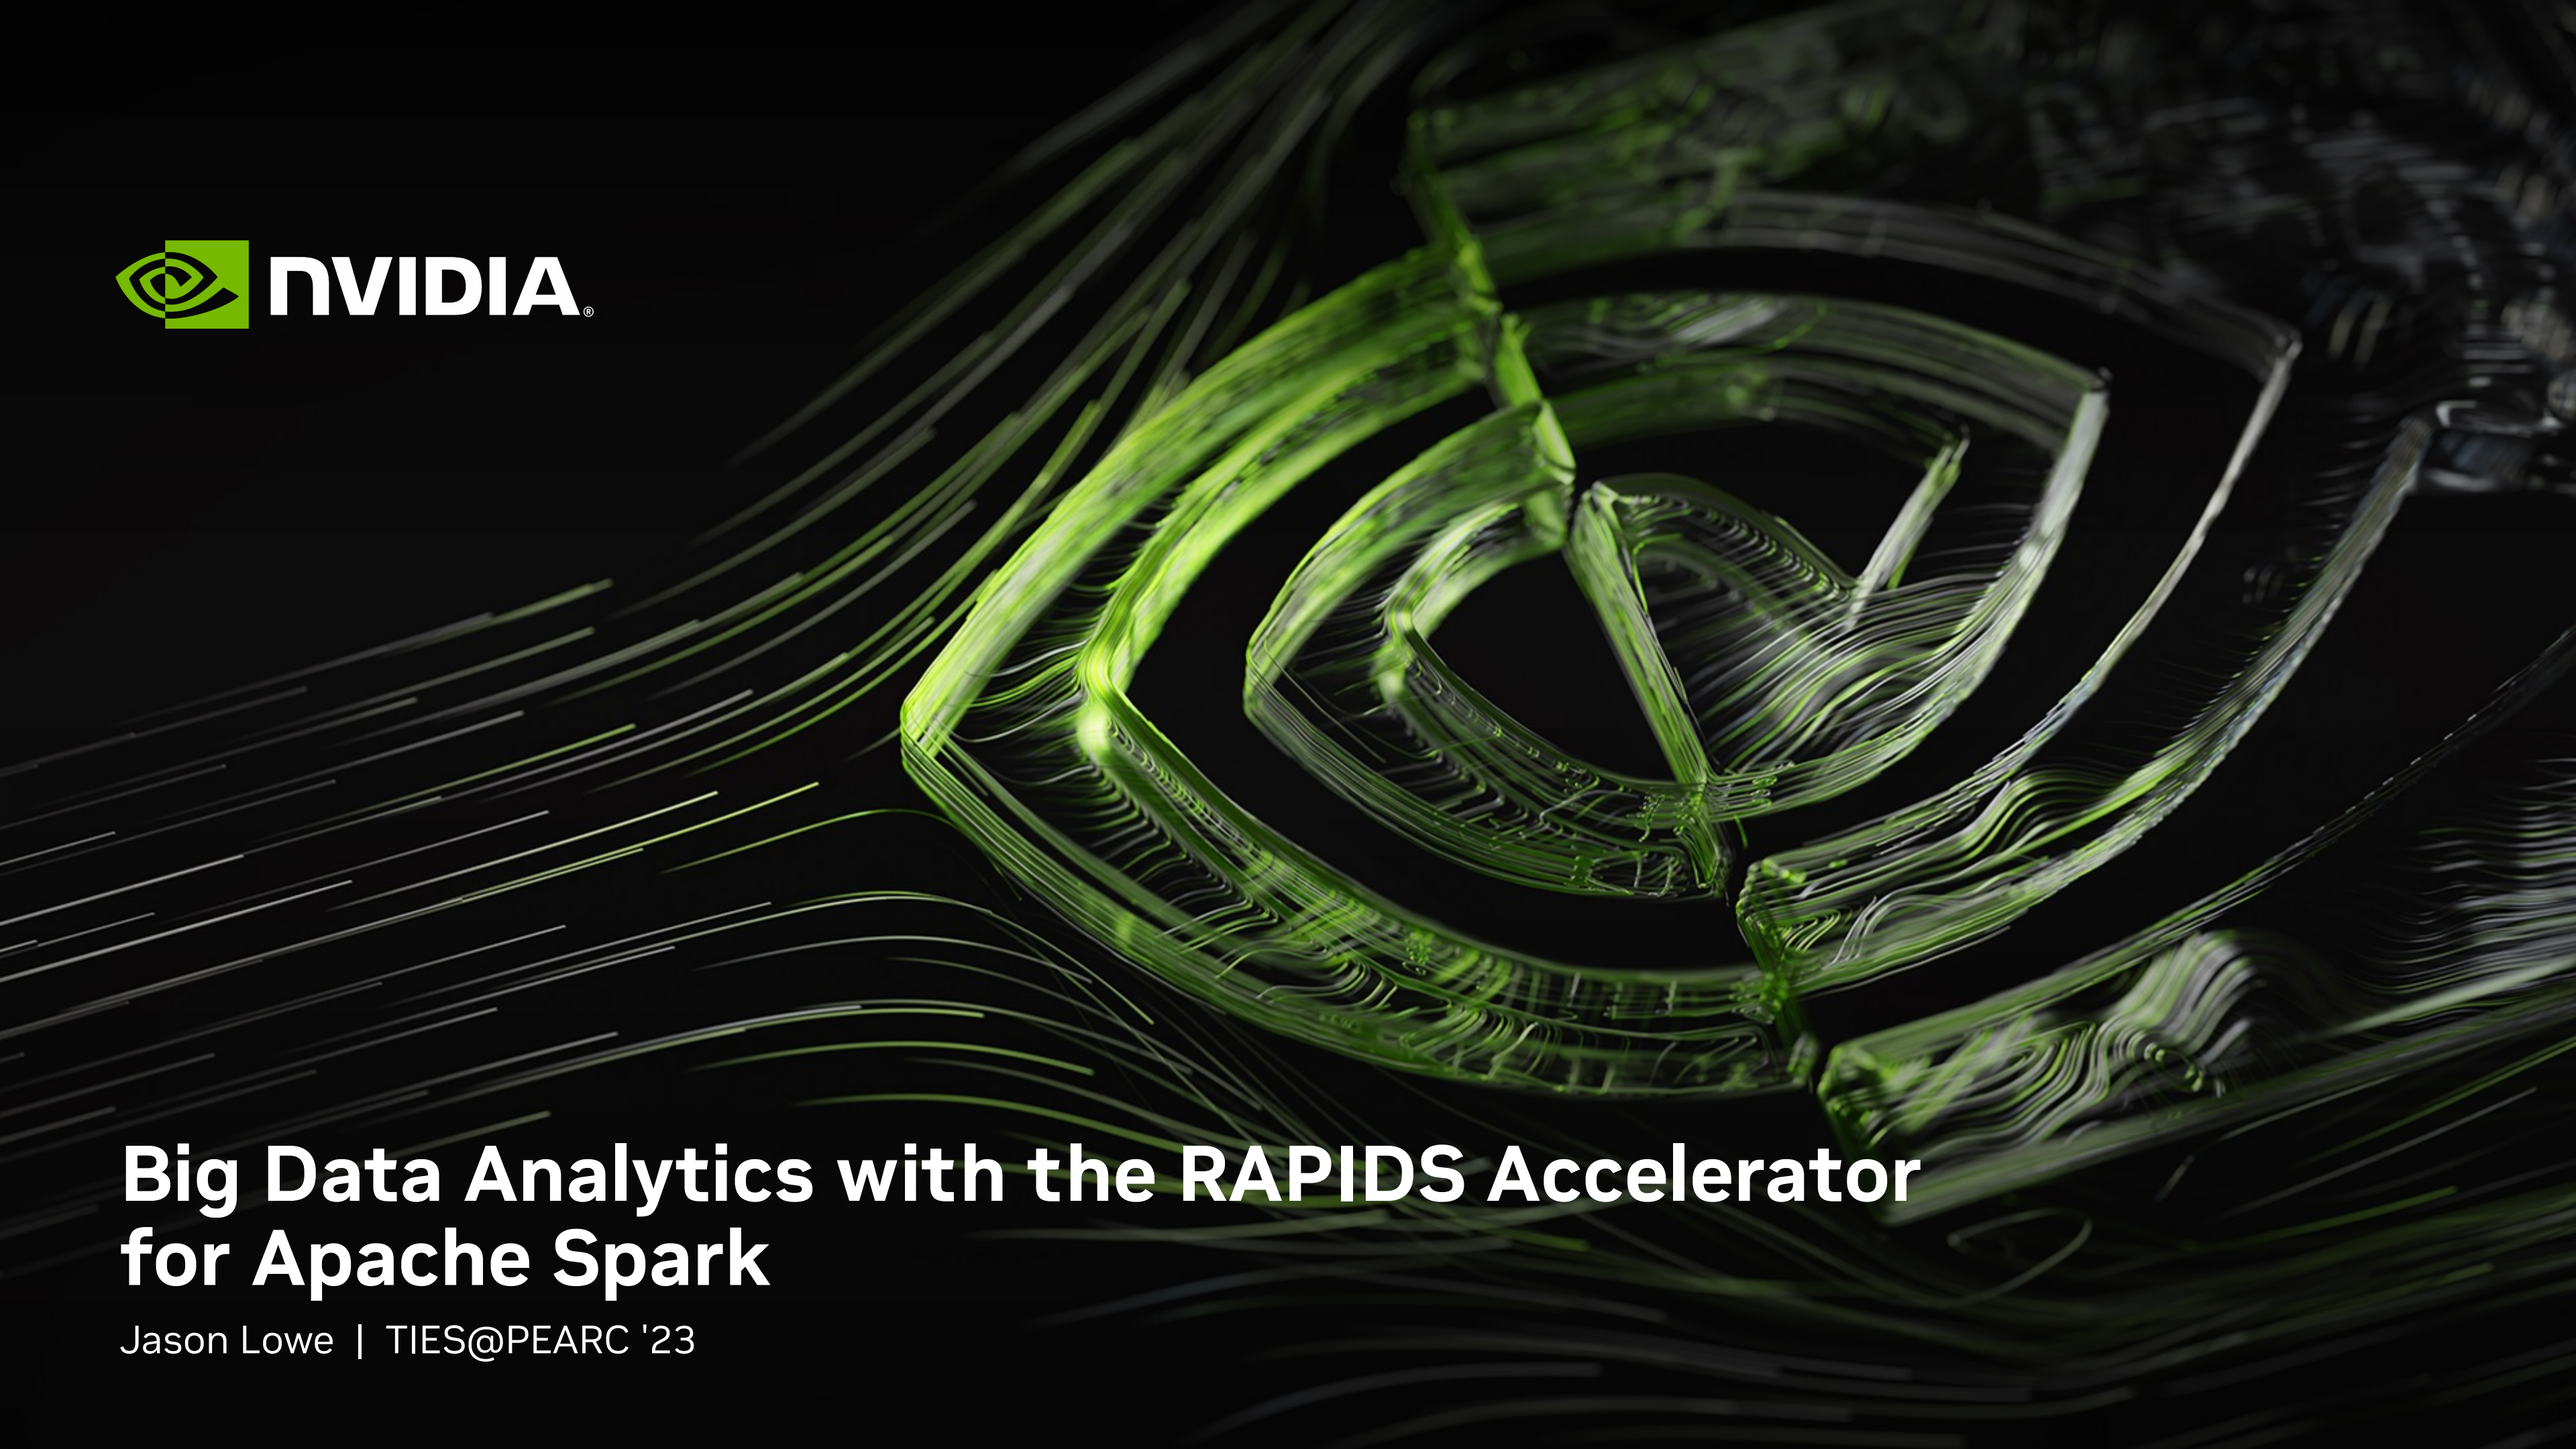

Big Data Analytics with the RAPIDS Accelerator for Apache Spark

Jason Lowe | TIES@PEARC '23

221 Zettabytes of Data Generated by 2026

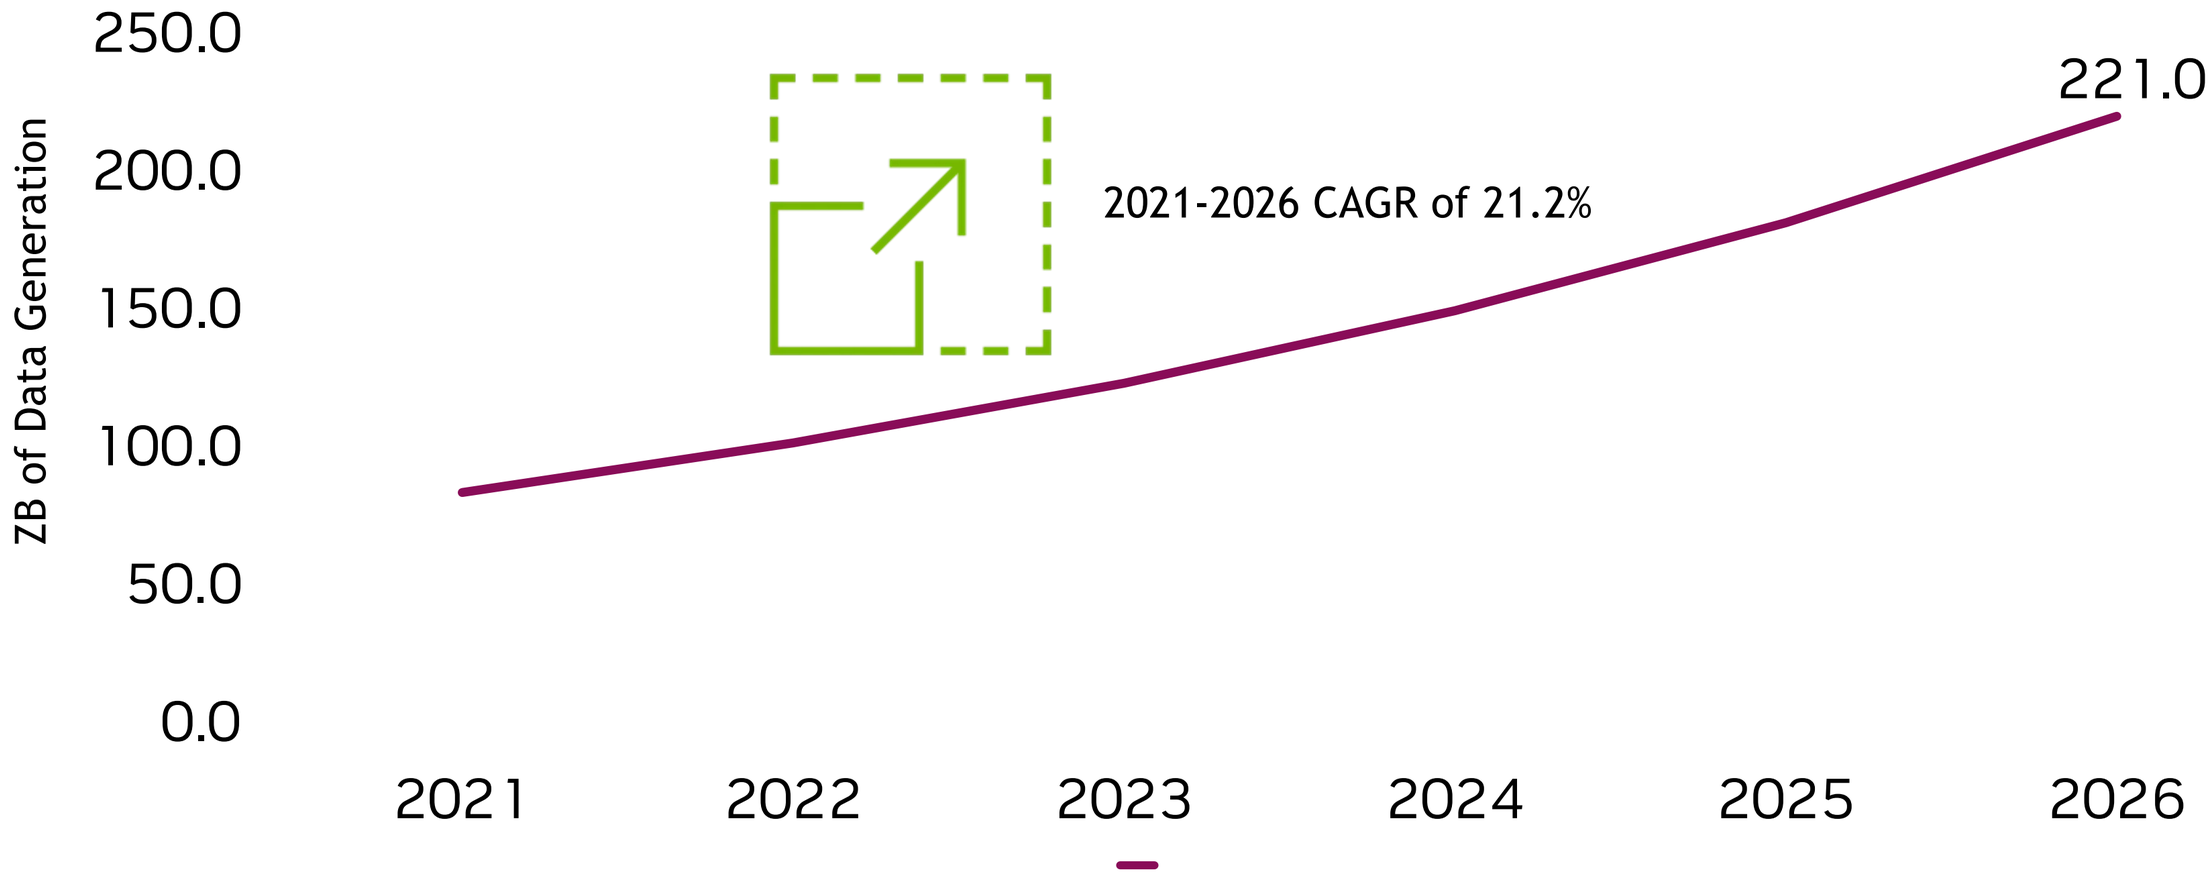

Source: IDC, Global DataSphere Forecast, 2022-2026: Enterprise Organizations Driving Most of the Growth, May 2022

How to Deal With Data Growth?

Increase Time

Accept higher latency as data volumes grow

How to Deal With Data Growth?

Increase Time

Accept higher latency as data volumes grow

Increase Cost

Increase compute capacity to meet growing data demands

How to Deal With Data Growth?

