Spark on RAPIDS

Powered by NVIDIA

Optimize Cost and Operation of your Spark workloads on AWS with NVIDIA RAPIDS

Legacy Platforms with Scaled Data Costs

Organizations running large-scale workloads on legacy platforms such as Cloudera, on-premise Hadoop, or other cloud providers face escalating operational costs, limited scalability, and inefficient job management. *With annual operating costs ranging from \$1M to \$50M*, these platforms hinder innovation and resource efficiency. This creates an urgent need for companies to transition to scalable and cost-effective platforms that support high performance and operational efficiency, allowing companies to remain competitive.

Scalable Acceleration for Spark Workloads

Our NVIDIA RAPIDS acceleration for Spark workloads AWS integration delivers a novel solution for enterprises. Our approach reduces operational costs, enhances scalability, and accelerates job execution, enabling organizations to process larger datasets faster. With seamless data migration to AWS and a metric-driven feasibility analysis, we establish a clear business case for adopting next-generation processing capabilities.

Key Activities

Business Case Development

Provide a business case for the technical conversion including a level of effort for deploying and migrating production Spark Jobs to RAPIDS on EC2

CAYLE

Environment Configuration

Deploy and configure a Spark workload with NVIDIA RAPIDS on AWS, using the customer's AWS account and non-production VPC.

Performance & Cost Analysis

Measure and analyze a POC workload, deploying a Spark job using an existing job script while gathering workload metrics running on RAPIDS.

Source: Caylent Cost Analysis 2024

Highlights

- Reduce job processing time with NVIDIA RAPIDS
- Reduce operational costs by optimizing compute resource usage
- Ensure reliable performance for massive datasets with AWS cloud infrastructure
- Metric-driven comparisons provide a transparent business case for migration

Deliverables

- Non-production Spark workload running on NVIDIA RAPIDS for evaluation
- Business case development with operational cost projections at scale
- Technical feasibility study comparing
 workload performance
- Migrated dataset to AWS for consistent
 platform benchmarking



