DATASHEET

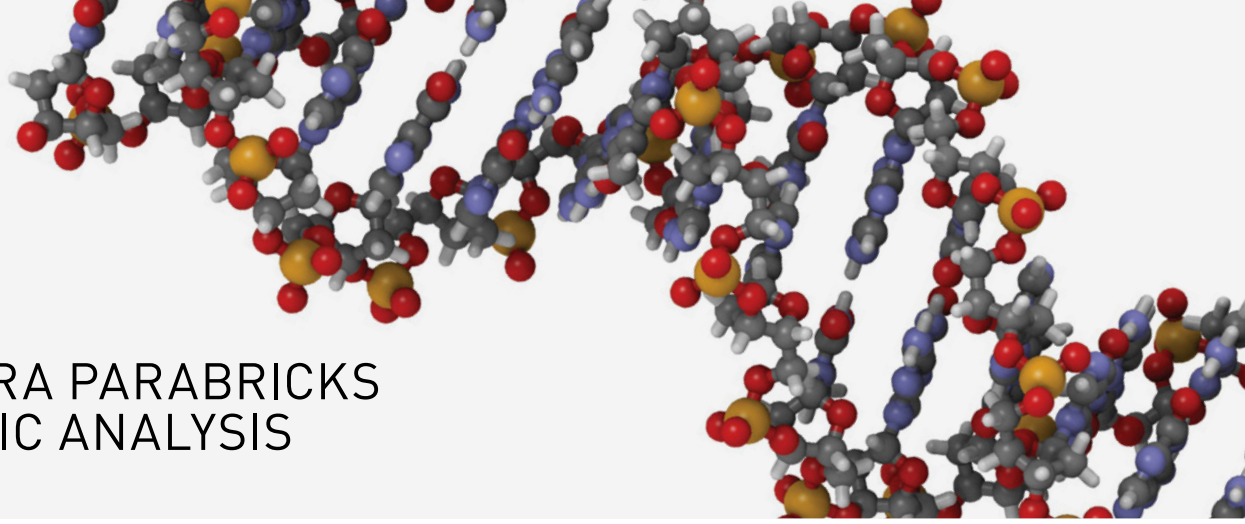

NVIDIA CLARA PARABRICKS FOR GENOMIC ANALYSIS

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 30-60X compared to industry-standard workflows, whole genomes with 30X coverage can be analyzed in 22 minutes, from FASTQ to Variant Call Format (VCF), using the Genome Analysis Toolkit's (GATK) HaplotypeCaller. Similarly, whole human exomes can be analyzed in four minutes from FASTQ to VCF with Google's DeepVariant. Running on a suite of NVIDIA GPU platforms, Clara Parabricks provides over 33 accelerated tools, including Burrows-Wheeler Aligner (BWA-MEM), Picard, Samtools, Spliced Transcripts Alignment to a Reference (STAR), and STAR-fusion, along with four somatic callers and two structural variant callers, generating the same results as native CPU instances.

Accelerating Genomic Analysis for More Insights

Starting with FASTQ files, Clara Parabricks uses GPUs to accelerate genomic data analysis, with the flexibility for researchers to deploy multiple analysis workflows that meet the specific needs of their project. Currently, for a whole human genome at 30X coverage, a server with 32 virtual CPUs (vCPUs) takes nearly 30 hours to generate a VCF. With Clara Parabricks, the compute time is approaching 22 minutes on eight NVIDIA A100 Tensor Core GPUs. The VCFs are equivalent between Parabricks and the native CPU. Unlike other acceleration platforms, Clara Parabricks scales to handle projects from just a few genomes to tens of thousands of genomes to over one million exomes, meeting the demands of any project.