Increase Time

Accept higher latency as data volumes grow

Increase Cost

Increase compute capacity to meet growing data demands

Reduce Fidelity

Down sample data to meet cost & SLA targets

Scaling ETL Processing

Growth in Requirement for Data Processing

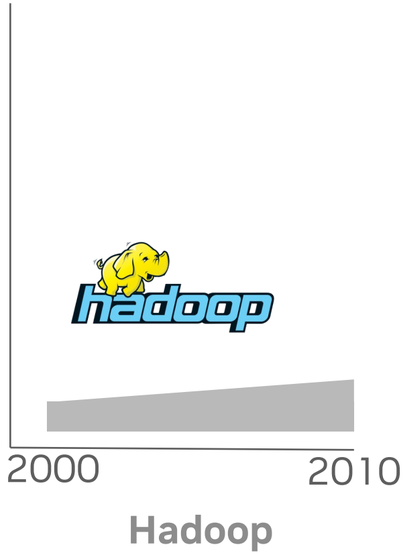

Scaling ETL Processing with Apache Spark

Growth in Requirement for Data Processing

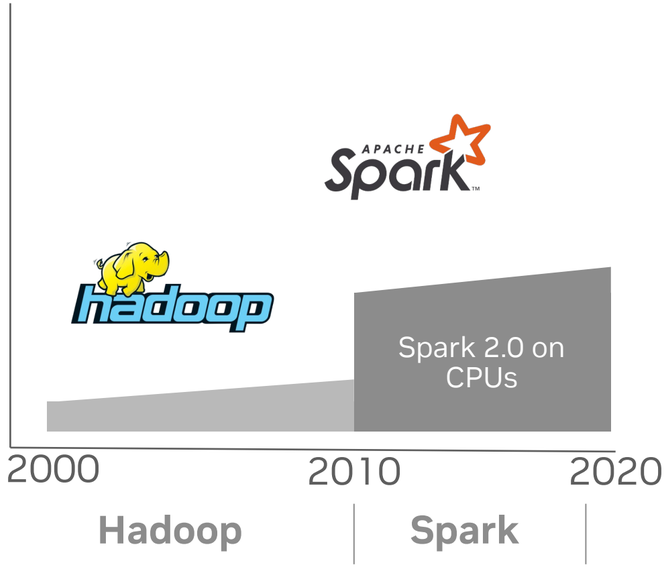

Scaling ETL Processing with Apache Spark with GPUs

RAPIDS Accelerator for Apache Spark

Growth in Requirement for Data Processing

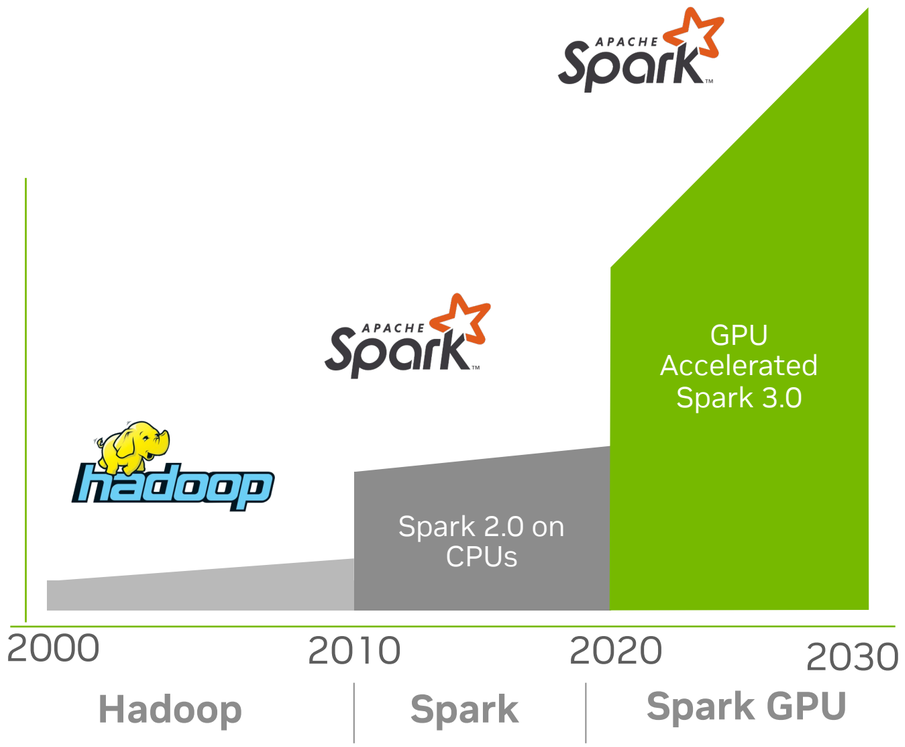

NVIDIA Decision Support Benchmark

NVIDIA Decision Support (NDS) is our adaptation of the TPC-DS benchmark often used by Spark customers and providers.

NDS consists of the same 100+ SQL queries as the industry standard benchmark but has modified parts for execution scripts.

The NDS benchmark is derived from the TPC-DS benchmark and as such is not comparable to published TPC-DS results, as the NDS results do not comply with the TPC-DS Specification.

https://github.com/nvidia/spark-rapids-benchmarks

NDS Benchmark Environment on Google Cloud Dataproc

Parquet data at scale factor 3000

CPU Cluster

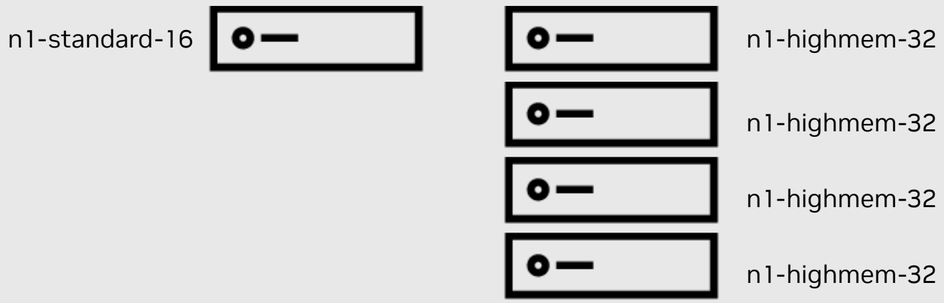

Nodes	1 driver, 4 workers
RAM	208 GB
Storage	Google Cloud Storage + 2 NVME / node
Software	Dataproc 2.1, Spark 3.3.0, YARN

Worker cost: $1.89/hr

GPU Cluster

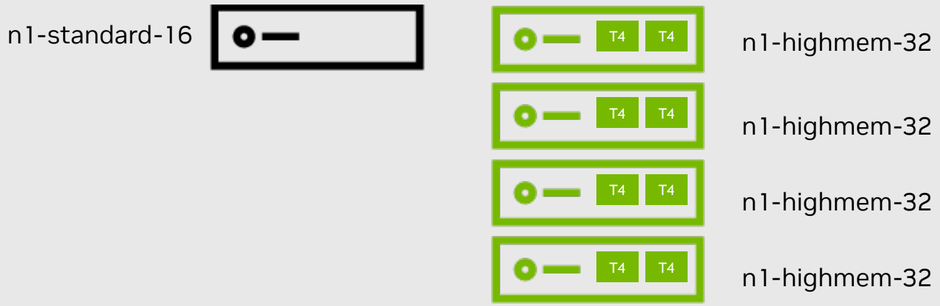

Nodes	1 driver, 4 workers
RAM	208 GB
Storage	Google Cloud Storage + 2 NVME / node
Software	Dataproc 2.1, Spark 3.3.0, YARN, **RAPIDS Accelerator plugin jar**
GPU	**2xT4 / node**

Worker cost: $2.59/hr
($1.89/hr + 2 x ($0.35)/hr/GPU)

Pricing: https://cloud.google.com/compute/vm-instance-pricing, us-central1, Feb 2023

NDS Benchmark Configuration

GCP Dataproc

Config	Dataproc CPU	Dataproc GPU	Group
spark.driver.memory	50gb	50gb	Resources
spark.driver.maxResultSize	1gb (default)	2gb	
spark.executor.cores	16	16	
spark.executor.memory	16gb	16gb	
spark.executor.instances	8	8	
spark.executor.memoryOverhead	1.6gb (default)	16gb	
spark.task.resource.gpu.amount		0.0625	
spark.scheduler.minRegisteredResourcesRatio	1.0	1.0	Scheduling
spark.locality.wait	0	0	
spark.sql.shuffle.partitions	128	200 (default)	Shuffle
spark.sql.files.maxPartitionBytes	128mb (default)	2gb	
spark.shuffle.manager		com.nvidia.spark.rapids.spark330.RapidsShuffleManager	
spark.rapids.shuffle.multiThreaded.writer.threads		16	
spark.rapids.shuffle.multiThreaded.reader.threads		16	
spark.rapids.sql.batchSizeBytes		1gb	GPU specific
spark.rapids.sql.concurrentGpuTasks		2	
spark.rapids.memory.host.spillStorageSize		32gb	
spark.rapids.memory.pinnedPool.size		8gb	

NVIDIA Decision Support Benchmark 3TB

RAPIDS Spark release 23.02

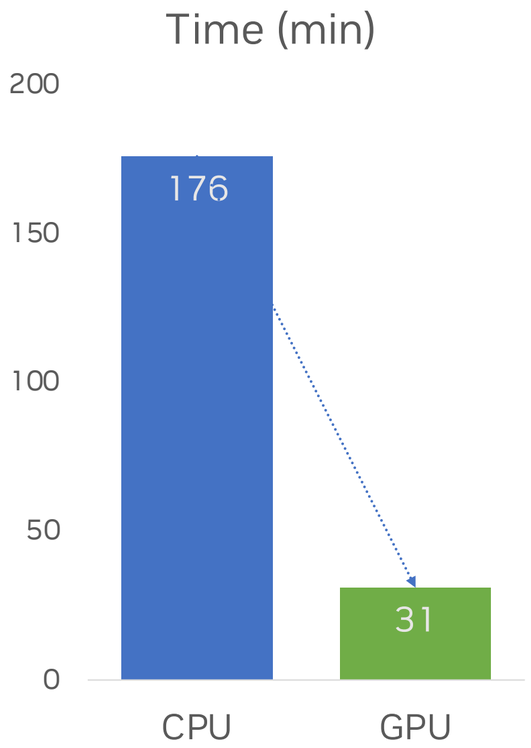

5.7x speedup

The NDS benchmark is derived from the TPCDS benchmark and as such is not comparable to published TPCDS results, as the NDS results do not comply with the TPCDS Specification.

