



NVIDIA RAPIDS Accelerator for Apache Spark

Faster results and lower cost for data analytics and machine learning in financial services.



Institutions across financial services are realizing significant ROI and cost savings using data analytics and machine learning. In consumer finance, some of the top use cases include application fraud, credit underwriting, backtesting and monitoring, marketing and propensity, and loan and mortgage pricing, while capital markets groups are deploying models for feature engineering, backtesting, and risk calculations.

Many enterprises use Apache Spark for key operations such as ingesting raw data into data lakes, business process analytics, loading data into data warehouses, and data preprocessing at the start of machine learning pipelines. However, growing workloads are being constrained by slow, CPU-based infrastructure—and slow processing costs time, money, and energy, resulting in a larger carbon footprint.

The Solution: Accelerate Apache Spark

The **NVIDIA RAPIDS™ Accelerator for Apache Spark** takes advantage of NVIDIA GPUs to accelerate Apache Spark workloads without code changes. It operates as a plug-in to popular Apache Spark platforms. The RAPIDS Accelerator speeds up selected Spark operations while allowing other operations to continue running on the CPU. As a result, processing time can be accelerated up to 5X or more, allowing the same work to be completed with 4X less infrastructure cost.


Who Needs the RAPIDS Accelerator?

Financial services institutions perform data analytics, and many Apache Spark workloads can benefit. Accelerating data analysis allows businesses to achieve faster time to results. Businesses can be more cost-effective while getting more done with the same investment.

Quickly Analyze Workloads for Cost and Performance

The RAPIDS Accelerator includes the **Accelerated Spark Analysis Tool**, which analyzes workloads and provides savings estimates. This makes it easy to determine which workloads would benefit most from GPU acceleration.

CPU vs. GPU: Example of Relative Cost



Based on NDS benchmark.
Does not include NVIDIA AI Enterprise license cost.

5X

Faster Execution Time

- > Move data in and out of data lakes more quickly
- > Take advantage of faster analytics
- > Accelerate AI pipelines

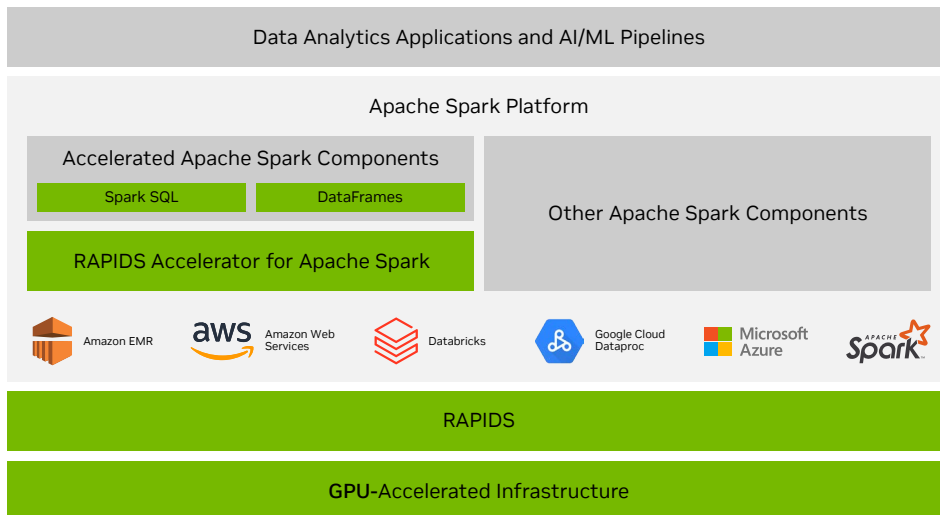
4X

Lower Costs

- > Save on cloud usage costs
- > Reduce power consumption and carbon footprint

How Does It Work?

The RAPIDS Accelerator supports data analytics workflows for GPU-accelerated infrastructure. By operating as a plug-in to Apache Spark platforms, it accelerates workflows without requiring code changes. It can work with Spark standalone, YARN clusters, and Kubernetes clusters. Additionally, advanced features such as Adaptive Query Execution for SQL and Dynamic Partitioning Pruning ensure consistent accelerated performance.



NVIDIA AI Enterprise Software Platform

The RAPIDS Accelerator is available with **NVIDIA AI Enterprise**, an end-to-end, secure, cloud-native AI software platform optimized to accelerate enterprises to the leading edge of AI. NVIDIA AI Enterprise delivers validation and integration for NVIDIA AI open-source software, access to AI solution workflows to speed time to production, certifications to deploy AI everywhere, and enterprise-grade support, security, and API stability to mitigate the potential risks of open-source software. NVIDIA AI Enterprise is optimized and tested for production deployments on all major Apache Spark platforms.

NVIDIA Enterprise Support

NVIDIA AI Enterprise includes **NVIDIA Enterprise Business-Standard Support**. With direct access to NVIDIA experts to support implementation and optimization, defined service-level agreements, and priority security notifications and fixes, this service ensures that projects stay on track.

Ready to Get Started?

To learn more about GPU-Accelerated Apache Spark, visit:

nvidia.com/spark-tool