NVIDIA Clara Parabricks Accelerated Tools

From FASTQ to BAM to VCF for DNA and RNA applications

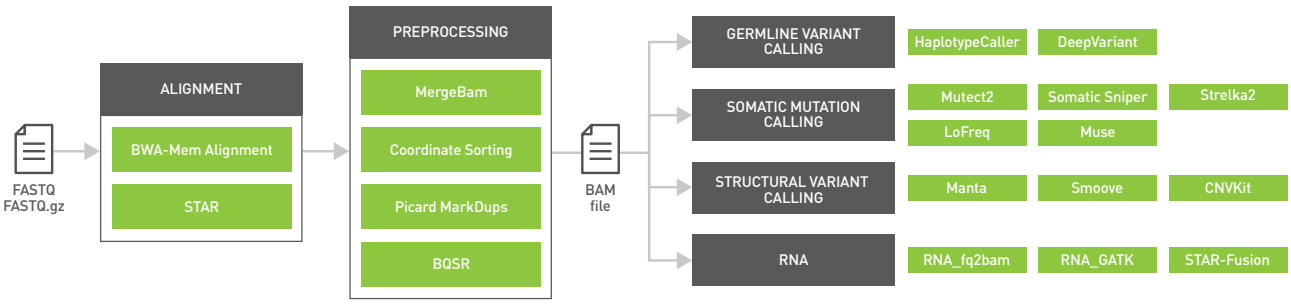

More tools available:

JOINT GENOTYPING	ImportGVCF, SelectVariants, GenotypeGVCF	VARIANT PROCESSING	VQSR, Variant Filtration, Select Variants, CNNScore Variants
QUALITY CHECKING	Alignment Summary, InsertSize, SequencingArtifact, GcBias, QualityScoreDistribution, WGSMetrics, RAWWGSMetrics, BaseDistributionByCycle, MeanQualityByCycle	VARIANT QC	Variants can be annotated, filtered, and QC'd using NVIDIA generated tools, including merging multiple VCFs

The Parabricks Advantage

HIGH THROUGHPUT
On a single server, the software can process more than 60 whole genomes per day.

30-60X FASTER
By running 30–60X faster, Clara Parabricks reduces computing costs up to 50 percent compared to CPU-only solutions.

BETTER ACCURACY
No need to sacrifice accuracy while reducing turnaround time. A suite of both germline and somatic callers is available.

DETERMINISTIC REPRODUCTIBLE
Any configuration of the software on any platform generates identical results every time for a given input.

FLEXIBLE PIPELINE
With over 33 tools, it's possible to create customized, accelerated pipelines for any specific pipeline or application need.

AI AND MACHINE LEARNING INTEGRATION
The GPU ecosystem is designed for deep learning and instantly integrates with state-of-the-art AI and machine learning frameworks and libraries.

Reduced 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 re-analysis of next-generation sequencing data, Parabricks is rapidly becoming the preferred solution, providing current pipelines with vast improvements in efficiency, while enabling user-driven customization.

For users on the cloud, computing costs are proportional to execution time. By reducing runtime by a factor of 30–60X, Parabricks reduces total computing costs by up to 50 percent compared to CPU-only solutions.

PERFORMANCE COMPARISON

Germline End-to-End Secondary Analysis

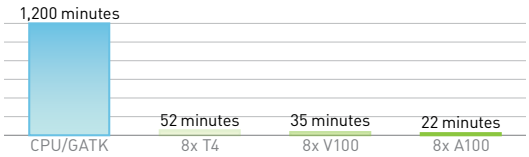

Data was generated using publicly available data (https://precision.fda.gov/challenges/truth) for NA12878, deprecating the data to 30X coverage. For the 22-minute runtime, DGX A100 with 320G memory was used. The native GATK4.1 numbers were generated using 32 vCPU (3.1 GHz Intel Xeon® Platinum 8175M) using 28Gb RAM.

High-Performance Solution

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 GATK4.1 variant caller is nondeterministic and can generate slightly different results based on certain parameters. The deterministic variant results generated by Parabricks is within 99.999 percent of the baseline results.

NVIDIA GPU RECOMMENDATIONS	
Optimized for scale-out performance	NVIDIA A30 Tensor Core GPU
Optimized for fastest turnaround time	NVIDIA A100 Tensor Core GPU

SUPPORT FOR AN ANNUAL PER-NODE LICENSE	
Full access to all workflows in the NVIDIA Clara Parabricks software suite	✓
No limitations on the number of genomes analyzed	✓