NVIDIA Decision Support Benchmark 3TB

RAPIDS Spark release 23.02

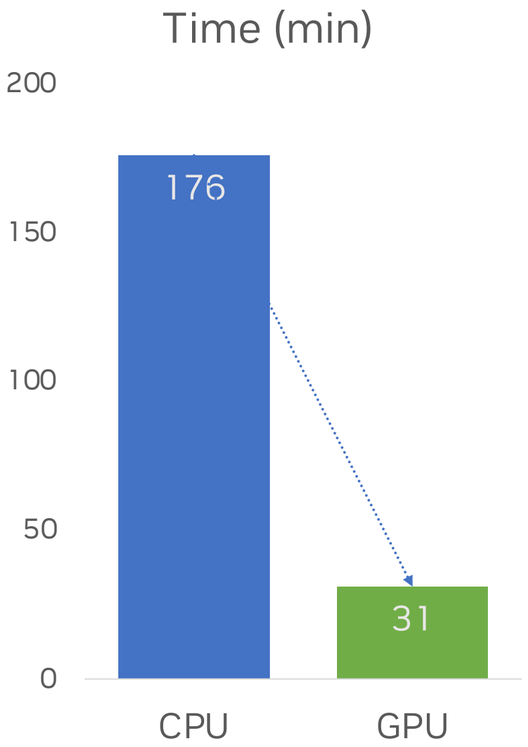

5.7x speedup

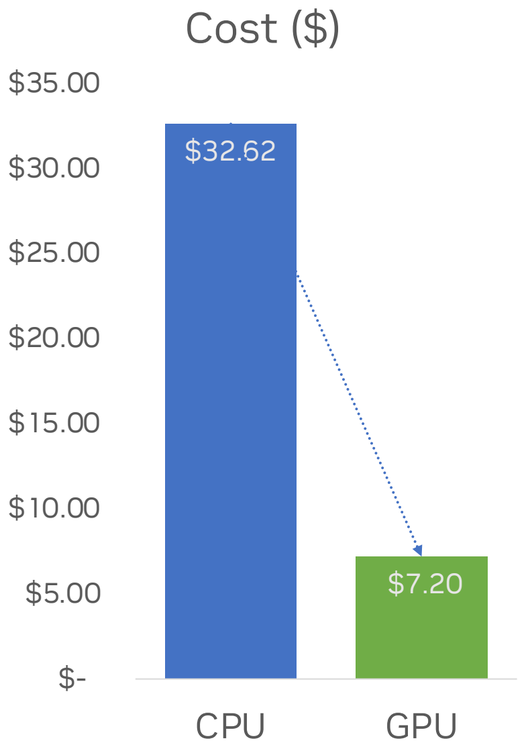

4.5x cost savings

The NDS benchmark is derived from the TPCDS benchmark and as such is not comparable to published TPCDS results, as the NDS results do not comply with the TPCDS Specification.

NVIDIA Decision Support Benchmark 3TB

RAPIDS Spark release 23.02

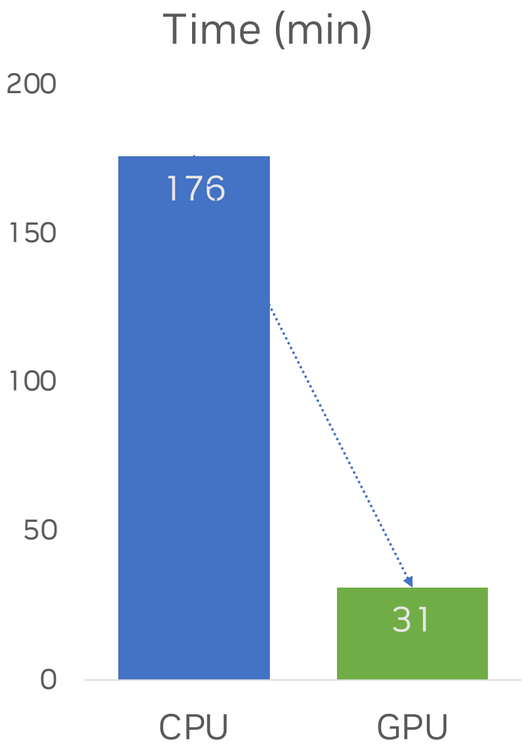

5.7x speedup

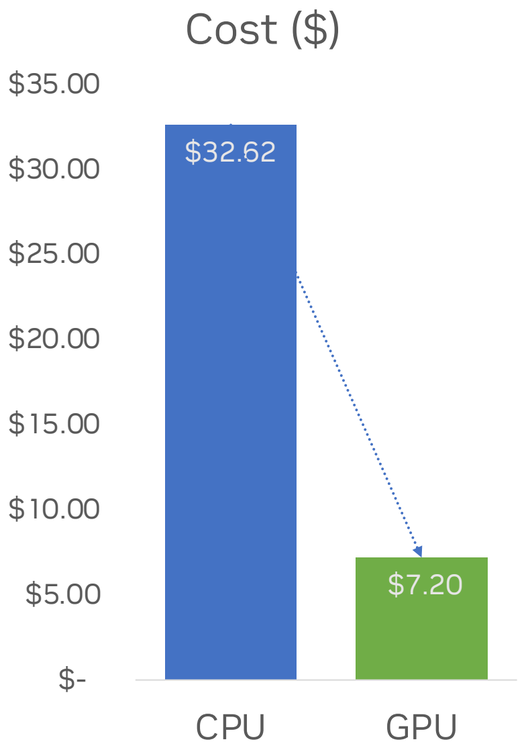

4.5x cost savings

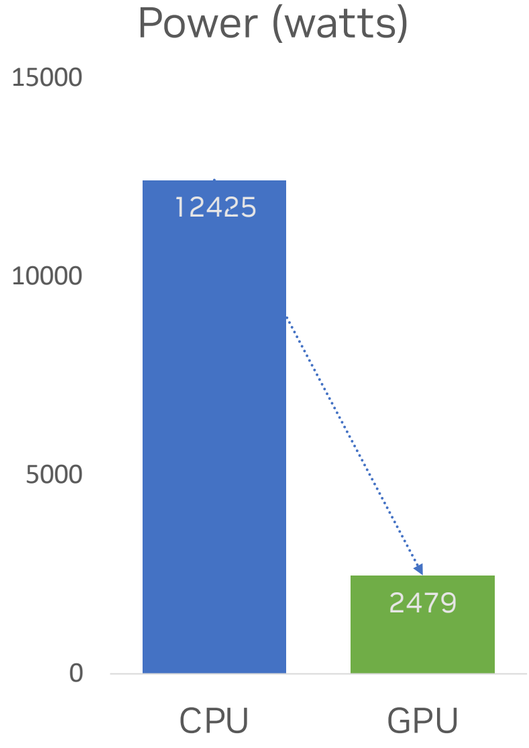

5x more efficient

The NDS benchmark is derived from the TPCDS benchmark and as such is not comparable to published TPCDS results, as the NDS results do not comply with the TPCDS Specification.

Cost Savings Across Clouds

NDS benchmark scale factor 3000

	GCP Dataproc 2.1
Cost Savings	78%
CPU	n1-highmem-32 (4)
GPU	n1-highmem-32 (4) + 2 x T4 / node

GCP pricing: https://cloud.google.com/compute/vm-instance-pricing, us-central1 , Feb 2023

Cost Savings Across Clouds

NDS benchmark scale factor 3000

	GCP Dataproc 2.1	AWS EMR 6.9
Cost Savings	78%	42%
CPU	n1-highmem-32 (4)	m6gd.8xlarge (8)
GPU	n1-highmem-32 (4) + 2 x T4 / node	g4dn.12xlarge (2) 4 x T4 / node

GCP Pricing: https://cloud.google.com/compute/vm-instance-pricing, us-central1 , Feb 2023
AWS Pricing: https://aws.amazon.com/emr/pricing/ , us-west-2, Feb 2023

Cost Savings Across Clouds

NDS benchmark scale factor 3000

	GCP Dataproc 2.1	AWS EMR 6.9	AWS Databricks Photon 10.4
Cost Savings	78%	42%	39%
CPU	n1-highmem-32 (4)	m6gd.8xlarge (8)	m6gd.8xlarge (8)
GPU	n1-highmem-32 (4) + 2 x T4 / node	g4dn.12xlarge (2) 4 x T4 / node	g5.8xlarge (4) 1 x A10 / node

GCP pricing: https://cloud.google.com/compute/vm-instance-pricing, us-central1 , Feb 2023
AWS pricing: https://aws.amazon.com/emr/pricing/, us-west-2, Feb 2023
Databricks AWS pricing: https://www.databricks.com/product/aws-pricing , Feb 2023

Cost Savings Across Clouds

NDS benchmark scale factor 3000

	GCP Dataproc 2.1	AWS EMR 6.9	AWS Databricks Photon 10.4	Azure Databricks Photon 10.4
Cost Savings	78%	42%	39%	34%
CPU	n1-highmem-32 (4)	m6gd.8xlarge (8)	m6gd.8xlarge (8)	Standard_E16ds_v4 (8)
GPU	n1-highmem-32 (4) + 2 x T4 / node	g4dn.12xlarge (2) 4 x T4 / node	g5.8xlarge (4) 1 x A10 / node	Standard_NC8as_T4_v3 (8) 1 x T4 / node

GCP pricing: https://cloud.google.com/compute/vm-instance-pricing, us-central1 , Feb 2023
AWS pricing: https://aws.amazon.com/emr/pricing/ , us-west-2, Feb 2023
Databricks AWS pricing: https://www.databricks.com/product/aws-pricing , Feb 2023
Databricks Azure pricing: https://azure.microsoft.com/en-us/pricing/details/databricks/#instance-type-support , Feb 2023

RAPIDS Accelerator Distribution Availability

Apache Spark 3.x Community Release	Amazon EMR	Cloudera CDP	Databricks Machine Learning Runtime	Google Cloud Dataproc	Microsoft Azure Synapse Analytics
Available Now	**Available Now**	**Available Now**	**Available Now**	**Available Now**	**Available Now**

