industry-standard workflows, 30X coverage of whole human genomes can be analyzed in 22 minutes on a single server. Similarly, whole human exomes can be analyzed in four minutes, generating variant call format (VCFs) files from germline or somatic variant callers.

Running on a suite of NVIDIA GPU platforms, Clara Parabricks provides accelerated versions of gold-standard tools for read alignment, processing and quality control (QC), and variant calling. Clara Parabricks can also be easily deployed in common bioinformatics workflow managers for improved scalability, customizability, and intertwining of GPU- and CPU-powered tasks. The combination of accelerated tools across the entire workflow means results can be generated in minutes as opposed to hours or days, producing the same results as the non-accelerated tools.

Access the Work ows >

Derive more insights with faster analysis

Starting with FASTQ files, Clara Parabricks uses the GPU architecture to accelerate genomic data analysis, delivering the flexibility researchers require to meet the specific needs of their projects. Currently, for a whole human genome at 30X coverage, a server with 32 virtual CPUs (vCPUs) takes approximately 20 hours to generate a VCF. With Clara Parabricks, the compute time is approximately 22 minutes on eight NVIDIA A100 Tensor Core GPUs. Whether using Haplotypecaller or DeepVariant, the VCFs are the same between Parabricks and the native CPU tools. Unlike other acceleration platforms, Clara Parabricks scales from handling projects where the turnaround time for analysis is paramount to tackling large-scale studies of millions of exomes or genomes.

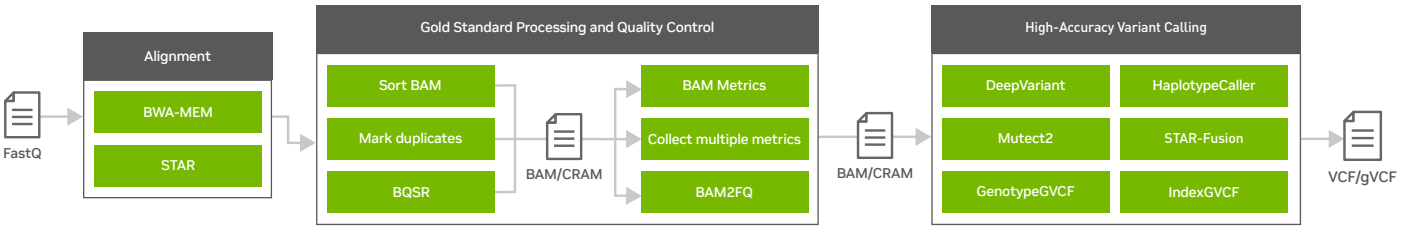

Explore the Parabricks advantage

KEY APPLICATIONS

Take advantage of accelerated tooling for gold-standard germline, somatic, and RNA analysis.

UP TO 80X FASTER

Experience up to 80X faster performance than CPU-only solutions, reducing computing costs by up to 50 percent.

BETTER ACCURACY

Bring the power of deep learning to your genomic analysis with Clara Parabricks and GPUs.

FLEXIBLE PIPELINE

Create powerful, customized workflows by configuring accelerated tools in Workflow Description Language (WDL) and NextFlow.

Deploy with flexibility

Clara Parabricks can be deployed in on-premises data centers, in public clouds, and on leading genomics and life science platforms. NVIDIA also works with sequencer developers to integrate Clara Parabricks into sequencers themselves, enabling real-time analysis at the edge. Whether you deploy in the cloud or on premises, you can choose from a variety of options to find the GPU solution best suited to your performance requirements.

Support for common bioinformatics workflow languages and managers, such as WDL, Nextflow, and Cromwell, allows for simple, scalable deployments no matter where your compute resides. Open-source reference workflows are available as blueprints to branch off of community-developed pipelines.

Reduce computing costs

With only one GPU server, Clara Parabricks can provide throughput comparable to more than 50 CPU servers, reducing IT management overhead and operating costs (including power and cooling). Whether it's analysis or reanalysis of next-generation sequencing data, Parabricks is rapidly becoming the preferred solution, providing vast improvements in efficiency, while enabling user-driven customization.

For users on the cloud, computing costs are proportional to the execution time. By reducing runtime by a factor of up to 80X, Parabricks cuts total computing costs by up to 50 percent compared to CPU-only solutions.

Performance Comparison

Germline End-To-End Secondary Analysis

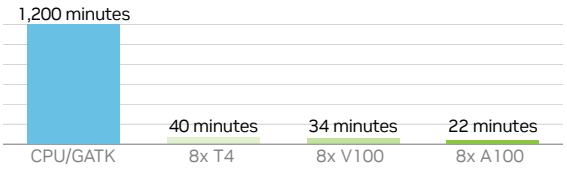

Data was generated using publicly available data (https://precision.fda.gov/challenges/truth) for NA12878, deprecating the data to 30X coverage. For the 22-minute runtime, DGX A100 with 320G memory was used. The native GATK4.1 numbers were generated using 32 vCPU (3.1 GHz Intel Xeon® Platinum 8175M) using 320Gb RAM.

Achieve the same results as CPU-only solutions

The output binary alignment map (BAM) file after alignment, sorting, marking duplicates, and applying base quality score recalibration (BQSR) is identical to the baseline CPU equivalent. The baseline GATK4.2 variant caller is nondeterministic and can generate slightly different results based on certain parameters. The deterministic variant results generated by Parabricks are within 99.999 percent of the baseline results.

NVIDIA GPU Recommendations	
Optimized for scale-out performance	NVIDIA A10 and T4 Tensor Core GPUs
Optimized for fastest turnaround time	NVIDIA A100 Tensor Core GPUs

Deploy with enterprise-grade support

Users wanting the assurance of full enterprise-grade support can purchase NVIDIA AI Enterprise, offered for deployment on NVIDIA-Certified servers and in the cloud. NVIDIA AI Enterprise includes full enterprise support for Clara Parabricks, as well as a suite of other data science tools and frameworks.

Features	Free	NVIDIA AI Enterprise
Accelerated tools	+	+
High-throughput workflows	+	+
NVIDIA Enterprise Support for Clara Parabricks		+
Full-stack support for other NVIDIA software		+