Apache Spark 3.x

Resource aware scheduling

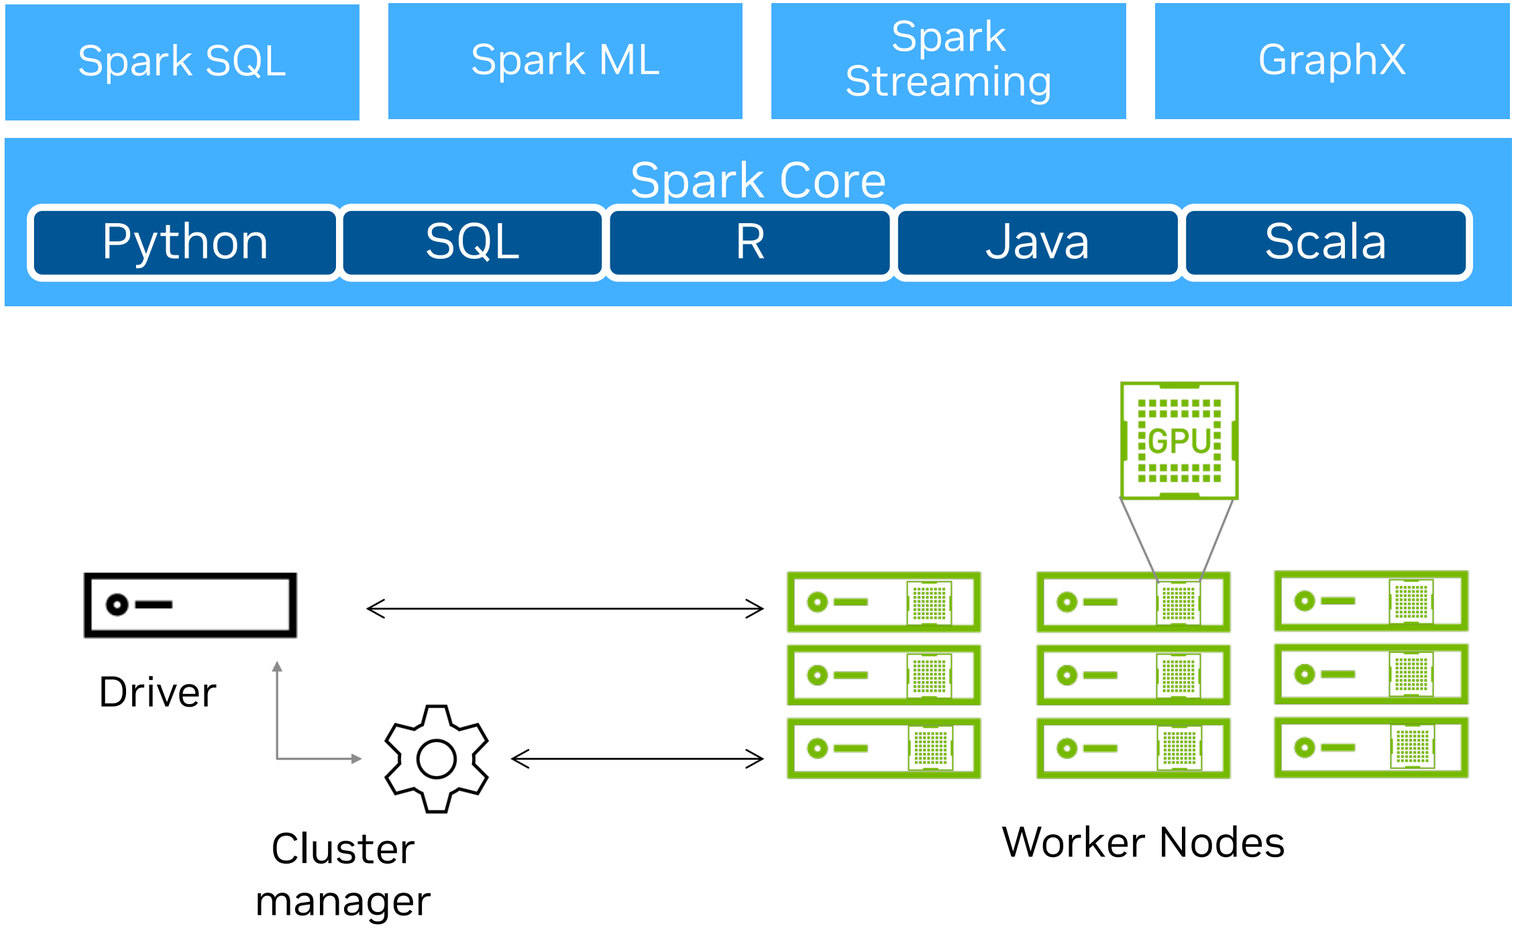

Apache Spark 3.x

Columnar processing

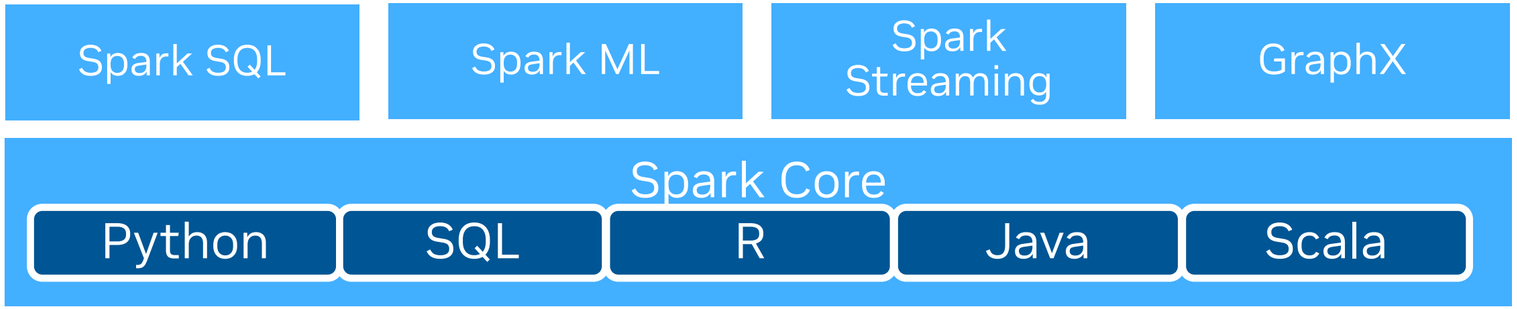

Row processing

Column processing

RAPIDS Accelerator for Apache Spark

Spark Plugin for GPU Acceleration

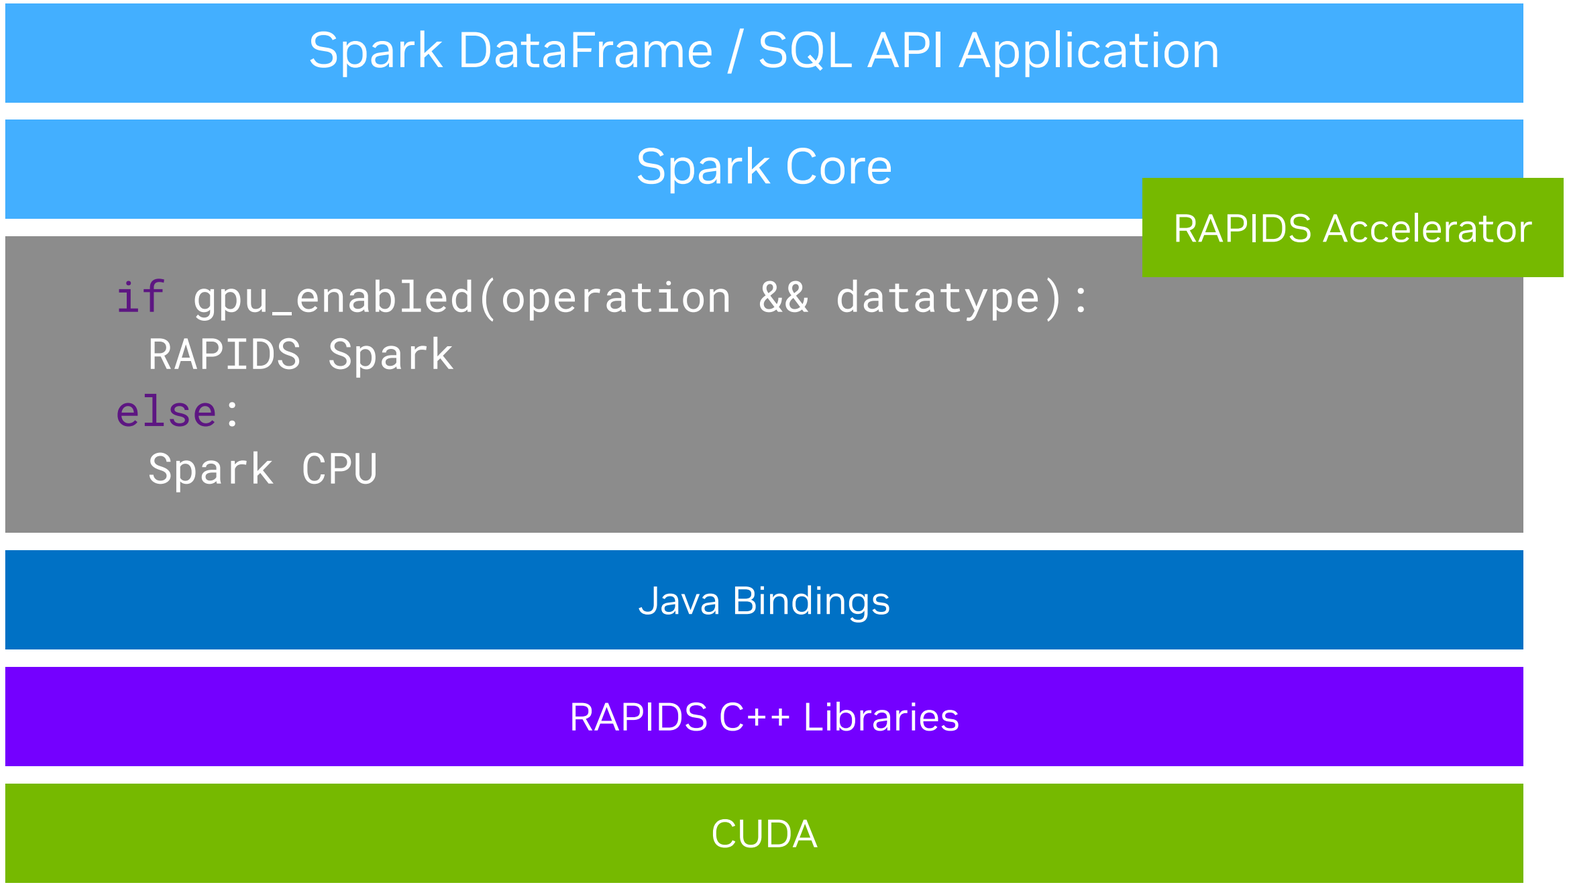

No Query Changes

- Add jars to classpath and set spark.plugins config
- Same SQL and DataFrame code
- Compatible with PySpark, SparkR, Koalas, and other DataFrame-based APIs
- Seamless fallback to CPU for unsupported operations

```
spark.sql( """

    SELECT
        o_order_priority
        count(*) as order_count
    FROM
        orders
    WHERE
        o_orderdate >= DATE '1993-07-01'
        AND o_orderdate < DATE '1993-07-01' + interval '3' month
        AND EXISTS (
            SELECT
                *
            FROM lineitem
            WHERE
                l_orderkey = o_orderkey
                AND l_commitdate < l_receiptdate
             )
    GROUP BY
        o_orderpriority ORDER BY o_orderpriority

""" ).show()
```

Example: SQL query text processed directly by Spark using ***spark.sql()*** API.
No changes are required to the query or any of the Spark code.

Spark SQL & DataFrame Query Execution

```
bar.groupBy(
    col(”product_id”),
    col(“ds”))
.agg(
    max(col(“price”)) -
    min(col(“price”)).alias(“range”))
```

```
SELECT product_id, ds,
   max(price) – min(price)
AS range
FROM bar
GROUP BY product_id, ds
```

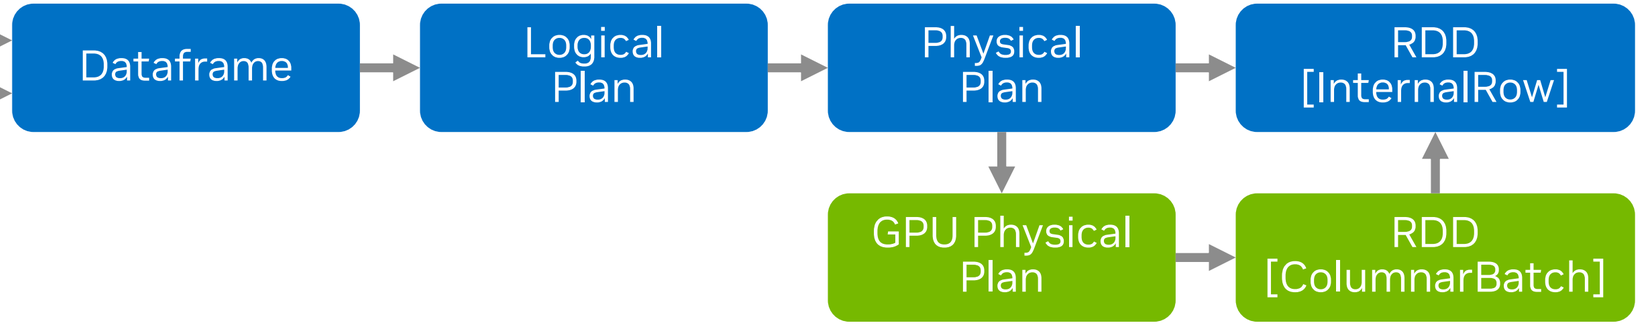

Spark SQL & DataFrame Query Execution

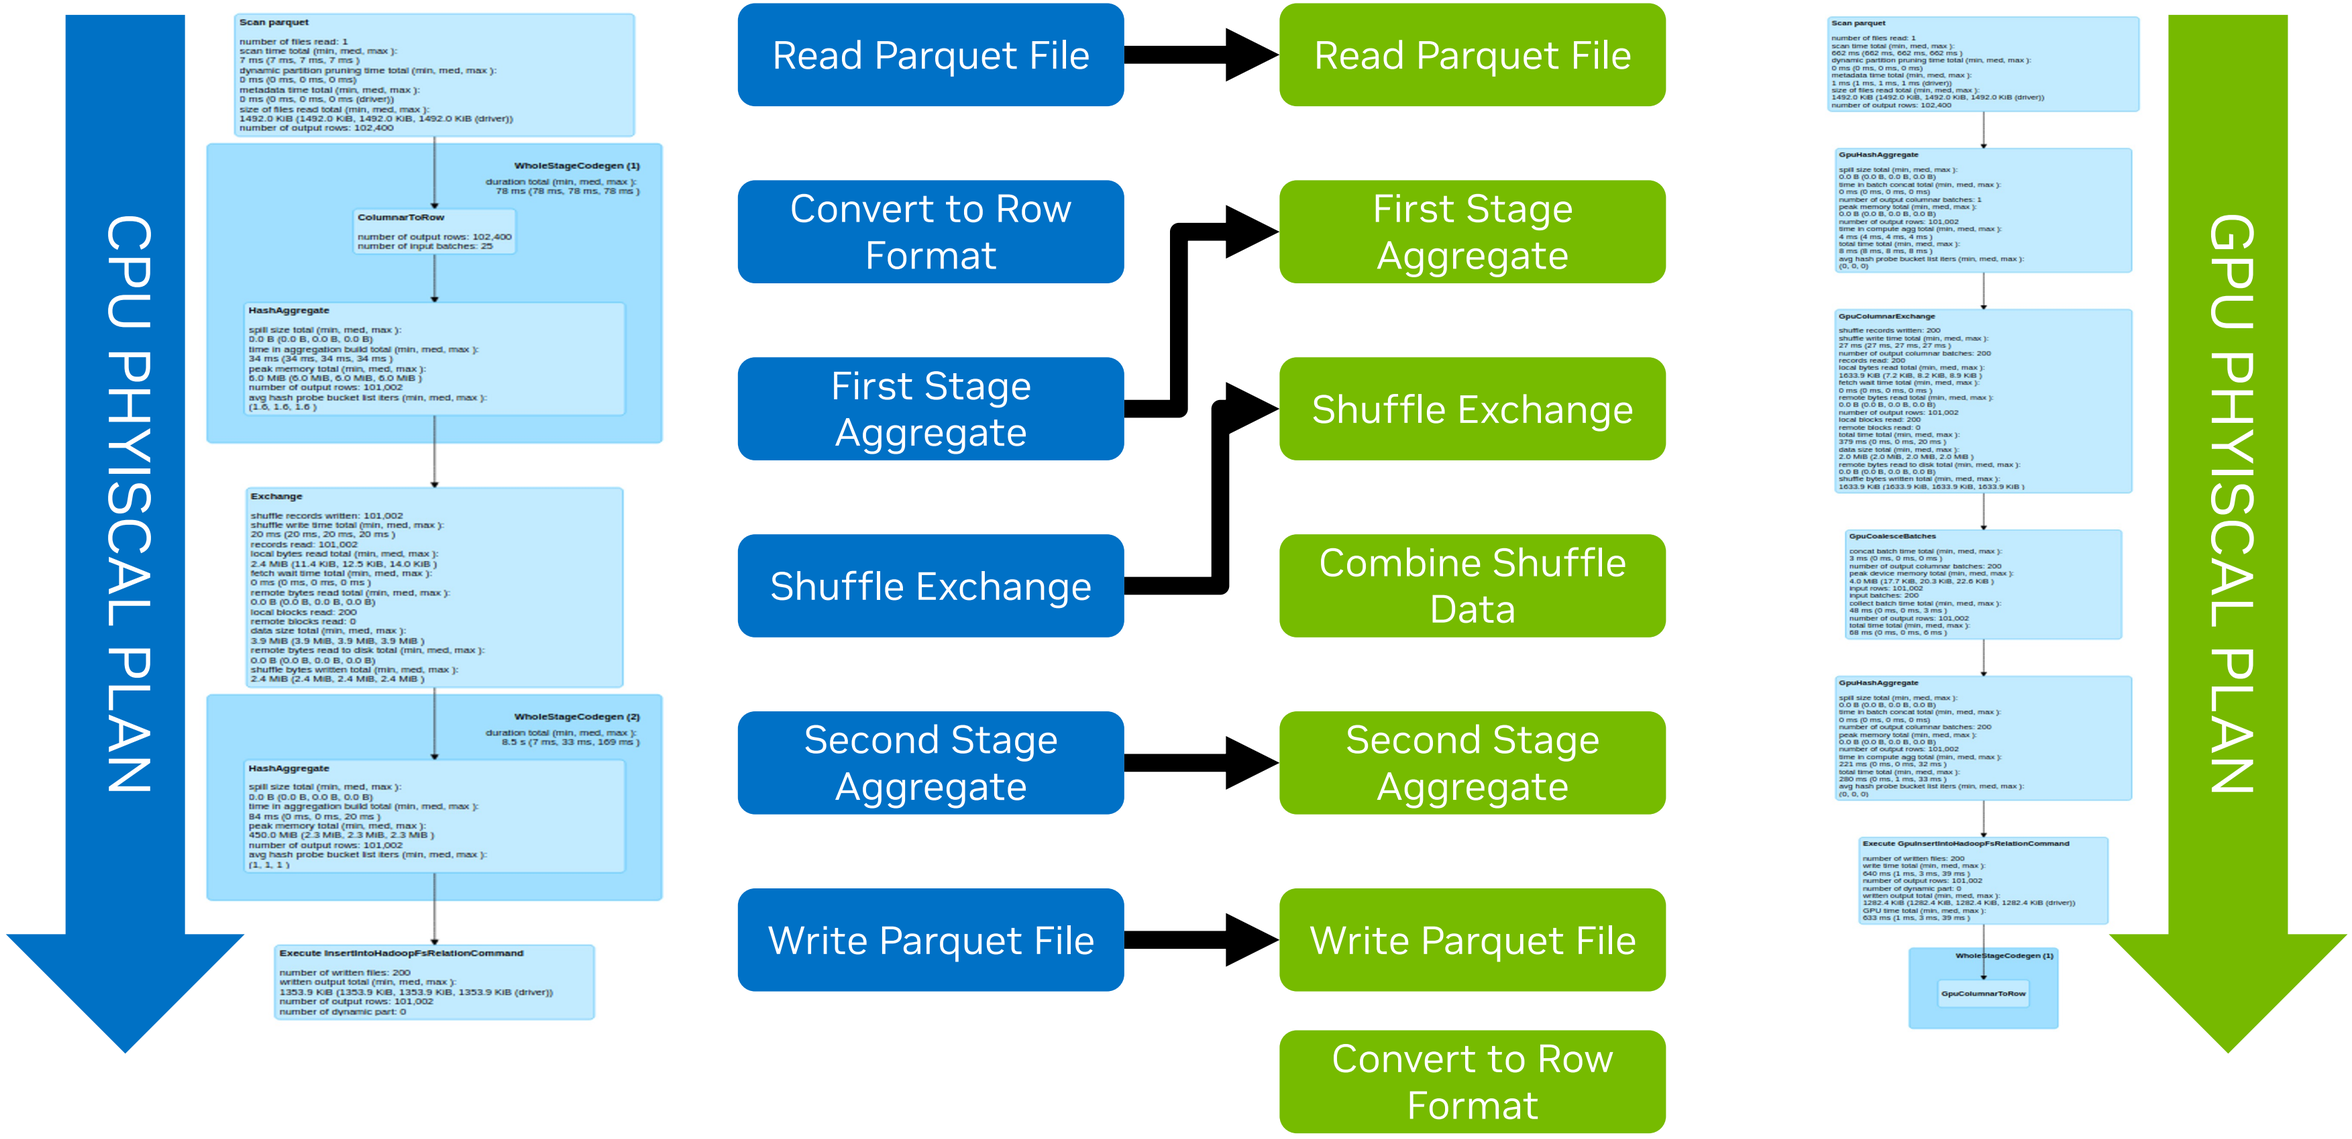

RAPIDS Accelerator Qualification Tool

Predicting the benefit of Spark + GPUs

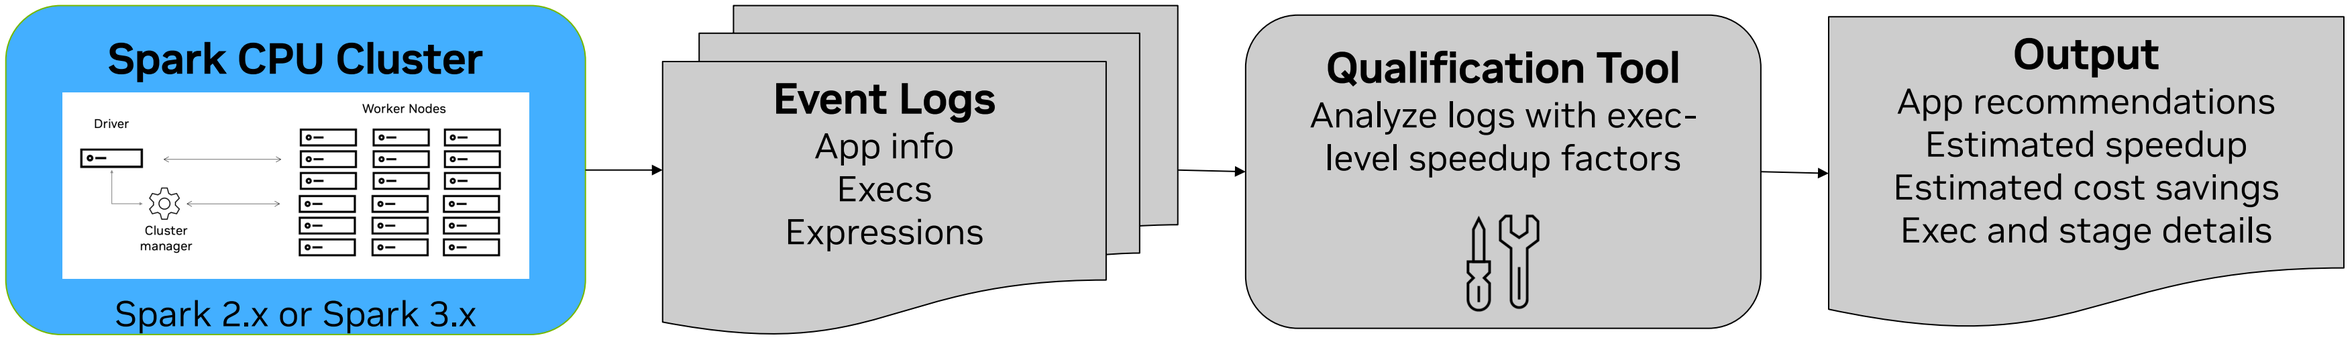


```
$ spark_rapids_dataproc qualification --cluster qualification-demo --region us-central1
```

Workload Qualification Output

Sample tool output

App Name	Recommendation	Estimated GPU Speedup	Estimated GPU Duration(s)	App Duration(s)	Estimated GPU Savings(%)
Customer App #1	Strongly Recommended	3.66	651.24	2384.32	64.04
Sales App #1	Strongly Recommended	3.14	89.61	281.62	58.11
Sales App #2	Strongly Recommended	3.12	300.39	939.21	57.89
Customer App #2	Strongly Recommended	2.55	698.16	1783.65	48.47

Report Summary:

Total applications	4
RAPIDS candidates	4
Overall estimated speedup	3.10
Overall estimated cost savings	57.50%

Qualification Tool

⚠ **Disclaimer!**

- Estimates provided by the Qualification tool are based on the currently supported "*SparkPlan*" or "*Executor Nodes*" used in the application. It currently does not look at the expressions or datatypes used.
- Please refer to the Supported Operators guide to check the types and expressions you are using are supported.

Total Applications

105

1.7 h Total Run Durations

RAPIDS Candidates

76

72.38% Fit for GPU acceleration

GPU Opportunity

1.1 h

1.2 h Total SqlDF Durations

94.16% Supported SQL DF Durations

GPU Recommendations Table

Export

Filters Active - 0

Collapse All | Show All | Clear All

Recommendations | Spark User

Search:

App Name	App ID	App Duration	Estimated Speed-up	Recommendation
TPC-DS Like Bench q14a	app-20220208013901-0189	5.5 min	2.7	Strongly Recommended
TPC-DS Like Bench q67	app-20220208023329-0246	13 min	2.6	Strongly Recommended
TPC-DS Like Bench q24b	app-20220208020315-0202	2.8 min	2.6	Strongly Recommended
TPC-DS Like Bench q24a	app-20220208020026-0201	2.8 min	2.6	Strongly Recommended
TPC-DS Like Bench q14b	app-20220208014431-0190	5.0 min	2.6	Strongly Recommended
TPC-DS Like Bench q4	app-20220208012938-0179	3.4 min	2.6	Strongly Recommended
TPC-DS Like Bench q23b	app-20220208015647-0200	3.6 min	2.5	Strongly Recommended

AT&T Optimizing Cost, Performance, and Pipeline Simplicity

Case Study

- **Challenges**: AT&T faced challenges managing costs and scaling a large multistage AI pipeline.
  - Perform feature engineering for about 3 trillion call records each month
  - Goal – improve speed, cost, and pipeline simplicity
- **Solution**: RAPIDS Accelerator for Apache Spark and GPUs
  - On Microsoft Azure, a GPU cluster was compared against an Apache Spark CPU cluster using Databricks' Photon engine
  - Cost and performance was measured across all five pipeline stages
- **Outcome**: Data preparation, model training, and optimization time was significantly reduced
  - GPU acceleration benefited all five pipeline stages
  - Executing the pipeline was simpler, cheaper, and faster

73%
Lower Cost (average)

68%
Faster Execution (average)

Use Cases

- Network planning and optimization
- Fraud
- Sales and marketing
- Taxes

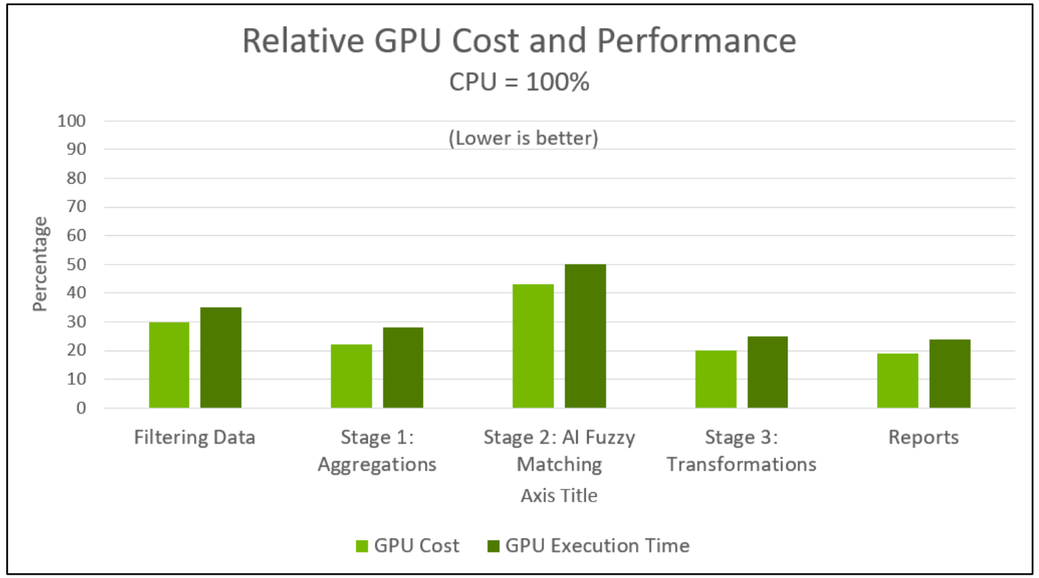

AT&T

NVIDIA

Taboola Optimizes Data Center Capacity and Cost

Case Study

- Most context-relevant webpage advertisements are served by Taboola's complex and compute-hungry data pipeline
- **Challenges**: Scaling capacity and minimizing cost for Apache Spark data pipelines
  - Frequent need to scale Apache Spark CPU cluster capacity to address constantly growing compute and storage requirements
  - Scaling CPU clusters was expensive
- **Solution**: RAPIDS Accelerator and A30 Tensor Core GPUs to accelerate data pipelines more cost-effectively than CPUs
- **Outcome**:
  - Greater scalability at lower cost
  - For some workloads, one A30 GPU sustained the same production load as 200 CPU cores, and with greater energy efficiency

20X GPU Speed-up

Average measured across multiple workloads on Intel CPUs

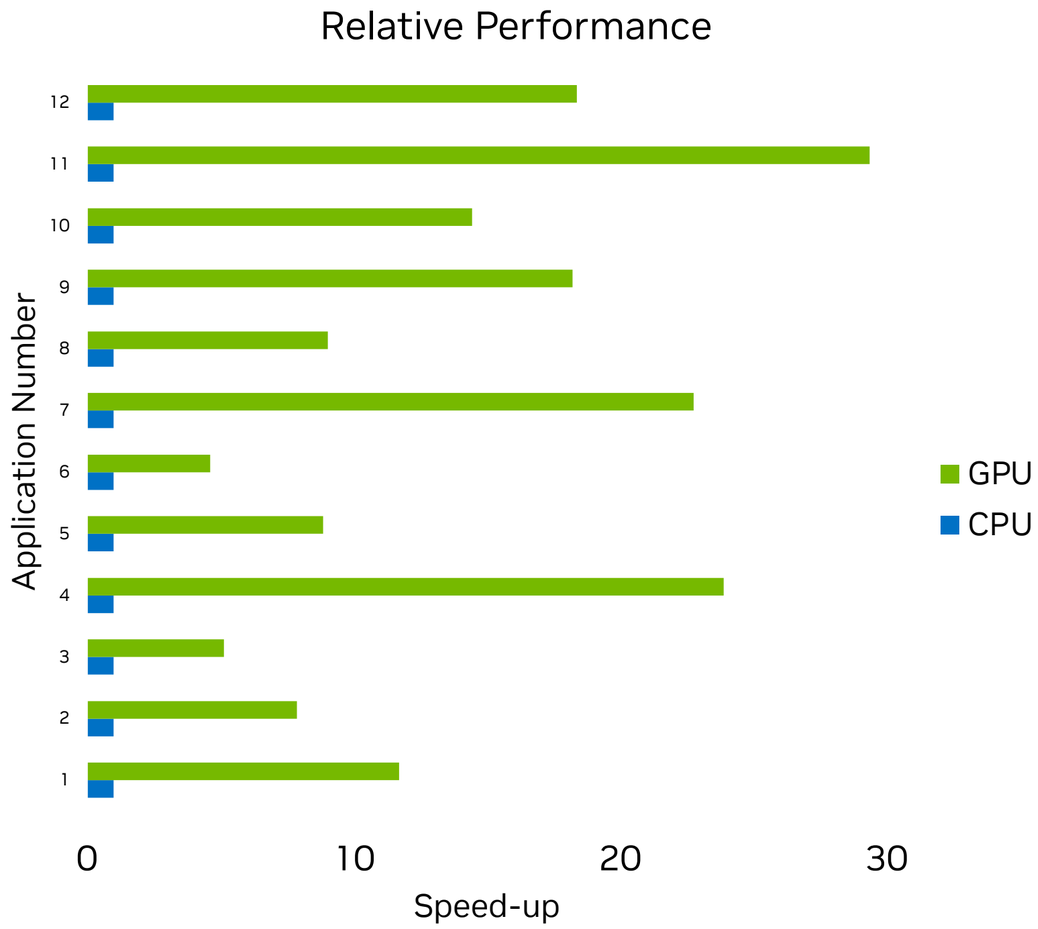

Taboola

NVIDIA

Saving Time and Money in E-Commerce

Case Study

Large Retailer

- **Challenges**: increase sales in an increasingly online market
  - Internal tool rearranges online shelves based on price, popularity and other constraints, using a multiple stage ML pipeline that starts with ETL
  - Tool generated more than $300M in incremental revenue once implemented on Google Dataproc but optimizing a single run to complete under an hour was expensive
- **Solution**: RAPIDS Accelerator reduced job time to below one hour, while saving 70% in infrastructure costs
- **Outcome**:
  - Greater than $150K/year savings for this tool alone
  - Many other applications have similar potential

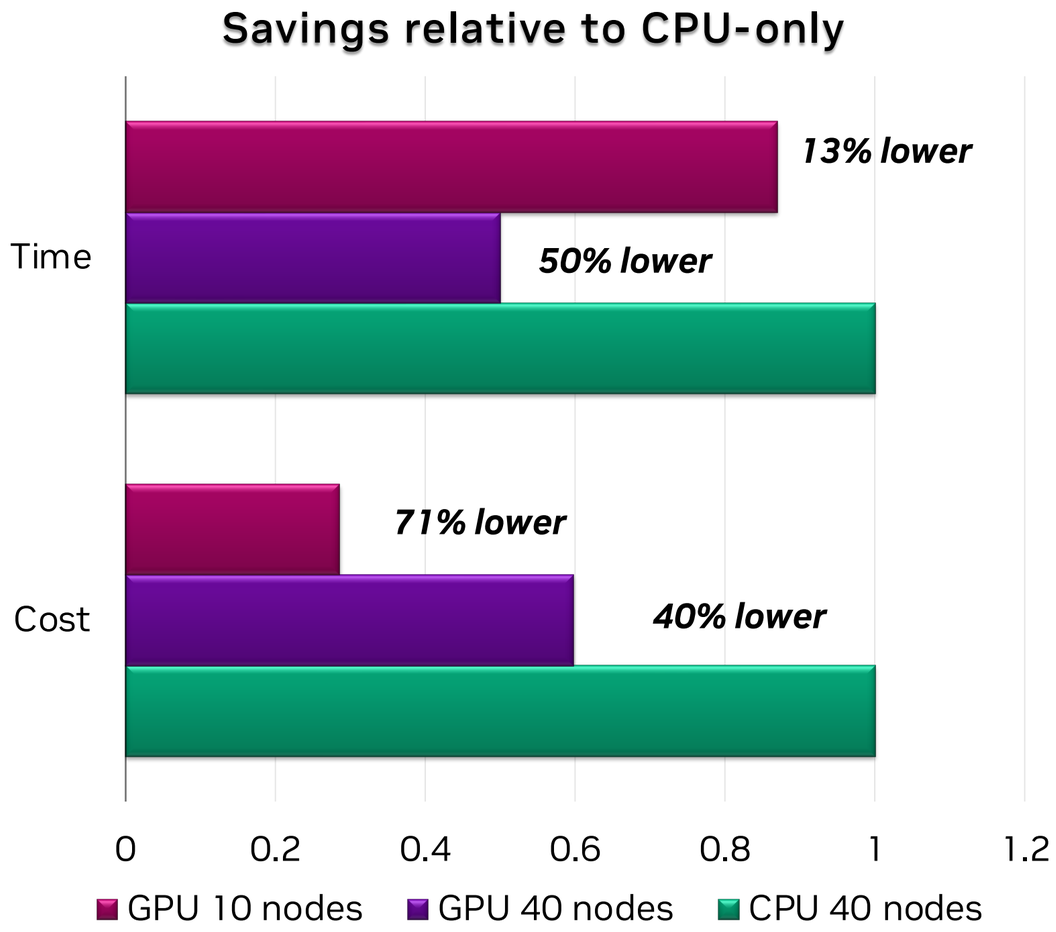

NVIDIA AI Platform

End-to-end open platform for production AI

Application Workflows

NVIDIA AI Enterprise

AI and Data Science Development and Deployment Tools

Accelerate Data Prep → Train AI at Scale → Optimize for Inference → Deploy at Scale

- NVIDIA RAPIDS
- NVIDIA TAO Toolkit
- PyTorch/ TensorFlow
- NVIDIA TensorRT
- NVIDIA Triton™ Inference Server

Cloud Native Management and Orchestration

- NVIDIA GPU Operator
- NVIDIA Network Operator
- Certified Partner Integrations
  MLOps I Container Orchestration I Virtualization

Infrastructure Optimization

- NVIDIA vGPU
- NVIDIA Magnum IO™
- NVIDIA CUDA-X AI™

Accelerated Infrastructure

NVIDIA LaunchPad

Hands-on Labs

Spark 3.x on NVIDIA GPUs

RAPIDS Accelerator for Apache Spark

Reduce Cost

Complete jobs faster with less hardware

Save on-prem and in the cloud

Speed Execution

Accelerate data preparation

Quickly move to next stages of the pipeline

Focus on most-critical activities

More Efficient

Reduce power consumption

Do more with less

Next Steps

RAPIDS Accelerator Overview: https://www.nvidia.com/spark

RAPIDS Accelerator User Docs: https://nvidia.github.io/spark-rapids

GitHub: https://github.com/NVIDIA/spark-rapids/

Contact us: spark-rapids-support@nvidia.com


NVIDIA